

Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework*

Manuel Arellano[†] Richard Blundell[‡] Stéphane Bonhomme[§]

Revised version: January 21, 2017

Abstract

We develop a new quantile-based panel data framework to study the nature of income persistence and the transmission of income shocks to consumption. Log-earnings are the sum of a general Markovian persistent component and a transitory innovation. The persistence of past shocks to earnings is allowed to vary according to the size and sign of the current shock. Consumption is modeled as an age-dependent nonlinear function of assets, unobservable tastes and the two earnings components. We establish the nonparametric identification of the nonlinear earnings process and of the consumption policy rule. Exploiting the enhanced consumption and asset data in recent waves of the Panel Study of Income Dynamics, we find that the earnings process features nonlinear persistence and conditional skewness. We confirm these results using population register data from Norway. We then show that the impact of earnings shocks varies substantially across earnings histories, and that this nonlinearity drives heterogeneous consumption responses. The framework provides new empirical measures of partial insurance in which the transmission of income shocks to consumption varies systematically with assets, the level of the shock and the history of past shocks.

JEL CODE: C23, D31, D91.

KEYWORDS: Earnings dynamics, consumption, nonlinear persistence, partial insurance, panel data, quantile regression, latent variables.

*This paper was the basis for Arellano's Presidential Address to the Econometric Society in 2014. We are grateful to the co-editor and three anonymous referees for detailed comments. We also thank participants in regional meetings and seminars, especially to Xiaohong Chen, Mariacristina De Nardi, Jordi Gali, Fatih Guvenen, Lars Hansen, Jim Heckman, Yingyao Hu, Josep Pijoan, Enrique Sentana, and Kjetil Storesletten for their comments. We are particularly grateful to Luigi Pistaferri and Itay Saporta-Eksten for help with the PSID data, Magne Mogstad and Michael Graber for providing the estimations using the Norwegian population register data as part of the project on 'Labour Income Dynamics and the Insurance from Taxes, Transfers and the Family', and Ran Gu and Raffaele Saggio for excellent research assistance. Arellano acknowledges research funding from the Ministerio de Economía y Competitividad, Grant ECO2011-26342. Blundell would like to thank the ESRC Centre CPP at IFS and the ERC MicroConLab project for financial assistance. Bonhomme acknowledges support from the European Research Council/ ERC grant agreement n^o 263107.

[†]CEMFI, Madrid.

[‡]University College London and Institute for Fiscal Studies.

[§]University of Chicago.

1 Introduction

Consumption decisions and earnings dynamics are intricately linked. Together with the net value of assets, the size and durability of any income shock dictates how much consumption will need to adjust to ensure a reasonable standard of living in future periods of the life-cycle.¹ Understanding the persistence of earnings is therefore of key interest not only because it affects the permanent or transitory nature of inequality, but also because it drives much of the variation in consumption.² The precise nature of labor income dynamics and the distribution of idiosyncratic shocks also plays a central role in the design of optimal social insurance and taxation.³ This paper proposes a new nonlinear framework to study the persistence of earnings *and* the impact of earnings shocks on consumption.

With some notable exceptions (see the discussion and references in Meghir and Pistaferri, 2011), the literature on earnings dynamics has focused on linear models. The random walk permanent/transitory model is a popular example (Abowd and Card, 1989). Linear models have the property that all shocks are associated with the same persistence, irrespective of the household’s earnings history. Linearity is a convenient assumption, as it allows one to study identification and estimation using standard covariance techniques. However, by definition linear models rule out nonlinear transmission of shocks, while nonlinearities in income dynamics are likely to have a first-order impact on consumption and saving choices.

The existing literature on earnings shocks and consumption follows two main approaches. One approach is to take a stand on the precise mechanisms that households use to smooth consumption, for example saving and borrowing or labor supply, and to take a fully-specified life-cycle model to the data, see Gourinchas and Parker (2002), Guvenen and Smith (2014), or Kaplan and Violante (2014), for example. Except in very special cases (as in Hall and Mishkin, 1982) the consumption function is generally a complex nonlinear function of earnings components.⁴ Another approach is to estimate the degree of “partial insurance” from the data without precisely specifying the insurance mechanisms, see Blundell, Pistaferri and

¹See, for example, Jappelli and Pistaferri (2010) and references therein.

²See Deaton and Paxson (1994) for key initial insights on consumption inequality, and the subsequent literature reviewed in Blundell (2014).

³Golosov and Tsyvinski (2016) provide a recent review. In a dynamic Mirrlees tax design setting, optimal labor distortions for unexpectedly high shocks are determined mainly by the need to provide intertemporal insurance. Golosov *et al.* (2014) show that deviations from log normality can have serious repercussions for optimal capital and labor taxation.

⁴Interesting recent exceptions are Heathcote, Storesletten and Violante (2014) and the semi-structural approach in Alan, Browning and Ejrnaes (2014).

Preston (2008) for example. Linear approximations to optimality conditions from the optimization problem deliver tractable estimating equations. However, linear approximations may not always be accurate (Kaplan and Violante, 2010). Moreover, some aspects of consumption smoothing such as precautionary savings or asset accumulation in the presence of borrowing constraints and nonlinear persistence are complex in nature, making a linear framework less attractive. In this paper we develop a comprehensive new approach to study the nonlinear relationship between shocks to household earnings and consumption over the life cycle.

Our *first* contribution is to specify and estimate a nonlinear earnings process. In this framework, log-earnings are the sum of a general Markovian persistent component and a transitory innovation. Our interest mainly centers on the *conditional* distribution of the persistent component given its past. This is a comprehensive measure of the earnings risk faced by households. Conditional (or predictive) earnings distributions are a main feature of many models of consumption responses to income shocks, as in Kaplan and Violante (2014) for example, and also play a central role in the optimal design of social insurance, as in Golosov, Troshkin, and Tsyvinski (2014). Using quantile methods on both US and Norwegian data, we show that the conditional distribution of the persistent component of earnings exhibits important asymmetries. Our setup provides a tractable framework for incorporating such properties into structural decision models.

Our modeling approach to earnings dynamics captures the intuition that, unlike in linear models, different shocks may be associated with different persistence. The notion of persistence we propose is one of *persistence of histories*. In a Markovian setup, this is conveniently summarized using a derivative effect which measures by how much earnings at period t vary with the earnings component at $t - 1$, when hit by a particular shock at time t . This approach provides a new dimension of persistence where the impact of past shocks on current earnings may be altered by the size and sign of new shocks. For example, our framework allows for “unusual” shocks to wipe out the memory of past shocks.

Allowing for nonlinear persistence, and more generally for flexible models of conditional earnings distributions given past earnings, has both theoretical and empirical appeal. The main features of our nonlinear framework for earnings dynamics relate directly to existing structural labor market models. Consider a worker that accumulates occupational or industry specific skills, and these skills are permanently learned, perhaps with a small decay rate,

if used in the same occupation. A change of occupation or industry where such skills are unused makes them less valuable and, possibly, makes them depreciate much faster. In job ladder models earnings risk is asymmetric, job-loss risk affecting workers at the top of the ladder while workers at the bottom face opportunities to move up (Lise, 2013). In a recent model of earnings losses upon displacement, Huckfeldt (2016) includes an unskilled sector that tends to absorb laid-off workers. With low probability, workers have the opportunity to escape this sector and move to the skilled sector that features a job ladder. The escape from the unskilled sector is the event that wipes out the memory of the past bad shocks. From an empirical perspective, “unusual” shocks could correspond to job losses, changes of career, or health shocks. If such life-changing events are occasionally experienced by households, one would expect their predictive probability distributions over future income to feature nonlinear dynamic asymmetries.

Consider large, negative “unusual” income shocks, which not only have a direct effect but also cancel out the persistence of a good income history. For example, a worker hit by an adverse occupation-specific shock might find her skills less valuable. In that case her previous earnings history may matter much less after the shock. Using a parallel with the macroeconomic literature on disaster risk, these shocks could be called “microeconomic disasters”. While macroeconomic disasters could have potentially large effects on saving behavior (Rietz, 1988, Barro, 2006), they are so unlikely that they are statistically elusive events. In contrast, disasters at the micro level happen all the time to some individuals and therefore their dynamic consequences may have clear-cut empirical content.⁵

Such features are prominent in the empirical results that we report in this paper, and they are all at odds with linear models commonly used in the earnings dynamics literature. Moreover, despite recent advances on models of distributional earnings dynamics (for example Meghir and Pistaferri, 2004, or Botosaru and Sasaki, 2015), existing models do not seem well-suited to capture the nonlinear transmission of income shocks that we uncover.

Our *second* contribution is to develop an estimation framework to assess how consumption responds to earnings shocks in the data. In the baseline analysis we model the consumption policy rule as an age-dependent nonlinear function of assets, unobserved tastes and the

⁵The notion of “micro disasters” is also related to Castañeda, Díaz-Giménez and Ríos-Rull (2003), who argue that allowing for a substantial probability of downward risk for high-income households may help explain wealth inequality. See also Constantinides and Gosh (2016) and Schmidt (2015), who emphasize the asset pricing implications of income risk asymmetries.

persistent and transitory earnings components. We motivate our specification using a standard life-cycle model of consumption and saving with incomplete markets (as in Huggett, 1993, for example). In this model, as we illustrate through a small simulation exercise, a nonlinear earnings process with dynamic skewness of the type we uncover in our conditional quantile analysis will have qualitatively different implications for the level and distribution of consumption and assets over the life cycle in comparison to a linear earnings model.

The empirical consumption rule we develop is nonlinear, thus allowing for age-specific interactions between asset holdings and the earnings components. However, unlike fully specified structural approaches we model the consumption rule nonparametrically, leaving functional forms unrestricted. This modeling approach allows capturing an array of response coefficients that provides a rich picture of the extent of consumption insurance in the data. Moreover, there is no need for approximation arguments as we directly estimate the nonlinear consumption rule. Our consumption rule allows for unobserved household heterogeneity. We also show how to extend the framework to allow for advance information on earnings shocks and habits in consumption.

A virtue of our consumption framework is its ability to produce new empirical quantities, such as nonparametric marginal propensities to consume, that narrow the gap between policy-relevant evidence and structural modeling. At the same time, in the absence of further assumptions the model can be thought of as semi-structural and cannot generally be used to perform policy counterfactuals. As an example, in order to assess the impact of a change in the earnings process on consumption dynamics, one would need to take a precise stand on preferences and expectations, among other factors. Although given our goal to document nonlinear effects we do not impose such assumptions in this paper, economic structure could be added to our framework in order to conduct policy evaluation exercises.

Beyond the traditional covariance methods that have dominated the literature, new econometric techniques are needed to study our nonlinear model of earnings and consumption. Nonparametric identification can be established in our setup by building on a recent literature on nonlinear models with latent variables. Identification of the earnings process follows Hu and Schennach (2008) and Wilhelm (2015).⁶ Identification of the consumption rule relies on novel arguments, which extend standard instrumental-variables methods (as in Blundell *et al.*, 2008, for example) to our nonparametric setup.

⁶Lochner and Shin (2014) rely on related techniques to establish identification of a different nonlinear model of earnings.

To devise a tractable estimation strategy, we rely on the approach introduced by Arellano and Bonhomme (2016), adapted to a setup with time-varying latent variables. Wei and Carroll (2009) introduced a related estimation strategy in a cross-sectional context. The approach combines quantile regression methods, which are well-suited to capture nonlinear effects of earnings shocks, with regression methods on bases of functions, which are well-suited to flexibly model conditional distributions.⁷ To deal with the presence of the latent earnings components, we use a sequential estimation algorithm that consists in iterating between quantile regression estimation, and draws from the posterior distribution of the latent persistent components of earnings.

We take the earnings and consumption model to data from the Panel Study of Income Dynamics (PSID) for 1999-2009 and focus on working age families. Unlike earlier waves of the PSID, these data contain enhanced information on asset holdings and consumption expenditures in addition to labor earnings, see Blundell, Pistaferri and Saporta-Eksten (2016), for example. This is the first household panel to include detailed information on consumption and assets across the life-cycle for a representative sample of households. Although we abstract from labor supply, our modeling and estimation approach makes full use of the availability of panel information on earnings, consumption and assets. In addition, the quantile regression specifications that we use allow us to obtain rather precise estimates, despite the flexibility of the model and the moderate sample size.

Our empirical results show that the impact of earnings shocks varies substantially across households' earnings histories, and that this nonlinearity is a driver of heterogeneous consumption responses. Earnings data show the presence of nonlinear persistence, where “unusual” positive shocks for low earnings households, and negative shocks for high earnings households, are associated with lower persistence than other shocks. That is, such shocks have a higher propensity to wipe out the impact of the previous earnings history. Related to this, we find that conditional log-earnings distributions are asymmetric, skewed to the right (respectively, left) for households at the bottom (resp., top) of the income distribution. Although most of our results are based on PSID data, we show that similar empirical patterns hold in Norwegian administrative earnings data.

Regarding consumption, we find a significant degree of insurability of shocks to the persistent earnings components. We also uncover asymmetries in consumption responses to

⁷Misra and Surico (2014) use quantile methods to document the consumption responses associated with tax rebates.

earnings shocks that hit households at different points of the income distribution. Lastly, we find that assets play a role in the insurability of earnings shocks.

The literature on earnings dynamics is vast. Recent work has focused on non-Gaussianity (Geweke and Keane, 2000, Bonhomme and Robin, 2010) and heterogeneity (Alvarez, Browning and Ejrnaes, 2010). The nonlinear earnings persistence that we uncover is consistent with findings on US administrative tax records, such as Guvenen, Ozcan and Song (2014) and especially recent independent work by Guvenen *et al.* (2015). Relative to this growing body of research, the fact that our quantile-based methods are able to uncover previously unknown results in PSID survey data, and that these results also hold in administrative “big data” sets, is important because PSID uniquely provides joint longitudinal data on wealth, income and expenditures at household level. This allows us to conduct a joint empirical analysis of earnings and consumption patterns.⁸

The outline of the paper is as follows: In the next section we describe the earnings process and develop our measure of nonlinear persistence. Section 3 lays out the consumption model and defines a general representation of partial insurance to earnings shocks. In Section 4 we establish identification of the model. Section 5 describes our estimation strategy and the panel dataset. In Section 6 we present our empirical results. Section 7 concludes with a summary and some directions for future research. The Supplementary Appendix contains additional results.

2 Model (I): Earnings process

We start by describing our nonlinear model of earnings dynamics. In the next section we will present the consumption model.

2.1 The model

We consider a cohort of households, $i = 1, \dots, N$, and denote as t the age of the household head (relative to $t = 1$). Let Y_{it} be the pre-tax labor earnings of household i at age t , and

⁸Our nonlinear earnings model featuring asymmetric persistence is also related to time-series regime-switching models that are popular to analyze business cycle dynamics, see for example Evans and Watchel (1993)’s model of inflation uncertainty and Teräsvirta (1994) on smooth transition autoregressive models.

let y_{it} denote log- Y_{it} , net of a full set of age dummies. We decompose y_{it} as follows:⁹

$$y_{it} = \eta_{it} + \varepsilon_{it}, \quad i = 1, \dots, N, \quad t = 1, \dots, T, \quad (1)$$

where the probability distributions of η 's and ε 's are absolutely continuous.

The first, *persistent* component η_{it} is assumed to follow a general first-order Markov process. We denote the τ th conditional quantile of η_{it} given $\eta_{i,t-1}$ as $Q_t(\eta_{i,t-1}, \tau)$, for each $\tau \in (0, 1)$. The following representation is then without loss of generality:

$$\eta_{it} = Q_t(\eta_{i,t-1}, u_{it}), \quad (u_{it} | \eta_{i,t-1}, \eta_{i,t-2}, \dots) \sim \text{Uniform}(0, 1), \quad t = 2, \dots, T. \quad (2)$$

The dependence structure of the η process is not restricted beyond the first-order Markov assumption. The identification assumptions will only require η 's to be dependent over time, without specifying a particular (parametric) form of dependence.

The second, *transitory* component ε_{it} is assumed to have zero mean, to be independent over time and independent of η_{is} for all s . Even though more general moving average representations are commonly used in the literature, the biennial nature of the PSID data makes this assumption more plausible. Model (1)-(2) is intended as a representation of the uncertainty about persistent and transitory labor income in future periods that households face when deciding how much to spend and save. Our approach can be extended to allow for a moving average ε component, provided additional time periods are available, and for an unobserved time-invariant household-specific effect in addition to the two latent time-varying components η and ε (as done in the Supplementary Appendix).

Survey data like the PSID are often contaminated with errors (Bound *et al.*, 2001). In the absence of additional information, it is not possible to disentangle the transitory innovation from classical measurement error. Thus, an interpretation of our estimated distribution of ε_{it} is that it represents a mixture of transitory shocks and measurement error.¹⁰

Both earnings components are assumed mean independent of age t . However, the conditional quantile functions Q_t , and the marginal distributions of ε_{it} , may all depend on t . For a given cohort of households, age and calendar time are perfectly collinear, so this dependence

⁹Model (1) is additive in η and ε . Given our nonlinear approach, it is in principle possible to allow for interactions between the two earnings components, for example in $y_{it} = H_t(\eta_{it}, \varepsilon_{it})$, where identification could be established along the lines of Hu and Shum (2012).

¹⁰If additional information were available and the marginal distribution of a classical measurement error were known, one could recover the distribution of ε_{it} using a deconvolution argument. The estimation algorithm can be modified to deal with this case.

may capture age effects as well as aggregate shocks. The distribution of the initial condition η_{i1} is left unrestricted.

An important special case of model (1)-(2) is obtained when

$$y_{it} = \eta_{it} + \varepsilon_{it}, \quad \eta_{it} = \eta_{i,t-1} + v_{it}, \quad (3)$$

that is, when η_{it} follows a random walk. When v_{it} is independent of $\eta_{i,t-1}$ and has cumulative distribution function F_t , (2) becomes: $\eta_{it} = \eta_{i,t-1} + F_t^{-1}(u_{it})$. We will refer to the random walk plus independent shock as the *canonical model* of earnings dynamics.

2.2 Nonlinear dynamics

Model (1)-(2) allows for nonlinear dynamics of earnings. Here we focus on the ability of this specification to capture nonlinear persistence, and general forms of conditional heteroskedasticity.

Nonlinear persistence. We introduce the following quantities

$$\rho_t(\eta_{i,t-1}, \tau) = \frac{\partial Q_t(\eta_{i,t-1}, \tau)}{\partial \eta}, \quad \rho_t(\tau) = \mathbb{E} \left[\frac{\partial Q_t(\eta_{i,t-1}, \tau)}{\partial \eta} \right], \quad (4)$$

where $\partial Q_t / \partial \eta$ denotes the partial derivative of Q_t with respect to its first component and the expectation is taken with respect to the distribution of $\eta_{i,t-1}$.

The ρ 's in (4) are measures of nonlinear persistence of the η component.¹¹ $\rho_t(\eta_{i,t-1}, \tau)$ measures the persistence of $\eta_{i,t-1}$ when it is hit by a current shock u_{it} that has rank τ . This quantity depends on the lagged component $\eta_{i,t-1}$, and on the percentile of the shock τ . Average persistence across η values is $\rho_t(\tau)$. Note that, while the shocks u_{it} are i.i.d. by construction, they may greatly differ in the earnings persistence associated with them. The ρ 's are thus measures of *persistence of earnings histories*.

In the canonical model of earnings dynamics (3) where η_{it} is a random walk, $\rho_t(\eta_{i,t-1}, \tau) = 1$ irrespective of $\eta_{i,t-1}$ and τ . In contrast, in model (2) the persistence of $\eta_{i,t-1}$ may depend on the magnitude and direction of the shock u_{it} . As a result, the persistence of a shock to $\eta_{i,t-1}$ depends on the size and sign of current and future shocks $u_{it}, u_{i,t+1} \dots$. In particular, our

¹¹Note that $\rho_t(\eta_{i,t-1}, \tau)$ may be positive or negative, and may exceed 1 in absolute value. As a simple illustration, if $\ln(\eta_{it})$ is a random walk with standard Gaussian innovations, η_{it} itself is a multiplicative random walk, for which derivative measures of persistence in (4) do not vary with lagged η but vary with the value of the shock. For example, at the median shock the derivative is 1, but at the bottom quartile shock the derivative is around 0.5, and at the top quartile it is around 2.

model allows particular shocks to wipe out the memory of past shocks. As reviewed in the introduction, labor market models of the job ladder and occupational mobility can involve workers facing an increasing risk of a large fall in earnings, while those recently laid-off have a small probability of a positive shock that takes them into a skilled job where they can advance along the ladder. The interaction between the shock u_{it} and the lagged persistent component $\eta_{i,t-1}$ is a central feature of our nonlinear approach and, as we show below, it has substantive implications for consumption decisions.

It is useful to consider the following specification of the quantile function

$$Q_t(\eta_{i,t-1}, \tau) = \alpha_t(\tau) + \beta_t(\tau)'h(\eta_{i,t-1}), \quad (5)$$

where h is a multi-valued function. Our empirical specification will be based on (5), taking the components of h in a polynomial basis of functions. Persistence and average persistence in (5) are, respectively,

$$\rho_t(\eta_{i,t-1}, \tau) = \beta_t(\tau)' \frac{\partial h(\eta_{i,t-1})}{\partial \eta}, \quad \rho_t(\tau) = \beta_t(\tau)' \mathbb{E} \left[\frac{\partial h(\eta_{i,t-1})}{\partial \eta} \right],$$

thus allowing shocks to affect the persistence of $\eta_{i,t-1}$ in a flexible way. This measure is related to quantile autoregression parameters, as in Koenker and Xiao (2006).

Conditional heteroskedasticity. As model (2) does not restrict the form of the conditional distribution of η_{it} given $\eta_{i,t-1}$, it allows for general forms of heteroskedasticity. In particular, a measure of period- t uncertainty generated by the presence of shocks to the persistent earnings component is, for some $\tau \in (1/2, 1)$, $\sigma_t(\eta_{i,t-1}, \tau) = Q_t(\eta_{i,t-1}, \tau) - Q_t(\eta_{i,t-1}, 1 - \tau)$. For example, in the canonical model (3) with $v_{it} \sim \mathcal{N}(0, \sigma_{v_t}^2)$, we have $\sigma_t(\eta_{i,t-1}, \tau) = 2\sigma_{v_t} \Phi^{-1}(\tau)$.¹²

In addition, the model allows for conditional skewness and kurtosis in η_{it} . Along the lines of the skewness measure proposed by Kim and White (2004), one can consider, for some $\tau \in (1/2, 1)$,¹³

$$sk_t(\eta_{i,t-1}, \tau) = \frac{Q_t(\eta_{i,t-1}, \tau) + Q_t(\eta_{i,t-1}, 1 - \tau) - 2Q_t(\eta_{i,t-1}, \frac{1}{2})}{Q_t(\eta_{i,t-1}, \tau) - Q_t(\eta_{i,t-1}, 1 - \tau)}. \quad (6)$$

¹²The shock u_{it} is a rank. A persistent shock of a magnitude comparable to η_{it} can be constructed, among other ways, as $\zeta_{it} = Q_t(m_t, u_{it})$ where m_t is the median of η_{it} .


¹³Similarly, a measure of conditional kurtosis is, for some $\alpha < 1 - \tau$,

$$kur_t(\eta_{i,t-1}, \tau, \alpha) = \frac{Q_t(\eta_{i,t-1}, 1 - \alpha) - Q_t(\eta_{i,t-1}, \alpha)}{Q_t(\eta_{i,t-1}, \tau) - Q_t(\eta_{i,t-1}, 1 - \tau)}.$$

Figure 1: Quantile autoregressions of log-earnings

(a) PSID data

(b) Norwegian administrative data



Note: Residuals y_{it} of log pre-tax household labor earnings, Age 25-60 1999-2009 (US), Age 25-60 2005-2006 (Norway). See Section 6 and Appendix C for the list of controls. Estimates of the average derivative of the conditional quantile function of y_{it} given $y_{i,t-1}$ with respect to $y_{i,t-1}$. Quantile functions are specified as third-order Hermite polynomials.

Source: The Norwegian results are part of the project on ‘Labour Income Dynamics and the Insurance from Taxes, Transfers and the Family’. See Appendix C.

The empirical estimates below suggest that conditional skewness is a feature of the earnings process.

Preliminary evidence on nonlinear persistence. Suggestive evidence of nonlinearity in the persistence of earnings can be seen from Figure 1. This figure plots estimates of the average derivative, with respect to last period income $y_{i,t-1}$, of the conditional quantile function of current income y_{it} given $y_{i,t-1}$. This average derivative effect is a measure of persistence analogous to ρ_t in (4), except that here we use residuals y_{it} of log pre-tax household labor earnings on a set of demographics (including education and a polynomial in the age of the household head) as outcome variables. Given the nature of the PSID sample, panel (a) features biennial persistence estimates. On the two horizontal axes we report the percentile of $y_{i,t-1}$ (“ τ_{init} ”), and the percentile of the innovation of the quantile process (“ τ_{shock} ”). For estimation we use a series quantile specification, as in (5), based on a third-order Hermite polynomial.

This simple descriptive analysis not only shows the similarity in the patterns of the

nonlinearity of household earnings in both the PSID household panel survey data and in the annual population register data from Norway. It also suggests differences in the impact of an innovation to the quantile process (τ_{shock}) according to both the direction and magnitude of τ_{shock} and the percentile of the past level of income τ_{init} . Persistence of earnings history is highest when high earnings households (that is, high τ_{init}) are hit by a good shock (high τ_{shock}), and when low earnings households (that is, low τ_{init}) are hit by a bad shock (low τ_{shock}). In both cases, estimated persistence is close to $.9 - 1$. In contrast, bad shocks hitting high-earnings households, and good shocks hitting low-earnings ones, are associated with much lower persistence of earnings history, as low as $.3 - .4$. In Section 6 we will see that our nonlinear model that separates transitory shocks from the persistent component, reproduces the nonlinear persistence patterns of Figure 1 for the PSID panel survey data and the Norwegian population register data.

3 Model (II): Consumption rule

In order to motivate our empirical specification of the consumption function, we start by describing a standard stochastic life-cycle consumption model. We then use this setup to derive the form of the policy rule for household consumption, and describe the empirical consumption model that we will take to the data.

3.1 The consumption rule in a simple life-cycle model

We consider a theoretical framework where households act as single agents. Throughout their lifetime households have access to a single risk-free, one-period bond whose constant return is $1 + r$, and, at age t , face a period-to-period budget constraint

$$A_{it} = (1 + r)A_{i,t-1} + Y_{i,t-1} - C_{i,t-1}, \quad (7)$$

where A_{it} , Y_{it} and C_{it} denote assets, income and consumption, respectively. Family log-earnings are given by $\ln Y_{it} = \kappa_t + \eta_{it} + \varepsilon_{it}$, where κ_t is a deterministic age profile, and η_{it} and ε_{it} are persistent and transitory earnings components, respectively. At age t agents know η_{it} , ε_{it} and their past values, but not $\eta_{i,t+1}$ or $\varepsilon_{i,t+1}$, so there is no advance information. All distributions are known to households, and there is no aggregate uncertainty. In each period t in the life-cycle, the optimization problem is represented by the Bellman equation

$$V_t(A_{it}, \eta_{it}, \varepsilon_{it}) = \max_{C_{it}} u(C_{it}) + \beta \mathbb{E}_t [V_{t+1}(A_{i,t+1}, \eta_{i,t+1}, \varepsilon_{i,t+1})], \quad (8)$$

where $u(\cdot)$ is agents' utility, and β is the discount factor. An important element in (8) is the *conditional* distribution of the Markov component $\eta_{i,t+1}$ given η_{it} , which enters the expectation.

In a nonlinear earnings model such as (1)-(2), the presence of “unusual” shocks to earnings may lead to precautionary motives that induce high-income households to save more than they would do under a linear earnings model. Even with certainty equivalent preferences, under model (1)-(2) the discounting applied to persistent shocks will be state-dependent. In Section S1 of the Supplementary Appendix we illustrate these theoretical mechanisms in a two-period version of the model. In Section S2 we report the results of a simulation exercise based on the life-cycle model. We compare the canonical linear earnings model with an earnings model that features the presence of positive and negative “unusual” shocks. In the simulated economy, we find that an implication of the nonlinear earnings process is to reduce consumption among households on higher incomes. In particular, a negative shock for those on higher incomes reduces the persistence of the past and consequently is more damaging in terms of expected future incomes, thus inducing higher saving, higher wealth accumulation and lower consumption.¹⁴

In such a life-cycle model with uncertainty, the consumption rule will have the form

$$C_{it} = G_t(A_{it}, \eta_{it}, \varepsilon_{it}), \quad (9)$$

for some age-dependent function G_t . We will base our empirical specification on (9). The consumption rule at age t will be of this nonparametric form provided the state variables at t are period- t assets and the latent earnings components.¹⁵

In documenting dynamic patterns of consumption and earnings, one strategy is to take a stand on the functional form of the utility function and the distributions of the shocks, and to calibrate or estimate the model's parameters by comparing the model's predictions with the data. Another strategy is to linearize the Euler equation, with the help of the budget constraint; with a linear approximated problem at hand, standard covariance-based methods may be used for estimation. Our approach differs from those strategies as we

¹⁴In addition, in Section S2 of the Supplementary Appendix we report simulation results based on the same life-cycle model, using as input the earnings process we estimate on the PSID.

¹⁵Our approach may be extended to allow for habits or advance information, through simple modifications of the vector of state variables. There could also be additional borrowing constraints in each period. In that case, the nonparametric consumption rule in (9) would no longer be differentiable. The derivative effects defined below require differentiability of g_t .

directly estimate the nonlinear consumption rule (9). Doing so, we avoid linearized first-order conditions, and we estimate a flexible rule that is consistent with the above life-cycle consumption model. This approach allows documenting a rich set of derivative effects, thus shedding light on the patterns of consumption responses in the data.

3.2 Empirical consumption rule

Consider a cohort of households. Let c_{it} denote log-consumption net of age dummies. Similarly, let a_{it} denote log-assets net of age dummies. Our empirical specifications are based on

$$c_{it} = g_t(a_{it}, \eta_{it}, \varepsilon_{it}, \nu_{it}), \quad t = 1, \dots, T. \quad (10)$$

The ν_{it} are unobserved arguments of the consumption function, in addition to assets and the latent earnings components. In the specification without unobserved individual heterogeneity, ν_{it} are independent across periods and independent of $(a_{it}, \eta_{it}, \varepsilon_{it})$, and g_t is monotone in ν . An economic interpretation for ν is as a taste shifter that increases marginal utility. In the single-asset life-cycle model of Subsection 3.1 monotonicity is implied by the Bellman equation, provided $\frac{\partial u(C, \nu')}{\partial C} > \frac{\partial u(C, \nu)}{\partial C}$ for all C if $\nu' > \nu$, where the household's utility function is $u(C, \nu)$. Without loss of generality we normalize the marginal distribution of ν_{it} to be standard uniform in each period. From an empirical perspective the presence of the taste shifters ν_{it} in the consumption rule (10) may also partly capture measurement error in consumption expenditures. In the specification with unobserved individual heterogeneity, ν_{it} comprises two components: a time-invariant latent household factor, and an i.i.d. uniform component independent of the latter. Consumption is monotone in the second component, while being fully nonlinear in the first.

Clearly, the net assets variable a_{it} is not exogenous. In order to ensure identification, it suffices to specify a_{it} as sequentially exogenous (or “predetermined”). We will specify assets as a function of lagged assets, consumption, earnings, the persistent earnings component η , and age, as follows:

$$a_{it} = h_t(a_{i,t-1}, c_{i,t-1}, y_{i,t-1}, \eta_{i,t-1}, v_{it}), \quad (11)$$

where h_t is an age-specific function which is increasing in its last argument, and v_{it} are i.i.d. uniform independent of the other arguments. Note that the standard linear asset rule (7) is a special case of (11), so imposing (7) is not needed for identification or estimation. Taking

a stand on the budget constraint will however be required in order to simulate the impact of earnings shocks over the life cycle.

Derivative effects. Average consumption, for given values of asset holdings and earnings components, is

$$\mathbb{E}[c_{it} | a_{it} = a, \eta_{it} = \eta, \varepsilon_{it} = \varepsilon] = \mathbb{E}[g_t(a, \eta, \varepsilon, \nu_{it})].$$

Our framework allows us to document how average consumption varies as a function of assets and the two earnings components, and over the life cycle. In particular, the average derivative of consumption with respect to η is

$$\phi_t(a, \eta, \varepsilon) = \mathbb{E}\left[\frac{\partial g_t(a, \eta, \varepsilon, \nu_{it})}{\partial \eta}\right]. \quad (12)$$

The parameter $\phi_t(a, \eta, \varepsilon)$ reflects the degree of insurability of shocks to the persistent earnings component.¹⁶ We will document how this new *measure of partial insurance* varies over the life cycle, and how it depends on households' asset holdings, by reporting estimates of the average derivative effect $\bar{\phi}_t(a) = \mathbb{E}[\phi_t(a, \eta_{it}, \varepsilon_{it})]$. The quantity $1 - \bar{\phi}_t(a)$ is then a general measure of the degree of consumption insurability of shocks to the persistent earnings component, as a function of age and assets.

Dynamic effects of earnings shocks on consumption. Other measures of interest are the effects of an earnings shock u_{it} to the η component on consumption profile $c_{i,t+s}$, $s \geq 0$. For example, the contemporaneous effect can be computed, using the chain rule and equation (12), as

$$\mathbb{E}\left[\frac{\partial}{\partial u}\Big|_{u=\tau} g_t(a, Q_t(\eta, u), \varepsilon, \nu_{it})\right] = \phi_t(a, Q_t(\eta, \tau), \varepsilon) \frac{\partial Q_t(\eta, \tau)}{\partial u}.$$

This derivative effect depends on η through the insurance coefficient ϕ_t , but also through the quantity $\frac{\partial Q_t(\eta, \tau)}{\partial u}$ as the earnings model allows for general forms of conditional heteroskedasticity and skewness. The quantity $\frac{\partial Q_t(\eta, \tau)}{\partial u}$ measures the responsiveness of earnings to a shock u_{it} on impact. Note that its derivative with respect to η is equal to $\frac{\partial \rho_t(\eta, \tau)}{\partial \tau}$, where $\rho_t(\eta, \tau)$ is our persistence measure. In the empirical analysis we will report finite-difference counterparts to these derivative effects (“impulse responses”), and find an asymmetric impact of earnings shocks at different points of the income distribution.

¹⁶Likewise, the average derivative with respect to ε is $\psi_t(a, \eta, \varepsilon) = \mathbb{E}\left[\frac{\partial g_t(a, \eta, \varepsilon, \nu_{it})}{\partial \varepsilon}\right]$.

Multiple assets. Our approach can be generalized to represent the consumption policy function of a model with multiple assets differing in the stochastic properties of their returns. An example is a model that distinguishes between a risky asset and a risk-free asset, as often used in studies of household portfolios (e.g., Alan, 2012). A version of the consumption rule (10) with two assets would be consistent with this type of model as long as excess returns of the risky asset are not heterogeneous across households. In the presence of kinks induced by participation or transaction costs, our empirical consumption rule would capture a smoothed approximation.

The consumption rule (10) can also be extended to the Kaplan and Violante (2014) model of wealthy “hand-to-mouth” consumers. In that framework, the consumption policy rule is a nonlinear function of assets disaggregated into liquid and illiquid parts. Access to the illiquid, higher return asset involves a transaction cost. The separate assets interact in nonlinear ways with earnings. In this dynamic choice environment, the nonlinearity in the consumption model (10) incorporating the two separate assets can provide a smoothed approximation to the complex paths involved in the Kaplan and Violante (2014) composite consumption function.¹⁷

4 Identification

The earnings and consumption models take the form of nonlinear state-space models. A series of recent papers (notably Hu and Schennach, 2008, and Hu and Shum, 2012) has established conditions under which nonlinear models with latent variables are nonparametrically identified under conditional independence restrictions. Techniques developed in this literature can be used in order to establish identification of the models we consider.

4.1 Earnings process

Consider model (1)-(2), where η_{it} is a Markovian persistent component and ε_{it} are independent over time and independent of the η 's. We assume that the data contain T consecutive periods, $t = 1, \dots, T$. So, for a given cohort of households, $t = 1$ corresponds to the age at which the household head enters the sample, and $t = T$ corresponds to the last period of

¹⁷A separate question of interest, but one that is beyond the scope of this paper, concerns the identification of the latent consumption policy functions associated with the agent's accessing or not accessing the illiquid asset. Such a question could be posed if we observed a time-varying indicator of whether or not consumers access their illiquid assets.

observation.¹⁸ For that cohort, our aim is to identify the joint distributions of $(\eta_{i1}, \dots, \eta_{iT})$ and $(\varepsilon_{i1}, \dots, \varepsilon_{iT})$ given i.i.d. data from (y_{i1}, \dots, y_{iT}) . Four periods are needed in order to identify at least one Markov transition Q_t .

Conditions for the nonparametric identification of the earnings process are direct consequences of the analysis in Hu and Schennach (2008) and Wilhelm (2015). We provide such conditions in Appendix A. Identification is derived under several high-level assumptions. In particular, the distributions of $(y_{it}|y_{i,t-1})$ and $(\eta_{it}|y_{i,t-1})$ both need to satisfy *completeness* conditions. For example, the first condition requires that the only function h (in a suitable functional space) satisfying $\mathbb{E}[h(y_{it}) | y_{i,t-1}] = 0$ be $h = 0$. This requires that $\eta_{i,t-1}$ and η_{it} be statistically dependent, albeit without specifying the form of that dependence. An intuition for this is that if η 's were independent over time there would be no way to distinguish them from the transitory ε 's. Completeness is commonly assumed in nonparametric instrumental variables problems (Newey and Powell, 2003).

4.2 Consumption rule

Let us now turn to the identification of the consumption rule (10), starting with the case without unobserved heterogeneity. We make the following assumptions, where we denote $z_i^t = (z_{i1}, \dots, z_{it})$.

Assumption 1 For all $t \geq 1$,

- i*) $u_{i,t+s}$ and $\varepsilon_{i,t+s}$, for all $s \geq 0$, are independent of a_i^t, η_i^{t-1} , and y_i^{t-1} . ε_{i1} is independent of a_{i1} and η_{i1} .
- ii*) $a_{i,t+1}$ is independent of $(a_i^{t-1}, c_i^{t-1}, y_i^{t-1}, \eta_i^{t-1})$ conditional on $(a_{it}, c_{it}, y_{it}, \eta_{it})$.
- iii*) the taste shifter ν_{it} in (10) is independent of η_{i1} , $(u_{is}, \varepsilon_{is})$ for all s , ν_{is} for all $s \neq t$, and a_i^t .

Part *i*) in Assumption 1 requires current and future earnings shocks, which are independent of past components of earnings, to be independent of current and past asset holdings as well. At the same time, we let η_{i1} and a_{i1} be arbitrarily dependent. This is important, because asset accumulation upon entry in the sample may be correlated with past earnings shocks. Part *ii*) is a first-order Markov condition on asset accumulation. It is satisfied in a standard life-cycle model with one single risk-less asset, see equation (7). The assumption

¹⁸We consider a balanced panel for simplicity but our arguments can be extended to unbalanced panels.

also holds in such a model when the interest rate r_t is time-varying and known to households. More generally, the assumption allows the latent components of earnings η_{it} and ε_{it} to affect asset holdings separately, as in (11). Lastly, part *iii*) requires taste shifters to be independent over time, independent of earnings components, and independent of current and past assets. In particular, this rules out the presence of unobserved heterogeneity in consumption. We will relax this condition in the next subsection.

The identification argument proceeds in a sequential way. Starting with the first period, letting $y_i = (y_{i1}, \dots, y_{iT})$, and using f as a generic notation for a density function, we have

$$f(a_1|y) = \int f(a_1|\eta_1)f(\eta_1|y)d\eta_1, \quad (13)$$

where we have used that, by Assumption 1*i*), $f(a_1|\eta_1, y)$ and $f(a_1|\eta_1)$ coincide. We can rewrite (13) as

$$f(a_1|y) = \mathbb{E}[f(a_1|\eta_{i1}) | y_i = y], \quad (14)$$

where the expectation is taken with respect to the density of η_{i1} given y_i , for a fixed value a_1 . Hence, provided the distribution of $(\eta_{i1}|y_i)$ (which is identified from the earnings process, see above) is complete, the density $f(a_1|\eta_1)$ is identified from (14).¹⁹ Note that the density $f(a_1, \eta_1|y) = f(a_1|\eta_1)f(\eta_1|y)$ is also identified.

We then have, using the consumption rule and Assumption 1*iii*),

$$f(c_1|a_1, y) = \int f(c_1|a_1, \eta_1, y_1)f(\eta_1|a_1, y)d\eta_1, \quad (15)$$

or equivalently

$$f(c_1|a_1, y) = \mathbb{E}[f(c_1|a_{i1}, \eta_{i1}, y_{i1}) | a_{i1} = a_1, y_i = y], \quad (16)$$

where the conditional expectation is taken at fixed c_1 . Under completeness in (y_{i2}, \dots, y_{iT}) of the distribution of $(\eta_{i1}|a_{i1}, y_i)$ (which is identified from the previous paragraph),²⁰ the densities $f(c_1|a_1, \eta_1, y_1)$ and $f(c_1, \eta_1|a_1, y)$ are thus identified. Identification of the consumption function (10) for $t = 1$ follows since g_1 is the conditional quantile function of c_1 given a_1, η_1 , and $\varepsilon_1 = y_1 - \eta_1$.

¹⁹In fact, given that we are working with bounded density functions, it is sufficient that the distribution of $(\eta_{i1}|y_i)$ be *boundedly complete*; see Blundell, Chen and Kristensen (2007) for analysis and discussion.

²⁰Here by completeness in y_{i2} of the distribution of $(y_{i1}|y_{i2}, x_i)$ we mean that the only solution to $\mathbb{E}[h(y_{i1}, x_i)|y_{i2}, x_i] = 0$ is $h = 0$. This is the same as $(y_{i1}, x_i)|(y_{i2}, x_i)$ being complete. Note that, similarly as before, the weaker condition of bounded completeness suffices.

Second period's assets. Turning to period 2 we have, using Assumption 1*i*) and 1*iii*),

$$f(a_2|c_1, a_1, y) = \int f(a_2|c_1, a_1, \eta_1, y_1) f(\eta_1|c_1, a_1, y) d\eta_1, \quad (17)$$

from which it follows that the density $f(a_2|c_1, a_1, \eta_1, y_1)$ is identified, provided the distribution of $(\eta_{i1}|c_{i1}, a_{i1}, y_i)$ (which is identified from above) is complete in (y_{i2}, \dots, y_{iT}) .

In addition, using Bayes' rule and Assumption 1*i*) and 1*iii*),

$$f(\eta_2|a_2, c_1, a_1, y) = \int \frac{f(y|\eta_1, \eta_2, y_1) f(\eta_1, \eta_2|a_2, c_1, a_1, y_1)}{f(y|a_2, c_1, a_1, y_1)} d\eta_1.$$

So, as the density $f(\eta_1|a_2, c_1, a_1, y_1)$ is identified from above, and by Assumption 1

$$f(\eta_1, \eta_2|a_2, c_1, a_1, y_1) = f(\eta_1|a_2, c_1, a_1, y_1) f(\eta_2|\eta_1),$$

it follows that $f(\eta_2|a_2, c_1, a_1, y)$ is identified.

Subsequent periods. To see how the argument extends to subsequent periods, consider second period's consumption. We have, using Assumption 1*iii*),

$$f(c_2|a_2, c_1, a_1, y) = \int f(c_2|a_2, \eta_2, y_2) f(\eta_2|a_2, c_1, a_1, y) d\eta_2. \quad (18)$$

Provided the distribution of $(\eta_{i2}|a_{i2}, c_{i1}, a_{i1}, y_i)$ (which is identified from the previous paragraph) is complete in $(c_{i1}, a_{i1}, y_{i1}, y_{i3}, \dots, y_{iT})$, the density $f(c_2|a_2, \eta_2, y_2)$ is identified.

By induction, using in addition Assumption 1*ii*) from the third period onward, the joint density of η 's, consumption, assets, and earnings is identified provided, for all $t \geq 1$, the distributions of $(\eta_{it}|c_i^t, a_i^t, y_i)$ and $(\eta_{it}|c_i^{t-1}, a_i^t, y_i)$ are complete in $(c_i^{t-1}, a_i^{t-1}, y_i^{t-1}, y_{i,t+1}, \dots, y_{iT})$.

Discussion. The identification arguments depend on completeness conditions, which relate to the relevance, in a nonparametric sense, of the instruments. To illustrate this, consider the completeness of the distribution of $(\eta_{i1}|y_i)$ in (y_{i2}, \dots, y_{iT}) , which we use to show the identification of the consumption rule in the first period, see (16). Here we abstract from assets for simplicity. The completeness condition then depends on the properties of the earnings process. As an example, consider the case where $T = 2$, and $(\eta_{i1}, y_{i1}, y_{i2})$ follows a multivariate normal distribution with zero mean. Then $\eta_{i1} = \alpha y_{i1} + \beta y_{i2} + \zeta_i$, where ζ_i is normal $(0, \sigma^2)$, independent of (y_{i1}, y_{i2}) . It can be easily shown that $\beta \neq 0$ if $\text{Cov}(\eta_{i1}, \eta_{i2}) \neq 0$, in which case the distribution of $(\eta_{i1}|y_{i1}, y_{i2})$ is complete in y_{i2} . As in the identification of the earnings process, identification of the consumption rule thus relies on η 's being dependent

over time. Beyond the settings studied so far in the literature (such as in D’Haultfoeuille, 2011, Andrews, 2011, or Hu and Shiu, 2012), it would be of great interest to provide primitive conditions for completeness in the nonlinear models we focus on here.

An intuitive explanation for the identification argument comes from the link to nonparametric instrumental variables (NPIV). In period 1, for a fixed a_1 , (14) is analogous to an NPIV problem where η_{i1} is the endogenous regressor and $y_i = (y_{i1}, \dots, y_{iT})$ is the vector of instruments. Likewise, conditional on (a_{i1}, y_{i1}) , (y_{i2}, \dots, y_{iT}) are the “excluded instruments” for η_{i1} in (16). In subsequent periods, lagged consumption and assets are used as instruments, together with lags and leads of earnings. Using leads of log-earnings for identifying consumption responses is a common strategy in linear models, see for example Hall and Mishkin (1982) and Blundell *et al.* (2008).

4.3 Household unobserved heterogeneity

Accounting for unobserved heterogeneity in preferences or discounting, for example, may be empirically important. Heterogeneity in discount factors is also a popular mechanism in quantitative macroeconomic models to generate realistic wealth inequality, see for example Krusell and Smith (1998) and Krueger, Mitman and Perri (2015). With unobserved heterogeneity the consumption rule takes the form

$$c_{it} = g_t(a_{it}, \eta_{it}, \varepsilon_{it}, \xi_i, \tilde{\nu}_{it}), \quad t = 1, \dots, T, \quad (19)$$

where ξ_i is a household-specific effect and $\tilde{\nu}_{it}$ are i.i.d. standard uniform. The distribution of $(\xi_i, \eta_{i1}, a_{i1})$ is left unrestricted. Therefore, ξ_i is treated as a “fixed effect”. Even if a fully unstructured distinction between unobserved heterogeneity and individual dynamics in a finite horizon panel is not possible, finite-dimensional fixed effects can be included nonparametrically in the consumption (and earnings) equations as long as T is sufficiently large. In Appendix A we provide conditions for identification of this model, by relying on results from Hu and Schennach (2008). For simplicity we consider scalar heterogeneity ξ_i . Depending on the number of available time periods, a vector of unobserved heterogeneity could be allowed for.

Lastly, in Section S3 of the Supplementary Appendix we consider several extensions of the model. We show how to allow for unobserved heterogeneity in earnings, as in the following

specification that we will take to the data:

$$y_{it} = \eta_{it} + \zeta_i + \varepsilon_{it}, \quad (20)$$

where $\eta_{it} = Q_t(\eta_{i,t-1}, u_{it})$ is first-order Markov. We also describe how to allow for dependence in transitory shocks, advance earnings information, and consumption habits.

5 Estimation strategy

5.1 Empirical specification

Earnings components. The earnings model depends on the Markovian transitions of the persistent component $Q_t(\cdot, \cdot)$, the marginal distributions of ε_{it} , and the marginal distribution of the initial persistent component η_{i1} . We now explain how we empirically specify these three components.

Let φ_k , for $k = 0, 1, \dots$, denote a dictionary of bivariate functions, with $\varphi_0 = 1$. Letting age_{it} denote the age of the head of household i in period t , we specify

$$\begin{aligned} Q_t(\eta_{i,t-1}, \tau) &= Q(\eta_{i,t-1}, age_{it}, \tau) \\ &= \sum_{k=0}^K a_k^Q(\tau) \varphi_k(\eta_{i,t-1}, age_{it}). \end{aligned} \quad (21)$$

In practice we use low-order products of Hermite polynomials for φ_k . We specify the quantile function of ε_{it} (for $t = 1, \dots, T$) given age_{it} , and that of η_{i1} given age at the start of the period age_{i1} , in a similar way. Specifically, we set

$$\begin{aligned} Q_\varepsilon(age_{it}, \tau) &= \sum_{k=0}^K a_k^\varepsilon(\tau) \varphi_k(age_{it}), \\ Q_{\eta_1}(age_{i1}, \tau) &= \sum_{k=0}^K a_k^{\eta_1}(\tau) \varphi_k(age_{i1}), \end{aligned}$$

with outcome-specific choices for K and φ_k .

The quantile model (21) provides a flexible specification of the conditional distribution of η_{it} given $\eta_{i,t-1}$ and age. Similarly, our quantile specifications flexibly model how ε_{it} and η_{i1} depend on age, at every quantile. We include the age of the household head as a control, while ruling out dependence on calendar time. This choice is motivated by our desire to model life-cycle evolution, as well as by the relative stationarity of the earnings distributions (conditional on age) during the 1999-2009 period and the relatively small sample size. On

larger samples, an interesting avenue will be to allow for variation in both age and calendar time within our framework.²¹

Consumption rule. We specify the conditional distribution of consumption given current assets and earnings components as follows:

$$\begin{aligned} g_t(a_{it}, \eta_{it}, \varepsilon_{it}, \tau) &= g(a_{it}, \eta_{it}, \varepsilon_{it}, age_{it}, \tau) \\ &= \sum_{k=1}^K b_k^g \tilde{\varphi}_k(a_{it}, \eta_{it}, \varepsilon_{it}, age_{it}) + b_0^g(\tau), \end{aligned} \quad (22)$$

where $\tilde{\varphi}_k$ is a dictionary of functions (in practice, another product of Hermite polynomials).

Equation (22) is a nonlinear regression model. In contrast with (21), the consumption model is additive in τ . It would be conceptually straightforward to let all coefficients b_k^g depend on τ , although this would lead to a less parsimonious specification. Below we augment the consumption function to also depend nonlinearly on a household-specific effect.

Assets evolution. We specify the distribution of initial assets a_{i1} conditional on the initial persistent component η_{i1} and the age at the start of the period age_{i1} as

$$Q_a(\eta_{i1}, age_{i1}, \tau) = \sum_{k=0}^K b_k^a(\tau) \tilde{\varphi}_k(\eta_{i1}, age_{i1}), \quad (23)$$

for different choices for K and $\tilde{\varphi}_k$. We then specify how assets evolve as a function of lagged assets, consumption, earnings, the persistent earnings component η , and age, using (11), where

$$\begin{aligned} h_t(a_{i,t-1}, c_{i,t-1}, y_{i,t-1}, \eta_{i,t-1}, \tau) &= h(a_{i,t-1}, c_{i,t-1}, y_{i,t-1}, \eta_{i,t-1}, age_{it}, \tau) \\ &= \sum_{k=1}^K b_k^h \tilde{\varphi}_k(a_{i,t-1}, c_{i,t-1}, y_{i,t-1}, \eta_{i,t-1}, age_{it}) + b_0^h(\tau), \end{aligned} \quad (24)$$

for some K and $\tilde{\varphi}_k$.²²

²¹The functional form in (21) does not enforce monotonicity in τ but our estimation method will produce an automatic rearrangement of quantiles if needed.

²²In a previous version of the paper we estimated the model imposing that $\eta_{i,t-1}$ does not enter (24), which is still consistent with the budget constraint (7) and avoids the need to model predetermined assets. We obtained qualitatively similar empirical results.

Implementation. The functions a_k^Q , a_k^ε and $a_k^{\eta_1}$ are indexed by a finite-dimensional parameter vector θ . Likewise, the functions b_0^g , b_0^h , and b_k^a are indexed by a parameter vector μ that also contains $b_1^g, \dots, b_K^g, b_1^h, \dots, b_K^h$. We base our implementation on Wei and Carroll (2009) and Arellano and Bonhomme (2016). As in these papers we model the functions a_k^Q as piecewise-polynomial interpolating splines on a grid $[\tau_1, \tau_2], [\tau_2, \tau_3], \dots, [\tau_{L-1}, \tau_L]$, contained in the unit interval. We extend the specification of the intercept coefficient a_0^Q on $(0, \tau_1]$ and $[\tau_L, 1)$ using a parametric model indexed by λ^Q . All a_k^Q for $k \geq 1$ are constant on $[0, \tau_1]$ and $[\tau_L, 1]$, respectively. Hence, denoting $a_{k\ell}^Q = a_k^Q(\tau_\ell)$, the functions a_k^Q depend on $\{a_{11}^Q, \dots, a_{KL}^Q, \lambda^Q\}$.

Unlike in an ordinary quantile regression, the dependence of the parameters on the percentiles τ needs to be specified because some of our regressors are latent variables. In practice, we take $L = 11$ and $\tau_\ell = \ell/(L + 1)$. The functions a_k^Q are taken as piecewise-linear on $[\tau_1, \tau_L]$. An advantage of this specification is that the likelihood function is available in closed form. In addition, we specify a_0^Q as the quantile of an exponential distribution on $(0, \tau_1]$ (with parameter λ_-^Q) and $[\tau_L, 1)$ (with parameter λ_+^Q).²³

We proceed similarly to model a_k^ε , $a_k^{\eta_1}$, and b_k^a . Moreover, as our data show little evidence against consumption being log-normal, we set $b_0^g(\tau)$ to $\alpha + \sigma\Phi^{-1}(\tau)$, where (α, σ) are parameters to be estimated. We proceed similarly for $b_0^h(\tau)$. We also estimated two different versions of the model with more flexible specifications for $b_0^g(\tau)$ and $b_0^h(\tau)$: based on quantiles on a grid with $L = 11$ knots, and allowing for an age effect in the variance of the consumption innovation. In both cases we found very similar results to the ones we report below. We use tensor products of Hermite polynomials for φ_k and $\tilde{\varphi}_k$, each component of the product taking as argument a standardized variable.²⁴

²³As a result, we have

$$a_0^Q(\tau) = \frac{1}{\lambda_-^Q} \log\left(\frac{\tau}{\tau_1}\right) \mathbf{1}\{0 < \tau < \tau_1\} + \sum_{\ell=1}^{L-1} \left(a_{k\ell}^Q + \frac{a_{k,\ell+1}^Q - a_{k\ell}^Q}{\tau_{\ell+1} - \tau_\ell} (\tau - \tau_\ell) \right) \mathbf{1}\{\tau_\ell \leq \tau < \tau_{\ell+1}\} - \frac{1}{\lambda_+^Q} \log\left(\frac{1-\tau}{1-\tau_L}\right) \mathbf{1}\{\tau_L \leq \tau < 1\}.$$

²⁴For example, $a_t/std(a)$, $\eta_t/std(y)$, $\varepsilon_t/std(y)$, and $(age_t - mean(age))/std(age)$ are used as arguments of the consumption rule.

Household unobserved heterogeneity. When allowing for household unobserved heterogeneity in consumption/assets, we model log-consumption as

$$c_{it} = g(a_{it}, \eta_{it}, \varepsilon_{it}, age_{it}, \xi_i, \tilde{\nu}_{it}), \quad (25)$$

which we specify similarly as in (22), with parameters \tilde{b}_k^g . To fix the scale of the function we impose that

$$\sum_{k=1}^K \tilde{b}_k^g \tilde{\varphi}_k(0, 0, 0, \overline{age}, \xi) = \xi, \quad \text{for all } \xi, \quad (26)$$

where \overline{age} denotes the mean value of age in the sample. Likewise, we model assets as

$$a_{it} = h(a_{i,t-1}, c_{i,t-1}, y_{i,t-1}, \eta_{i,t-1}, age_{it}, \xi_i, \tilde{\nu}_{it}), \quad (27)$$

with a similar specification as in (24). Lastly, we specify $\xi_i = q(a_{i1}, \eta_{i1}, age_{i1}, \omega_i)$, with ω_i uniform on $(0, 1)$, using a quantile modeling as in (23).²⁵

5.2 Overview of the estimation algorithm

The algorithm is an adaptation of techniques developed in Arellano and Bonhomme (2016) to a setting with time-varying latent variables. The first estimation step recovers estimates of the earnings parameters θ . The second step recovers estimates of the consumption and assets parameters μ , given a previous estimate of θ . Our choice of a sequential estimation strategy, rather than joint estimation of (θ, μ) , is motivated by the fact that θ is identified from the earnings process alone. In contrast, in a joint estimation approach, estimates of the earnings process would be partly driven by the consumption model. Here we describe the estimation of the earnings parameters θ . Estimation of the consumption parameters μ is similar. The model’s restrictions are described in detail in Appendix B.

A compact notation for the restrictions implied by the earnings model is

$$\bar{\theta} = \underset{\theta}{\operatorname{argmin}} \mathbb{E} \left[\int R(y_i, \eta; \theta) f_i(\eta; \bar{\theta}) d\eta \right],$$

where R is a known function, $\bar{\theta}$ denotes the true value of θ , and $f_i(\cdot; \bar{\theta}) = f(\cdot | y_i^T, age_i^T; \bar{\theta})$ denotes the posterior density of $(\eta_{i1}, \dots, \eta_{iT})$ given the earnings data. The estimation algorithm is closely related to the “stochastic EM” algorithm (Celeux and Diebolt, 1993). Stochastic

²⁵We proceed analogously when allowing for an additive household-specific effect ζ_i in log-earnings $y_{it} = \eta_{it} + \zeta_i + \varepsilon_{it}$, where η_{it} is given by (2). There we allow for flexible dependence between η_{i1} , ζ_i and age_{i1} through another series quantile model.

EM is a simulated version of the classical EM algorithm of Dempster *et al.* (1977), where new draws from η are computed in every iteration of the algorithm. One difference is that, unlike in EM, our problem is not likelihood-based. Instead, we exploit the computational convenience of quantile regression and replace likelihood maximization by a sequence of quantile regressions in each M-step of the algorithm.

Starting with a parameter vector $\widehat{\theta}^{(0)}$, we iterate the following two steps on $s = 0, 1, 2, \dots$ until convergence of the $\widehat{\theta}^{(s)}$ process.

1. *Stochastic E-step*: Draw $\eta_i^{(m)} = (\eta_{i1}^{(m)}, \dots, \eta_{iT}^{(m)})$ for $m = 1, \dots, M$ from $f_i(\cdot; \widehat{\theta}^{(s)})$.
2. *M-step*: Compute

$$\widehat{\theta}^{(s+1)} = \underset{\theta}{\operatorname{argmin}} \sum_{i=1}^N \sum_{m=1}^M R(y_i, \eta_i^{(m)}; \theta).$$

Note that, as the likelihood function is available in closed form, the E-step is straightforward. In practice we use a random-walk Metropolis-Hastings sampler for this purpose, targeting an acceptance rate of approximately 30%. The M-step consists of a number of quantile regressions. For example, the parameters $a_{k\ell}^Q$ are updated as

$$\min_{(a_{0\ell}^Q, \dots, a_{K\ell}^Q)} \sum_{i=1}^N \sum_{t=2}^T \sum_{m=1}^M \rho_{\tau_\ell} \left(\eta_{it}^{(m)} - \sum_{k=0}^K a_{k\ell}^Q \varphi_k(\eta_{i,t-1}^{(m)}, \operatorname{age}_{it}) \right), \quad \ell = 1, \dots, L,$$

where $\rho_\tau(u) = u(\tau - \mathbf{1}\{u \leq 0\})$ is the ‘‘check’’ function. This is a set of standard quantile regressions, associated with convex objective functions. We proceed in a similar way to update all other parameters, see Appendix B for details.

In practice we first estimate the effect of age on mean log-earnings by regressing them on a quartic in age. We then impose in each iteration of the algorithm that ε_{it} and age are uncorrelated (although we allow for age effects on the variance and quantiles of ε_{it}). We take $M = 1$, stop the chain after a large number of iterations, and report an average across the last \widetilde{S} values $\widehat{\theta} = \frac{1}{\widetilde{S}} \sum_{s=S-\widetilde{S}+1}^S \widehat{\theta}^{(s)}$, and similarly for consumption-related parameters $\widehat{\mu}$.²⁶ The results for the earnings parameters are based on $S = 500$ iterations, with 200 Metropolis-Hastings draws in each iteration. Consumption-related parameters are estimated using 200 iterations with 200 draws per iteration. In both cases we take $\widetilde{S} = S/2$. In our experiments we observed that the algorithm may get ‘‘stuck’’ on what appears to be a local regime of the

²⁶As an alternative to taking a large M , taking $M = 1$ and averaging the parameter draws is computationally convenient in our setting.

Markov chain. We started the algorithm from a large number of initial parameter values, and selected the estimates yielding the highest average log-likelihood over iterations. The non-selected values tended to give very similar pictures to the ones we report below.

Properties. Nielsen (2000) studies the statistical properties of the stochastic EM algorithm in a likelihood case. He provides conditions under which the Markov Chain $\hat{\theta}^{(s)}$ is ergodic, for a fixed sample size. He also characterizes the asymptotic distribution of $\hat{\theta}$ as the sample size N tends to infinity. Arellano and Bonhomme (2016) characterize the asymptotic distribution of $\hat{\theta}$ in a case where the optimization step is not likelihood-based but relies on quantile-based estimating equations. The estimator $\hat{\theta}$ is root- N consistent and asymptotically normal under correct specification of the parametric model, for K and L fixed. Note that an alternative, nonparametric approach, would be to let K and L increase with N at an appropriate rate so as to let the approximation bias tend to zero.²⁷ Studying inference in our problem as (N, K, L) jointly tend to infinity is an interesting avenue for future work.

6 Empirical results on earnings and consumption

In this section we present our empirical results. We start by describing our main data source, the Panel Study of Income Dynamics (PSID). We then show how earnings and consumption respond to income shocks. We also corroborate our findings for the nonlinear earnings process using administrative data on household earnings from the Norwegian population register. Finally, we report simulation exercises based on the estimated model.

6.1 Panel Data

Panel data on household consumption, income and assets are rare. The PSID began the collection of detailed data on consumption expenditures and asset holdings in 1999, in addition to household earnings and demographics. An annual wave is available every other year. We use data for the 1999-2009 period (six waves).

Earnings Y_{it} are total pre-tax household labor earnings. We construct y_{it} as residuals from regressing log household earnings on a set of demographics, which include cohort interacted with education categories for both household members, race, state and large-city dummies, a

²⁷See Belloni *et al.* (2016) for an analysis of inference for series quantile regression, and Arellano and Bonhomme (2016) for a consistency analysis in a panel data model closely related to the one we consider here.

family size indicator, number of kids, a dummy for income recipient other than husband and wife, and a dummy for kids out of the household. Controls for family size and composition are included so as to equalize household earnings (likewise for consumption and assets below). Education, race and geographic dummies are included in an attempt to capture individual heterogeneity beyond cohort effects and the initial persistent component of earnings η_{i1} . Removing demographic-specific means in a preliminary step has been the standard practice in the empirical analysis of earnings dynamics. A more satisfactory approach would integrate both steps, especially given our emphasis on nonlinearities. However, except for age, we did not attempt a richer conditioning in light of sample size.

We use data on consumption C_{it} of nondurables and services. The panel data contain information on health expenditures, utilities, car-related expenditures and transportation, education, child care, and food expenditures. Recreation, alcohol, tobacco and clothing (the latter available from 2005) are the main missing items. Rent information is available for renters, but not for home owners. We follow Blundell, Pistaferri and Saporta-Eksten (2016) and impute rent expenditures for home owners.²⁸ In total, approximately 67% of consumption expenditures on nondurables and services are covered. We construct c_{it} as residuals of log total consumption on the same set of demographics as for earnings.

Asset holdings A_{it} are constructed as the sum of financial assets (including cash, stocks and bonds), real estate value, pension funds, and car value, net of mortgages and other debt. We construct residuals a_{it} by regressing log-assets on the same set of demographics as for earnings and consumption. These log-assets residuals will enter as arguments of the nonlinear consumption rule (10).

To select the sample we follow Blundell *et al.* (2016) and focus on a sample of participating and married male heads aged between 25 and 60. We drop all observations for which data on earnings, consumption, or assets, either in levels or log-residuals, are missing. See Appendix C for further details. In the analysis we focus on a balanced subsample of $N = 792$ households. Table C1 in Appendix C shows mean total earnings, consumption and asset holdings, by year. Compared to Blundell *et al.* (2016), households in our balanced sample have higher assets, and to a less extent higher earnings and consumption. The table also shows a large and increasing dispersion of assets across households. The evolution of assets may partly reflect the housing boom and bust, including the effect of the Great recession at

²⁸Note that, as a result, consumption responds automatically to variations in house prices. An alternative would be to exclude rents and imputed rents from consumption expenditures.

the end of the sample. Although our framework could be used to document distributional dynamics along the business cycle, we abstract from business cycle effects in this paper.

Lastly, the sample that we use is relatively homogeneous. Including households with less stable employment histories would be interesting, but it would require extending our framework. We return to this point in the conclusion.

6.2 Earnings

We next comment on the empirical estimates of the earnings process. Figure 2 (a) reproduces Figure 1 (a). It shows estimates of the average derivative of the conditional quantile function of log-earnings residuals y_{it} given $y_{i,t-1}$ with respect to $y_{i,t-1}$ in the PSID sample. The figure suggests the presence of nonlinear persistence, which depends on both the percentile of past income (τ_{init}) and the percentile of the quantile innovation (τ_{shock}). This empirical pattern is also present for male wages, see Figure S1 of the Supplementary Appendix. We then estimate the earnings model,²⁹ and given the estimated parameters we simulate the model.³⁰ Figure 2 (b), which is based on simulated data, shows that our nonlinear model reproduces the patterns of nonlinear persistence well. In contrast, standard models have difficulty fitting this empirical pattern. For example, we estimated a simple version of the canonical earnings dynamics model (3) with a random walk component and independent transitory shocks.³¹ Figure 2 (c) shows that the average derivative of the quantile function is nearly constant (up to simulation error) with respect to τ_{shock} and τ_{init} . This linear specification without interaction effects between earnings shocks and past earnings components stands in contrast with the data.³²

Figure 2 (d) then shows the estimated persistence of the earnings component η_{it} . Specifically, the graph shows $\rho_t(\eta_{i,t-1}, \tau)$ from equation (4), evaluated at percentiles τ_{init} and τ_{shock} and at the mean age in the sample (47.5 years).³³ Persistence in η 's is higher than persis-

²⁹We use tensor products of Hermite polynomials of degrees (3, 2) for the conditional quantile function of η_{it} given $\eta_{i,t-1}$ and age, and second-order polynomials for ε_{it} and η_{i1} as a function of age.


³⁰We draw 20 earnings values per household. In the simulation we impose that the support of simulated η draws be less than 3 times the empirical support of log-earnings residuals. This affects very few observations.

³¹Estimation is based on equally-weighted minimum distance using the covariance structure predicted by the canonical model.

³²In order to assess the sensitivity to the polynomial functional form that we use, in Figure S2 of the Supplementary Appendix we report persistence estimates based on piecewise-linear specifications based on 9 or 25 pieces. The characteristic shape of Figure 2 (a) remains present.

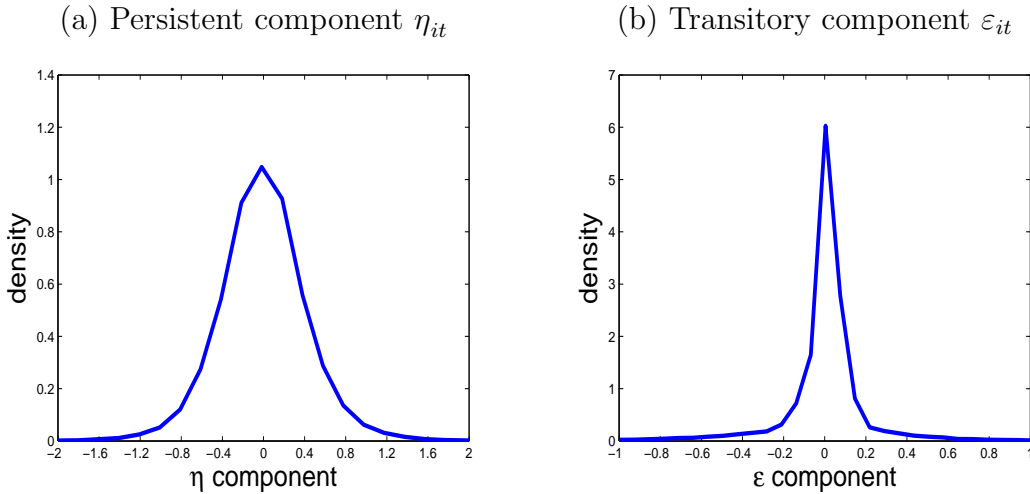
³³The estimated persistence is similar when averaging over age, see Figure S3 of the Supplementary Appendix.

Figure 2: Nonlinear persistence



Note: Graphs (a), (b), and (c) show estimates of the average derivative of the conditional quantile function of y_{it} given $y_{i,t-1}$ with respect to $y_{i,t-1}$, evaluated at percentile τ_{shock} and at a value of $y_{i,t-1}$ that corresponds to the τ_{init} percentile of the distribution of $y_{i,t-1}$. Graph (a) is based on the PSID data, graph (b) is based on data simulated according to our nonlinear earnings model with parameters set to their estimated values, and graph (c) is based on data simulated according to the canonical random walk earnings model (3). Graph (d) shows estimates of the average derivative of the conditional quantile function of η_{it} on $\eta_{i,t-1}$ with respect to $\eta_{i,t-1}$, based on estimates from the nonlinear earnings model.

Figure 3: Densities of persistent and transitory earnings components



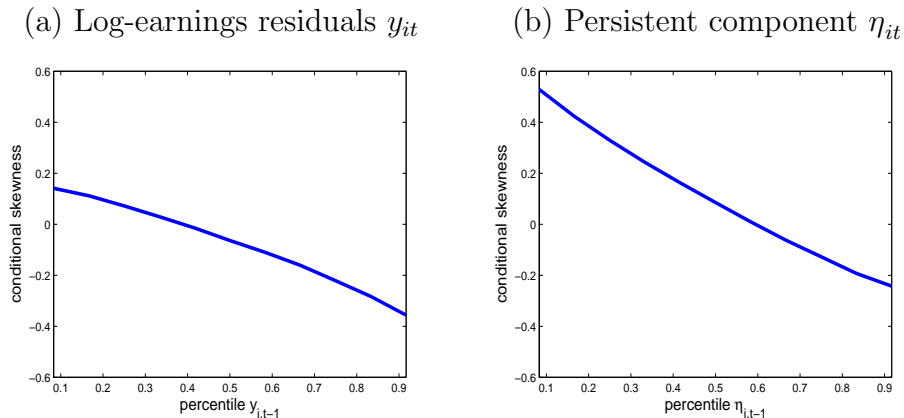
Note: Nonparametric estimates of densities based on simulated data according to the nonlinear model, using a Gaussian kernel.

tence in log-earnings residuals, consistently with the fact that Figure 2 (d) is net of transitory shocks. Persistence is close to 1 for high earnings households hit by good shocks, and for low earnings households hit by bad shocks. At the same time, persistence is lower, down to .6 – .8, when bad shocks hit high-earnings households or good shocks hit low-earnings ones.

Densities and moments. Figure 3 shows estimates of the marginal distributions of the persistent and transitory earnings components at mean age. While the persistent component η_{it} shows small departures from Gaussianity, the density of ε_{it} is clearly non-normal and presents high kurtosis and fat tails. These results are qualitatively consistent with empirical estimates of non-Gaussian linear models in Horowitz and Markatou (1996) and Bonhomme and Robin (2010).

In Figure 4 we report the measure of conditional skewness in (6), for $\tau = 11/12$, for both log-earnings residuals (left graph) and the η component (right). Panel (b) shows that η_{it} is positively skewed for low values of $\eta_{i,t-1}$, and negatively skewed for high values of $\eta_{i,t-1}$. This is in line with the nonlinear persistence reported in Figure 2 (d): when low- η households are hit by an unusually positive shock, dependence of η_{it} on $\eta_{i,t-1}$ is low with the result that they have a relatively large probability of outcomes far to the right from the central part of the distribution. Likewise, high- η households have a relatively large probability of getting outcomes far to the left of their distribution associated with low persistence episodes.

Figure 4: Conditional skewness of log-earnings residuals and η component



Note: Conditional skewness $sk(y, \tau)$ and $sk(\eta, \tau)$, see equation (6), for $\tau = 11/12$. Log-earnings residuals (data, left) and η component (right). The x-axis shows the conditioning variable, the y-axis shows the corresponding value of the conditional skewness measure.

Panel (a) similarly suggests the presence of conditional asymmetry in log-earnings residuals, although the evidence seems less strong than for η .

In addition, in Figures S4 to S8 of the Supplementary Appendix we report several measures of fit of the model. We show quantile-based estimates of conditional dispersion and conditional skewness. We also report estimates of the skewness, kurtosis and densities of log-earnings residuals growth at various horizons, from 2 to 10 years. The data suggests the presence of ARCH effects (as in Meghir and Pistaferri, 2004). It also shows that log-earnings growth is non-Gaussian, displaying negative skewness and high kurtosis. Guvenen *et al.* (2015) document similar features on US administrative data. This shows both the qualitative similarity between the PSID and the administrative US data in terms of higher moments of log-earnings growth, and the ability of our nonlinear model to fit these features. Note that, in our model, skewness and excess kurtosis of log-earnings growth at long horizons are mostly due to the non-Gaussianity of the *transitory* component ε .

Confidence bands. We compare two different methods to compute confidence intervals for these estimates. The first method is the nonparametric bootstrap, clustered at the household level. The second method is the parametric bootstrap. While the latter requires correct specification of the parametric model, the former may still be consistent under misspecification and allows for unrestricted serial correlation (however we are not aware of a formal justification for it in this setting). In Figures S9 to S14 of the Supplementary Appendix we

report pointwise 95% confidence bands based on both methods, and we also show uniform confidence bands based on the nonparametric bootstrap. In all cases, the main findings on nonlinear persistence and conditional skewness seem rather precisely estimated. At the same time the uniform confidence bands are wider, especially for the η component.

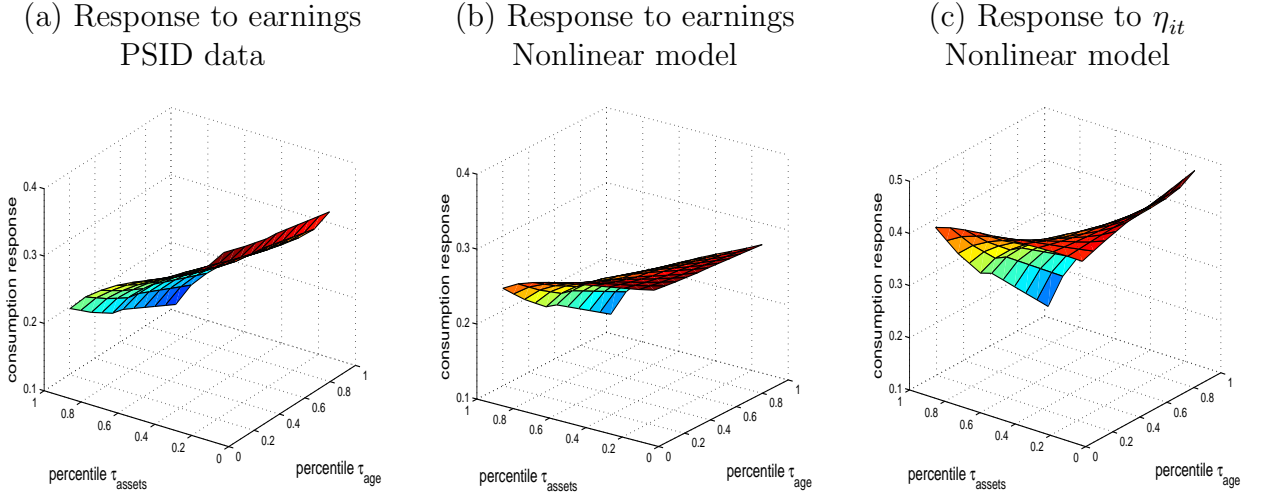
Household unobserved heterogeneity. In Figure S15 of the Supplementary Appendix we report the nonlinear persistence and conditional skewness of the η component in model (20), which allows in addition for an additive household-specific effect. Compared to Figure 2, allowing for a household effect reduces persistence. Moreover, the nonlinear pattern is more pronounced than in the homogeneous case. Persistence is close to 1 for values of τ_{init} and τ_{shock} that are close to each other, but it is substantially lower when a large positive (respectively negative) shock hits a low-earnings (resp. high-earnings) household.

Norwegian population register data. The above results suggest that nonlinear persistence and conditional skewness are features of earnings processes in the PSID. In order to corroborate these findings using a different, larger data set, we estimated the earnings process using a balanced subsample of 2,873 households from the 2000-2005 Norwegian administrative data. The estimates are shown in Figures S16 to S19 of the Supplementary Appendix. Like the PSID, the Norwegian population register data presents a similar pattern of conditional skewness. Moreover, the Norwegian data shows similar nonlinear persistence in the persistent component η as the PSID. At the same time, the dispersion of the transitory component ε is much smaller in the Norwegian data, suggesting either the presence of large measurement error in the PSID or smaller true transitory innovations in Norway. In order to shed more light on the differences, it would be very interesting to also estimate our nonlinear model on a large administrative data set for the US.

6.3 Consumption

We next turn to consumption. We estimated two main specifications: with and without household unobserved heterogeneity. Since due to the small sample size allowing for unobserved heterogeneity is somewhat challenging on these data, here we focus on the results without unobserved heterogeneity. The results with unobserved heterogeneity are reported in the Supplementary Appendix.

Figure 5: Consumption responses to earnings shocks, by assets and age, model without household-specific unobserved heterogeneity



Note: Graphs (a) and (b) show estimates of the average derivative of the conditional mean of c_{it} , with respect to y_{it} , given y_{it} , a_{it} and age_{it} , evaluated at values of a_{it} and age_{it} that corresponds to their τ_{assets} and τ_{age} percentiles, and averaged over the values of y_{it} . Graph (a) is based on the PSID data, and graph (b) is based on data simulated according to our nonlinear model with parameters set to their estimated values. Graph (c) shows estimates of the average consumption responses $\bar{\phi}_t(a)$ to variations in η_{it} , evaluated at τ_{assets} and τ_{age} .

Figure 5 (a) shows estimates of the average derivative, with respect to y_{it} , of the conditional mean of c_{it} given y_{it} , a_{it} and age_{it} . The function is evaluated at percentiles of log-assets and age (τ_{assets} and τ_{age} , respectively), and averaged over y_{it} .³⁴ The derivative effects lie between .2 and .3. Moreover, the results indicate that consumption of older households, and of households with higher assets, is less correlated to variations in earnings. Figure 5 (b) shows the same response surface based on simulated data from our full nonlinear model of earnings and consumption. The fit of the model, though not perfect, seems reasonable. In particular, the model reproduces the main pattern of correlation with age and assets.^{35,36}

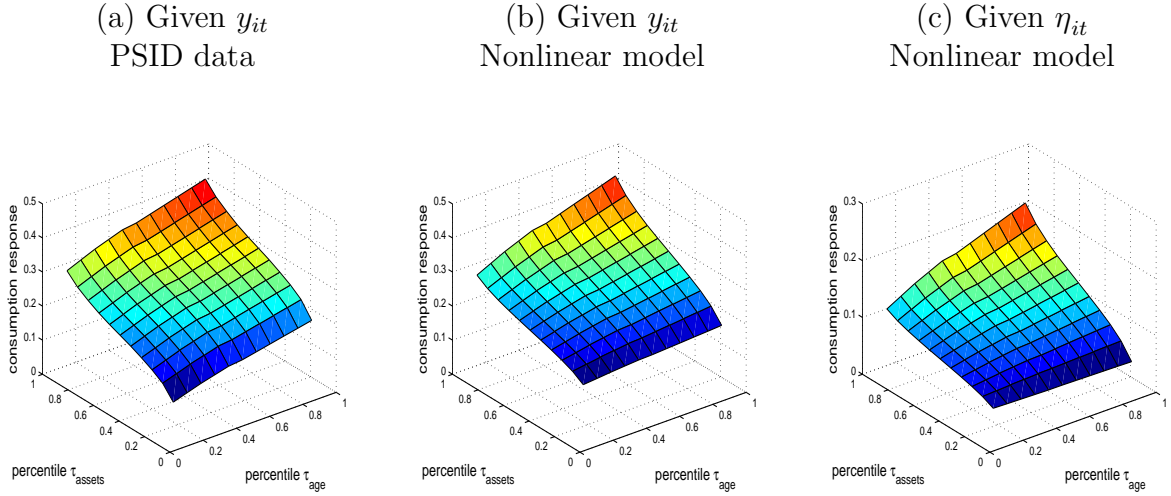
Figure 5 (c) shows estimates of the average consumption response $\bar{\phi}_t(a)$ to variations in the persistent component of earnings. As described in Section 3, $1 - \bar{\phi}_t(a)$ can be regarded

³⁴We use tensor products of Hermite polynomials with degrees (2, 2, 1) in the estimation of the consumption rule.

³⁵In Figure S20 of the Supplementary Appendix we show that the model fit to the density of log-consumption is also good.

³⁶While the covariances between log-earnings and log-consumption residuals are well reproduced, the baseline model does not perform as well in fitting the dynamics of consumption, as it systematically underestimates the autocorrelations between log-consumption residuals. The specification with household unobserved heterogeneity improves the fit to consumption dynamics.

Figure 6: Consumption responses to assets



Note: Estimates of the average derivative of the conditional mean of c_{it} , with respect to a_{it} , given y_{it} (respectively, given η_{it} and ε_{it} in graph (c)), a_{it} , and age_{it} , evaluated at values of a_{it} and age_{it} that corresponds to their τ_{assets} and τ_{age} percentiles, and averaged over the values of y_{it} (resp., over the values of η_{it} and ε_{it} in graph (c)). Model without unobserved heterogeneity.

as a measure of the degree of consumption insurability of shocks to the persistent earnings component, as a function of age and assets. On average the estimated $\bar{\phi}_t(a)$ parameter lies between .3 and .4, suggesting that more than half of pre-tax household earnings fluctuations is effectively insured. Moreover, variation in assets and age suggests the presence of an interaction effect. In particular, older households with high assets seem better insured against earnings fluctuations.³⁷

In Figures S22 and S23 of the Supplementary Appendix we report 95% confidence bands for $\bar{\phi}_t(a)$ based on both parametric bootstrap and nonparametric bootstrap. The findings on insurability of shocks to the persistent earnings component seem quite precisely estimated. In Figure S24 of the Supplementary Appendix we report estimates of the model with household unobserved heterogeneity in consumption, see (19). Estimated consumption responses are quite similar to the ones without unobserved heterogeneity, although the nonlinearity with respect to assets and age seems more pronounced.

Consumption responses to assets. In addition to consumption responses to earnings shocks, our nonlinear framework can be used to document derivative effects with respect to

³⁷Consumption responses to transitory shocks are shown in Figure S21 of the Supplementary Appendix.

assets. Such quantities are often of great interest, for example when studying the implications of tax reforms. Graph (c) in Figure 6 shows estimated average derivatives, in a model without unobserved heterogeneity in consumption.³⁸ The quantile polynomial specifications are the same as in Figure 5. We see that the responses range between .05 and .2, and that the derivative effects seem to increase with age and assets.

6.4 Simulating the impact of persistent earnings shocks

In this last subsection, we simulate life-cycle earnings and consumption according to our nonlinear model, and show the evolution of earnings and consumption following a persistent earnings shock. In Figure 7 we report the difference between the age-specific medians of log-earnings of two types of households: households that are hit, at the same age 37, by either a large negative shock to the persistent earnings component ($\tau_{shock} = .10$), or by a large positive shock ($\tau_{shock} = .90$), and households that are hit by a median shock $\tau = .50$ to the persistent component.³⁹ We report age-specific medians across 100,000 simulations of the model. At the start of the simulation (that is, age 35) all households have the same persistent component indicated by the percentile τ_{init} . With some abuse of terminology we refer to the resulting earnings and consumption paths as “impulse responses”.⁴⁰

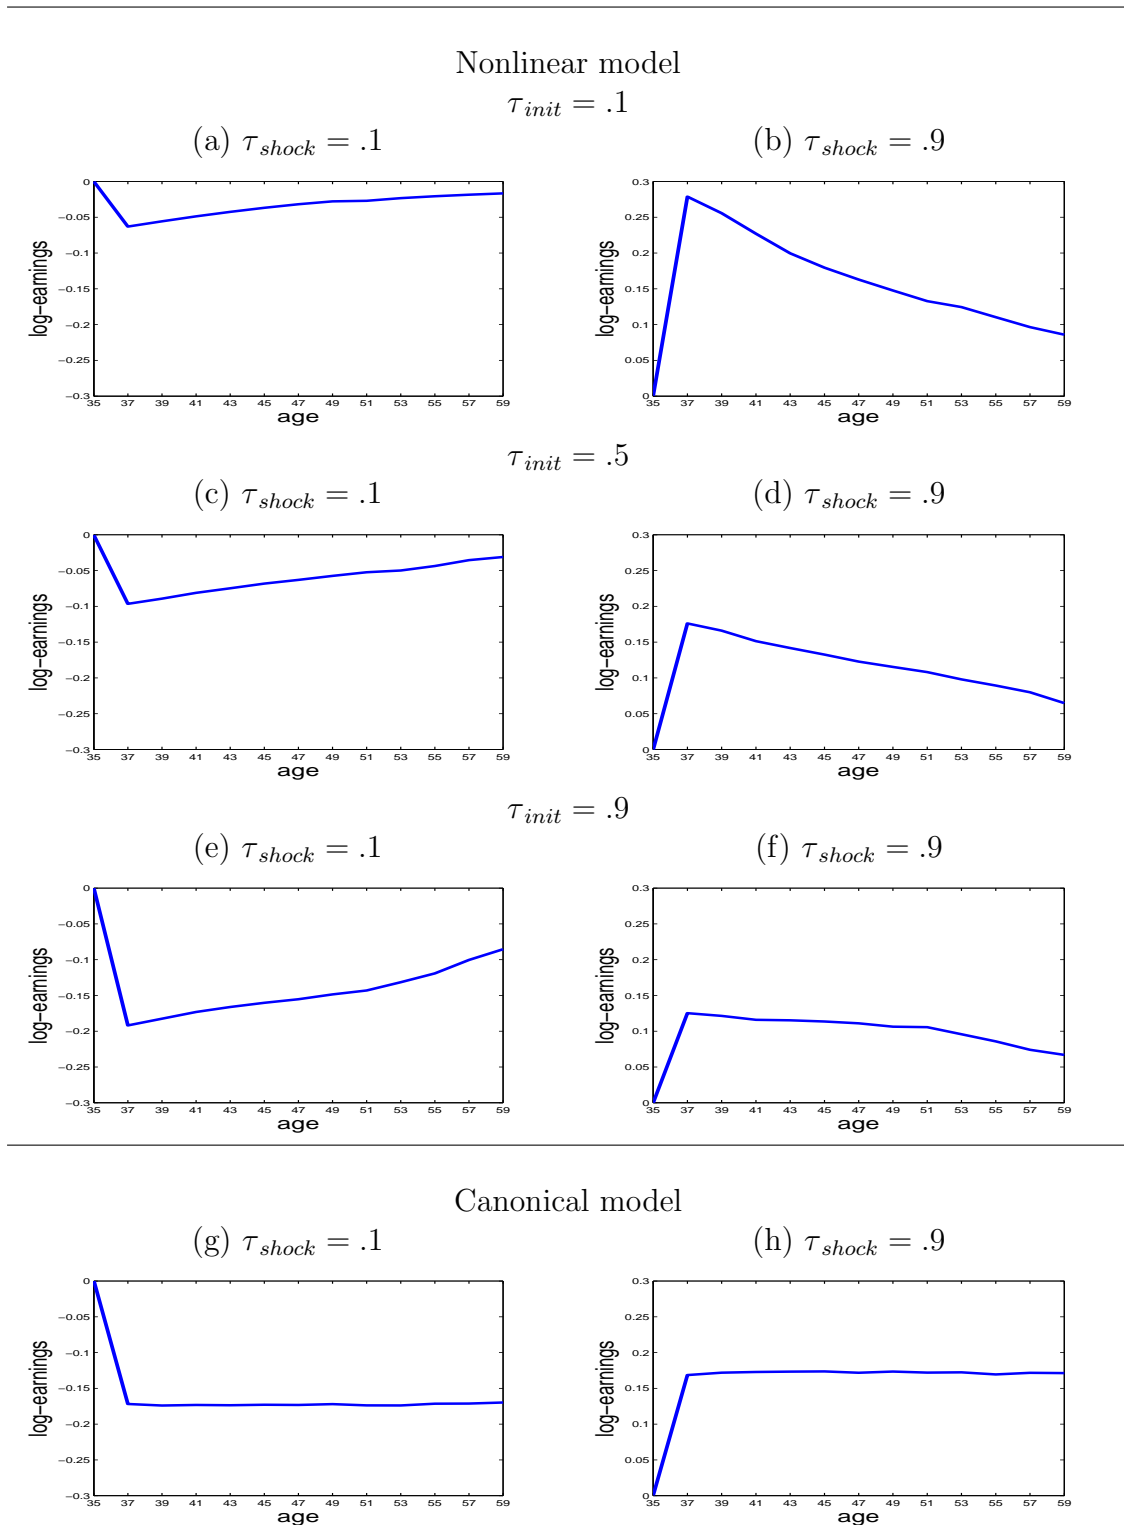
Earnings responses reported in Figure 7 are consistent with the presence of interaction effects between the rank in the distribution of earnings component (τ_{init}) and the sign and size of the shock to the persistent component (τ_{shock}). While a large negative shock ($\tau_{shock} = .10$) is associated with a 7% drop in earnings for low earnings households ($\tau_{init} = .10$), a similar shock is associated with a 19% drop for high-earnings households ($\tau_{init} = .90$). We also find interaction effects in the response to large positive shocks ($\tau_{shock} = .90$). This suggests the presence of asymmetries in the persistence of earnings histories, depending on the previous earnings history of the household and the size and magnitude of the shock. Moreover, the long-run impact of these shocks over the life cycle also depends on the initial condition. For example, Figure 7 (e) shows a very slow recovery from a negative earnings shock when starting from a high-earnings position, while graph (a) shows a quicker recovery.

³⁸Graphs (a) and (b) in Figure 6 show that the model replicates well the empirical relationship between consumption and assets conditional on earnings and age. See Figure S25 of the Supplementary Appendix for the results in a model with unobserved heterogeneity.

³⁹Note that such positive or negative shocks being “large” are relative statements, given that they correspond to ranks of different conditional distributions.


⁴⁰See for example Gallant *et al.* (1993) and Koop *et al.* (1996) for work on impulse response functions in nonlinear models.

Figure 7: Impulse responses, earnings




Note: Persistent component at percentile τ_{init} at age 35. The graphs show the difference between a household hit by a shock τ_{shock} at age 37, and a household hit by a .5 shock at the same age. Age-specific medians across 100,000 simulations. Graphs (a) to (f) correspond to the nonlinear model. Graphs (g) and (h) correspond to the canonical model (3) of earnings dynamics.

Figure 8: Impulse responses, consumption



Note: See notes to Figure 7. Graphs (a) to (f) correspond to the nonlinear model. Graphs (g) and (h) correspond to the canonical model of earnings dynamics (3) and a linear consumption rule. Linear assets accumulation rule (7), $r = 3\%$. $a_{it} \geq 0$. Model without household unobserved heterogeneity in consumption.

Figure 9: Impulse responses by age and initial assets



Note: See notes to Figure 8. Initial assets at age 35 (for “young” households) or 51 (for “old” households) are at percentile .10 (dashed curves) and .90 (solid curves). Linear assets accumulation rule (τ), $r = 3\%$. $a_{it} \geq 0$. Model without household unobserved heterogeneity in consumption.

These earnings impulse responses on impact are in line with the shape of the persistence function shown in Figure 2 (d). In Figure 7, moving from panel (a) to panel (c) shows little difference in the impact of a shock, while moving to panel (e) shows a larger impact. This is consistent with the persistence function $\rho_t(\eta_{i,t-1}, \tau)$ being steeper in τ for high income households who receive a large negative shock, compared to other households hit by such shock. Conversely, the earnings response is quite similar when moving between panels (f) and (d), while it is higher in panel (b), again in line with the shape of the persistence function.

In graphs (g) and (h) of Figure 7 we report results based on the “canonical model” of earnings dynamics where η is a random walk, see equation (3). In this model, there are by assumption no interaction effects between income shocks and the ranks of households in the income distribution. The implications of the nonlinear earnings model thus differ markedly from those of standard linear models.

In Figure 8 we report the results of a similar exercise to Figure 7, but we now focus on consumption responses in a model without unobserved heterogeneity in consumption. In order to simulate consumption paths, one needs to take a stand on the rule of assets accumulation. We use the linear assets accumulation rule (7), with a constant (biennial) interest rate $r = 3\%$. In the simulation we impose that $a_{it} \geq 0$. We observed little sensitivity to varying the floor on assets.

We see that the nonlinearities observed in the earnings response matter for consumption too. For example, while a large negative shock ($\tau_{shock} = .10$) is associated with a 2% drop in consumption for low earnings households, it is associated with an 8% drop for high-earnings households. Conversely, a large positive shock is associated with a 5% increase in consumption for high-earnings households, and with an 11% increase for low-earnings households. In Figures S26 to S29 of the Supplementary Appendix we report bootstrap confidence bands for earnings and consumption responses, which suggest that the results are relatively precisely estimated. In addition, graphs (g) and (h) of Figure 8 report results based on the canonical earnings model with a linear log-consumption rule.⁴¹ The fact that the canonical model assumes away the presence of interaction effects between income shocks and households' positions in the income distribution appears at odds with the data.

The results with unobserved heterogeneity in consumption are reported in Figure S30 of the Supplementary Appendix. They show smaller consumption responses to variations in earnings compared to the case without unobserved heterogeneity. For example, a large positive shock ($\tau_{shock} = .90$) is now associated with a 7% increase in consumption for low earnings households. In addition, effects on consumption seem to revert more quickly towards the median in the model with heterogeneity.⁴²

In Figure 9 we perform similar exercises, while varying the timing of shocks and the asset holdings that households possess. Graphs (a) to (d) suggest that a negative shock ($\tau_{shock} = .10$) for high-earnings households has a higher impact on earnings at later ages: the earnings drop is 40% when the shock hits at age 53, compared to 20% when a similar shock hits at age 37. The impact of a positive shock on low-earnings individuals seems to vary less with age. Graphs (e) to (h) in Figure 9 show the consumption responses in the

⁴¹Specifically, c_{it} is modeled as a linear function of η_{it} , ε_{it} , and an independent additive error term i.i.d. over time. The model is estimated by equally-weighted minimum distance based on covariance restrictions.

⁴²In Figures S31 and S32 of the Supplementary Appendix we show the results from the nonlinear assets rule we have estimated, see (11). The results do not differ markedly compared to the baseline specification.

model without heterogeneity. The results suggest that, while the presence of asset holdings does not seem to affect the insurability of positive earnings shocks, it does seem to attenuate the consumption response to negative shocks, particularly for households who are hit later in the life cycle. Figure S33 of the Supplementary Appendix shows similar patterns when allowing for unobserved heterogeneity.⁴³

7 Conclusion

In this paper we have developed a nonlinear framework for modeling persistence that sheds new light on the nonlinear transmission of income shocks and the nature of consumption insurance. In this framework, household income is the sum of a first-order Markov persistent component and a transitory component. The consumption policy rule is an age-dependent, nonlinear function of assets, unobserved heterogeneity, persistent income and transitory income. The model reveals asymmetric persistence patterns, where “unusual” earnings shocks are associated with a drop in persistence. It also leads to new empirical measures of partial insurance.

We provide conditions under which the model is nonparametrically identified, and we develop a tractable simulation-based sequential quantile regression method for estimation. These methods open the way to identify and estimate nonlinear models of earnings and consumption dynamics. They also provide new tools to assess the suitability of existing life-cycle models of consumption and savings, and potentially help guide the development of new structural models.

Our results suggest that nonlinear persistence and conditional skewness are important features of earnings processes. These features, which are present in both the PSID and in Norwegian population register data, are not easy to capture using existing models of earnings dynamics, motivating the use of new econometric methods to document distributional dynamics. Estimating models that allow for persistent and transitory components of income on a relatively homogeneous sample of households from the PSID, we find the presence of nonlinear persistence and conditional asymmetries in earnings.

The nonlinearities observed in the earnings responses are shown to impact consumption choices. For example, we found that while a large negative shock is associated with a

⁴³The results for the estimated nonlinear assets rule reported in Figure S34 of the Supplementary Appendix show some differences compared to Figure 9, particularly for the responses to positive earnings shocks.

relatively small drop in consumption for low earnings households, it is associated with a sizable drop for high-earnings households. We also identified differences in persistence across different demographic groups. The results suggest that, while the presence of asset holdings seems not to affect the insurability of positive earnings shocks, it appears to attenuate consumption responses to negative shocks, particularly for households who are hit later in the life cycle. Standard linear models, which assume away the presence of interaction effects between income shocks and the position in the income distribution, deliver qualitatively different predictions that appear at odds with the data.

A natural next step is to combine the framework introduced in this paper with more structural approaches. It is in fact easy to take our estimated Markovian earnings components to simulate or estimate fully-specified life-cycle models of consumption and savings. We provide an illustration of a simple life-cycle simulation model using the nonlinear dynamic quantile specification for earnings in Section S2 of the Supplementary Appendix. Moreover, the nonlinear model can be generalized to allow for other states and choices. For example, we could extend the analysis to incorporate intensive and extensive margins of labor supply, as in Low, Meghir and Pistaferri (2010), and family labor supply, as in Blundell *et al.* (2016).

Lastly, in this paper we have abstracted from the role of business cycle fluctuations. In a recent paper on US Social Security Data for 1978-2010, Guvenen, Ozcan and Song (2014) find that the left-skewness of earnings shocks is counter-cyclical. In future work it will be interesting to apply our framework to document distributional dynamics over the business cycle.

References

- [1] Abowd, J., and D. Card (1989): “On the Covariance Structure of Earnings and Hours Changes,” *Econometrica*, 57(2): 411–45.
- [2] Alan, S. (2012): “Do Disaster Expectations Explain Household Portfolios?,” *Quantitative Economics*, 3, 1–28.
- [3] Alan, S., M. Browning, and M. Ejrnaes (2014): “Income and Consumption: a Micro Semi-Structural Analysis with Pervasive Heterogeneity,” mimeo University of Essex, November.
- [4] Alvarez, J., M. Browning, and M. Ejrnaes (2010): “Modelling Income Processes with Lots of Heterogeneity,” *Review of Economic Studies*, 77(4), 1353–1381.
- [5] Andrews, D. (2011): “Examples of L^2 -Complete and Boundedly-Complete Distributions,” unpublished manuscript.
- [6] Arellano, M., and S. Bonhomme (2016): “Nonlinear Panel Data Estimation via Quantile Regressions,” *Econometrics Journal*, 19, C61-C94.
- [7] Barro, R. (2006): “Rare Disasters and Asset Markets in the Twentieth Century,” *Quarterly Journal of Economics*, 121 (3), 823–866.
- [8] Belloni, A., Chernozhukov, V., Chetverikov, D., and I. Fernández-Val (2016): “Conditional Quantile Processes based on Series or Many Regressors,” unpublished manuscript.
- [9] Blundell, R. (2014): “Income Dynamics and Life-Cycle Inequality: Mechanisms and Controversies,” *Economic Journal*, 124 (2), 289–318.
- [10] Blundell, R., X. Chen, and D. Kristensen (2007): “Semi-Nonparametric IV Estimation of Shape-Invariant Engel Curves” *Econometrica*, 75(6): 1613–1669.
- [11] Blundell, R., M. Graber and M. Mogstad (2015): “Labor Income Dynamics and the Insurance from Taxes, Transfers, and the Family,” *Journal of Public Economics*, 127, 58–73.
- [12] Blundell, R., L. Pistaferri, and I. Preston (2008): “Consumption Inequality and Partial Insurance,” *American Economic Review*, 98(5): 1887–1921.
- [13] Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016): “Consumption Smoothing and Family Labor Supply,” *American Economic Review*, 106(2), 387–435.
- [14] Bonhomme, S., and J. M. Robin (2010): “Generalized Nonparametric Deconvolution with an Application to Earnings Dynamics,” *Review of Economic Studies*, 77(2), 491–533.
- [15] Bound, J., C. Brown, and N. Mathiowetz (2001): “Measurement Error in Survey Data,” in *Handbook of Econometrics*, Vol. 5, ed. by J. Heckman and E. Leamer. Amsterdam: North Holland, Elsevier Science.
- [16] Botosaru, I., and Y. Sasaki (2015): “Nonparametric Heteroskedasticity in Persistent Panel Processes: An Application to Earnings Dynamics,” unpublished manuscript.

- [17] Castañeda, A., J. Díaz-Giménez, and J. V. Ríos-Rull (2003): “Accounting for the U.S. Earnings and Wealth Inequality,” *Journal of Political Economy*, 111(4), 818–857.
- [18] Celeux, G., and J. Diebolt (1993): “Asymptotic Properties of a Stochastic EM Algorithm for Estimating Mixing Proportions,” *Comm. Statist. Stochastic Models*, 9, 599–613.
- [19] Chernozhukov, V., I. Fernández-Val, and A. Galichon (2010): “Quantile and Probability Curves without Crossing ,” *Econometrica*, 78(3), 1093–1125.
- [20] Constandinides, G., and A. Gosh (2016): “Asset Pricing with Countercyclical Household Consumption Risk ,” to appear in the *Journal of Finance*.
- [21] Deaton, A., and C. Paxson (1994): “Intertemporal Choice and Inequality, ” *Journal of Political Economy*, 102(3), 437–467.
- [22] Dempster, A. P., N. M. Laird, and D. B. Rubin (1977): “Maximum Likelihood from Incomplete Data via the EM Algorithm,” *Journal of the Royal Statistical Society, B*, 39, 1–38.
- [23] D’Haultfoeulle, X. (2011): “On the Completeness Condition for Nonparametric Instrumental Problems,” *Econometric Theory*, 27, 460–471
- [24] Evans, M., and P. Wachtel (1993): “Inflation Regimes and the Sources of Inflation Uncertainty,” *Journal of Money, Credit, and Banking*, 25, 475–511.
- [25] Evdokimov, K., and H. White (2012): “Some Extensions of a Lemma of Kotlarski,” *Econometric Theory*, 4, 925–932.
- [26] Gallant, A. R., P. E. Rossi,, and G. Tauchen (1993): “Nonlinear Dynamic Structures,” *Econometrica*, 871–907.
- [27] Geweke, J., and M. Keane (2000): “An Empirical Analysis of Earnings Dynamics Among Men in the PSID: 1968-1989,” *Journal of Econometrics*, 96(2), 293–356.
- [28] Golosov, M., and A. Tsyvinski (2015): “Policy Implications of Dynamic Public Finance,” *Annual Reviews of Economics*, 7, pp.147-171
- [29] Golosov, M., M. Troshkin, and A. Tsyvinski (2016): “Redistribution and social insurance, ” *American Economic Review*, 106(2): 359-386.
- [30] Gourinchas, P.O. and J. A. Parker (2002): “Consumption over the Life Cycle.” *Econometrica*, 70, 47-91.
- [31] Guvenen, F., F. Karahan, S. Ozcan, and J. Song (2015): “What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk?” mimeo University of Minnesota, March.
- [32] Guvenen, F., S. Ozcan, and J. Song (2014): “The Nature of Countercyclical Income Risk,” *Journal of Political Economy*, 122(3), 621–660.
- [33] Guvenen, F., and A. Smith (2014): “Inferring Labor Income Risk from Economic Choices: An Indirect Inference Approach,” *Econometrica*, November, 82(6), 2085–2129.

- [34] Hall, R., and F. Mishkin (1982): “The sensitivity of Consumption to Transitory Income: Estimates from Panel Data of Households,” *Econometrica*, 50(2): 261–81.
- [35] Heathcote, J., K. Storesletten, and G. Violante (2014): “Consumption and Labor Supply with Partial Insurance: An Analytical Framework,” *American Economic Review*, July, 104(7), 2075–2126.
- [36] Horowitz, J. L., and M. Markatou (1996): “Semiparametric Estimation of Regression Models for Panel Data”, *Review of Economic Studies*, 63, 145–168.
- [37] Hu, Y. (2015): “Microeconomic Models with Latent Variables: Applications of Measurement Error Models in Empirical Industrial Organization and Labor Economics,” Technical report, Cemmap Working Papers, CWP03/15.
- [38] Hu, Y. and S. M. Schennach (2008): “Instrumental Variable Treatment of Nonclassical Measurement Error Models,” *Econometrica*, 76, 195–216.
- [39] Hu, Y. and J.-L. Shiu (2012): “Nonparametric Identification Using Instrumental Variables: Sufficient Conditions for Completeness,” unpublished manuscript.
- [40] Hu, Y. and M. Shum (2012): “Nonparametric Identification of Dynamic Models with Unobserved State Variables,” *Journal of Econometrics*, 171, 32–44.
- [41] Huckfeldt, C. (2016): “Understanding the Scarring Effect of Recessions”, mimeo, Economics Department, Cornell University, March.
- [42] Huggett, M. (1993), “The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies”. *Journal of Economic Dynamics and Control*, 17, 953–969.
- [43] Jappelli, T., and L. Pistaferri (2010): “The Consumption Response to Income Changes” *Annual Reviews in Economics*, 2, 479–506.
- [44] Kaplan, G., and G. Violante (2010): “How Much Consumption Insurance Beyond Self-Insurance” *American Economic Journal*, 2(4), 53–87.
- [45] Kaplan, G., and G. Violante (2014): “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, 82(4), 1199-1239.
- [46] Kim, T.H., and H. White (2004): “On More Robust Estimation of Skewness and Kurtosis,” *Finance Research Letters*, 1, 56–73.
- [47] Koenker, R. and G. J. Bassett (1978): “Regression quantiles,” *Econometrica*, 46, 33–50.
- [48] Koenker, R., and Z. Xiao (2006): “Quantile Autoregression,” *Journal of the American Statistical Association*, 101(475), 980–990.
- [49] Koop, G., M. H. Pesaran, and S. M. Potter (1996): “Impulse Response Analysis in Nonlinear Multivariate Models,” *Journal of Econometrics*, 74(1), 119–147.
- [50] Krueger, D., K. Mitman, and F. Perri (2015): “Macroeconomics and Heterogeneity, Including Inequality,” unpublished manuscript.
- [51] Krusell, P., and A. Smith (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106, 867–896.

- [52] Lise, J. (2013): “On-the-Job Search and Precautionary Savings,” *Review of Economic Studies*, 80(3), 1086–1113.
- [53] Lochner, L., and Y. Shin (2014): “Understanding Earnings Dynamics: Identifying and Estimating the Changing Roles of Unobserved Ability, Permanent and Transitory Shocks,” unpublished manuscript.
- [54] Low, H. L., C. Meghir, and L. Pistaferri (2010): “Wage Risk and Employment Risk over the Life-Cycle,” *American Economic Review*, 100, 1432–1467.
- [55] Meghir, C., and L. Pistaferri (2004): “Income Variance Dynamics and Heterogeneity,” *Econometrica*, 72(1): 1–32.
- [56] Meghir, C., and L. Pistaferri (2011): “Earnings, Consumption and Life Cycle Choices,” *Handbook of Labor Economics*, Elsevier.
- [57] Misra, K., and P. Surico (2014): “Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs,” *A EJ: Macroeconomics*, 6(4): 84–106.
- [58] Newey, W. , and J. Powell (2003): “Instrumental Variable Estimation of Nonparametric Models,” *Econometrica*.
- [59] Nielsen, S. F. (2000): “The Stochastic EM Algorithm: Estimation and Asymptotic Results,” *Bernoulli*, 6(3): 457–489.
- [60] Rietz, Th. (1988): “The Equity Risk Premium: A Solution,” *Journal of Monetary Economics*, 22 (1), 117–131.
- [61] Schmidt, L. (2015): “Climbing and Falling Off the Ladder: Asset Pricing Implications of Labor Market Event Risk,” unpublished manuscript.
- [62] Teräsvirta, T. (1994): “Specification, Estimation, and Evaluation of Smooth Transition Autoregressive Models,” *Journal of the American Statistical Association*, 89, 208–218.
- [63] Wei, Y. and R. J. Carroll (2009): “Quantile Regression with Measurement Error,” *Journal of the American Statistical Association*, 104, 1129–1143.
- [64] Wilhelm, D. (2015): “Identification and Estimation of Nonparametric Panel Data Regressions with Measurement Error,” Cemmap working paper CWP34/15.

APPENDIX

A Identification

A.1 Earnings process

In the following, all conditional and marginal densities are assumed to be bounded away from zero and infinity on their supports. With some abuse of notation, in the absence of ambiguity we use $f(a|b)$ as a generic notation for the conditional density $f_{A|B}(a|b)$, and for simplicity we omit the i index in density arguments.

Operator injectivity. The identification arguments below rely on the concept of operator injectivity, which we now formally define. A linear operator \mathcal{L} is a linear mapping from a functional space \mathcal{H}_1 to another functional space \mathcal{H}_2 . \mathcal{L} is *injective* if the only solution $h \in \mathcal{H}_1$ to the equation $\mathcal{L}h = 0$ is $h = 0$.

One special case of operator injectivity (“deconvolution”) obtains when $Y_{i2} = Y_{i1} + \epsilon_{i1}$, with Y_{i1} independent of ϵ_{i1} , and $[\mathcal{L}h](y_2) = \int h(y_1)f_{\epsilon_1}(y_2 - y_1)dy_1$. \mathcal{L} is then injective if the characteristic function of ϵ_{i1} has no zeros on the real line. The normal and many other standard distributions satisfy this property.⁴⁴ If the marginal distributions f_{Y_2} and f_{ϵ_1} are known, injectivity implies that $h = f_{Y_1}$ is the only solution to the functional equation $\int h(y_1)f_{\epsilon_1}(y_2 - y_1)dy_1 = f_{Y_2}(y_2)$. In other words, f_{Y_1} is identified from the knowledge of f_{Y_2} and f_{ϵ_1} .

Another, important special case of operator injectivity (“completeness”) is obtained when \mathcal{L} is the conditional expectation operator associated with the distribution of $(Y_{i1}|Y_{i2})$, in which case $[\mathcal{L}h](y_2) = \mathbb{E}[h(Y_{i1})|Y_{i2} = y_2]$. \mathcal{L} being injective is then equivalent to the distribution of $(Y_{i1}|Y_{i2})$ being *complete*.

Building block for identification. To establish nonparametric identification of the earnings process, we rely on results from Hu and Schennach (2008) and Wilhelm (2015). In the context of a panel data model with measurement error, Wilhelm (2015) provides conditions under which the marginal distribution of ϵ_{i2} is identified, given three periods of observations (y_{i1}, y_{i2}, y_{i3}) . We provide a brief summary of the identification argument used by Wilhelm in Section S4 of the Supplementary Appendix.

The key condition that underlies identification in this context is the fact that, in the earnings model with $T = 3$, log-earnings (y_{i1}, y_{i2}, y_{i3}) are conditionally independent given η_{i2} .⁴⁵ This “Hidden Markov” structure fits into the general setup considered in Hu and Schennach (2008). Hu (2015) provides a recent survey of applications of this line of work.

Identification of the earnings process. Returning to the earnings dynamics model (1)-(2), let now $T \geq 3$. Suppose that the conditions in Wilhelm (2015) are satisfied on each of the three-year subpanels $t \in \{1, 2, 3\}$ to $t \in \{T - 2, T - 1, T\}$. It follows from Wilhelm’s result that the marginal distributions of ϵ_{it} are identified for all $t \in \{2, 3, \dots, T - 1\}$. By serial independence of the ϵ ’s, the joint distribution of $(\epsilon_{i2}, \epsilon_{i3}, \dots, \epsilon_{i,T-1})$ is thus also identified.

Hence, if the characteristic functions of ϵ_{it} do not vanish on the real line, then by a deconvolution argument the joint distribution of $(\eta_{i2}, \eta_{i3}, \dots, \eta_{i,T-1})$ is identified. As a result, all Markov transitions $f_{\eta_t|\eta_{t-1}}$ are identified for $t = 3, \dots, T - 1$, and the marginal distribution of η_{i2} is identified

⁴⁴Injectivity also holds if the zeros of the characteristic function of ϵ_{i1} are isolated. See Evdokimov and White (2012).

⁴⁵Indeed, $f(y_1, y_2, y_3|\eta_2) = f(y_1|\eta_2)f(y_2|\eta_2, y_1)f(y_3|\eta_2, y_2, y_1) = f(y_1|\eta_2)f(y_2|\eta_2)f(y_3|\eta_2)$.

as well (so we need $T \geq 4$ to identify at least one Markov transition). Moreover, it is easy to show that the conditional distributions of $\eta_{i2}|y_{i1}$ and $y_{iT}|\eta_{i,T-1}$ are identified.⁴⁶

Note that, in the case where $\varepsilon_{i1}, \dots, \varepsilon_{iT}$ have the same marginal distribution, then the distributions of the initial and terminal components $\varepsilon_{i1}, \eta_{i1}$, and $\varepsilon_{iT}, \eta_{iT}$ are also identified. However, the first and last-period distributions are generally not identified in a fully non-stationary setting. In the empirical analysis we impose time-stationarity restrictions, and pool different cohorts of households together in order to identify the distributions of η 's and ε 's at all age.⁴⁷

A.2 Consumption rule with unobserved heterogeneity

We make the following assumption.

Assumption A1

i) $u_{i,t+s}$ and $\varepsilon_{i,t+s}$, for all $s \geq 0$, are independent of $a_i^t, \eta_i^{t-1}, y_i^{t-1}$, and ξ_i . ε_{i1} is independent of a_{i1}, η_{i1} and ξ_i .

ii) $a_{i,t+1}$ is independent of $(a_i^{t-1}, c_i^{t-1}, y_i^{t-1}, \eta_i^{t-1})$ conditional on $(a_{it}, c_{it}, y_{it}, \eta_{it}, \xi_i)$.

iii) the taste shifter ν_{it} in (19) is independent of $\eta_{i1}, (u_{is}, \varepsilon_{is})$ for all s , ν_{is} for all $s \neq t$, a_i^t , and ξ_i .

The identification strategy proceeds in two steps. First we have, by Assumption A1i) and iii), for all $t \geq 1$,

$$f(c^t, a^t|y) = \int f(c^t, a^t|\eta^t, y^t)f(\eta^t|y)d\eta^t,$$

or, equivalently,

$$f(c^t, a^t|y) = \mathbb{E} [f(c^t, a^t|\eta_i^t, y_i^t) | y_i = y],$$

where the expectation is taken for fixed (c^t, a^t) . Let $t = 3$. $f(c^3, a^3|\eta^3, y^3)$ is thus identified, provided the distribution of $(\eta_i^3|y_i)$ is boundedly complete in (y_{i4}, \dots, y_{iT}) . In particular, this argument requires that $T \geq 6$.

For the second step, we note that, by Assumption A1,

$$f(c^3, a^3|\eta^3, y^3) = \int f(a_1, c_1, a_2|\eta_1, y_1, \xi)f(c_2, a_3|a_2, \eta_2, y_2, \xi)f(c_3|a_3, \eta_3, y_3, \xi)f(\xi|\eta^3, y^3)d\xi. \quad (\text{A1})$$

For fixed (a^3, η^3, y^3) , equation (A1) is formally analogous to the nonlinear instrumental variables set-up of Hu and Schennach (2008). Hence the consumption rules, the asset evolution distributions, and the distribution of the latent heterogeneous component ξ_i , will all be nonparametrically identified under the conditions of Hu and Schennach's main theorem. These conditions include injectivity/completeness conditions analogous to the ones we have used in the baseline model, as well as a "scaling" condition. For example, in a consumption model that is additive in ν_{it} (as in our empirical application), a possible scaling condition (and the one we use) is that the mean of

⁴⁶Indeed we have $f_{y_2|y_1}(y_2|y_1) = \int f_{\varepsilon_2}(y_2 - \eta_2)f_{\eta_2|y_1}(\eta_2|y_1)d\eta_2$. Hence, since the characteristic function of ε_{i2} is non-vanishing, $f_{\eta_2|y_1}(\cdot|y_1)$ is identified for given y_1 . A similar argument shows that $f_{y_T|\eta_{T-1}}(y_T|\cdot)$ is identified for given y_T .

⁴⁷Specifically, the above arguments allow to nonparametrically recover, for each cohort entering the sample at age j , the distributions of ε at ages $j + 2, j + 4, j + 6$, and $j + 8$ (based on biennial data). In our dataset, j belongs to $\{25, \dots, 50\}$. Pooling across cohorts, we obtain that the distributions of ε are nonparametrically identified at all ages between 27 and 58 years. In turn, the joint distribution of η 's is nonparametrically identified in this age range. Identification at ages 25, 26 and 59, 60 intuitively comes from parametric extrapolation using the quantile models.

c_{i3} , conditional on ξ_i and some values of $(a_{i3}, \eta_{i3}, y_{i3})$, is increasing in ξ_i . In that case identification is to be understood up to an increasing transformation of ξ_i .⁴⁸ Consumption rules and asset distributions for $t \geq 4$ can then be identified by relying on additional periods or, alternatively, under time-stationarity assumptions by pooling information from different cohorts.

B Estimation

B.1 Model's restrictions

Let $\rho_\tau(u) = u(\tau - \mathbf{1}\{u \leq 0\})$ denote the ‘‘check’’ function of quantile regression (Koenker and Bassett, 1978). Let also $\bar{\theta}$ denote the true value of θ , and let

$$f_i(\eta_i^T; \bar{\theta}) = f(\eta_i^T | y_i^T, age_i^T; \bar{\theta})$$

denote the posterior density of $\eta_i^T = (\eta_{i1}, \dots, \eta_{iT})$ given the earnings data. As the earnings model is fully specified, f_i is a known function of $\bar{\theta}$.

We start by noting that, for all $\ell \in \{1, \dots, L\}$,

$$(\bar{a}_{0\ell}^Q, \dots, \bar{a}_{K\ell}^Q) = \underset{(a_{0\ell}^Q, \dots, a_{K\ell}^Q)}{\operatorname{argmin}} \sum_{t=2}^T \mathbb{E} \left[\int \rho_{\tau_\ell} \left(\eta_{it} - \sum_{k=0}^K a_{k\ell}^Q \varphi_k(\eta_{i,t-1}, age_{it}) \right) f_i(\eta_i^T; \bar{\theta}) d\eta_i^T \right], \quad (\text{B2})$$

where $\bar{a}_{k\ell}^Q$ denotes the true value of $a_{k\ell}^Q = a_k^Q(\tau_\ell)$, and the expectation is taken with respect to the distribution of (y_i^T, age_i^T) . To see that (B2) holds, note that the objective function is smooth (due to the presence of the integrals) and convex (because of the ‘‘check’’ function). The first-order conditions of (B2) are satisfied at true parameter values as, by (21), for all $k \in \{0, \dots, K\}$, $\ell \in \{1, \dots, L\}$, and $t \geq 2$,

$$\mathbb{E} \left[\mathbf{1} \left\{ \eta_{it} \leq \sum_{k=0}^K \bar{a}_{k\ell}^Q \varphi_k(\eta_{i,t-1}, age_{it}) \right\} \middle| \eta_i^{t-1}, age_i^T \right] = \tau_\ell.$$

Likewise, we have, for all ℓ ,

$$(\bar{a}_{0\ell}^\varepsilon, \dots, \bar{a}_{K\ell}^\varepsilon) = \underset{(a_{0\ell}^\varepsilon, \dots, a_{K\ell}^\varepsilon)}{\operatorname{argmin}} \sum_{t=1}^T \mathbb{E} \left[\int \rho_{\tau_\ell} \left(y_{it} - \eta_{it} - \sum_{k=0}^K a_{k\ell}^\varepsilon \varphi_k(age_{it}) \right) f_i(\eta_i^T; \bar{\theta}) d\eta_i^T \right], \quad (\text{B3})$$

and, for all ℓ ,

$$(\bar{a}_{0\ell}^{\eta_1}, \dots, \bar{a}_{K\ell}^{\eta_1}) = \underset{(a_{0\ell}^{\eta_1}, \dots, a_{K\ell}^{\eta_1})}{\operatorname{argmin}} \mathbb{E} \left[\int \rho_{\tau_\ell} \left(\eta_{i1} - \sum_{k=0}^K a_{k\ell}^{\eta_1} \varphi_k(age_{i1}) \right) f_i(\eta_i^T; \bar{\theta}) d\eta_i^T \right]. \quad (\text{B4})$$

In addition to (B2)-(B3)-(B4), the model implies other restrictions on the tail parameters λ , which are given in the next subsection. All the restrictions depend on the posterior density f_i . Given the use of piecewise-linear interpolating splines, the joint likelihood function of $(\eta_i^T, y_i^T | age_i^T; \bar{\theta})$ is available in closed form, and we provide an explicit expression in the next subsection. In practice, this means that it is easy to simulate from f_i . We take advantage of this feature in our estimation algorithm.

⁴⁸Arellano and Bonhomme (2016) apply Hu and Schennach (2008)’s results to a class of nonlinear panel data models.

Turning to consumption we have

$$\begin{aligned} (\bar{\alpha}, \bar{b}_1^g, \dots, \bar{b}_K^g) = \underset{(\alpha, b_1^g, \dots, b_K^g)}{\operatorname{argmin}} \sum_{t=1}^T \mathbb{E} \left[\int \left(c_{it} - \alpha - \sum_{k=1}^K b_k^g \tilde{\varphi}_k(a_{it}, \eta_{it}, y_{it} - \eta_{it}, age_{it}) \right) \right. \\ \left. \dots \times g_i(\eta_i^T; \bar{\theta}, \bar{\mu}) d\eta_i^T \right], \end{aligned}$$

where

$$g_i(\eta_i^T; \bar{\theta}, \bar{\mu}) = f(\eta_i^T | c_i^T, a_i^T, y_i^T, age_i^T; \bar{\theta}, \bar{\mu})$$

denotes the posterior density of $(\eta_{i1}, \dots, \eta_{iT})$ given the earnings, consumption, and asset data.

Moreover, the variance of taste shifters satisfies

$$\bar{\sigma}^2 = \frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[\int \left(c_{it} - \bar{\alpha} - \sum_{k=1}^K \bar{b}_k^g \tilde{\varphi}_k(a_{it}, \eta_{it}, y_{it} - \eta_{it}, age_{it}) \right)^2 g_i(\eta_i^T; \bar{\theta}, \bar{\mu}) d\eta_i^T \right]. \quad (\text{B5})$$

Likewise, for assets we have

$$\begin{aligned} (\bar{\alpha}^h, \bar{b}_1^h, \dots, \bar{b}_K^h) = \underset{(\alpha^h, b_1^h, \dots, b_K^h)}{\operatorname{argmin}} \sum_{t=2}^T \mathbb{E} \left[\int \left(a_{it} - \alpha^h - \sum_{k=1}^K b_k^h \tilde{\varphi}_k(a_{i,t-1}, c_{i,t-1}, y_{i,t-1}, \eta_{i,t-1}, age_{it}) \right) \right. \\ \left. \dots \times g_i(\eta_i^T; \bar{\theta}, \bar{\mu}) d\eta_i^T \right], \end{aligned}$$

with a similar expression for the variance of $b_0^h(v_{it})$ as in (B5).

Lastly we have, for all ℓ ,

$$(\bar{b}_{0\ell}^a, \dots, \bar{b}_{K\ell}^a) = \underset{(b_{0\ell}^a, \dots, b_{K\ell}^a)}{\operatorname{argmin}} \mathbb{E} \left[\int \rho_{\tau_\ell} \left(a_{i1} - \sum_{k=0}^K b_{k\ell}^a \tilde{\varphi}_k(\eta_{i1}, age_{i1}) \right) g_i(\eta_i^T; \bar{\theta}, \bar{\mu}) d\eta_i^T \right],$$

with additional restrictions characterizing tail parameters given in the next subsection.

B.2 Estimation algorithm

Additional model restrictions. The tail parameters λ satisfy simple moment restrictions. For example, we have

$$\bar{\lambda}_-^Q = - \frac{\sum_{t=2}^T \mathbb{E} \left[\int \mathbf{1} \left\{ \eta_{it} \leq \sum_{k=0}^K \bar{a}_{k1}^Q \varphi_k(\eta_{i,t-1}, age_{it}) \right\} f_i(\eta_i^T; \bar{\theta}) d\eta_i^T \right]}{\sum_{t=2}^T \mathbb{E} \left[\int \left(\eta_{it} - \sum_{k=0}^K \bar{a}_{k1}^Q \varphi_k(\eta_{i,t-1}, age_{it}) \right) \mathbf{1} \left\{ \eta_{it} \leq \sum_{k=0}^K \bar{a}_{k1}^Q \varphi_k(\eta_{i,t-1}, age_{it}) \right\} f_i(\eta_i^T; \bar{\theta}) d\eta_i^T \right]}, \quad (\text{B6})$$

and

$$\bar{\lambda}_+^Q = \frac{\sum_{t=2}^T \mathbb{E} \left[\int \mathbf{1} \left\{ \eta_{it} \geq \sum_{k=0}^K \bar{a}_{kL}^Q \varphi_k(\eta_{i,t-1}, age_{it}) \right\} f_i(\eta_i^T; \bar{\theta}) d\eta_i^T \right]}{\sum_{t=2}^T \mathbb{E} \left[\int \left(\eta_{it} - \sum_{k=0}^K \bar{a}_{kL}^Q \varphi_k(\eta_{i,t-1}, age_{it}) \right) \mathbf{1} \left\{ \eta_{it} \geq \sum_{k=0}^K \bar{a}_{kL}^Q \varphi_k(\eta_{i,t-1}, age_{it}) \right\} f_i(\eta_i^T; \bar{\theta}) d\eta_i^T \right]}, \quad (\text{B7})$$

with similar equations for the other tail parameters.

Likelihood function. The likelihood function is (omitting the conditioning on age for conciseness)

$$f(y_i^T, c_i^T, a_i^T, \eta_i^T; \theta, \mu) = \prod_{t=1}^T f(y_{it}|\eta_{it}; \theta) \prod_{t=1}^T f(c_{it}|a_{it}, \eta_{it}, y_{it}; \mu) \prod_{t=2}^T f(a_{it}|a_{i,t-1}, y_{i,t-1}, c_{i,t-1}, \eta_{i,t-1}; \mu) \\ \times \prod_{t=2}^T f(\eta_{it}|\eta_{i,t-1}; \theta) f(a_{i1}|\eta_{i1}; \mu) f(\eta_{i1}; \theta). \quad (\text{B8})$$

The likelihood function is fully specified and available in closed form. For example, we have

$$f(y_{it}|\eta_{it}; \theta) = \mathbf{1}\{y_{it} - \eta_{it} < A_{it}^\varepsilon(1)\} \tau_1 \lambda_-^\varepsilon \exp[\lambda_-^\varepsilon (y_{it} - \eta_{it} - A_{it}^\varepsilon(1))] \\ + \sum_{\ell=1}^{L-1} \mathbf{1}\{A_{it}^\varepsilon(\ell) \leq y_{it} - \eta_{it} < A_{it}^\varepsilon(\ell+1)\} \frac{\tau_{\ell+1} - \tau_\ell}{A_{it}^\varepsilon(\ell+1) - A_{it}^\varepsilon(\ell)} \\ + \mathbf{1}\{A_{it}^\varepsilon(L) \leq y_{it} - \eta_{it}\} (1 - \tau_L) \lambda_+^\varepsilon \exp[-\lambda_+^\varepsilon (y_{it} - \eta_{it} - A_{it}^\varepsilon(L))],$$

where $A_{it}^\varepsilon(\ell) \equiv \sum_{k=0}^K a_{k\ell}^\varepsilon \varphi_k(\text{age}_{it})$ for all (i, t, ℓ) . Note that the likelihood function is non-negative by construction. In particular, drawing from the posterior density of η automatically produces rearrangement of the various quantile curves (Chernozhukov, Galichon and Fernandez-Val, 2010).

Estimation algorithm: earnings. Start with $\hat{\theta}^{(0)}$. Iterate on $s = 0, 1, 2, \dots$ the two following steps.

Stochastic E-step: Draw M values $\eta_i^{(m)} = (\eta_{i1}^{(m)}, \dots, \eta_{iT}^{(m)})$ from

$$f(\eta_i^T | y_i^T; \hat{\theta}^{(s)}) \propto \prod_{t=1}^T f(y_{it}|\eta_{it}; \hat{\theta}^{(s)}) f(\eta_{i1}; \hat{\theta}^{(s)}) \prod_{t=2}^T f(\eta_{it}|\eta_{i,t-1}; \hat{\theta}^{(s)}),$$

where $a \propto b$ means that a and b are equal up to a proportionality factor independent of η .

M-step: Compute,⁴⁹ for $\ell = 1, \dots, L$,

$$\left(\hat{a}_{0\ell}^{Q,(s+1)}, \dots, \hat{a}_{K\ell}^{Q,(s+1)} \right) = \underset{(a_{0\ell}^Q, \dots, a_{K\ell}^Q)}{\operatorname{argmin}} \sum_{i=1}^N \sum_{t=2}^T \sum_{m=1}^M \rho_{\tau_\ell} \left(\eta_{it}^{(m)} - \sum_{k=0}^K a_{k\ell}^Q \varphi_k(\eta_{i,t-1}^{(m)}, \text{age}_{it}) \right), \\ \left(\hat{a}_{0\ell}^{\varepsilon,(s+1)}, \dots, \hat{a}_{K\ell}^{\varepsilon,(s+1)} \right) = \underset{(a_{0\ell}^\varepsilon, \dots, a_{K\ell}^\varepsilon)}{\operatorname{argmin}} \sum_{i=1}^N \sum_{t=1}^T \sum_{m=1}^M \rho_{\tau_\ell} \left(y_{it} - \eta_{it}^{(m)} - \sum_{k=0}^K a_{k\ell}^\varepsilon \varphi_k(\text{age}_{it}) \right), \\ \left(\hat{a}_{0\ell}^{\eta_1,(s+1)}, \dots, \hat{a}_{K\ell}^{\eta_1,(s+1)} \right) = \underset{(a_{0\ell}^{\eta_1}, \dots, a_{K\ell}^{\eta_1})}{\operatorname{argmin}} \sum_{i=1}^N \sum_{m=1}^M \rho_{\tau_\ell} \left(\eta_{i1}^{(m)} - \sum_{k=0}^K a_{k\ell}^{\eta_1} \varphi_k(\text{age}_{i1}) \right),$$

and compute

$$\hat{\lambda}_-^{Q,(s+1)} = - \frac{\sum_{i=1}^N \sum_{t=2}^T \sum_{m=1}^M \mathbf{1}\{\eta_{it}^{(m)} \leq \hat{A}_{itm}^{Q,(s+1)}\}}{\sum_{i=1}^N \sum_{t=2}^T \sum_{m=1}^M \left(\eta_{it}^{(m)} - \hat{A}_{itm}^{Q,(s+1)} \right) \mathbf{1}\{\eta_{it}^{(m)} \leq \hat{A}_{itm}^{Q,(s+1)}\}},$$

⁴⁹In practice, we used gradient descent algorithms in all M-step computations. This provided the fastest alternative in our repeated estimation, without any noticeable loss in accuracy.

where

$$\widehat{A}_{itm}^{Q,(s+1)} \equiv \sum_{k=0}^K \widehat{a}_{k1}^{Q,(s+1)} \varphi_k(\eta_{i,t-1}^{(m)}, age_{it}),$$

with similar updating rules for $\widehat{\lambda}_+^{Q,(s+1)}$, $\widehat{\lambda}_-^{\varepsilon,(s+1)}$, $\widehat{\lambda}_+^{\varepsilon,(s+1)}$, $\widehat{\lambda}_-^{\eta_1,(s+1)}$, and $\widehat{\lambda}_+^{\eta_1,(s+1)}$.

In practice, we start the algorithm with different choices for $\widehat{\theta}^{(0)}$. For example, for the initial values of the quantile parameters in η_{it} we run quantile regressions of log-earnings on lagged log-earnings and age. We proceed similarly to set other starting parameter values. We experimented with a number of other choices, and selected the parameter values corresponding to the highest average log-likelihood over iterations. We observed some effect of starting values on estimates of tail parameters although most results were stable.

Estimation algorithm: consumption. Similar to the earnings case. One difference is that in the stochastic E-step we draw $\eta_i^{(m)}$ from

$$\begin{aligned} f(\eta_i^T | y_i^T, c_i^T, a_i^T; \widehat{\theta}, \widehat{\mu}^{(s)}) &\propto \prod_{t=1}^T f(y_{it} | \eta_{it}; \widehat{\theta}) f(\eta_{i1}; \widehat{\theta}) \prod_{t=2}^T f(\eta_{it} | \eta_{i,t-1}; \widehat{\theta}) \\ &\quad \times f(a_{i1} | \eta_{i1}; \widehat{\mu}^{(s)}) \prod_{t=2}^T f(a_{it} | a_{i,t-1}, c_{i,t-1}, y_{i,t-1}, \eta_{i,t-1}; \widehat{\mu}^{(s)}) \\ &\quad \times \prod_{t=1}^T f(c_{it} | a_{it}, \eta_{it}, y_{it}; \widehat{\mu}^{(s)}). \end{aligned}$$

C Data appendix

C.1 PSID data

We use the 1999-2009 Panel Study of Income Dynamics (PSID) to estimate the model. The PSID started in 1968 collecting information on a sample of roughly 5,000 households. Of these, about 3,000 were representative of the US population as a whole (the core sample), and about 2,000 were low-income families (the Census Bureau's SEO sample). Thereafter, both the original families and their split-offs (children of the original family forming a family of their own) have been followed. The PSID data were collected annually until 1996 and biennially starting in 1997. A great advantage of PSID after 1999 is that, in addition to income data and demographics, it collects data about detailed asset holdings and consumption expenditures in each wave. To the best of our knowledge this makes the PSID the only representative large scale US panel to include income, hours, consumption, and assets data. Since we need both consumption and assets data, we focus on the 1999-2009 sample period.

We focus on non-SEO households with participating and married male household heads aged between 25 and 60, and with non missing information on key demographics (age, education, and state of residence). To reduce the influence of measurement error, we also drop observations with extremely high asset values (20 millions or more), as well as observations with total transfers more than twice the size of total household earnings. When calculating the relevant consumption, hourly wage and earnings moments, we do not use data displaying extreme "jumps" from one year to the next (most likely due to measurement error). Furthermore, we do not use earnings and wage data when the implied hourly wage is below one-half the state minimum wage. See Blundell, Pistaferri and Saporta-Eksten (2016) for further details of the sample selection.

C.2 Norwegian population register data

Sample selection. The Norwegian results are provided as part of the Blundell, Graber and Mogstad project on ‘Labour Income Dynamics and the Insurance from Taxes, Transfers and the Family’, see Blundell, Graber and Mogstad (2015) for details. For Figure 1 we use Norwegian register data provided under that project for the years 2005 and 2006 only.

We select a balanced panel of households where the male head is Norwegian, resident in Norway. We restrict the sample to include male, non-immigrant residents between the age 30 and 60 and their spouse (if they have one), with non-missing information on all key demographic variables. We choose a balanced panel of continuously married males, where household disposable income (that is, pooled labor income after tax and transfers of the spouses) is above the threshold of substantial gainful activity (one basic amount, 14,000 USD in 2014) and where the total income from self-employment is below one basic amount.

The residual income measure is obtained by regressing the log of household disposable income on a set of demographics including cohort interacted with education category of both spouses, dummies for children and region. Our measure of household disposable income pools the individual disposable income of the spouses (if the male has a spouse). In each year, we regress the log of household disposable income on dummies for region, marital status, number of children, education, and a fourth-order polynomial in age and the interaction of the latter two to obtain the residual income. To produce Figure 1 we use an equidistant grid of 11 quantiles and a third-order Hermite polynomial.

Table C1: Descriptive statistics

	1999	2001	2003	2005	2007	2009
	Earnings					
Mean	87,120	93,777	96,289	98,475	103,442	102,893
10%	34,863	37,532	36,278	35,005	35,533	31,992
25%	50,709	53,000	52,975	54,696	53,813	52,451
50%	73,423	77,000	76,576	78,944	80,292	79,181
75%	102,211	106,000	105,292	109,391	113,604	112,607
90%	145,789	152,000	150,280	154,971	171,688	163,879
	Consumption					
Mean	30,761	34,784	37,553	43,199	44,511	40,598
10%	15,804	17,477	18,026	20,365	21,634	20,008
25%	20,263	21,786	22,834	26,322	28,341	26,167
50%	26,864	29,366	31,924	37,381	38,704	34,570
75%	36,887	41,030	45,071	51,529	53,239	47,300
90%	48,977	53,870	62,864	73,338	73,715	67,012
	Net worth					
Mean	224,127	283,539	311,664	387,830	447,323	406,290
10%	19,016	26,100	28,494	38,287	41,854	33,592
25%	48,095	59,600	69,397	83,137	101,005	85,179
50%	114,096	137,500	159,230	191,663	217,599	188,354
75%	248,000	301,750	345,549	413,955	489,224	384,625
90%	535,827	586,000	654,437	830,462	939,583	867,786

Notes: Balanced subsample from PSID, 1999-2009. $N = 792$, $T = 6$. In 2001 dollars.

Supplementary Appendix to

“Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework”

Manuel Arellano, Richard Blundell and Stéphane Bonhomme

January 21, 2017

S1 Consumption responses in a two-period model

Consider a standard two-period setup, with a single risk-free asset. Let A_t denote beginning-of-period- t assets, and assume that $A_3 = 0$. Agents have CRRA utility. The Euler equation (assuming $\beta(1+r) = 1$ for simplicity) is

$$C_1^{-\gamma} = \mathbb{E}_1 \left[((1+r)A_2 + Y_2)^{-\gamma} \right],$$

where γ denotes risk aversion and the expectation is conditional on period-1 information. Here we have used the budget constraint $A_3 = (1+r)A_2 + Y_2 - C_2 = 0$. Equivalently,

$$C_1^{-\gamma} = \mathbb{E}_1 \left[((1+r)^2 A_1 + (1+r)Y_1 - (1+r)C_1 + Y_2)^{-\gamma} \right]. \quad (\text{S1})$$

Let $X_1 = (1+r)A_1 + Y_1$ denote “cash on hand” (as in Deaton, 1991). Let also $Y_2 = \mathbb{E}_1(Y_2) + \sigma W$. We will expand the Euler equation as $\sigma \rightarrow 0$. We denote the certainty equivalent consumption level as

$$\bar{C}_1 = \frac{(1+r)X_1 + \mathbb{E}_1(Y_2)}{2+r}.$$

Expanding in orders of magnitude of σ we have

$$C_1 \approx \bar{C}_1 + a\sigma + b\sigma^2 + c\sigma^3. \quad (\text{S2})$$

It is easy to see that $a = 0$, since $\mathbb{E}_1(W) = 0$. Hence,

$$C_1^{-\gamma} \approx \bar{C}_1^{-\gamma} \left(1 - \frac{\gamma}{\bar{C}_1} b\sigma^2 - \frac{\gamma}{\bar{C}_1} c\sigma^3 \right). \quad (\text{S3})$$

Moreover, by (S1) and (S2),

$$C_1^{-\gamma} \approx \mathbb{E}_1 \left[(\bar{C}_1 - (1+r)b\sigma^2 - (1+r)c\sigma^3 + \sigma W)^{-\gamma} \right],$$

from which it follows that

$$\begin{aligned}
C_1^{-\gamma} \approx \bar{C}_1^{-\gamma} \mathbb{E}_1 \left[1 + \frac{\gamma}{\bar{C}_1} (1+r)b\sigma^2 + \frac{\gamma}{\bar{C}_1} (1+r)c\sigma^3 - \frac{\gamma}{\bar{C}_1} \sigma W \right. \\
\left. + \frac{\gamma(\gamma+1)}{2} \left(\frac{1}{\bar{C}_1} \right)^2 \sigma^2 W^2 - \gamma(\gamma+1) \left(\frac{1}{\bar{C}_1} \right)^2 (1+r)b\sigma^3 W \right. \\
\left. - \frac{\gamma(\gamma+1)(\gamma+2)}{6} \left(\frac{1}{\bar{C}_1} \right)^3 \sigma^3 W^3 \right]. \tag{S4}
\end{aligned}$$

Finally, equating the coefficients of σ^2 and σ^3 in (S3) and (S4), using that $\mathbb{E}_1(W) = 0$, and denoting as $R = (1+r)X_1 + \mathbb{E}_1(Y_2)$ the expected period-2 resources, we obtain

$$b = -\frac{\gamma+1}{2R} \mathbb{E}_1(W^2), \quad c = \frac{(2+r)(\gamma+1)(\gamma+2)}{6R^2} \mathbb{E}_1(W^3).$$

This yields the following expression for period-1 consumption

$$\begin{aligned}
C_1 \approx \underbrace{\frac{(1+r)X_1 + \mathbb{E}_1(Y_2)}{2+r}}_{\text{certainty equivalent}} - \underbrace{\frac{\gamma+1}{2R} \mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^2)}_{\text{precautionary (variance)}} + \underbrace{\frac{(2+r)(\gamma+1)(\gamma+2)}{6R^2} \mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^3)}_{\text{precautionary (skewness)}}. \tag{S5}
\end{aligned}$$

Note that $\mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^2)$ is the conditional variance of Y_2 , and $\mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^3)$ is its conditional third-order moment.

Example: a simple nonlinear earnings process. To illustrate the effect of earnings shocks on consumption in this model, we consider the following simple earnings process (in levels):

$$Y_2 = Y_2^D + \rho(Y_1^P, V_2)Y_1^P + V_2 + Y_2^T,$$

where Y_2^D is the deterministic component, $Y_2^P = \rho(Y_1^P, V_2)Y_1^P + V_2$ is the persistent component, and Y_2^T is the transitory component. We set $\rho(Y_1^P, V_2) = 1 - \delta$ if $(Y_1^P < -c, V_2 > b)$ or $(Y_1^P > c, V_2 < -b)$, and $\rho(Y_1^P, V_2) = 1$ otherwise. Moreover, $\Pr(V_2 > b) = \Pr(V_2 < -b) = \tau$, with $\tau < 1/2$, and we assume that V_2 and Y_2^T are symmetrically distributed with zero mean. This earnings process has the following properties:

- If $|Y_1^P| \leq c$, then the process coincides with the “canonical” earnings model (in levels). So $\mathbb{E}_1(Y_2) = Y_2^D + Y_1^P$, $\mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^2) = \text{Var}(V_2) + \text{Var}(Y_2^T)$, and $\mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^3) = 0$.
- If $|Y_1^P| > c$, $\mathbb{E}_1(Y_2) = Y_2^D + (1 - \delta\tau)Y_1^P$ (“state-dependent persistence”).
- If $|Y_1^P| > c$, $\mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^2) = \tau(1 - \tau)\delta^2(Y_1^P)^2 + 2\tau\delta\mathbb{E}(V_2|V_2 > b)|Y_1^P| + \text{Var}(V_2) + \text{Var}(Y_2^T)$ (“state-dependent risk”).
- Lastly, if $Y_1^P < -c$, $\mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^3) > 0$, and if $Y_1^P > c$, $\mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^3) < 0$ (“state-dependent skewness”). For example, if $Y_1^P < -c$ we have

$$\begin{aligned}
\mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^3) = -\tau(1 - 2\tau)\delta Y_1^P \left[(1 - \tau)\delta^2(Y_1^P)^2 - 3\delta Y_1^P \mathbb{E}(V_2|V_2 > b) \right. \\
\left. + 3(\mathbb{E}(V_2^2|V_2 > b) - \mathbb{E}(V_2^2||V_2| \leq b)) \right] > 0.
\end{aligned}$$

Discussion. *State-dependent persistence* implies that low and high earnings households respond less to variations in Y_1^P than middle-earnings households. Low-earnings households save less than in the canonical linear model, while high-earnings households save more.

State-dependent risk implies that both low and high earnings households save more than in the canonical model because of higher variability of earnings. As shown by (S5), the effect is increasing in risk aversion and higher for low assets households. Note that the effect is scaled by expected resources R . Compared to the canonical linear earnings model, this effect will tend to increase savings for high earnings households and decrease savings for low earnings households.

Lastly, *state-dependent skewness* implies that, compared to the canonical model, high earnings households save more, and low earnings households save less.

Overall, the comparative statics for high earnings households are unambiguous, while the combined effect for low-earnings households is ambiguous.¹

S2 Simulations of a life-cycle model

Here we present an illustrative simulation to show some possible implications on consumption and assets of nonlinearity in income in the standard model outlined in Section 3. We start by comparing the linear canonical earnings model with a simple nonlinear earnings model which features the presence of “unusual” earnings shocks. We then report simulations based on the nonlinear earnings process we estimated on the PSID.

S2.1 Simulation exercise

To simulate the model we follow Kaplan and Violante (2010). Each household enters the labor market at age 25, works until 60, and dies with certainty at age 95. After retirement, families receive social security transfers Y_i^{ss} from the government, which are functions of the entire realizations of labor income. Income is assumed not to be subject to risk during retirement. Agents’ utility is CRRA with risk aversion $\gamma = 2$. The interest rate is $r = 3\%$ and the discount factor is $\beta = 1/(1+r) \approx .97$. We consider the following process for η_{it} :

$$\eta_{it} = \rho_t(\eta_{i,t-1}, v_{it})\eta_{i,t-1} + v_{it}, \quad (\text{S6})$$

and we compare two specifications. In the first specification, $\rho_t = 1$ (and v_{it} is normally distributed), which corresponds to the “canonical” earnings model used by Kaplan and Violante.² In the second specification, nonlinear persistence in income is approximated through a simple switching process:

$$\rho_t(\eta_{i,t-1}, v_{it}) = 1 - \delta \left(\mathbf{1}\{\eta_{i,t-1} < -d_{t-1}\} \mathbf{1}\{v_{it} > b_t\} + \mathbf{1}\{\eta_{i,t-1} > d_{t-1}\} \mathbf{1}\{v_{it} < -b_t\} \right), \quad (\text{S7})$$

where, at each age t , d_t is set so that $|\eta_{it}| > d_t$ with probability π , and b_t is set so that $|v_{it}| > b_t$ with probability π . In model (S6)-(S7), the persistence of the η process is equal to one unless an “unusual” positive shock v hits a low income household or an “unusual” negative shock v hits a high income household, leading persistence to drop to $1 - \delta = .8$. The latter happens with probability $\pi = .15$ in every period. Details on the simulation are provided below.

¹Note that here A_1 is taken as exogenous. In a complete model of the life cycle, household assets will be different when facing a nonlinear or a linear (“canonical”) earnings process.

²Kaplan and Violante (2010) also consider a more general AR(1) log-earnings process.

The simple parametric process (S7) is designed to roughly approximate the earnings process that we estimate on PSID data, see Section 6. It is worth noting that, since our flexible, quantile-based process is first-order Markov, it is easy to take the estimated η process to simulate, and possibly estimate, life-cycle models of consumption and saving such as the one we focus on here. In Subsection S2.3 of this Supplementary Appendix we show the results of using the nonlinear dynamic quantile earnings model we estimated on PSID data as an input in this simple life-cycle simulation model.

The simulation results are presented in Figure S35. In the simulation we use a natural borrowing limit.³ Graphs (a)-(b) show that a qualitative implication of the nonlinear earnings process is to reduce consumption among those on higher incomes. A negative shock for those on higher incomes reduces the persistence of the past and consequently is more damaging in terms of expected future incomes. This induces higher saving and lower consumption at younger ages. Conversely, we see that consumption is (slightly) higher for the nonlinear process for those on lower income. Graphs (c)-(d) show that the nonlinear model also results in a higher consumption variance among older households, and steeper accumulation and subsequent decumulation of assets over the life cycle. In addition, in graph (e) we report estimates of the average derivative of the conditional mean of log-consumption with respect to log-earnings, holding assets and age fixed at percentiles indicated on the two horizontal axes. We see that in the simulated economy consumption responses to changes in earnings tend to decrease with age and, to a lesser extent, the presence of assets. These simulation results provide further motivation for the use of a nonlinear earnings model to study consumption dynamics.

S2.2 Details on the simulations

Agents live for T periods, and work until age T_{ret} , where both T and T_{ret} are exogenous and fixed. ξ_t is the unconditional probability of surviving to age t , where $\xi_t = 1$ before retirement, and $\xi_t < 1$ after retirement. Households have expected life-time utility

$$\mathbb{E}_0 \sum_{t=1}^T \beta^{t-1} \xi_t u(C_{it}).$$

During working years $1 \leq t < T_{ret}$, agents receive after-tax labor income Y_{it} , which is decomposed into a deterministic experience profile κ_t , a permanent component η_{it} , and a transitory component ε_{it} :

$$\ln Y_{it} = \kappa_t + y_{it},$$

$$y_{it} = \eta_{it} + \varepsilon_{it}.$$

We consider the following process for η_{it}

$$\eta_{it} = \rho_t(\eta_{i,t-1}, v_{it}) \eta_{i,t-1} + v_{it},$$

where η_{i0} is drawn from an initial normal distribution with mean zero and variance $\sigma_{\eta_0}^2$. The shocks ε_{it} and v_{it} have mean zero, are normally distributed with variances σ_{ε}^2 and $\sigma_{v_t}^2$. The persistence of the η process is approximated as (S7) where, at each age t , d_t is set so that $|\eta_{it}| > d_t$ with probability π , and b_t is set so that $|v_{it}| > b_t$ with probability π . We set $\pi = .15$ and $1 - \delta = .8$.

³Results for a zero borrowing limit are given in Figure S36.

Define gross labor income as \tilde{Y}_{it} , with $\tilde{Y}_{it} = G(Y_{it})$, where G function is the inverse of a tax function estimated by Gouveia and Strauss (1994).⁴ After retirement, agents receive after-tax social security transfers Y^{ss} , which are a function of lifetime average individual gross earnings $P\left(\frac{1}{T_{ret}-1} \sum_{t=1}^{T_{ret}-1} \tilde{Y}_{it}\right)$.

Lastly, throughout their lifetime households have access to a single risk-free, one-period bond whose constant return is $1 + r$, and face a period-to-period budget constraint

$$A_{i,t+1} = (1 + r) A_{it} + Y_{it} - C_{it}, \quad \text{if } t < T_{ret},$$

$$\left(\frac{\xi_t}{\xi_{t+1}}\right) A_{i,t+1} = (1 + r) A_{it} + Y^{ss} - C_{it}, \quad \text{if } t \geq T_{ret}.$$

For the calibration we use Kaplan and Violante's (2010, KV hereafter) preferred parameters. Specifically we use the following:

Demographics. The model period is one year. Agents enter the labor market at age 25, retire at age 60, and die with certainty at age 95. So we set $T_{ret} = 35$ and $T = 70$. The survival rates ξ_t are obtained from the National Center for Health Statistics (1992).

Preferences. We assume the utility function is of the CRRA form $u(C) = C^{1-\gamma}/(1-\gamma)$, where the risk aversion parameter is set to $\gamma = 2$.

Discount factor and interest rate. The interest rate r is assumed to equal 3%. We set the discount factor β to match an aggregate wealth-income ratio of 2.5, which is the average wealth to average income ratio computed from the 1989 and 1992 Survey of Consumer Finances.

Income process. We use KV's choice of deterministic age profile κ_t , which is estimated from the PSID. The estimated profile peaks after 21 years of labor market experience at roughly twice the initial value, and then it slowly declines to about 80% of the peak value. For the stochastic components of the income process, we set the initial variance of the permanent shocks $\sigma_{\eta_0}^2 = .15$ to match the dispersion of household earnings at age 25. We set the variance of transitory shocks $\sigma_{\varepsilon}^2 = .05$, which is the value estimated by Blundell, Pistaferri and Preston (2008). The permanent component follows the nonlinear switching process (S7). In order to make its variance comparable with the random walk process in KV, we calibrate $\sigma_{v_t}^2$ to match the dispersion of age-specific variance of η_{it} in KV.

Initial wealth and borrowing limit. Households' initial assets are set to 0. We impose two alternative borrowing limits: either a natural borrowing limit, in which the agent cannot die with debt, or a zero borrowing limit, in which the agent's net worth can not fall below zero.

⁴The tax function is

$$\tau(\tilde{Y}_{it}) = 0.258 \times \left[\tilde{Y}_{it} - \left(\tilde{Y}_{it}^{-0.768} + \tau^s \right)^{-\frac{1}{0.768}} \right],$$

where τ^s is chosen so that the ratio of total personal current tax receipts on labor income (not including social security contributions) to total labor income is the same as for the US economy in 1990, i.e. roughly 25%.

Social security benefits. This setup follows KV exactly. Social security benefits are a function of lifetime average individual gross earnings $Y^{ss} = P\left(\frac{1}{T_{ret}-1} \sum_{t=1}^{T_{ret}-1} \tilde{Y}_{it}\right)$. The P function is designed to mimic the actual US system. This is achieved by specifying that benefits are equal to 90% of average past earnings up to a given bend point, 32% from this first bend point to a second bend point, and 15% beyond that. The two bend points are set at, respectively, 0.18 and 1.10 times cross-sectional average gross earnings, based on the US legislation and individual earnings data for 1990. Benefits are then scaled proportionately so that a worker earning average labor income each year is entitled to a replacement rate of 45%.

Discretization of the earnings process. In order to use the switching process (S7) in the life-cycle model, we compute age-specific Markov transition probabilities using a simulation approach, as follows.

1. Based on the process (S7), we simulate one million households throughout their working years.
2. At each age, we rank households by their earnings and put them into 40 different bins. We then count the transitions between any two bins from two neighboring ages, and estimate transition probabilities.

S2.3 Simulations based on the estimated nonlinear earnings model

In this subsection we report the results of a simulation exercise closely related to the one shown in the first part of this section, except that it is based on the nonlinear quantile-based earnings process that we estimated on the PSID. Calibration and simulation details are similar to the ones above, with a few differences. In order to ensure comparability with the data, we let each household enter the model at age 25, work until 61, and die at age 93. As the PSID is biennial, we set each period in the model to be equal to two years. There is one risk-free asset with a constant interest rate $r = 1.03^2 - 1 \approx 6\%$, and the discount factor is set to $\beta = (1 + r)^{-1}$. We impose a natural borrowing limit in the simulation.

We simulate life-cycle profiles for one million households. When discretizing persistent and transitory earnings components, we use 100 bins for the former and 80 bins for the latter. We checked that the discretized process fits the nonlinear persistence in Figure 1 well. We remove age-specific means in the persistent component (which may be different from zero for the particular cohort for which we are performing the simulation). In order to ensure comparability with a canonical linear earnings process, in the nonlinear process we compute age-specific variances of transitory shocks and target the age profile of the variance of the persistent component, and we set the parameters of the linear process (that is, a random walk plus an independent shock) using the resulting values.

The results of the simulation are presented in Figure S37. As in Figure S35, we see that households on higher incomes tend to consume less when exposed to the nonlinear process than to the linear one. Conversely, consumption is slightly higher for the nonlinear process for those on lower income. Consumption and asset variance are lower for the nonlinear process, similarly as in Figure S35, although the differences between the two processes are stronger in Figure S37, which is based on the process estimated on PSID.⁵ Lastly, the marginal propensity to consume out of

⁵Differences in scale with Figure S35 are due to the different period considered in this paper relative to

earnings in panel (e) tends to decrease with age and the presence of assets, similarly to what we found in the PSID.

S3 Extensions

Here we consider four extensions of the model: to allow for unobserved heterogeneity in earnings, dependent ε , advance earnings information, and consumption habits, respectively. The estimation strategy can be modified to handle each of these extensions.

S3.1 Unobserved heterogeneity in earnings

It is possible to allow for unobserved heterogeneity in earnings, in addition to heterogeneity in the initial condition η_{i1} . Specifically, let η_{it} be a first-order Markov process conditional on another latent component ζ_i :

$$\eta_{it} = Q_t(\eta_{i,t-1}, \zeta_i, u_{it}), \quad (\text{S8})$$

where u_{it} is i.i.d. standard uniform, independent of η_i^{t-1} and ζ_i . ε_{it} is independent over time, independent of η_{is} for all s , and independent of ζ_i .

With a vector-valued ζ_i , (S8) would nest linear earnings models with slope heterogeneity as in Guvenen (2007) and Guvenen and Smith (2014), for example. A simpler case is our baseline model (1)-(2) augmented with a household-specific fixed-effect, that is equation (20) in the paper.

Consider the scalar- ζ_i case for concreteness, and take $T = 5$. In this model, (y_{i1}, y_{i2}) , y_{i3} , and (y_{i4}, y_{i5}) are conditionally independent given (η_{i3}, ζ_i) .⁶ By Hu and Schennach (2008)'s theorem, for bivariate latent (η_{i3}, ζ_i) , and under suitable injectivity conditions, the marginal distribution of ε_{i3} is thus identified given five periods of earnings data. As a result, the joint density of η 's is identified by a similar argument as in Section 4. Identification of the densities of ζ_i and of η_{it} given $(\eta_{i,t-1}, \zeta_i)$ can then be shown along the lines of Hu and Shum (2012), under a suitable ‘‘scaling’’ condition. A scaling condition is implicit in equation (20), which is the model we implement.

S3.2 Dependence in the transitory earnings component

In the baseline model ε_{it} are independent over time. It is possible to allow for serial dependence while maintaining identification. To see this, consider the setup where ε_{it} is an m -dependent process with $m = 1$ (for example, an $MA(1)$ process), and consider a panel with $T \geq 5$ periods. Then it is easy to see that y_{i1} , y_{i3} and y_{i5} are conditionally independent given η_{i3} . As a result, identification arguments based on ‘‘Hidden Markov’’ structures (Hu and Schennach, 2008, Wilhelm, 2015) can be applied.

Kaplan and Violante (2010) and different parameter choices. Also, note that consumption data in the PSID waves which we use only capture a share of consumption expenditures, and similarly for assets. This can explain the differences between the levels in Figure S37 and the descriptive statistics reported in Table C1.

⁶Indeed,

$$\begin{aligned} f(y_1, y_2, y_3, y_4, y_5 | \eta_3, \zeta) &= f(y_1, y_2 | \eta_3, \zeta) f(y_3 | \eta_3, \zeta, y_1, y_2) f(y_4, y_5 | \eta_3, \zeta, y_3, y_2, y_1) \\ &= f(y_1, y_2 | \eta_3, \zeta) f(y_3 | \eta_3, \zeta) f(y_4, y_5 | \eta_3, \zeta). \end{aligned}$$

S3.3 Advance information

If households have advance information about future earnings shocks, the consumption rule (10) takes future earnings components as additional arguments, see Blundell *et al.* (2008). For example, consider a model where households know the realization of the one-period-ahead persistent component, in which case

$$c_{it} = g_t(a_{it}, \eta_{it}, \eta_{i,t+1}, \varepsilon_{it}, \nu_{it}), \quad t = 1, \dots, T-1. \quad (\text{S9})$$

Identification can be established using similar arguments as in the baseline model. To see this, consider first period's consumption. We have

$$f(c_1|a_1, y) = \int \int f(c_1|a_1, \eta_1, \eta_2, y_1) f(\eta_1, \eta_2|a_1, y) d\eta_1 d\eta_2.$$

It can be shown that $f(\eta_1, \eta_2|a_1, y)$ is identified under completeness in (y_{i2}, \dots, y_{iT}) of the distribution of $(\eta_{i1}, \eta_{i2}|y_i)$, using the earnings process and first period's assets. If the distribution of $(\eta_{i1}, \eta_{i2}|a_{i1}, y_i)$ is complete in (y_{i2}, \dots, y_{iT}) it thus follows that $f(c_1|a_1, \eta_1, \eta_2, y_1)$ is identified. In this case we need at least two "excluded instruments" in y_i for (η_{i1}, η_{i2}) . The other steps in the identification arguments of Section 4 can be similarly adapted.

Lastly, similar arguments can be used to show identification in models where households have advance information about future transitory shocks $\varepsilon_{i,t+s}$, as well as in models where the consumption rule depends on lags of η 's or ε 's, for example in models where η_{it} follows a higher-order Markov process.

S3.4 Consumption habits

In the presence of habits, the consumption rule takes the form

$$c_{it} = g_t(c_{i,t-1}, a_{it}, \eta_{it}, \varepsilon_{it}, \nu_{it}), \quad t = 2, \dots, T. \quad (\text{S10})$$

Identification can be shown under similar conditions as in Section 4. For example, in the second period equation (18) becomes

$$f(c_2|c_1, a_2, a_1, y) = \int f(c_2|c_1, a_2, \eta_2, y_2) f(\eta_2|c_1, a_2, a_1, y) d\eta_2.$$

Provided the distribution of $(\eta_{i2}|c_{i1}, a_{i2}, a_{i1}, y_i)$ is identified, and complete in $(a_{i1}, y_{i1}, y_{i3}, \dots, y_{iT})$, it thus follows that the density $f(c_2|c_1, a_2, \eta_2, y_2)$ is identified. Intuitively, in the presence of habits the first lag of consumption cannot be used as an "excluded instrument" since it affects second period's consumption directly.

S4 Summary of the argument in Wilhelm (2015)

We consider model (1)-(2) with $T = 3$. We omit i subscripts for conciseness. Let $L^2(f)$ denote the set of squared-integrable functions with respect to a weight function f . We define $\mathcal{L}_{y_2|y_1}$ as the linear operator such that $\mathcal{L}_{y_2|y_1} h(a) = \mathbb{E}[h(y_2)|y_1 = a] \in L^2(f_{y_1})$ for every function $h \in L^2(f_{y_2})$. Similarly, let $\mathcal{L}_{\eta_2|y_1}$ be such that $\mathcal{L}_{\eta_2|y_1} h(a) = \mathbb{E}[h(\eta_2)|y_1 = a] \in L^2(f_{y_1})$ for every function $h \in L^2(f_{\eta_2})$. We denote as $\mathcal{R}(\mathcal{L}_{y_2|y_1})$ the range of $\mathcal{L}_{y_2|y_1}$, that is

$$\mathcal{R}(\mathcal{L}_{y_2|y_1}) = \{k \in L^2(f_{y_1}), \text{ s.t. } k = \mathcal{L}_{y_2|y_1} h \text{ for some } h \in L^2(f_{y_2})\}.$$

We assume the following, in addition to the Markovian and independence assumptions on η 's and ε 's.

Assumption S1

- (i) $\mathcal{L}_{y_2|y_1}$ and $\mathcal{L}_{\eta_2|y_1}$ are injective.
- (ii) There exists a function $h \in L^2(f_{y_3})$ such that

$$\mathbb{E}[h(y_3)|y_1 = \cdot] \in \mathcal{R}(\mathcal{L}_{y_2|y_1}), \text{ and} \tag{S11}$$

$$\mathbb{E}[y_2 h(y_3)|y_1 = \cdot] \in \mathcal{R}(\mathcal{L}_{y_2|y_1}). \tag{S12}$$

Thus, there exist s_1 and s_2 in $L^2(f_{y_2})$ such that

$$\mathbb{E}[h(y_3)|y_1 = \cdot] = \mathcal{L}_{y_2|y_1} s_1, \text{ and } \mathbb{E}[y_2 h(y_3)|y_1 = \cdot] = \mathcal{L}_{y_2|y_1} s_2.$$

- (iii) Let $\tilde{s}_1(y) = y s_1(y)$. The Fourier transforms $\mathcal{F}(s_1)$, $\mathcal{F}(\tilde{s}_1)$, and $\mathcal{F}(s_2)$ (where $\mathcal{F}(h)(u) = \int h(x)e^{iux} dx$) are ordinary functions. Moreover, $\mathcal{F}(s_1)(u) \neq 0$ for all $u \in \mathbb{R}$.

Part (i) is an injectivity/completeness condition. Part (ii) is not standard. It is related to the existence problem in nonparametric instrumental variables. Horowitz (2009) proposes a test for (S11) in the case where $\mathcal{L}_{y_2|y_1}$ is a compact operator. Part (iii) is a high-level assumption. See Wilhelm (2015) for more primitive conditions.

By Assumption S1-(ii) we have, almost surely in y_1 ,

$$\begin{aligned} \mathbb{E}[h(y_3)|y_1] &= \mathbb{E}[s_1(y_2)|y_1], \\ \mathbb{E}[y_2 h(y_3)|y_1] &= \mathbb{E}[s_2(y_2)|y_1]. \end{aligned}$$

Moreover, s_1 and s_2 are the unique solutions to these equations by Assumption S1-(i).

Hence, given the model's assumptions

$$\mathbb{E}[\mathbb{E}(h(y_3)|\eta_2) | y_1] = \mathbb{E}[\mathbb{E}(s_1(y_2)|\eta_2) | y_1] \text{ a.s.}$$

It thus follows from the injectivity of $\mathcal{L}_{\eta_2|y_1}$ in Assumption S1-(i) that, almost surely in η_2 ,

$$\mathbb{E}[h(y_3)|\eta_2] = \mathbb{E}[s_1(y_2)|\eta_2]. \tag{S13}$$

Likewise, $\mathbb{E}[y_2 h(y_3)|\eta_2] = \mathbb{E}[s_2(y_2)|\eta_2]$. Hence

$$\eta_2 \mathbb{E}[h(y_3)|\eta_2] = \mathbb{E}[s_2(y_2)|\eta_2] \text{ a.s.} \tag{S14}$$

Combining (S13) and (S14), we obtain

$$\eta_2 \mathbb{E}[s_1(y_2)|\eta_2] = \mathbb{E}[s_2(y_2)|\eta_2] \text{ a.s.}$$

That is, almost surely in η_2 ,

$$\eta_2 \int s_1(y) f_{\varepsilon_2}(y - \eta_2) dy = \int s_2(y) f_{\varepsilon_2}(y - \eta_2) dy. \tag{S15}$$

The functional equation (S15) depends on s_1 and s_2 , which are both uniquely determined given the data generating process, and on the unknown f_{ε_2} . By Assumption S1-(iii) we can take Fourier transforms and obtain

$$i\mathcal{F}(s_1)(u)\frac{d\psi_{\varepsilon_2}(-u)}{du} + \mathcal{F}(\tilde{s}_1)(u)\psi_{\varepsilon_2}(-u) = \mathcal{F}(s_2)(u)\psi_{\varepsilon_2}(-u), \quad (\text{S16})$$

where $\psi_{\varepsilon_2}(u) = \mathcal{F}(f_{\varepsilon_2})(u)$ is the characteristic function of ε_2 .

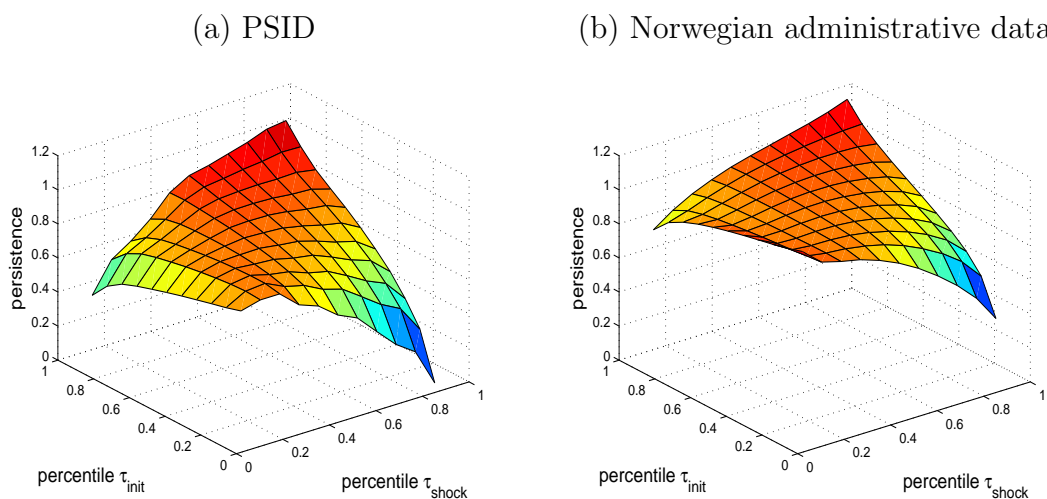
Noting that $\psi_{\varepsilon_2}(0) = 1$, (S16) can be solved in closed form for $\psi_{\varepsilon_2}(\cdot)$, because $\mathcal{F}(s_1)(u) \neq 0$ for all u by Assumption S1-(iii). This shows that the characteristic function of ε_2 , and hence its distribution function, are identified.

References

- [1] Deaton, A. (1991): “Saving and Liquidity Constraints,” *Econometrica*, 59(5), 1221–1248.
- [2] Gouveia, M., and R. P. Strauss (1994): “Effective Federal Individual Income Tax Functions: An Exploratory Empirical Analysis,” *National Tax Journal*, 317–339.
- [3] Guvenen, F. (2007): “Learning your Earning: Are Labor Income Shocks Really Very Persistent?” *American Economic Review*, 97(3), 687–712.
- [4] Guvenen, F., and A. Smith (2014): “Inferring Labor Income Risk from Economic Choices: An Indirect Inference Approach,” *Econometrica*, November, 82(6), 2085–2129.
- [5] Horowitz, J. (2009): “Specification Testing in Nonparametric Instrumental Variables Estimation,” unpublished manuscript.
- [6] Kaplan, G., and G. Violante (2014): “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, 82(4), 1199–1239.
- [7] Wilhelm, D. (2015): “Identification and Estimation of Nonparametric Panel Data Regressions with Measurement Error,” Cemmap working paper CWP34/15.

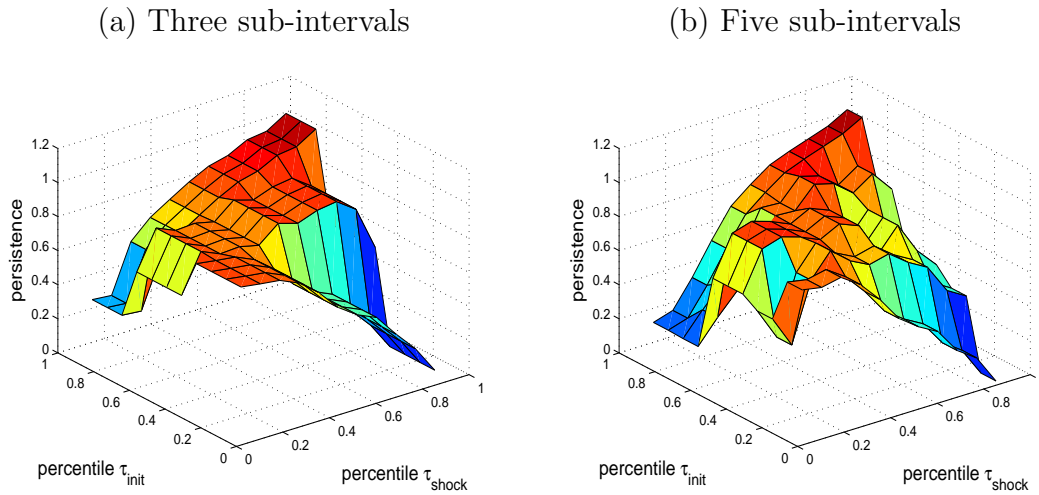
S5 Additional figures

Figure S1: Nonlinear earnings persistence in log-wages (PSID, males) and individual income (Norwegian population register data)



Note: Estimates of the average derivative of the conditional quantile function of y_{it} given $y_{i,t-1}$ with respect to $y_{i,t-1}$, evaluated at percentile τ_{shock} and at a value of $y_{i,t-1}$ that corresponds to the τ_{init} percentile of the distribution of $y_{i,t-1}$. Age 25-60. Left: PSID, male log hourly wage residuals, 1999-2009; right: Results from the Norwegian population register data, individual log-earnings residuals, years 2005-2006, see Appendix C; are provided as part of the project on 'Labour Income Dynamics and the Insurance from Taxes, Transfers and the Family'.


Figure S2: Nonlinear persistence, piecewise-linear specification quantile regression



Note: See the notes to Figure 2. Piecewise-linear quantile regression with 3 and 5 equally-spaced sub-intervals.

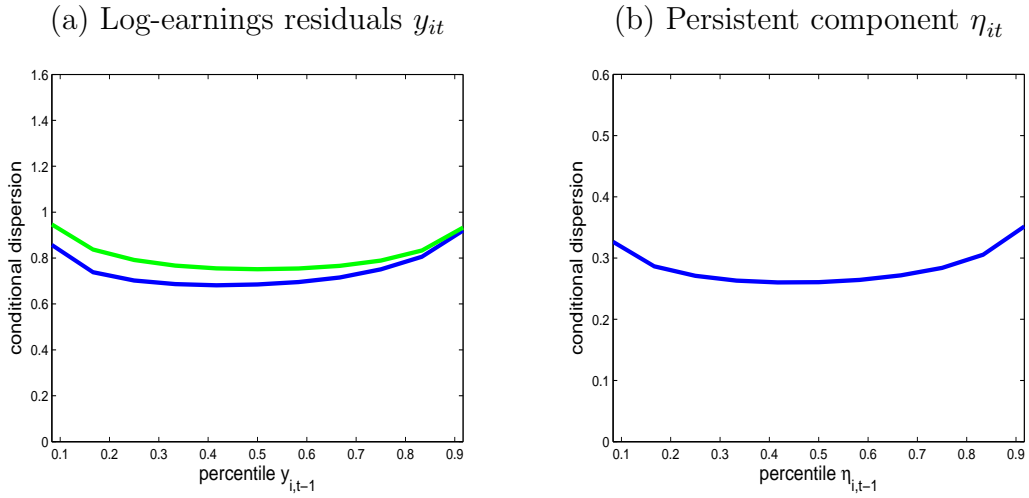
Figure S3: Nonlinear persistence of η_{it} , averaged over ages

(a) Persistent component η_{it} , nonlinear model



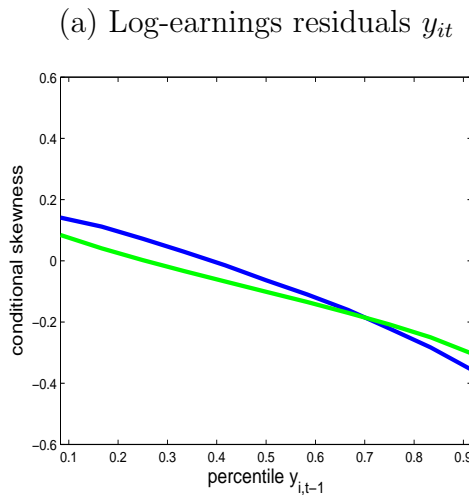
Note: Estimates of the average derivative of the conditional quantile function of η_{it} on $\eta_{i,t-1}$ with respect to $\eta_{i,t-1}$, based on estimates from the nonlinear earnings model, averaged over all ages of household heads.

Figure S4: Conditional dispersion of log-earnings residuals (fit) and η component




Note: Conditional dispersion $\sigma(y, \tau) = Q(\tau|y_{i,t-1} = y) - Q(1 - \tau|y_{i,t-1} = y)$ and $\sigma(\eta, \tau)$, for $\tau = 11/12$. Log-earnings residuals (left) and η component (right). On the left graph, dark is data and light is model fit. The x -axis shows the conditioning variable, the y -axis shows the corresponding value of the conditional dispersion measure.

Figure S5: Conditional skewness of log-earnings residuals, fit




Note: Conditional skewness $sk(y, \tau)$, see equation (6), for $\tau = 11/12$. Log-earnings residuals. The x -axis shows the conditioning variable, the y -axis shows the corresponding value of the conditional skewness measure. Dark is PSID data, light is nonlinear model.

Figure S6: Skewness of log-earnings residuals growth at various horizons, fit



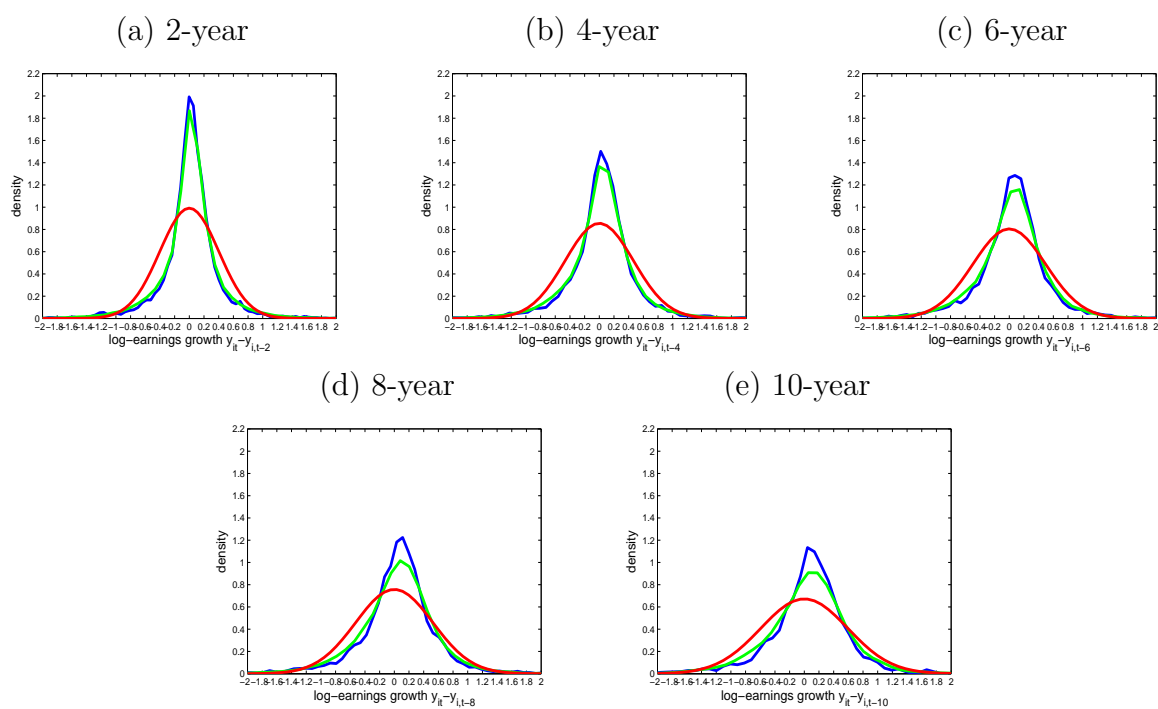
Note: Robust skewness estimate of $y_{it} - y_{i,t-s}$, at various horizons s (reported on the x-axis), computed according to the formula $sk(\tau, \alpha) = \frac{Q(\tau) + Q(1-\tau) - 2Q(1/2)}{Q(\tau) - Q(1-\tau)}$, where Q denote unconditional quantiles, and $\tau = 11/12$. Dark is PSID data, light is nonlinear model.

Figure S7: Kurtosis of log-earnings residuals growth at various horizons, fit



Note: Robust kurtosis estimate of $y_{it} - y_{i,t-s}$, at various horizons s (reported on the x-axis), computed according to the formula $kur(\tau, \alpha) = \frac{Q(1-\alpha) - Q(\alpha)}{Q(\tau) - Q(1-\tau)}$, where Q denote unconditional quantiles, $\tau = 10/12$, and $\alpha = 1/12$. Dark is PSID data, light is nonlinear model, and the horizontal line denotes the values of $kur(\tau, \alpha)$ for a normal distribution.

Figure S8: Densities of log-earnings growth at various horizons





Note: Estimated density of $y_{it} - y_{i,t-s}$, at various horizons s . Dark is PSID data, light is nonlinear model. Added to each graph is the Gaussian density with zero mean and the same variance as in the data.

Figure S9: Nonlinear persistence in earnings, 95% pointwise confidence bands



(a) PSID data

(b) Nonlinear model

Parametric bootstrap

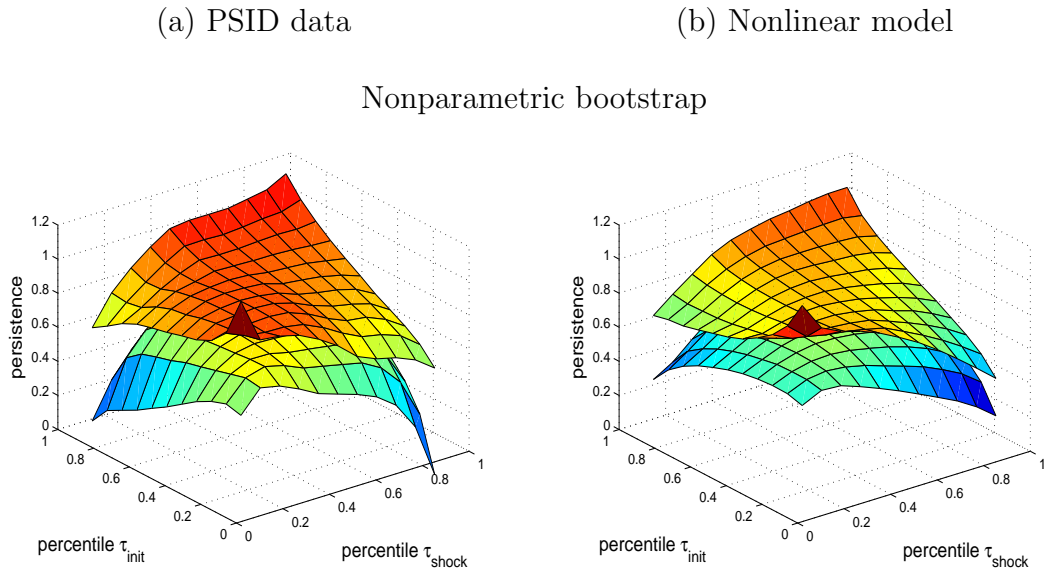


Nonparametric bootstrap



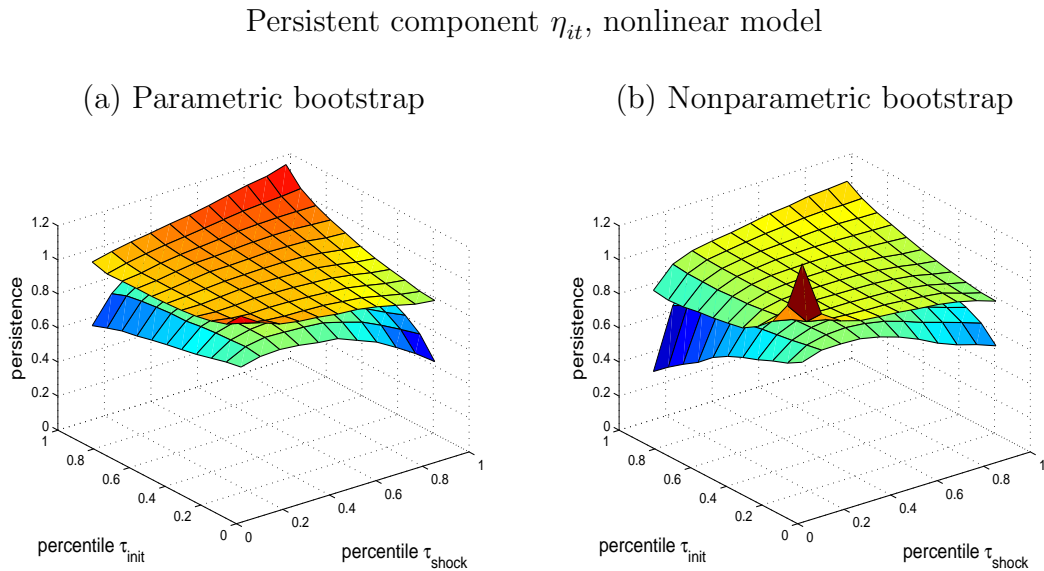
Note: See notes to Figure 2. Pointwise 95% confidence bands. 500 replications. Parametric bootstrap is based on the point estimates. Nonparametric bootstrap is clustered at the household level.

Figure S10: Nonlinear persistence in earnings, 95% uniform confidence bands



Note: See notes to Figure 2. Nonparametric bootstrap clustered at the household level. Uniform 95% confidence bands. 500 replications.

Figure S11: Nonlinear persistence in η_{it} , 95% pointwise confidence bands




Note: See notes to Figure 2. Pointwise 95% confidence bands. 500 replications. Parametric bootstrap is based on the point estimates. Nonparametric bootstrap is clustered at the household level.

Figure S12: Nonlinear persistence in η_{it} , 95% uniform confidence bands

Persistent component η_{it} , nonlinear model

Nonparametric bootstrap





Note: See notes to Figure 2. Nonparametric bootstrap clustered at the household level. Uniform 95% confidence bands. 500 replications.

Figure S13: Conditional skewness of log-earnings residuals and η component, 95% pointwise confidence bands



(a) Log-earnings residuals y_{it}

(b) Persistent component η_{it}

Parametric bootstrap

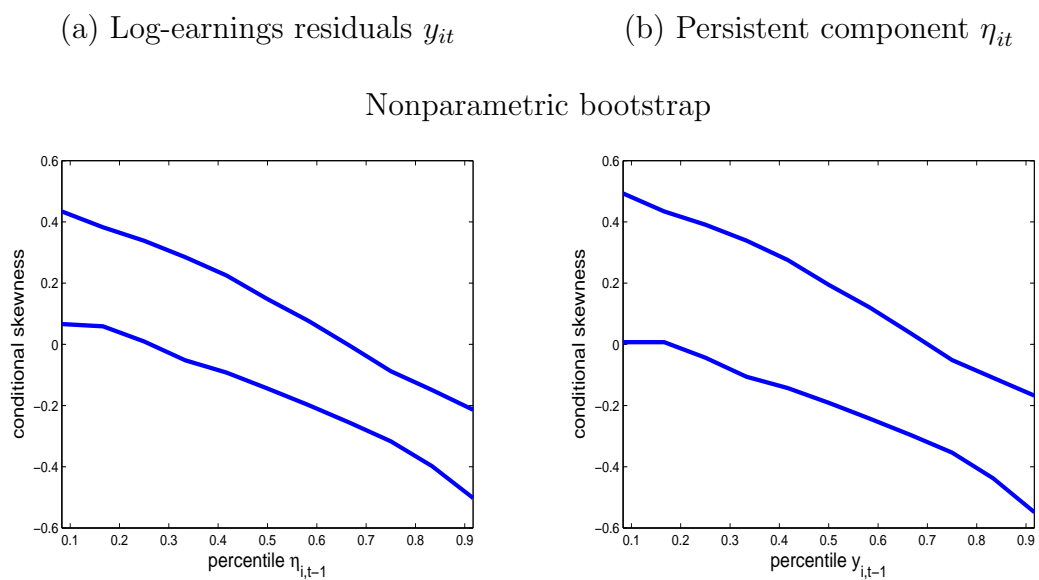


Nonparametric bootstrap



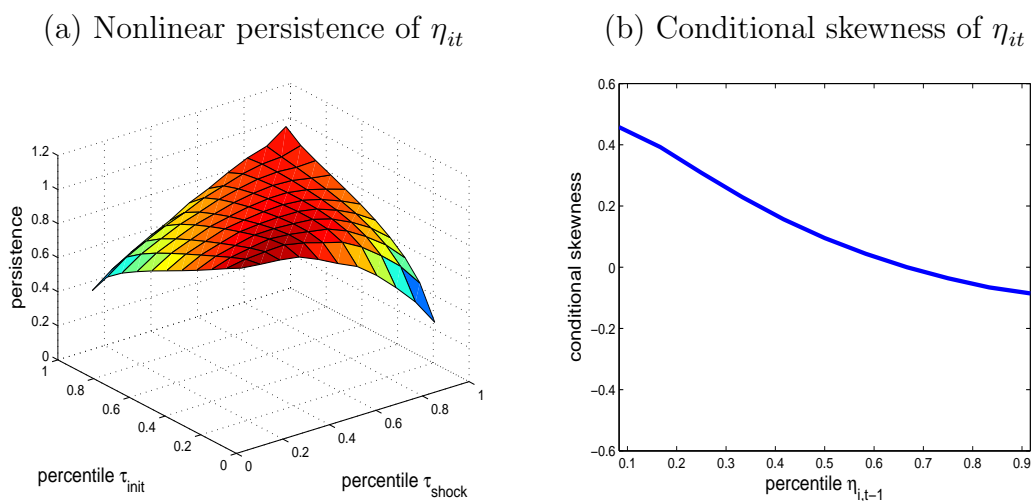
Note: See notes to Figure 4. Pointwise 95% confidence bands. 500 replications. Parametric bootstrap is based on the point estimates. Nonparametric bootstrap is clustered at the household level.

Figure S14: Conditional skewness of log-earnings residuals and η component, 95% uniform confidence bands



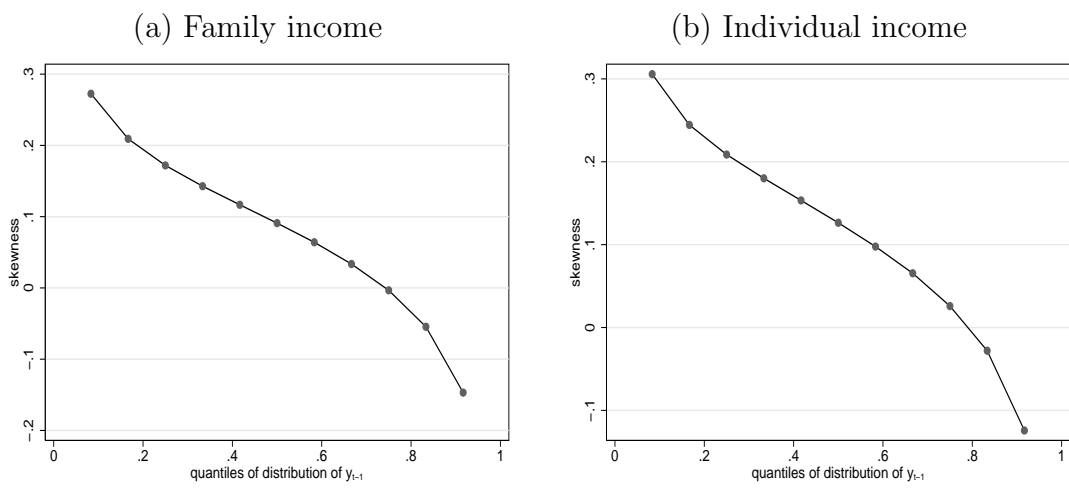
Note: See notes to Figure 2. Nonparametric bootstrap clustered at the household level. Uniform 95% confidence bands. 500 replications.

Figure S15: Household heterogeneity in earnings



Note: (a) Estimates of the average derivative of the conditional quantile function of η_{it} on $\eta_{i,t-1}$ with respect to $\eta_{i,t-1}$, based on estimates from the nonlinear earnings model with an additive household-specific effect. (b) Conditional skewness $sk(\eta, \tau)$, see equation (6), for $\tau = 11/12$, based on the same model.


Figure S16: Conditional skewness, Norwegian administrative data




Note: Conditional skewness of log-earnings measured as in (6) for $\tau = 1/10$. Age 25-60, years 2005-2006.

Figure S17: Nonlinear persistence, Norwegian data


(a) Earnings, Norwegian data



(b) Earnings, nonlinear model

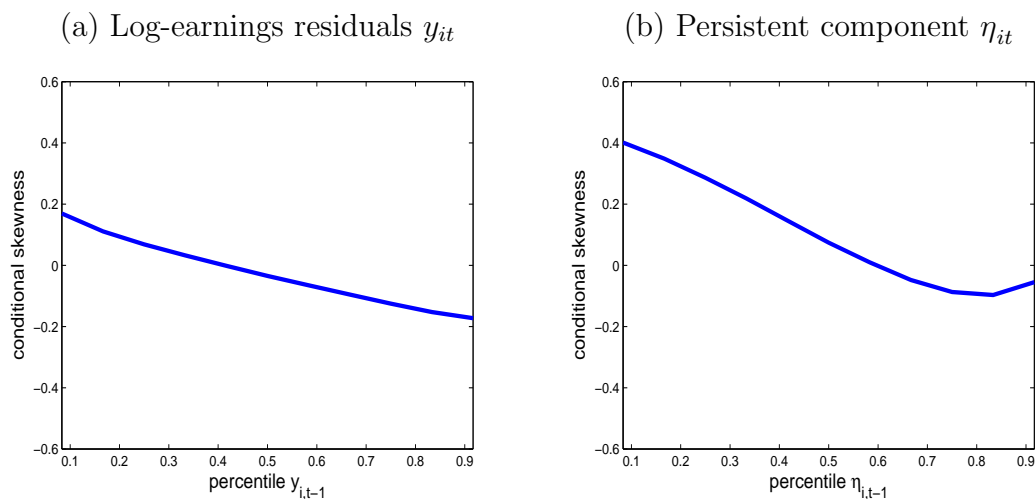


(c) Persistent component η_{it} , nonlinear model




Note: See the notes to Figure 2. Random subsample of 2,873 households, from 2000 – 2005 Norwegian administrative data, non-immigrant residents, age 30 to 60.

Figure S18: Conditional skewness of log-earnings residuals and η component, Norwegian data




Note: See the notes to Figure 4. Random subsample of 2,873 households, from 2000 – 2005 Norwegian administrative data, non-immigrant residents, age 30 to 60.

Figure S19: Densities of persistent and transitory earnings components, Norwegian data



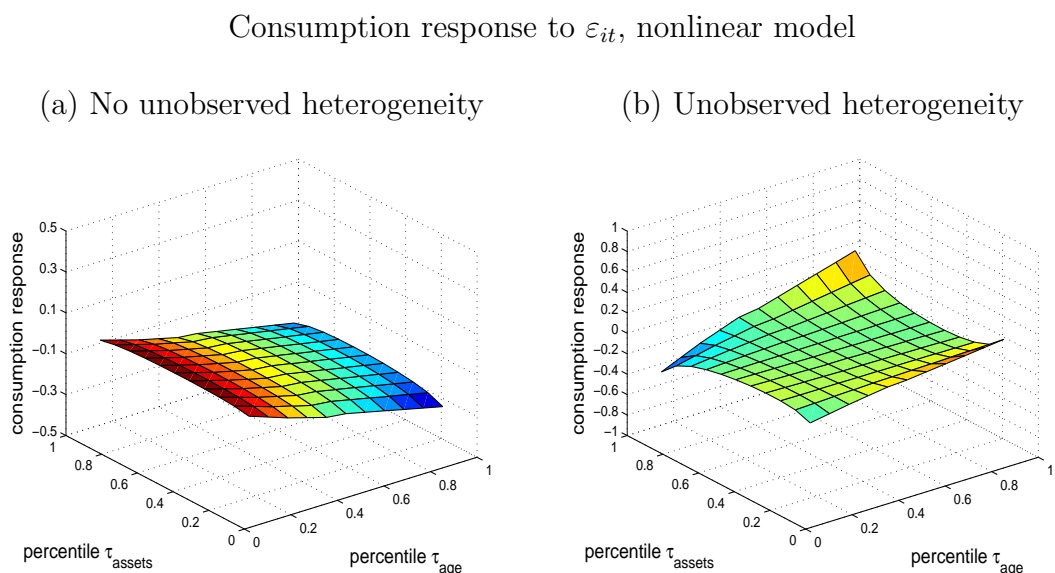
Note: See the notes to Figure 3. Random subsample of 2,873 households, from 2000 – 2005 Norwegian administrative data, non-immigrant residents, age 30 to 60.

Figure S20: Density of log-consumption in the data and model, with and without unobserved heterogeneity




Note: Source PSID and nonlinear consumption model. Kernel density estimates. Dashed is Gaussian fit.

Figure S21: Consumption responses to transitory earnings shocks, by assets and age



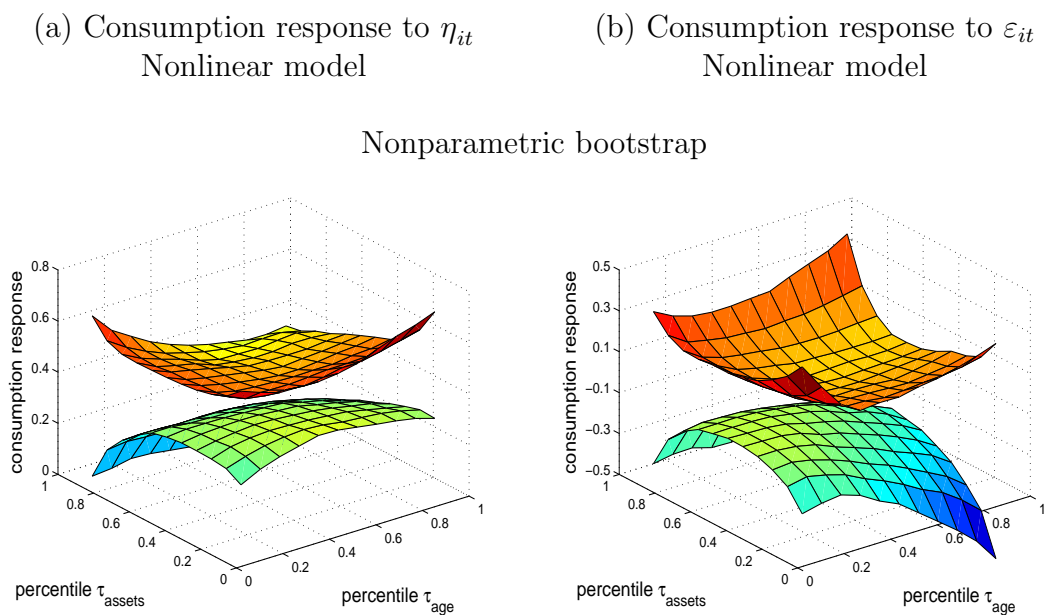
Note: Estimates of the average consumption responses to variations in ε_{it} , evaluated at τ_{assets} and τ_{age} .

Figure S22: Consumption responses to earnings shocks, by assets and age, 95% pointwise confidence bands



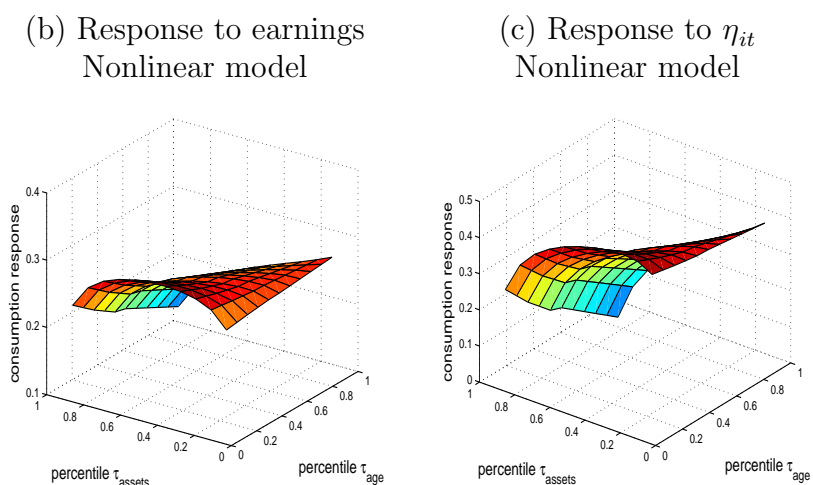
Note: See notes to Figures 5 and S21. Pointwise 95% confidence bands. 200 replications. Parametric bootstrap is based on the point estimates. Nonparametric bootstrap is clustered at the household level.

Figure S23: Consumption responses to earnings shocks, by assets and age, 95% uniform confidence bands



Note: See notes to Figures 5 and S21. Nonparametric bootstrap clustered at the household level. Uniform 95% confidence bands. 200 replications.


Figure S24: Consumption responses to earnings shocks, by assets and age, model with household-specific unobserved heterogeneity




Note: See the notes to Figure 5. Consumption and assets model with household-specific unobserved heterogeneity.

Figure S25: Consumption responses to assets, model with unobserved heterogeneity

(b) Given y_{it}
Nonlinear model




(c) Given η_{it}
Nonlinear model



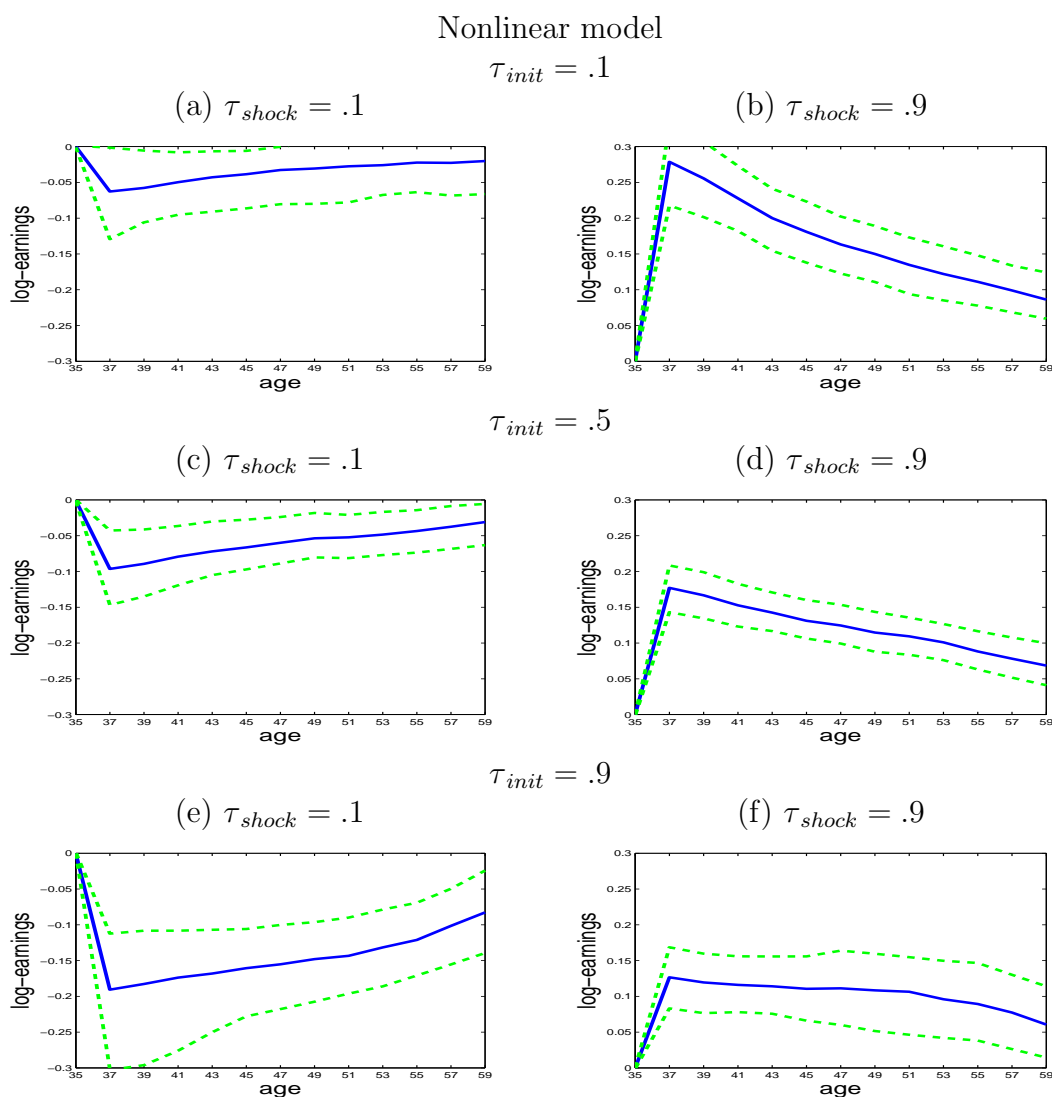
Note: See the notes to Figure 6. Model with unobserved heterogeneity.

Figure S26: Impulse responses, earnings, 95% pointwise confidence bands



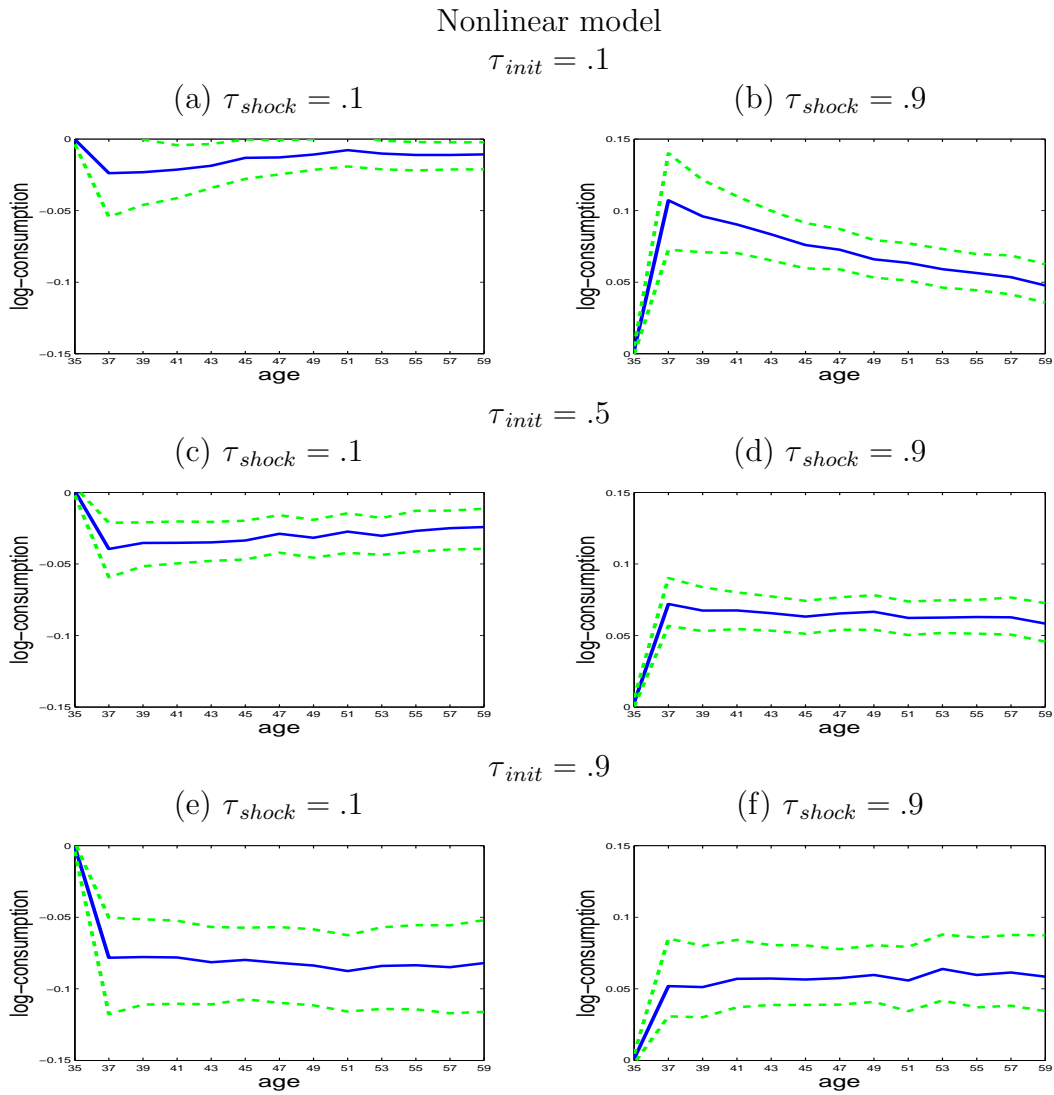
Note: See notes to Figure 7. Point estimates and pointwise 95% confidence bands from re-centered nonparametric bootstrap, clustered at the household level. 300 replications.

Figure S27: Impulse responses, earnings, 95% uniform confidence bands



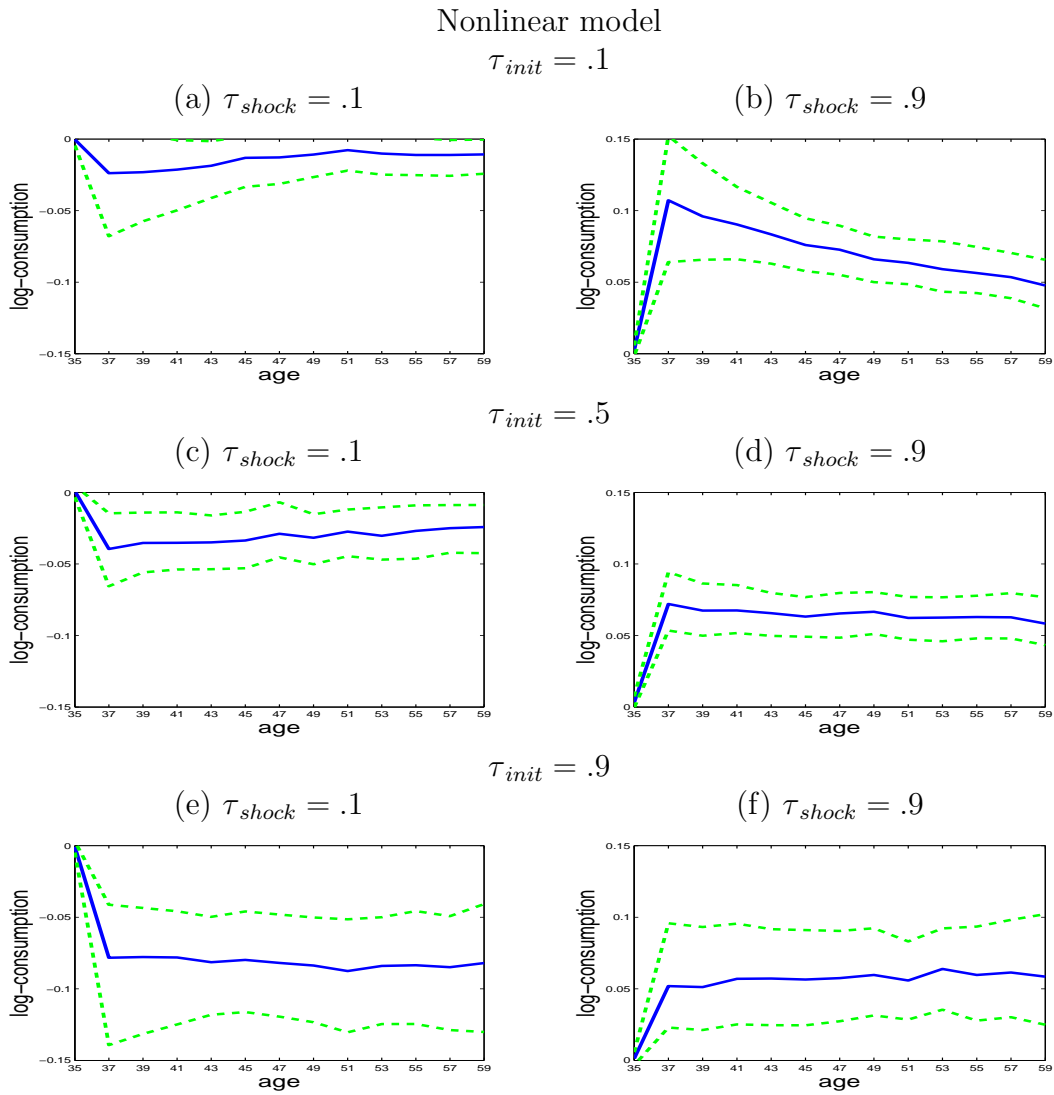
Note: See notes to Figure 7. Point estimates and uniform 95% confidence bands from re-centered nonparametric bootstrap, clustered at the household level. 300 replications.

Figure S28: Impulse responses, consumption, 95% pointwise confidence bands



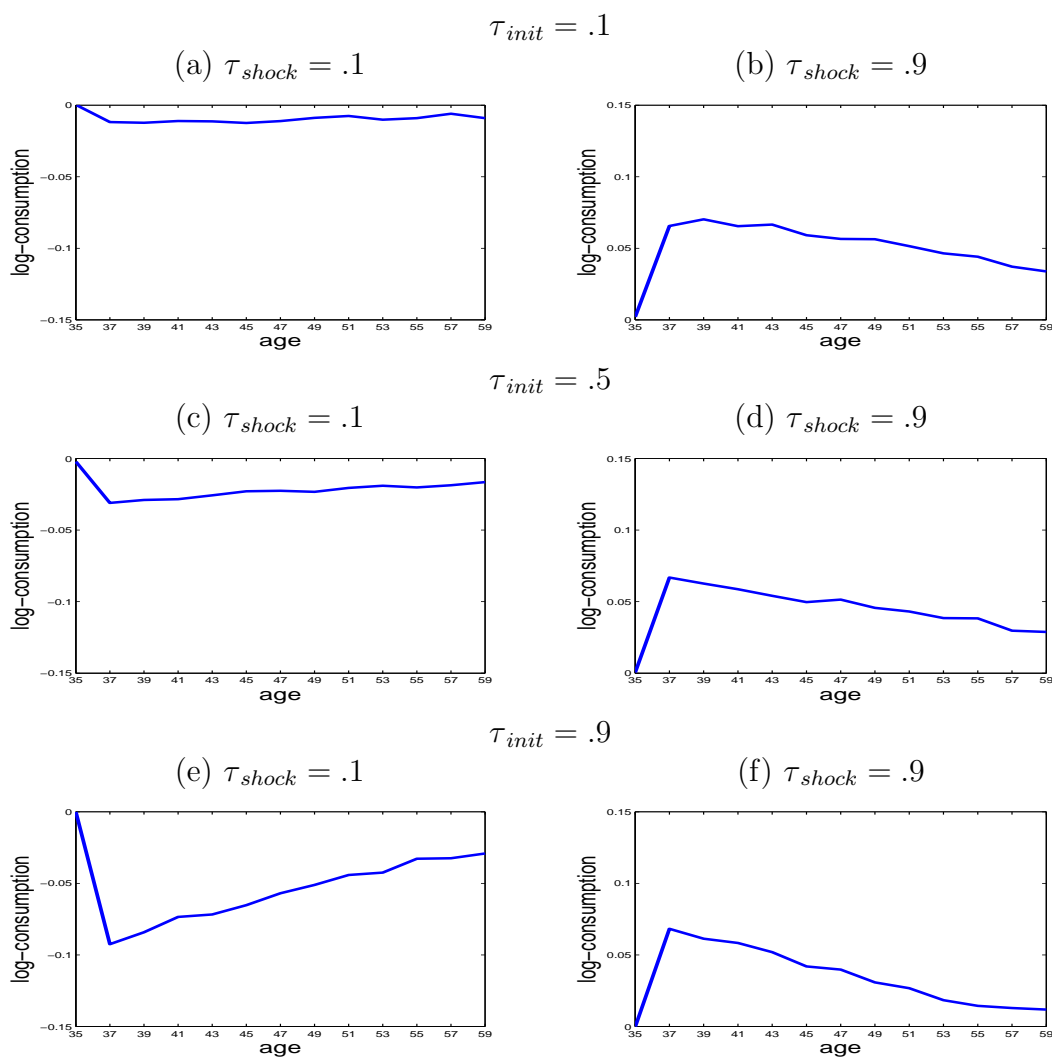
Note: See notes to Figure 8. Linear assets accumulation rule (γ), $r = 3\%$. $a_{it} \geq 0$. Point estimates and pointwise 95% confidence bands from re-centered nonparametric bootstrap, clustered at the household level. 300 replications.

Figure S29: Impulse responses, consumption, 95% uniform confidence bands



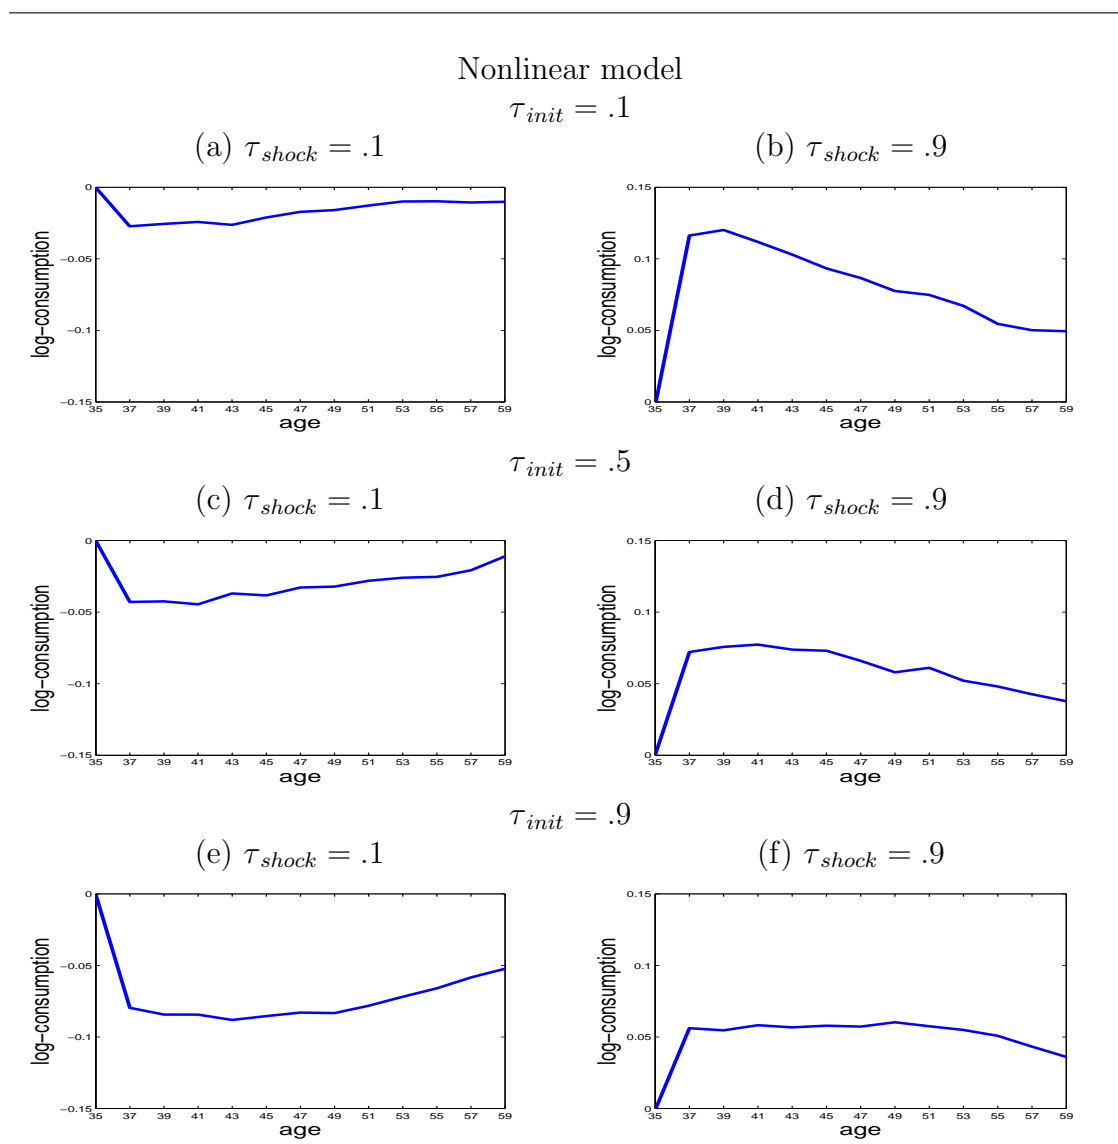
Note: See notes to Figure 8. Linear assets accumulation rule (γ), $r = 3\%$. $a_{it} \geq 0$. Point estimates and uniform 95% confidence bands from re-centered nonparametric bootstrap, clustered at the household level. 300 replications.

Figure S30: Impulse responses, consumption, model with household unobserved heterogeneity in consumption



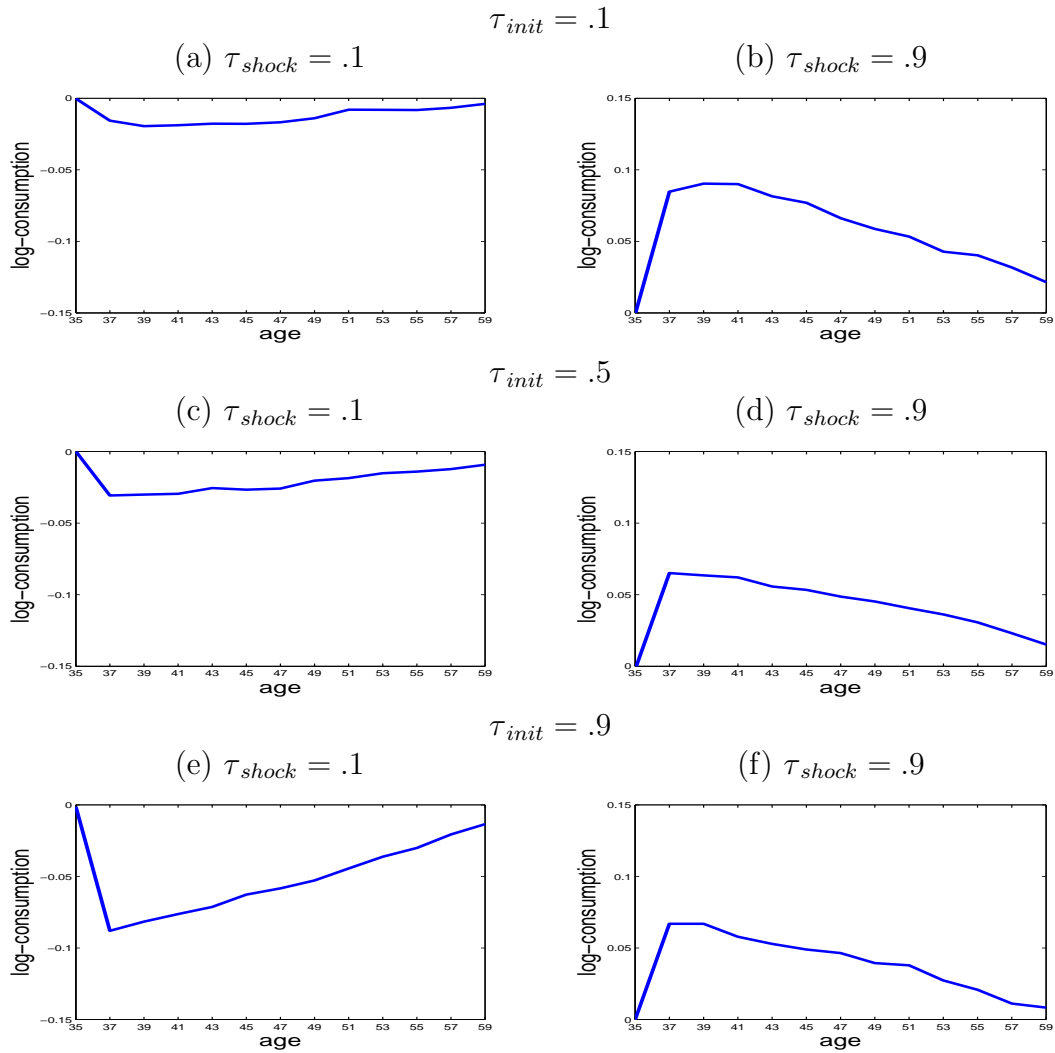
Note: See notes to Figure 8. Nonlinear model with household unobserved heterogeneity in consumption. Linear assets accumulation rule (7), $r = 3\%$. $a_{it} \geq 0$.

Figure S31: Impulse responses, consumption, model without household unobserved heterogeneity in consumption, nonlinear assets rule



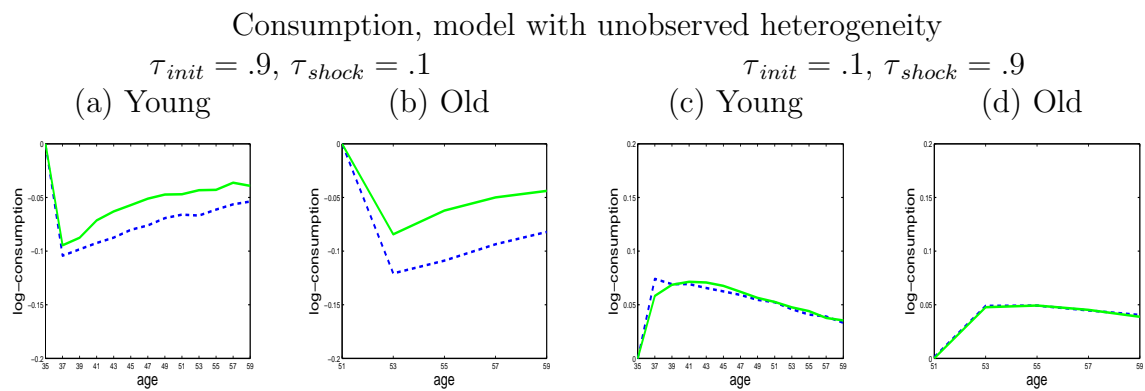
Note: See the notes to Figure 8. Estimated nonlinear assets rule.

Figure S32: Impulse responses, consumption, model with household unobserved heterogeneity in consumption, nonlinear assets rule




Note: See notes to Figure S31. Nonlinear model with household unobserved heterogeneity in consumption. Estimated nonlinear assets rule.

Figure S33: Impulse responses by age and initial assets




Note: See notes to Figures 8 and S30. Initial assets at age 35 (for “young” households) or 51 (for “old” households) are at percentile .10 (dashed curves) and .90 (solid curves). Linear assets accumulation rule (7), $r = 3\%$. $a_{it} \geq 0$. In the simulation of the model with unobserved heterogeneity ξ_i is set to zero.

Figure S34: Impulse responses by age and initial assets, nonlinear assets rule



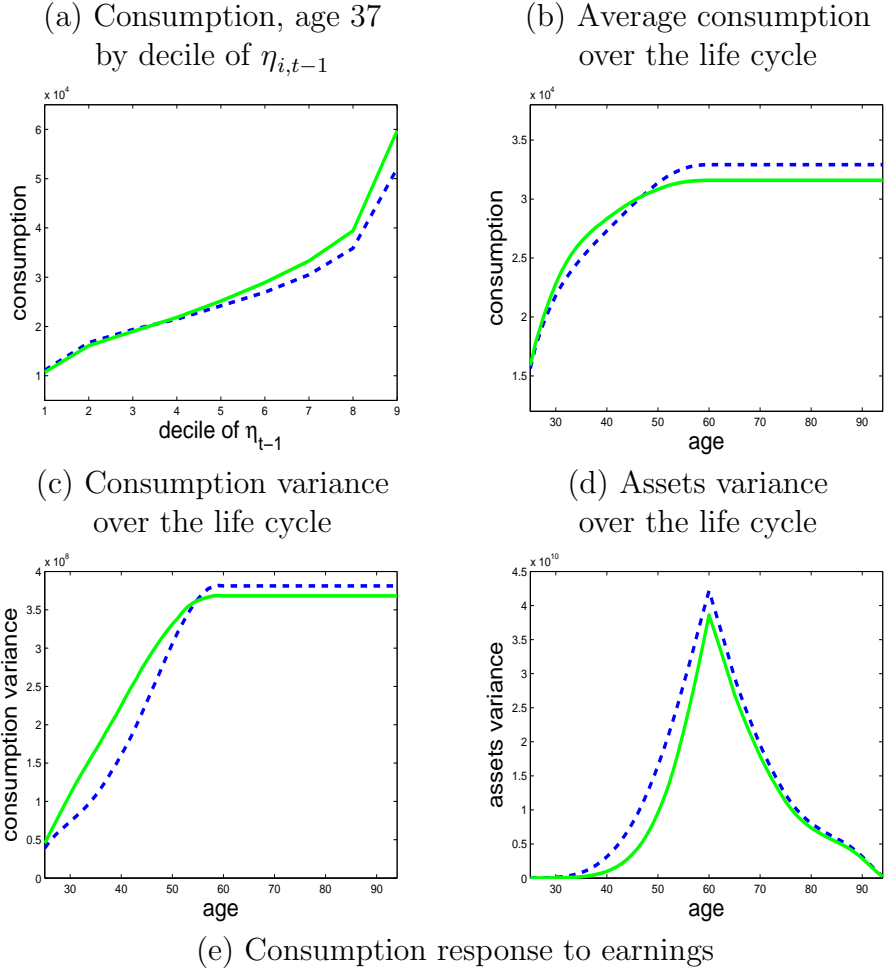
Note: See notes to Figure S31. Initial assets at age 35 (for “young” households) or 51 (for “old” households) are at percentile .10 (dashed curves) and .90 (solid curves). Estimated nonlinear assets rule. In the simulation of the model with unobserved heterogeneity ξ_i is set to zero.

Figure S35: Simulation exercise




Notes: In the top four panels, dashed is based on the nonlinear earnings process (S6)-(S7), solid is based on the canonical earnings process (3). Panel (e): estimate of the average derivative of the conditional mean of log-consumption with respect to log-earnings, given earnings, assets and age, evaluated at values of assets and age that corresponds to their τ_{assets} and τ_{age} percentiles, and averaged over the earnings values. Quantile regression on polynomials, see Section 6 for a description.

Figure S36: Simulation exercise, no borrowing



Notes: In the top four panels, dashed is based on the nonlinear earnings process (S6)-(S7), solid is based on the canonical earnings process (3). Panel (e): estimate of the average derivative of the conditional mean of log-consumption with respect to log-earnings, given earnings, assets and age, evaluated at values of assets and age that corresponds to their τ_{assets} and τ_{age} percentiles, and averaged over the earnings values. Assets are constrained to be non-negative.

Figure S37: Simulations based on the estimated nonlinear earnings model



Notes: In the top four panels, dashed is based on the nonlinear quantile-based earnings process estimated on the PSID, solid is based on a comparable canonical earnings process. Panel (e): estimate of the average derivative of the conditional mean of log-consumption with respect to log-earnings, given earnings, assets and age, evaluated at values of assets and age that corresponds to their τ_{assets} and τ_{age} percentiles, and averaged over the earnings values.