

Earnings and consumption dynamics: a nonlinear panel data framework

Manuel Arellano
Richard Blundell
Stéphane Bonhomme

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP53/15

Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework*

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

September 4, 2015

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 and the two earnings components. We establish the nonparametric identification of the nonlinear earnings process and the consumption policy rule. Exploiting the enhanced consumption and asset data in recent waves of the Panel Study of Income Dynamics, we find nonlinear persistence and conditional skewness to be key features of the earnings process. We show that the impact of earnings shocks varies substantially across earnings histories, and that this nonlinearity drives heterogeneous consumption responses. The transmission of shocks is found to vary systematically with assets.

JEL CODE: C23, D31, D91.

KEYWORDS: Earnings dynamics, consumption, 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 for comments from participants in regional meetings and seminars, especially to Xiaohong Chen, Jordi Gali, Fatih Guvenen, Lars Hansen, Jim Heckman, Yingyao Hu, Josep Pijoan, Enrique Sentana, and Kjetil Storesletten. We would also like to thank 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, 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. 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.²

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 to study identification and estimation using standard covariance techniques. However, by definition linear models rule out nonlinear transmission of shocks, and nonlinearities in income dynamics are likely to have a first-order impact on consumption 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 calibrate 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.³ Another approach is to estimate the degree of “partial insurance” from the data without precisely specifying the insurance mechanisms, see Blundell, Pistaferri and Preston (2008) for example. Linear approximations to equilibrium conditions from the optimization problem deliver tractable estimating equations. However, linear approxima-

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

²Golosov and Tsyvinski (2014) 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.* (2013) show that deviations from log normality can have serious repercussions for capital and labor taxation.

³Interesting recent exceptions are Heathcote, Storesletten and Violante (2014) and the semi-structural approach in Alan, Browning and Ejrnaes (2014).

tions 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 new framework to study the nonlinear relationship between shocks to household earnings and consumption over the life cycle. Our first contribution is to build and estimate a nonlinear earnings process. In our framework, log-earnings are the sum of a general Markovian persistent component and a transitory innovation. Our modeling approach allows to capture the intuition that, unlike in linear models, different shocks may be associated with different persistence. 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. In other words, the future persistence of a current shock depends on future shocks. For example, our framework allows for “unusual” shocks to wipe out the memory of past shocks. Moreover, in our model the densities of persistent and transitory income components are nonparametric and age-specific.

Allowing for nonlinear persistence, and more generally for flexible models of conditional earnings distributions given past earnings, has both theoretical and empirical appeal. Job ladder models suggest that 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). 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 for example large, negative “unusual” income shocks, which not only have a direct effect but also cancel out the persistence of a good income history. 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. The notion of “micro disasters” is also interestingly related to Castañeda, Díaz-Giménez and Ríos-Rull (2003), who find that allowing for a substantial probability of downward risk for high-income

households may help explain wealth inequality.⁴ Such features are prominent in the empirical results that we report in this paper. They are also consistent with some recent results independently obtained using administrative tax records (see Guvenen *et al.*, 2015, and references therein), 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 in this paper.

Our second contribution is to develop an estimation framework to assess how consumption responds to earnings shocks. In the baseline analysis we model the consumption policy rule as an age-dependent nonlinear function of assets and the 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 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 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 flexible modeling approach allows to capture 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. We also show how to extend the baseline specification to allow for household heterogeneity, and for advance information on earnings shocks and habits in consumption. A virtue of our framework is its ability to produce new empirical quantities, such as non-parametric marginal propensities to consume, that narrow the gap between policy-relevant evidence and structural modeling.

In contrast to linear models, our nonlinear model of earnings and consumption cannot be studied using standard techniques. New econometric methods are needed, beyond the traditional covariance methods that have dominated the literature. As a result, a large part of the paper is devoted to the econometric analysis. Nonparametric identification can be

⁴See also the recent work by Schmidt (2015), who emphasizes the asset pricing implications of income risk asymmetries.

established in our setup by building on a recent literature on nonlinear models with latent variables. Identification of the earnings process builds on Hu and Schennach (2008) and Wilhelm (2012).⁵ 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.

An important goal of this paper is to devise a tractable estimation approach. To achieve this, we combine quantile regression methods, which are well-suited to capture nonlinear effects of earnings shocks, with semiparametric methods based on series expansions in bases of functions, which are well-suited to model the dependence on conditioning variables (in our case, past earnings components and arguments of the consumption function). 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. This flexible approach builds on the methodology of Arellano and Bonhomme (2015), which we extend to a setup with time-varying latent variables. Wei and Carroll (2009) introduced a related estimation strategy in a cross-sectional context.

We take the 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 (2012), 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. 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 key driver of heterogeneous consumption responses. Earnings data show clear evidence 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. Related to this, we find that conditional log-earnings distributions are asymmetric, skewed to the right (respectively,

⁵Lochner and Shin (2014) rely on related techniques to establish identification of a different nonlinear model of earnings.

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 data. Regarding consumption, we find a high degree of insurability of shocks to the transitory and persistent earnings components. We also uncover interesting asymmetries in consumption responses to earnings shocks that hit households at different points of the income distribution. Lastly, we find that assets play an important role in the insurability of earnings shocks.

The fact that new (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.

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 baseline model and consider several extensions in Section 5. Section 6 describes our estimation strategy and the panel dataset. In Section 7 we present our empirical results. Section 8 concludes with a summary and some directions for future research.

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. Let Y_{it} be the pre-tax labor earnings of household i at age t , and let y_{it} denote $\log Y_{it}$, net of a full set of age dummies. We decompose y_{it} as follows:⁶

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

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

⁶Model (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 $y_{it} = H_t(\eta_{it}, \varepsilon_{it})$ subject to some scaling condition. Identification could then be established along the lines of Hu and Shum (2012).

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:

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

Note that, given that the PSID earnings data are recorded every other year, (2) is consistent with both first or second-order Markov assumptions at the yearly frequency. 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 the degree of dependence.

The second, *transitory* component ε_{it} is assumed to have zero mean, and 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 our 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.

In Section 5 we show how our approach can be extended to allow for a moving average ε component, provided additional time periods are available. We also show how to augment the model to allow for an unobserved time-invariant household-specific effect in addition to the two latent time-varying components η and ε .

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.⁷

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

⁷If additional information were available and the marginal distribution of the classical measurement error were known, one could recover the distribution of ε_{it} using a deconvolution argument. The estimation algorithm that we develop can be modified to deal with this case.

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. Let us consider 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. $\rho_t(\eta_{i,t-1}, \tau)$ 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 $\rho_t(\tau)$.

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 model allows particular shocks to wipe out the memory of past shocks. Just as in a job ladder model, an individual can face an increasing risk of a large fall in earnings, see Lise (2013) for example. The interaction between the shock u_{it} and the lagged persistent component $\eta_{i,t-1}$ is a key 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 capable of approximating any continuous function arbitrarily well as the number of polynomial terms increases. 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.⁸

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 $\tau \in (1/2, 1)$, $\sigma_t(\tau) = \mathbb{E} [Q_t(\eta_{i,t-1}, \tau) - Q_t(\eta_{i,t-1}, 1 - \tau)]$. For example, in the canonical model (3) with $v_{it} \sim \mathcal{N}(0, \sigma_{v_t}^2)$, we have $\sigma_t(\tau) = 2\sigma_{v_t} \Phi^{-1}(\tau)$. An analogous measure of uncertainty generated by the transitory shocks is $\sigma_{\varepsilon_t}(\tau) = F_{\varepsilon_t}^{-1}(\tau) - F_{\varepsilon_t}^{-1}(1 - \tau)$, where F_{ε_t} denotes the cumulative distribution function of ε_{it} .⁹

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 $\tau \in (1/2, 1)$,¹⁰

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

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

Preliminary evidence on nonlinear persistence. Prima-facie 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

⁸Regime-switching models that deliver asymmetric persistence are popular in the time series analysis of business cycles. See for example Evans and Watchel (1993)'s model of inflation uncertainty and Teräsvirta (1994) on smooth transition autoregressive models.

⁹The 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} .


¹⁰Similarly, 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 A 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: See Appendix A.

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 age) as outcome variables. 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 strong similarity in the patterns of the nonlinearity of household earnings in both the PSID household survey data and in the population register data from Norway. It also shows clear 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 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, as low as .3 – .4. In Section 7 we will show that our nonlinear earnings model that separates transitory shocks from the persistent component, estimated on the PSID, reproduces the nonlinear persistence patterns of Figure 1.

3 Model (II): Consumption rule

In order to motivate our empirical specification of the consumption function, we start by describing a standard life-cycle model of consumption and savings. In a simulated version of the model we outline some possible implications for consumption and asset accumulation of allowing for a nonlinear earnings specification. 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 A simple life-cycle model

We consider a simple theoretical framework where households act as single agents. Each household enters the labor market at age 25, works until 60, and dies with certainty at age 95. 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

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

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. In period 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. 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.

In each period t , the optimization problem is represented by

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

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

Figure 2: Simulation exercise



Notes: Dashed is based on the nonlinear earnings process (9)-(10); Solid is based on the canonical earnings process (3).

For 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 (“canonical”) earnings model. Even with certainty equivalent preferences, under model (1)-(2) the discounting applied to persistent shocks will be state dependent. In Appendix B we illustrate these theoretical insights in a two-period version of the model. Before describing how we empirically specify the consumption rule, we first present an illustrative simulation to outline some possible implications on consumption and assets of nonlinearity in income in this standard model.

Simulation exercise. To simulate the model we follow Kaplan and Violante (2010). 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} :

$$\eta_{it} = \rho_t(\eta_{i,t-1}, v_{it})\eta_{i,t-1} + v_{it}, \quad (9)$$

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. In the second specification, nonlinear persistence in income is approximated through a simple switching process:

$$\rho_t(\eta_{i,t-1}, v_{i,t}) = 1 - \delta \left(\mathbf{1} \{ \eta_{i,t-1} < -d_{t-1} \} \mathbf{1} \{ v_{it} > b_t \} + \mathbf{1} \{ \eta_{i,t-1} > d_{t-1} \} \mathbf{1} \{ v_{it} < -b_t \} \right), \quad (10)$$

where, at each age t , d_t is set so that $|\eta_{it}| > d_t$ with probability τ , and b_t is set so that $|v_{it}| > b_t$ with probability τ . In model (9)-(10), 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 $\tau = .15$ in every period. This simple parametric process is designed to roughly approximate the earnings process that we estimate on PSID data, see Section 7.

The simulation results are presented in Figure 2. The upper panel shows a clear 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. The lower panel shows 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. These simulation results provide further motivation for the use of a nonlinear earnings model to study consumption dynamics. In the next subsection we describe the consumption rule that we take to the data.

3.2 Deriving a consumption rule

In a life-cycle model with uncertainty such as the one outlined in the previous subsection, the consumption rule is of the form

$$C_{it} = G_t(A_{it}, \eta_{it}, \varepsilon_{it}), \quad (11)$$

for some age-dependent function G_t . We will base our empirical specification on (11). The consumption rule will be of this nonparametric form provided the state variables at time t are period- t assets and the latent earnings components. As we show in Section 5, our approach may be extended to allow for habits or advance information, through simple modifications of the vector of state variables.¹¹

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 directly estimate the nonlinear consumption rule (11). Doing so, we avoid linearized first-order conditions, and we estimate a flexible rule that is consistent with the life-cycle consumption model outlined in the previous subsection. This approach allows to document a rich set of derivative effects, thus shedding light on the patterns of consumption responses in the data.

An empirical consumption rule. Consider a cohort of households. Let c_{it} denote log-consumption net of a full set of age dummies. Similarly, let a_{it} denote assets net of age dummies. Our empirical specification is based on

$$c_{it} = g_t(a_{it}, \eta_{it}, \varepsilon_{it}, \nu_{it}), \quad t = 1, \dots, T, \quad (12)$$

where 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 $\nu' > \nu$. 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 (12) may also partly capture measurement error in consumption expenditures.

Clearly, the net assets variable a_{it} is not exogenous. In the next section we explain the stochastic assumptions under which g_t is identified.

¹¹There could also be additional borrowing constraints in each period. In that case, the nonparametric consumption rule in (11) would no longer be differentiable.

Derivative effects of persistent and transitory income. 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 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

$$\phi_t(a, \eta, \varepsilon) = \mathbb{E}\left[\frac{\partial g_t(a, \eta, \varepsilon, \nu_{it})}{\partial \eta}\right], \quad (13)$$

while the average derivative with respect to ε is

$$\psi_t(a, \eta, \varepsilon) = \mathbb{E}\left[\frac{\partial g_t(a, \eta, \varepsilon, \nu_{it})}{\partial \varepsilon}\right]. \quad (14)$$

The parameters $\phi_t(a, \eta, \varepsilon)$ and $\psi_t(a, \eta, \varepsilon)$ reflect the degree of insurability of shocks to the persistent and transitory earnings components, respectively. We will document how they vary over the life cycle, and how they depend on households' asset holdings, by reporting estimates of the following average derivative effects:

$$\bar{\phi}_t(a) = \mathbb{E}[\phi_t(a, \eta_{it}, \varepsilon_{it})], \quad \bar{\psi}_t(a) = \mathbb{E}[\psi_t(a, \eta_{it}, \varepsilon_{it})]. \quad (15)$$

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 (13), 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.

In the empirical analysis we will report finite-difference counterparts to these derivative effects (“impulse responses”), and document strong asymmetries in the effect of earnings shocks on consumption at different points of the income distribution.

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. Here we rely on techniques developed in this literature 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. 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 observation.¹² 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}) . In the following, all conditional and marginal densities are assumed to be bounded away from zero and infinity. 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 $Y_{i2} = Y_{i1} + \varepsilon_{i1}$, with Y_{i1} independent of ε_{i1} , and $[\mathcal{L}h](y_2) = \int h(y_1)f_{\varepsilon_1}(y_2 - y_1)dy_1$. \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.¹³ 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 $\int h(y_1)f_{\varepsilon_1}(y_2 - y_1)dy_1 = f_{Y_2}(y_2)$. In other words, f_{Y_1} is identified from the knowledge of f_{Y_2} and f_{ε_1} .

¹²We consider a balanced panel for simplicity but our arguments can be extended to unbalanced panels.

¹³Injectivity also holds if the zeros of the characteristic function of ε_{i1} are isolated. See Carrasco and Florens (2011) and Evdokimov and White (2012).

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*. Completeness is commonly assumed in nonparametric instrumental variables problems, see Newey and Powell (2003). While completeness is a high-level assumption, recent work provides primitive conditions for it in specific cases; see D’Haultfoeulle (2011) and Andrews (2011).

Building block for identification. To establish nonparametric identification of the earnings process, we rely on results from Hu and Schennach (2008) and Wilhelm (2012). In the context of a panel data model with measurement error, Wilhelm (2012) 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 Appendix C.

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} .¹⁴ 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.

The identification of the marginal distribution of ε_{i2} is derived under several high-level assumptions. In particular, it requires that the distributions of $(y_{i2}|y_{i1})$ and $(\eta_{i2}|y_{i1})$ both satisfy completeness conditions. This requires that η_{i1} and η_{i2} 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.

Identification of the earnings process. Returning to the earnings dynamics model (1)-(2), let now $T \geq 3$. Suppose that the conditions in Wilhelm (2012) 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 $(\varepsilon_{i2}, \varepsilon_{i3}, \dots, \varepsilon_{i,T-1})$ 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,

¹⁴Indeed, $f(y_1, y_2, y_3|\eta_2) = f(y_1|\eta_2)f(y_2|\eta_2, y_1)f(y_3|\eta_2, y_2, y_1) = f(y_1|\eta_2)f(y_2|\eta_2)f(y_3|\eta_2)$.

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 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.¹⁵

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 will impose time-stationary restrictions, and pool different cohorts of households together in order to identify the distributions of η 's and ε 's at all ages, see Section 6.

4.2 Consumption rule

Let us now turn to the identification of the consumption rule (12). 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 (12) is independent of $\eta_{i1}, (u_{is}, \varepsilon_{is})$ 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*) in Assumption 1 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 also holds in such a model when the interest rate r_t is time-varying and known

¹⁵Indeed 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, as 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 .

to households. More generally, the assumption allows the latent components of earnings η_{it} and ε_{it} to affect asset holdings separately.

Lastly, part *iii*) in Assumption 1 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 Section 5.

The identification argument proceeds in a sequential way, starting with period 1.

First period's assets. Let us start by analyzing period $t = 1$. Letting $y_i = (y_{i1}, \dots, y_{iT})$, we have

$$f(a_1|y) = \int f(a_1|\eta_1)f(\eta_1|y)d\eta_1, \quad (16)$$

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 (16) as

$$f(a_1|y) = \mathbb{E} [f(a_1|\eta_{i1}) | y_i = y], \quad (17)$$

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 Subsection 4.1) is complete, the density $f(a_1|\eta_1)$ is identified from (17). We will return below to the requirement that $f(\eta_1|y)$ be complete. In 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. Note also that, under bounded completeness, the density $f(a_1, \eta_1|y) = f(a_1|\eta_1)f(\eta_1|y)$ is identified.

First period's consumption. We have, using the consumption rule and Assumption 1*iii*),

$$f(c_1|a_1, y) = \int f(c_1|a_1, \eta_1, y_1)f(\eta_1|a_1, y)d\eta_1, \quad (18)$$

or equivalently

$$f(c_1|a_1, y) = \mathbb{E} [f(c_1|a_{i1}, \eta_{i1}, y_{i1}) | a_{i1} = a_1, y_i = y], \quad (19)$$

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),¹⁶ the

¹⁶Here 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.

densities $f(c_1|a_1, \eta_1, y_1)$ and $f(c_1, \eta_1|a_1, y)$ are thus identified.

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

$$f(a_2|c_1, a_1, y) = \int f(a_2|c_1, a_1, \eta_1, y_1)f(\eta_1|c_1, a_1, y)d\eta_1, \quad (20)$$

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 the previous paragraph) is complete in (y_{i2}, \dots, y_{iT}) .

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

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

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

$$f(c_2|a_2, c_1, a_1, y) = \int f(c_2|a_2, \eta_2, y_2)f(\eta_2|a_2, c_1, a_1, y)d\eta_2. \quad (21)$$

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. An intuitive explanation for the identification argument comes from the link to the nonparametric instrumental variables (NPIV) literature, see for example Newey and Powell (2003). In period 1, for a fixed a_1 , (17) 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 (19). 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). In subsequent periods, lagged consumption and assets are used as instruments, together with lags and leads

of earnings. Again, this strategy is a familiar one in linear models. Here we generalize it to deal with nonlinear, nonparametric models of earnings, consumption and assets.

The identification arguments depend on (bounded) 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 (19). 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} .¹⁷ As in the identification of the earnings process, identification of the consumption rule thus relies on η 's being dependent over time.

5 Extensions

In this section we introduce several extensions of the baseline model. We start by showing how to incorporate household unobserved heterogeneity. 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.

5.1 Household heterogeneity

In the baseline model, households differ *ex-ante* in their earnings due to heterogeneous initial conditions η_{i1} and level of assets a_{i1} . In contrast, the consumption rule is fully homogeneous. As accounting for unobserved heterogeneity in preferences or discounting, for example, may be empirically important, we now develop an extension of the model that allows for a household-specific effect ξ_i .¹⁸ With unobserved heterogeneity the consumption

¹⁷Denoting as ϕ the standard Gaussian density, taking Fourier transforms (for a fixed y_1 value) in

$$\int_{-\infty}^{\infty} g(\eta_1, y_1) \frac{1}{\sigma} \phi\left(\frac{\eta_1 - \alpha y_1 - \beta y_2}{\sigma}\right) d\eta_1 = 0$$

yields $g(\cdot, y_1) = 0$, provided $\beta \neq 0$.

¹⁸Heterogeneity in discount factors is also a popular mechanism in quantitative macro models to generate realistic wealth inequality. See for example Krusell and Smith (1998), and Krueger, Mitman and Perri (2015).

rule takes the form

$$c_{it} = g_t(a_{it}, \eta_{it}, \varepsilon_{it}, \xi_i, \nu_{it}), \quad t = 1, \dots, T. \quad (22)$$

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.

We make the following assumption.

Assumption 2

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 (22) 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 .

In a similar spirit as Assumption 1, Assumption 2 leaves the distribution of $(\xi_i, \eta_{i1}, a_{i1})$ unrestricted. Therefore, ξ_i is treated as a “fixed effect”.

The identification strategy proceeds in two steps. First we have, by Assumption 2*i*) 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 2,

$$f(c^3, a^3|\eta^3, y^3) = \int f(a_1, c_1, a_2|\eta_1, y_1, \xi)f(c_2, a_3|a_2, \eta_2, y_2, \xi)f(c_3|a_3, \eta_3, y_3, \xi)f(\xi|\eta^3, y^3)d\xi. \quad (23)$$

For fixed (a^3, η^3, y^3) , equation (23) 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 non-parametrically 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 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 .¹⁹ 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 (see Section 6).

Unobserved heterogeneity in earnings. It is possible to allow for unobserved heterogeneity in earnings as well, 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 :

$$\eta_{it} = Q_t(\eta_{i,t-1}, \zeta_i, u_{it}), \quad (24)$$

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 , (24) 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

$$y_{it} = \eta_{it} + \zeta_i + \varepsilon_{it}, \quad (25)$$

where $\eta_{it} + \zeta_i = Q_t(\eta_{i,t-1}, u_{it}) + \zeta_i$ is first-order Markov conditional on ζ_i .

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) .²⁰ 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

¹⁹Arellano and Bonhomme (2015) apply Hu and Schennach (2008)'s results to a class of nonlinear panel data models.

²⁰Indeed,

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

(2012), under a suitable scaling condition. A scaling condition is implicit in equation (25), which is the model we implement.

5.2 Additional extensions

Here we briefly consider extensions of the setup to allow for dependence in the ε 's, advance information, and habit formation. We focus on identification. The estimation strategy outlined in the next section can be modified to handle each of these extensions.

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, 2012) can be applied.

Advance information. If households have advance information about future earnings shocks, the consumption rule (12) 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

$$c_{it} = g_t(a_{it}, \eta_{it}, \eta_{i,t+1}, \varepsilon_{it}, \nu_{it}), \quad t = 1, \dots, T - 1. \quad (26)$$

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.

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

$$c_{it} = g_t(c_{i,t-1}, a_{it}, \eta_{it}, \varepsilon_{it}, \nu_{it}), \quad t = 2, \dots, T. \quad (27)$$

Identification can be shown under similar conditions as in Section 4. For example, in the second period equation (21) 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 that it is 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” as it affects period- t consumption directly.

6 Data and estimation strategy

6.1 Data

Panel data on 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 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} .²¹

²¹Removing 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, specially 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, and child care. 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 (2012) and impute rent expenditures for home owners.²² 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 (12).

To select the sample we follow Blundell *et al.* (2012) 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 A for further details. In the analysis we focus on a balanced subsample of $N = 792$ households.

Table 1 shows mean total earnings, consumption and asset holdings, by year. Compared to Blundell *et al.* (2012), households in our balanced sample have higher assets, and to a less extent higher earnings and consumption. We also can see 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 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.

²²Note 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.

Table 1: 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
	Assets					
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.
Source: See Appendix A.

6.2 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

$$\begin{aligned} Q_t(\eta_{i,t-1}, \tau) &= Q(\eta_{i,t-1}, age_{it}, \tau) \\ &= \sum_{k=0}^K a_k^Q(\tau) \varphi_k(\eta_{i,t-1}, age_{it}). \end{aligned} \quad (28)$$

In practice we use lower-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 series quantile model (28) 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 that we consider. The functional form in (28) does not enforce monotonicity in τ but our estimation method will produce an automatic rearrangement of quantiles if needed.

Note that the identification argument of Section 4.1 allows 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.

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

$$\begin{aligned} g_t(a_{it}, \eta_{it}, \varepsilon_{it}, \tau) &= g(a_{it}, \eta_{it}, \varepsilon_{it}, age_{it}, \tau) \\ &= \sum_{k=1}^K b_k^g \tilde{\varphi}_k(a_{it}, \eta_{it}, \varepsilon_{it}, age_{it}) + b_0^g(\tau), \end{aligned} \quad (29)$$

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

Equation (29) is a nonlinear regression model. In contrast with (28), 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.

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

$$Q_a(\eta_{i1}, age_{i1}, \tau) = \sum_{k=0}^K b_k^a(\tau) \tilde{\varphi}_k(\eta_{i1}, age_{i1}), \quad (30)$$

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, as follows:

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

where

$$\begin{aligned} h_t(a_{i,t-1}, c_{i,t-1}, y_{i,t-1}, \eta_{i,t-1}, \tau) &= h(a_{i,t-1}, c_{i,t-1}, y_{i,t-1}, \eta_{i,t-1}, age_{it}, \tau) \\ &= \sum_{k=1}^K b_k^h \tilde{\varphi}_k(a_{i,t-1}, c_{i,t-1}, y_{i,t-1}, \eta_{i,t-1}, age_{it}) + b_0^h(\tau), \end{aligned} \quad (31)$$

for some K and $\tilde{\varphi}_k$.²³

²³In a previous version of the paper we estimated the model imposing that $\eta_{i,t-1}$ does not enter (31), which is still consistent with the budget constraint (7) and avoids the modeling of 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 (2015). 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 $\{a_{11}^Q, \dots, a_{KL}^Q, \lambda^Q\}$.

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).²⁴

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 $b_0^h(\tau)$.²⁵ 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.²⁶

Extensions: adding household unobserved heterogeneity. We also implement two extensions of the model that allow for household unobserved heterogeneity in consumption/assets and earnings, respectively. In the first extension we model log-consumption as

$$c_{it} = g(a_{it}, \eta_{it}, \varepsilon_{it}, age_{it}, \xi_i, \nu_{it}), \quad (32)$$

²⁴As a result, we have

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

²⁵We also estimated a version 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. We found very similar results to the ones we report below.

²⁶For 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.

which we specify similarly as in (29), with parameters \tilde{b}_k^g . As a scaling condition we impose that

$$\sum_{k=1}^K \tilde{b}_k^g \tilde{\varphi}_k(0, 0, 0, \overline{age}, \xi) = \xi, \quad \text{for all } \xi, \quad (33)$$

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

$$a_{it} = h(a_{i,t-1}, c_{i,t-1}, y_{i,t-1}, \eta_{i,t-1}, age_{it}, \xi_i, v_{it}), \quad (34)$$

with a similar specification as in (31). Lastly, we specify

$$\xi_i = q(a_{i1}, \eta_{i1}, age_{i1}, \omega_i),$$

with ω_i uniform on $(0, 1)$, using a series quantile modeling as in (30).

In the second extension we allow for an additive household-specific effect ζ_i in log-earnings and model $y_{it} = \eta_{it} + \zeta_i + \varepsilon_{it}$, where η_{it} is given by (2). Here we allow for flexible dependence between η_{i1} , ζ_i and age_{i1} through another series quantile model.

6.3 Estimation algorithm

The algorithm is an adaptation of techniques developed in Arellano and Bonhomme (2015) 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 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.

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 $\ell \in \{1, \dots, L\}$,

$$\left(\bar{a}_{0\ell}^Q, \dots, \bar{a}_{K\ell}^Q \right) = \underset{(a_{0\ell}^Q, \dots, a_{K\ell}^Q)}{\operatorname{argmin}} \sum_{t=2}^T \mathbb{E} \left[\int \rho_{\tau_\ell} \left(\eta_{it} - \sum_{k=0}^K a_{k\ell}^Q \varphi_k(\eta_{i,t-1}, age_{it}) \right) f_i(\eta_i^T; \bar{\theta}) d\eta_i^T \right], \quad (35)$$

where $\bar{a}_{k\ell}^Q$ denotes the true value of $a_{k\ell}^Q = a_k^Q(\tau_\ell)$. To see that (35) 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 (35) are satisfied at true parameter values as, by (28), 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 ,

$$(\bar{a}_{0\ell}^\varepsilon, \dots, \bar{a}_{K\ell}^\varepsilon) = \underset{(a_{0\ell}^\varepsilon, \dots, a_{K\ell}^\varepsilon)}{\operatorname{argmin}} \sum_{t=1}^T \mathbb{E} \left[\int \rho_{\tau_\ell} \left(y_{it} - \eta_{it} - \sum_{k=0}^K a_{k\ell}^\varepsilon \varphi_k(age_{it}) \right) f_i(\eta_i^T; \bar{\theta}) d\eta_i^T \right], \quad (36)$$

and, for all ℓ ,

$$(\bar{a}_{0\ell}^{\eta_1}, \dots, \bar{a}_{K\ell}^{\eta_1}) = \underset{(a_{0\ell}^{\eta_1}, \dots, a_{K\ell}^{\eta_1})}{\operatorname{argmin}} \mathbb{E} \left[\int \rho_{\tau_\ell} \left(\eta_{i1} - \sum_{k=0}^K a_{k\ell}^{\eta_1} \varphi_k(age_{i1}) \right) f_i(\eta_i^T; \bar{\theta}) d\eta_i^T \right]. \quad (37)$$

In addition to (35)-(36)-(37), the model implies other restrictions on the tail parameters λ , which are given in Appendix C. 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 Appendix C. In practice, this means that it is easy to simulate from f_i . We take advantage of this feature in our estimation algorithm.

Turning to consumption we have

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

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

$$\bar{\sigma}^2 = \frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[\int \left(c_{it} - \bar{\alpha} - \sum_{k=1}^K \bar{b}_k^g \tilde{\varphi}_k(a_{it}, \eta_{it}, y_{it} - \eta_{it}, age_{it}) \right)^2 g_i(\eta_i^T; \bar{\theta}, \bar{\mu}) d\eta_i^T \right]. \quad (38)$$

Likewise, for assets we have

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

with a similar expression for the variance of $b_0^h(v_{it})$ as in (38).

Lastly we have, for all ℓ ,

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

with additional restrictions characterizing tail parameters given in Appendix C.

Overview of the algorithm. Here we describe the estimation of the earnings parameters θ . Estimation of the consumption parameters μ is similar. The estimation algorithms are described in more detail in Appendix C.

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 and $\bar{\theta}$ denotes the true value of θ .

Our estimation algorithm is closely related to the ‘‘stochastic EM’’ algorithm of Celeux and Diebolt (1993). Stochastic 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.²⁷ 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 $\hat{\theta}^{(0)}$, we iterate the following two steps on $s = 0, 1, 2, \dots$ until convergence of the $\hat{\theta}^{(s)}$ process:

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

2. *M-step:* Compute

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

²⁷Nielsen (2000b) compares the stochastic EM algorithm with the simulated EM algorithm of McFadden and Ruud (1994), where in contrast the same underlying uniform draws are re-used in every iteration.

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

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

which 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 C 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 \tilde{S} values $\hat{\theta} = \frac{1}{\tilde{S}} \sum_{s=S-\tilde{S}+1}^S \hat{\theta}^{(s)}$, and similarly for consumption-related parameters $\hat{\mu}$. 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 $\tilde{S} = S/2$. In our experiments we observed that the algorithm may sometimes get “stuck” on what appears to be a local regime of the 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 (2000a) 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 (2015) 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. We use the parametric bootstrap for inference.

Finally, 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.

See Belloni, Chernozhukov and Fernandez-Val (2011) for an analysis of inference for series quantile regression, and Arellano and Bonhomme (2015) for a consistency analysis in a panel data model closely related to the one we consider here. Studying inference in our problem as (N, K, L) jointly tend to infinity is an interesting avenue for future work.

7 Earnings and consumption in the PSID

In this section we present our empirical results. We start by describing how earnings and consumption respond to income shocks. We then report simulation exercises based on the estimated model. Lastly, we present the results of several additional specifications.

7.1 Earnings

We start by commenting on the empirical estimates of the earnings process. Figure 3 (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 shows clear evidence 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 D1 in Appendix D, where we also report similar patterns in the Norwegian data for individual income. We then estimate the earnings model,²⁸ and given the estimated parameters we simulate the model.²⁹ Figure 3 (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 evidence. For example, we estimated a simple version of the canonical earnings dynamics model (3) with a random walk component and independent transitory shocks.³⁰ Figure 3 (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 stands in sharp contrast with the data, and suggests that interaction effects between earnings shocks and past earnings components are key.


Figure 3 (d) then shows the estimated persistence of the earnings component η_{it} in model

²⁸We 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.

²⁹We 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.


³⁰Estimation is based on equally-weighted minimum distance using the covariance structure predicted by the canonical model.

Figure 3: 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 4: Densities of persistent and transitory earnings components



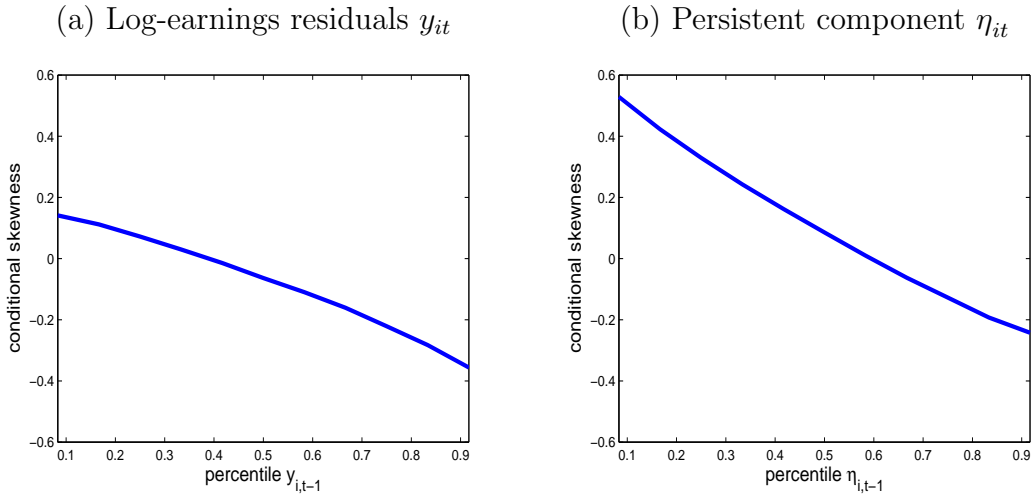
Note: Nonparametric estimates of densities based on simulated data according to the nonlinear model, using a Gaussian kernel.

(1)-(2). 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). Persistence in η 's is higher than persistence in log-earnings residuals, consistently with the fact that Figure 3 (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.

Component densities. Figure 4 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).

Lastly, in Figure 5 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 clear evidence 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 evidence of nonlinear persistence reported in Figure 3 (d): when low- η households are hit by an unusually positive shock, dependence of

Figure 5: 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 (left) and η component (right). The x-axis shows the conditioning variable, the y-axis shows the corresponding value of the conditional skewness measure.

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

These results suggest that conditional skewness is a key feature of earnings processes in the PSID. In addition, Figure D2 in Appendix D shows that the Norwegian administrative data presents a similar pattern of conditional skewness. This feature, and the related nonlinear persistence of earnings, are not easy to capture using existing models of earnings dynamics. As a notable example, models with variance dynamics such as Meghir and Pistaferri (2004) do not seem able to reproduce the nonlinear asymmetric effects apparent in Figures 3 and 5.

Inference. In Figures D3 and D4 in Appendix D we report 95% pointwise confidence bands for the persistence of log-earnings in the PSID, log-earnings simulated from our nonlinear model, and the η component, as well as conditional skewness of log-earnings and η . The bands are calculated using the parametric bootstrap. We see that the main evidence on nonlinear persistence and conditional skewness seems rather precisely estimated.

7.2 Consumption

We next turn to consumption-related parameters. Figure 6 (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} . We use tensor products of Hermite polynomials with degrees (2, 2, 1) in the estimation of the consumption rule. The derivative effects lie between .2 and .3. Moreover, the results indicate that consumption of older households, and of households with higher assets, is somewhat less correlated to variations in earnings. Figure 6 (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. While 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 (not shown). In Subsection 7.4 below we report the results of a specification allowing for household unobserved heterogeneity in consumption, where the fit to consumption dynamics is improved.



Figure 6 (c) shows estimates of the average consumption response $\bar{\phi}_t(a)$ to variations in the persistent component of earnings, see equation (15). $1 - \bar{\phi}_t(a)$ can be regarded 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 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.

In Figure 6 (d) we report estimates of $\bar{\psi}_t(a)$ to variations in the transitory component of earnings, see equation (15). The coefficient is negative, especially so for older individuals. While consistent with some of the empirical results reported in Blundell *et al.* (2012), a negative response coefficient may seem puzzling. Note, however, that ε 's are purely transitory so the η component is driving the main quantitative implications for the life-cycle evolution of earnings and consumption. The simulation exercises that we present next aim at illustrating the effects of earnings shocks on these life-cycle patterns.


Finally, in Figure D5 in Appendix D we report 95% pointwise confidence bands for $\bar{\phi}_t(a)$

Figure 6: Consumption responses to earnings shocks, by assets and age


(a) Consumption response to earnings PSID data (b) Consumption response to earnings Nonlinear model



(c) Consumption response to η_{it} Nonlinear model



(d) Consumption response to ε_{it} Nonlinear model



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. Graphs (c) and (d) show estimates of the average consumption responses $\bar{\phi}_t(a)$ and $\bar{\psi}_t(a)$ to variations in η_{it} and ε_{it} , respectively, evaluated at τ_{assets} and τ_{age} ; see equation (15).

and $\bar{\psi}_t(a)$ based on the parametric bootstrap. The evidence of insurability of shocks to the persistent earnings component is quite precisely estimated.

7.3 Model simulations

In this subsection we simulate life-cycle earnings and consumption according to our nonlinear model, and document the evolution of earnings and consumption following a persistent earnings shock. In Figure 7 we report the difference between the earnings paths of two types of households: households that are hit at 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. We report age-specific medians across 100,000 simulations of the model.³¹ At the start of the simulation (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”.

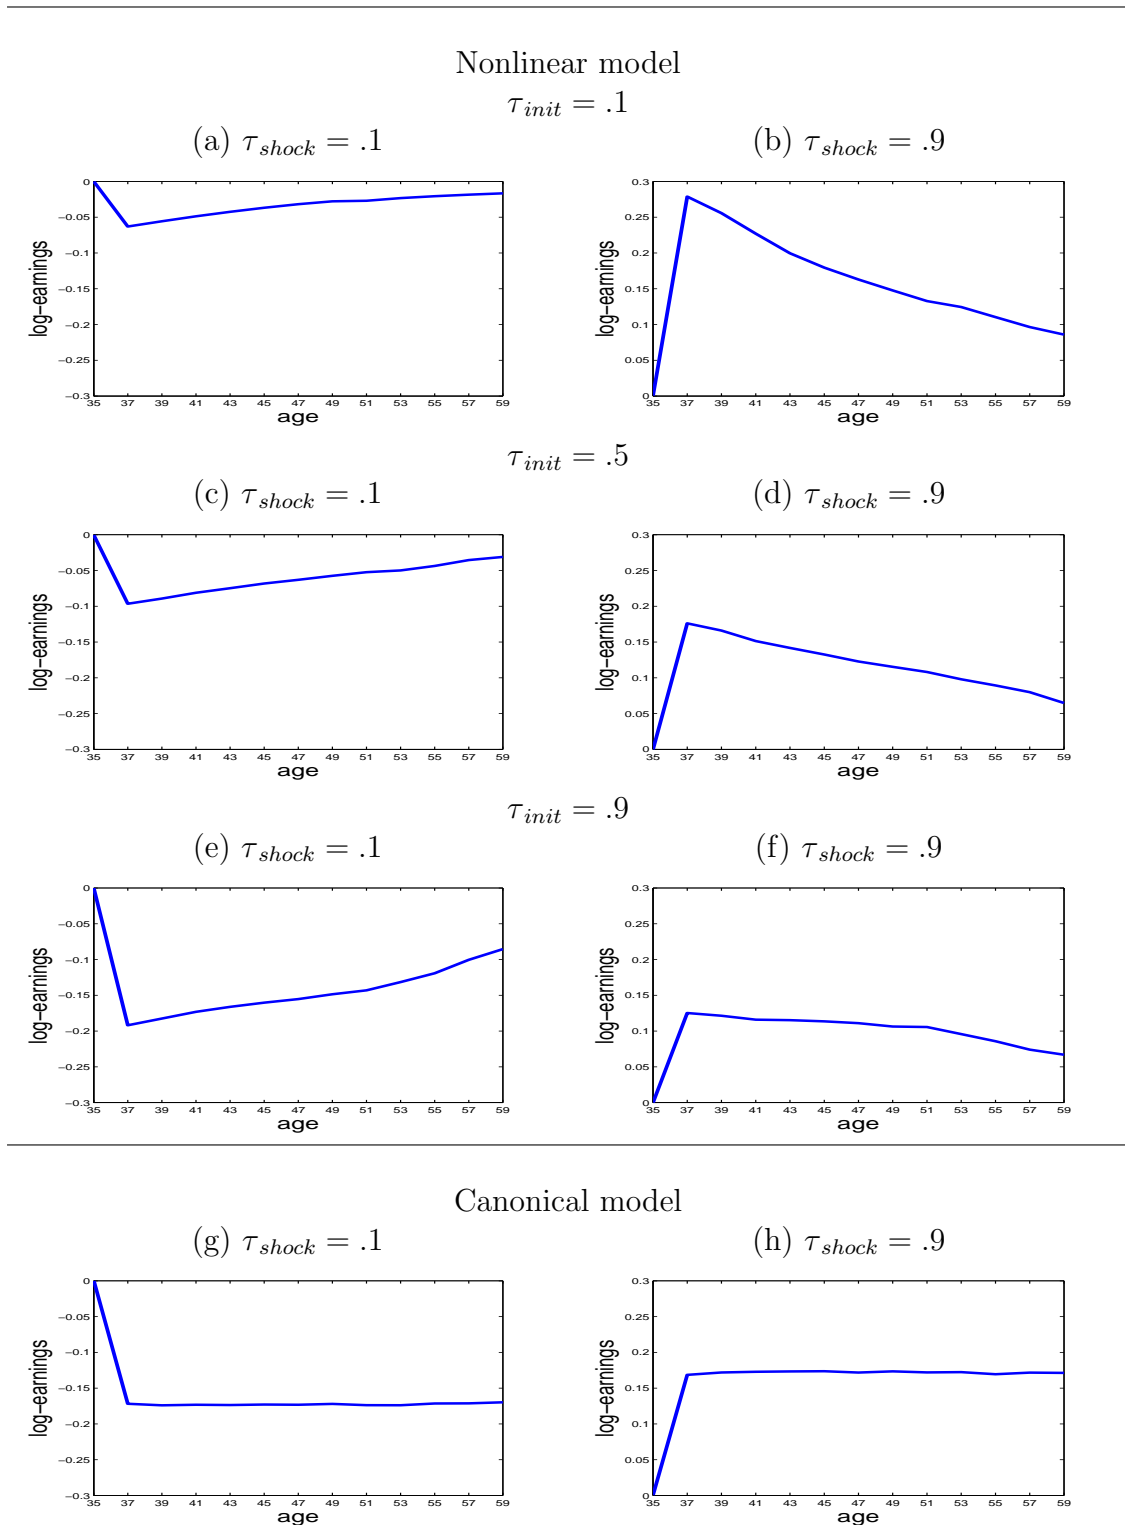
Earnings responses reported in Figure 7 are consistent with the presence of strong interaction effects between the rank in the distribution of earnings component (τ_{init}) and the sign and size of the shock to the persistent component (τ_{shock}). While a large negative shock ($\tau_{shock} = .10$) 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 strong interaction effects in the response to large positive shocks ($\tau_{shock} = .90$). Moreover, the persistence 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.

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. We see that the nonlinearities observed in the earnings response

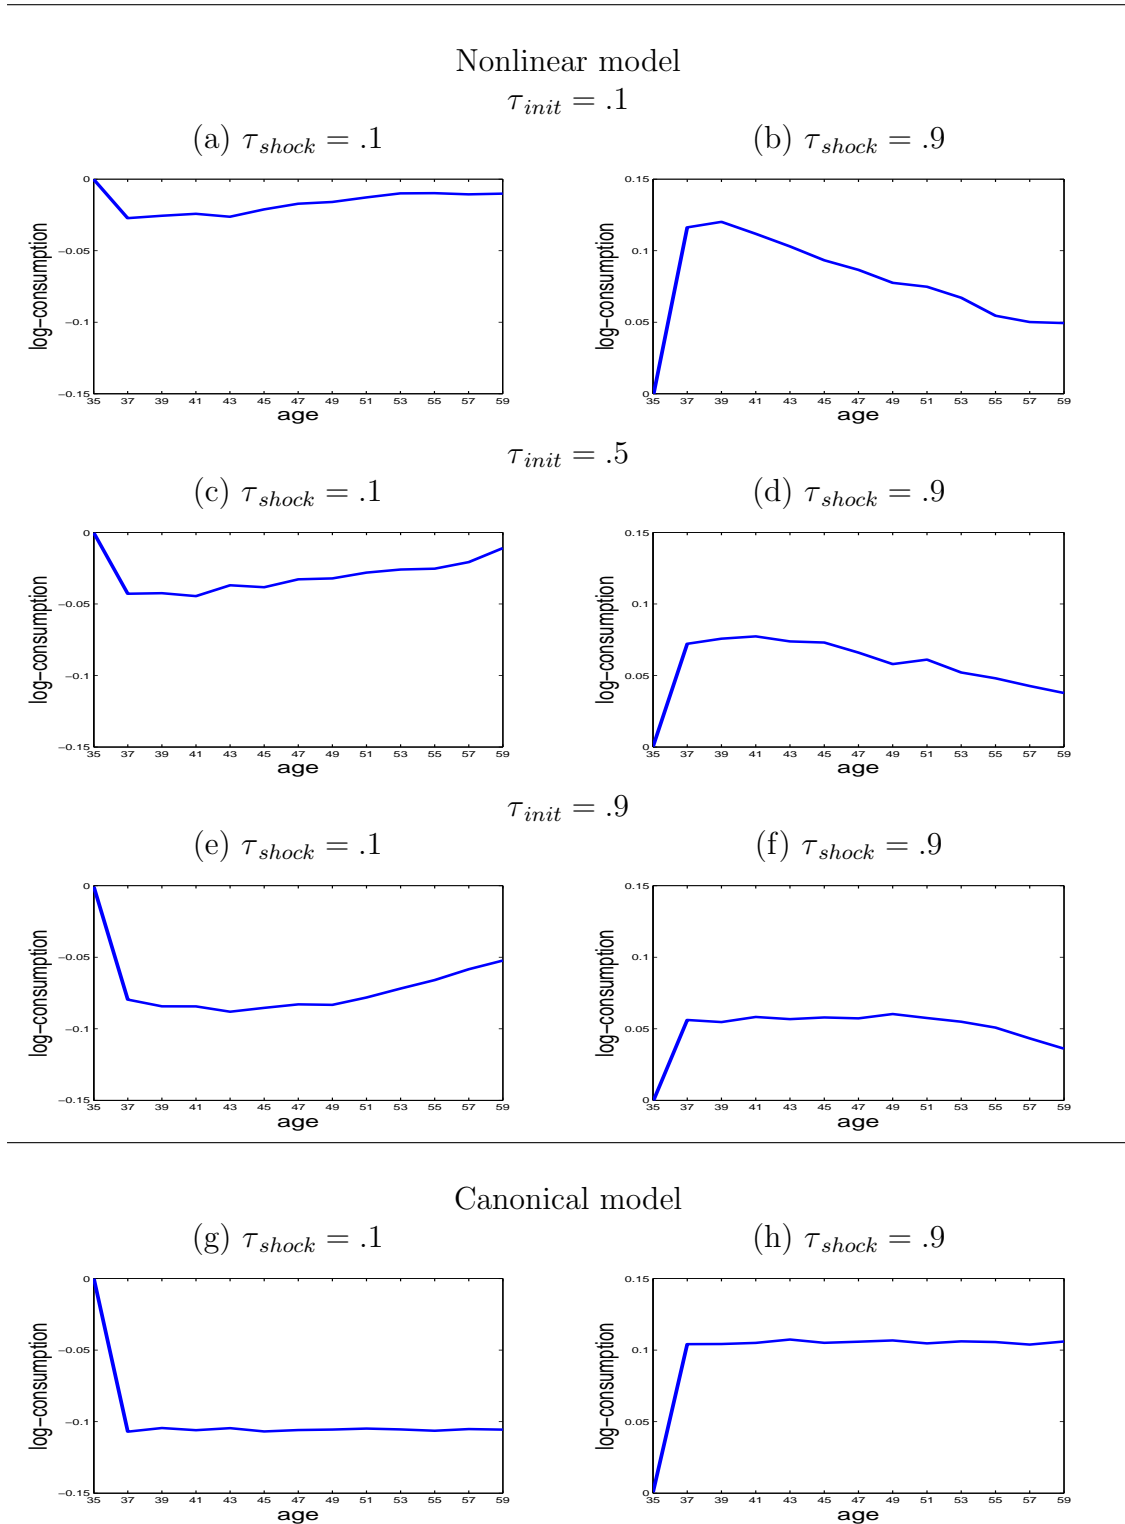
³¹We also computed age-specific quantiles across simulated households, but we do not report them here for brevity.

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.

Figure 9: Impulse responses by age and initial assets



Note: See notes to Figures 7 and 8. Initial assets at age 35 (for “young” households) or 51 (for “old” households) are at percentile .10 (dashed curves) and .90 (solid curves).

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. We also observe differences in persistence across the different scenarios.

In addition, graphs (g) and (h) of Figure 8 report results based on the canonical earnings model with a linear log-consumption rule.³² 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.

Lastly, 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 positive shocks for low earnings individuals seems to vary little with age.

Graphs (e) to (h) in Figure 9 show the consumption responses. The results suggest that, while the presence of asset holdings does not seem to affect the insurability of earnings shocks for younger households, it does seem to attenuate the consumption response for households who are hit later in the life cycle. These results are consistent with the estimates of the "partial insurance" coefficient $\bar{\phi}_t(a)$ as a function of assets and age reported in Figure 6.

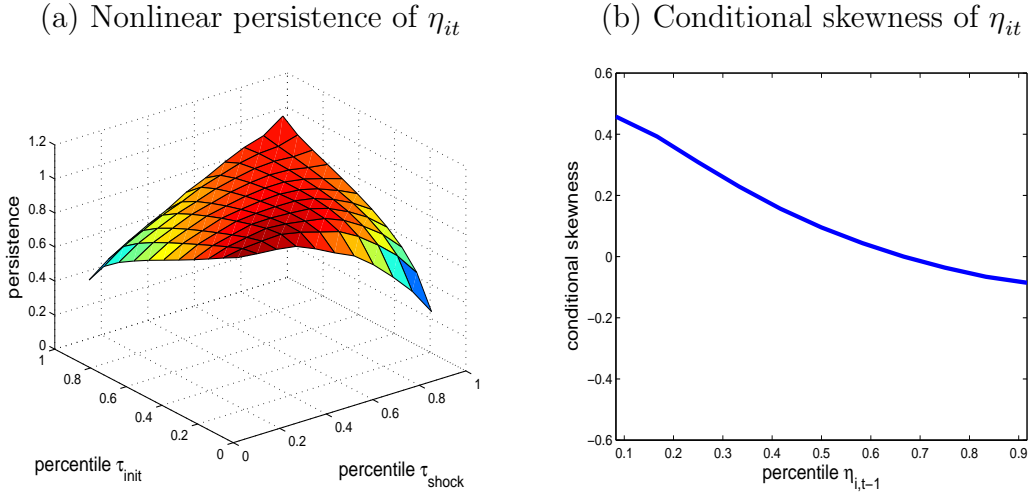
7.4 Additional empirical results

We turn to document how allowing for household unobserved heterogeneity in earnings and consumption affects the empirical results.

Household heterogeneity in earnings. In Figure 10 we report the nonlinear persistence and conditional skewness of the η component in model (25) that allows in addition for an additive household-specific effect. Graph (a) shows that, compared to Figure 3, allowing for a household effect reduces persistence. Interestingly, the nonlinear pattern here is more pronounced than in the homogeneous case. Persistence is close to one 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. Graph (b)

³²Specifically, c_{it} is modelled 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.

Figure 10: Household heterogeneity in earnings

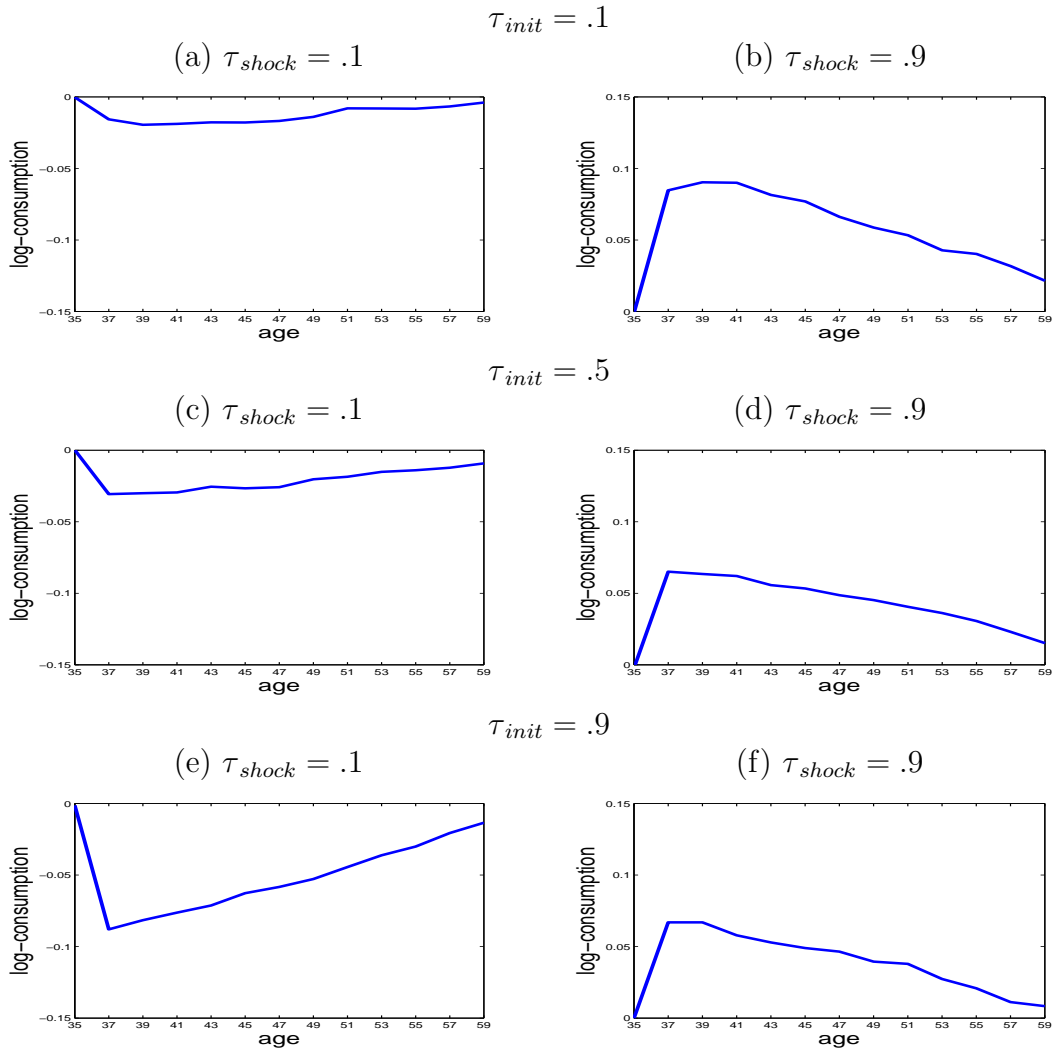


Note: (a) Estimates of the average derivative of the conditional quantile function of η_{it} on $\eta_{i,t-1}$ with respect to $\eta_{i,t-1}$, based on estimates from the nonlinear earnings model with an additive household-specific effect. (b) Conditional skewness $sk(\eta, \tau)$, see equation (6),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

shows that the η_{it} component is positively skewed when $\eta_{i,t-1}$ is low, but that it becomes increasingly less positively skewed as $\eta_{i,t-1}$ increases, similarly to the results from the baseline model.

Household unobserved heterogeneity in consumption. Estimating a model with household unobserved heterogeneity in consumption, as in (22), delivers interesting insights. In Figure 11 we report the life-cycle patterns of log-consumption of households subject to different earnings shocks at the same age (37 years old) while being at different points in the income distribution, averaged over ξ_i . Comparing the results with Figure 8 shows that allowing for unobserved heterogeneity tends to decrease the magnitude of consumption responses to variations in earnings. For example, a large positive shock ($\tau_{shock} = .90$) is associated with a 9% increase in consumption for low earnings households, compared to a 12% increase according to the baseline model without unobserved heterogeneity. We also see that effects on consumption seem to revert more quickly towards the median in the model with heterogeneity. In Appendix D, Figures D6 and D7 provide additional results based on this model.

Figure 11: Impulse responses, consumption, model with household unobserved heterogeneity in consumption



Note: See notes to Figures 7 and 8. Nonlinear model with household unobserved heterogeneity in consumption.

8 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, 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 the degree 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 models.

Our results suggest that nonlinear persistence and conditional skewness are key features of earnings processes. These features, which are present in both the PSID and in Norwegian administrative data, are not easy to capture using existing models of earnings dynamics, motivating the development 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 evidence of nonlinear persistence and conditional asymmetries in earnings, and that this nonlinearity has substantial effects on consumption. The results are robust to allowing for additional unobserved heterogeneity in earnings and consumption, and they are rather precisely estimated.

The nonlinearities observed in the earnings responses are shown to have a key role in consumption choices. For example, we found that while a large negative shock is associated with a relatively small drop in consumption for low earnings households, it is associated with a sizable drop for high-earnings households. We also identified clear differences in persistence across different demographic groups. The results suggest that, while the presence of asset holdings does not affect the insurability of earnings shocks for younger households, it does attenuate the consumption response 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.

The next step in our agenda will be to generalize the nonlinear model to allow for other states or choices, such as evolution of household size and both intensive and extensive margins of labor supply. 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 (2012) 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., M. Browning, and M. Ejrnaes (2014): “Income and Consumption: a Micro Semi-Structural Analysis with Pervasive Heterogeneity,” mimeo University of Essex, November.
- [3] Andrews, D. (2011): “Examples of L^2 -Complete and Boundedly-Complete Distributions,” unpublished manuscript.
- [4] Arellano, M., and S. Bonhomme (2015): “Nonlinear Panel Data Estimation via Quantile Regressions,” Cemmap Working Papers, CWP40/15.
- [5] Barro, R. (2006): “Rare Disasters and Asset Markets in the Twentieth Century,” *Quarterly Journal of Economics*, 121 (3), 823–866.
- [6] Belloni, A., Chernozhukov, V., and I. Fernández-Val (2011): “Conditional Quantile Processes based on Series or Many Regressors,” unpublished manuscript.
- [7] Blundell, R., X. Chen, and D. Kristensen (2007): “Semi-Nonparametric IV Estimation of Shape-Invariant Engel Curves” *Econometrica*, 75(6): 1613–1669.
- [8] 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.
- [9] Blundell, R., L. Pistaferri, and I. Preston (2008): “Consumption Inequality and Partial Insurance,” *American Economic Review*, 98(5): 1887–1921.
- [10] Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2012): “Consumption Smoothing and Family Labor Supply,” NBER Discussion Paper, November, to appear in the *American Economic Review*.

- [11] Bonhomme, S., and J. M. Robin (2010): “Generalized Nonparametric Deconvolution with an Application to Earnings Dynamics,” *Review of Economic Studies*, 77(2), 491–533.
- [12] 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.
- [13] Botosaru, I., and Y. Sasaki (2015): “Nonparametric Heteroskedasticity in Persistent Panel Processes: An Application to Earnings Dynamics,” unpublished manuscript.
- [14] Carrasco, M., and J. P. Florens (2011): “A Spectral Method for Deconvolving a Density,” *Econometric Theory*, 27, 546–581.
- [15] 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.
- [16] Celeux, G., and J. Diebolt (1993): “Asymptotic Properties of a Stochastic EM Algorithm for Estimating Mixing Proportions,” *Comm. Statist. Stochastic Models*, 9, 599–613.
- [17] Chernozhukov, V., I. Fernández-Val, and A. Galichon (2010): “Quantile and Probability Curves without Crossing,” *Econometrica*, 78(3), 1093–1125.
- [18] Deaton, A. (1991): “Saving and Liquidity Constraints,” *Econometrica*, 59(5), 1221–1248.
- [19] 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.
- [20] D’Haultfoeulle, X. (2011): “On the Completeness Condition for Nonparametric Instrumental Problems,” *Econometric Theory*, 27, 460–471
- [21] Evans, M., and P. Wachtel (1993): “Inflation Regimes and the Sources of Inflation Uncertainty,” *Journal of Money, Credit, and Banking*, 25, 475–511.
- [22] Evdokimov, K., and H. White (2012): “Some Extensions of a Lemma of Kotlarski,” *Econometric Theory*, 4, 925–932.
- [23] Golosov, M., and A. Tsyvinski (2014): “Policy Implications of Dynamic Public Finance,” to appear in *Annual Reviews of Economics*.
- [24] Golosov, M., M. Troshkin, and A. Tsyvinski (2013): “Redistribution and social insurance,” working paper, Princeton University.
- [25] Gourinchas, P.O. and J. A. Parker (2002): “Consumption over the Life Cycle.” *Econometrica*, 70, 47-91.
- [26] Guvenen, F. (2007): “Learning your Earning: Are Labor Income Shocks Really Very Persistent?” *American Economic Review*, 97(3), 687-712.

- [27] 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.
- [28] Guvenen, F., S. Ozcan, and J. Song (2012): “The Nature of Countercyclical Income Risk,” to appear in the *Journal of Political Economy*.
- [29] Guvenen, F., and A. Smith (2014): “Inferring Labor Income Risk from Economic Choices: An Indirect Inference Approach,” *Econometrica*, November, 82(6), 2085–2129.
- [30] Hall, R., and F. Mishkin (1982): “The sensitivity of Consumption to Transitory Income: Estimates from Panel Data of Households,” *Econometrica*, 50(2): 261–81.
- [31] 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.
- [32] Horowitz, J. (2009): “Specification Testing in Nonparametric Instrumental Variables Estimation,” unpublished manuscript.
- [33] Horowitz, J. L., and M. Markatou (1996): “Semiparametric Estimation of Regression Models for Panel Data”, *Review of Economic Studies*, 63, 145–168.
- [34] 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.
- [35] Hu, Y. and S. M. Schennach (2008): “Instrumental Variable Treatment of Nonclassical Measurement Error Models,” *Econometrica*, 76, 195–216.
- [36] Hu, Y. and M. Shum (2012): “Nonparametric Identification of Dynamic Models with Unobserved State Variables,” *Journal of Econometrics*, 171, 32–44.
- [37] Huggett, M. (1993), “The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies”. *Journal of Economic Dynamics and Control*, 17, 953–969.
- [38] Jappelli, T., and L. Pistaferri (2010): “The Consumption Response to Income Changes” *Annual Reviews in Economics*, 2, 479–506.
- [39] Kaplan, G., and G. Violante (2010): “How Much Consumption Insurance Beyond Self-Insurance” *American Economic Journal*, 2(4), 53–87.
- [40] Kaplan, G., and G. Violante (2014): “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, 82(4), 1199-1239.
- [41] Kim, T.H., and H. White (2004): “On More Robust Estimation of Skewness and Kurtosis,” *Finance Research Letters*, 1, 56–73.
- [42] Koenker, R. and G. J. Bassett (1978): “Regression quantiles,” *Econometrica*, 46, 33–50.
- [43] Krueger, D., K. Mitman, and F. Perri (2015): “Macroeconomics and Heterogeneity, Including Inequality,” unpublished manuscript.

- [44] Krusell, P., and A. Smith (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106, 867–896.
- [45] Lise, J. (2013): “On-the-Job Search and Precautionary Savings,” *Review of Economic Studies*, 80(3), 1086–1113.
- [46] Lochner, L., and Y. Shin (2014): “Understanding Earnings Dynamics: Identifying and Estimating the Changing Roles of Unobserved Ability, Permanent and Transitory Shocks,” unpublished manuscript.
- [47] McFadden, D.L., and P.A. Ruud (1994): “Estimation by Simulation,” *Review of Economics and Statistics*, 76, 591–608.
- [48] Meghir, C., and L. Pistaferri, (2004), “Income Variance Dynamics and Heterogeneity,” *Econometrica*, 72(1): 1–32.
- [49] Meghir, C., and L. Pistaferri, (2011): “Earnings, Consumption and Life Cycle Choices,” *Handbook of Labor Economics*, Elsevier.
- [50] Newey, W. , and J. Powell (2003): “Instrumental Variable Estimation of Nonparametric Models,” *Econometrica*.
- [51] Nielsen, S. F. (2000a): “The Stochastic EM Algorithm: Estimation and Asymptotic Results,” *Bernoulli*, 6(3): 457–489.
- [52] Nielsen, S. F. (2000b): “On Simulated EM Algorithms,” *Journal of Econometrics*, 96: 267–292.
- [53] Rietz, Th. (1988): “The Equity Risk Premium: A Solution,” *Journal of Monetary Economics*, 22 (1), 117–131.
- [54] Schmidt, L. (2015): “Climbing and Falling Off the Ladder: Asset Pricing Implications of Labor Market Event Risk,” unpublished manuscript.
- [55] Teräsvirta, T. (1994): “Specification, Estimation, and Evaluation of Smooth Transition Autoregressive Models,” *Journal of the American Statistical Association*, 89, 208–218.
- [56] Wei, Y. and R. J. Carroll (2009): “Quantile Regression with Measurement Error,” *Journal of the American Statistical Association*, 104, 1129–1143.
- [57] Wilhelm, D. (2012): “Identification and Estimation of Nonparametric Panel Data Regressions with Measurement Error,” unpublished manuscript.

APPENDIX

A Data appendix

A.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 (2012) for further details of the sample selection.

A.2 Norwegian register data

Sample selection. We use Norwegian register data for the years 2005 and 2006 only. We select a balanced panel of households where the male head is Norwegian, resident in Norway, age 25-60 and has no income from self-employment. Our measure of household disposable income pools the individual disposable income of the spouses (if the male has a spouse). We further select only households whose disposable income exceeds one basic amount in both years. This leaves us with a balanced sample of 789,982 households. See Blundell, Graber and Mogstad (2014) for details.

Residual log-income. In each year, we regress the log of household disposable income on dummies for region, marital status, number of children, education, and a 4th order polynomial in age and the interaction of the latter two to obtain the residual income.

Quantile regression. We use an equidistant grid of 11-quantiles and a 3rd degree Hermite polynomial.

B 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,

$$C_1^{-\gamma} = \mathbb{E}_1 \left[((1+r)^2 A_1 + (1+r)Y_1 - (1+r)C_1 + Y_2)^{-\gamma} \right]. \quad (\text{B1})$$

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

$$C_1 \approx \bar{C}_1 + a\sigma + b\sigma^2 + c\sigma^3. \quad (\text{B2})$$

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

$$C_1^{-\gamma} \approx \bar{C}_1^{-\gamma} \left(1 - \frac{\gamma}{\bar{C}_1} b\sigma^2 - \frac{\gamma}{\bar{C}_1} c\sigma^3 \right). \quad (\text{B3})$$

Moreover, by (B1) and (B2),

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

from which it follows that

$$\begin{aligned} C_1^{-\gamma} \approx & \bar{C}_1^{-\gamma} \mathbb{E}_1 \left[1 + \frac{\gamma}{\bar{C}_1} (1+r)b\sigma^2 + \frac{\gamma}{\bar{C}_1} (1+r)c\sigma^3 - \frac{\gamma}{\bar{C}_1} \sigma W \right. \\ & + \frac{\gamma(\gamma+1)}{2} \left(\frac{1}{\bar{C}_1} \right)^2 \sigma^2 W^2 - \gamma(\gamma+1) \left(\frac{1}{\bar{C}_1} \right)^2 (1+r)b\sigma^3 W \\ & \left. - \frac{\gamma(\gamma+1)(\gamma+2)}{6} \left(\frac{1}{\bar{C}_1} \right)^3 \sigma^3 W^3 \right]. \quad (\text{B4}) \end{aligned}$$

Finally, equating the coefficients of σ^2 and σ^3 in (B3) and (B4), 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

$$C_1 \approx \underbrace{\frac{(1+r)X_1 + \mathbb{E}_1(Y_2)}{2+r}}_{\text{certainty equivalent}} - \underbrace{\frac{\gamma+1}{2R} \mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^2)}_{\text{precautionary (variance)}} + \underbrace{\frac{(2+r)(\gamma+1)(\gamma+2)}{6R^2} \mathbb{E}_1((Y_2 - \mathbb{E}_1(Y_2))^3)}_{\text{precautionary (skewness)}}. \quad (\text{B5})$$

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

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

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 (B5), 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.³³

C Technical Appendix

C.1 Summary of the argument in Wilhelm (2012)

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

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

³³Note 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.

Assumption C1

- (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

$$\mathbb{E}[h(y_3)|y_1 = a] \in \mathcal{R}(\mathcal{L}_{y_2|y_1}), \text{ and} \quad (\text{C6})$$

$$\mathbb{E}[y_2 h(y_3)|y_1 = a] \in \mathcal{R}(\mathcal{L}_{y_2|y_1}). \quad (\text{C7})$$

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 (C6) in the case where $\mathcal{L}_{y_2|y_1}$ is a compact operator. Part (iii) is a high-level assumption; see Wilhelm (2012) for more primitive conditions.

By Assumption C1-(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 C1-(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 C1-(i) that, almost surely in η_2 ,

$$\mathbb{E}[h(y_3)|\eta_2] = \mathbb{E}[s_1(y_2)|\eta_2]. \quad (\text{C8})$$

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

$$\eta_2 \mathbb{E}[h(y_3)|\eta_2] = \mathbb{E}[s_2(y_2)|\eta_2] \text{ a.s.} \quad (\text{C9})$$

Combining (C8) and (C9), 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 ,

$$\eta_2 \int s_1(y) f_{\varepsilon_2}(y - \eta_2) dy = \int s_2(y) f_{\varepsilon_2}(y - \eta_2) dy. \quad (\text{C10})$$

The functional equation (C10) 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 C1-(iii) we can take Fourier transforms and obtain

$$i\mathcal{F}(s_1)(u) \frac{d\psi_{\varepsilon_2}(-u)}{du} + \mathcal{F}(\tilde{s}_1)(u) \psi_{\varepsilon_2}(-u) = \mathcal{F}(s_2)(u) \psi_{\varepsilon_2}(-u), \quad (\text{C11})$$

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

Noting that $\psi_{\varepsilon_2}(0) = 1$, (C11) can be solved in closed form for $\psi_{\varepsilon_2}(\cdot)$, because $\mathcal{F}(s_1)(u) \neq 0$ for all u by Assumption C1-(iii). This shows that the characteristic function of ε_2 , and hence its distribution function, are identified.

C.2 Estimation algorithm

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

$$\bar{\lambda}_-^Q = - \frac{\sum_{t=2}^T \mathbb{E} \left[\int \mathbf{1} \left\{ \eta_{it} \leq \sum_{k=0}^K \bar{a}_{k1}^Q \varphi_k(\eta_{i,t-1}, age_{it}) \right\} f_i(\eta_i^T; \bar{\theta}) d\eta_i^T \right]}{\sum_{t=2}^T \mathbb{E} \left[\int \left(\eta_{it} - \sum_{k=0}^K \bar{a}_{k1}^Q \varphi_k(\eta_{i,t-1}, age_{it}) \right) \mathbf{1} \left\{ \eta_{it} \leq \sum_{k=0}^K \bar{a}_{k1}^Q \varphi_k(\eta_{i,t-1}, age_{it}) \right\} f_i(\eta_i^T; \bar{\theta}) d\eta_i^T \right]}, \quad (\text{C12})$$

and

$$\bar{\lambda}_+^Q = \frac{\sum_{t=2}^T \mathbb{E} \left[\int \mathbf{1} \left\{ \eta_{it} \geq \sum_{k=0}^K \bar{a}_{kL}^Q \varphi_k(\eta_{i,t-1}, age_{it}) \right\} f_i(\eta_i^T; \bar{\theta}) d\eta_i^T \right]}{\sum_{t=2}^T \mathbb{E} \left[\int \left(\eta_{it} - \sum_{k=0}^K \bar{a}_{kL}^Q \varphi_k(\eta_{i,t-1}, age_{it}) \right) \mathbf{1} \left\{ \eta_{it} \geq \sum_{k=0}^K \bar{a}_{kL}^Q \varphi_k(\eta_{i,t-1}, age_{it}) \right\} f_i(\eta_i^T; \bar{\theta}) d\eta_i^T \right]}, \quad (\text{C13})$$

with similar equations for the other tail parameters.

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

$$\begin{aligned} f(y_i^T, c_i^T, a_i^T, \eta_i^T; \theta, \mu) &= \prod_{t=1}^T f(y_{it} | \eta_{it}; \theta) \prod_{t=1}^T f(c_{it} | a_{it}, \eta_{it}, y_{it}; \mu) \prod_{t=2}^T f(a_{it} | a_{i,t-1}, y_{i,t-1}, c_{i,t-1}, \eta_{i,t-1}; \mu) \\ &\quad \times \prod_{t=2}^T f(\eta_{it} | \eta_{i,t-1}; \theta) f(a_{i1} | \eta_{i1}; \mu) f(\eta_{i1}; \theta). \end{aligned} \quad (\text{C14})$$

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

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

where $A_{it}^\varepsilon(\ell) \equiv \sum_{k=0}^K a_{k\ell}^\varepsilon \varphi_k(age_{it})$ 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, for $\ell = 1, \dots, L$,

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

and compute

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

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 $\widehat{\theta}^{(0)}$, and we select the parameter values that correspond to the highest average log-likelihood over iterations.

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

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


D Additional Results

Figure D1: Nonlinear earnings persistence in log-wages (PSID, males) and individual income (Norwegian administrative 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 wages residuals, 1999-2009; right: Norwegian administrative data, individual log-earnings residuals, years 2005-2006.


Figure D2: 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 D3: Nonlinear persistence, 95% confidence bands


(a) Earnings, PSID data



(b) Earnings, nonlinear model




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



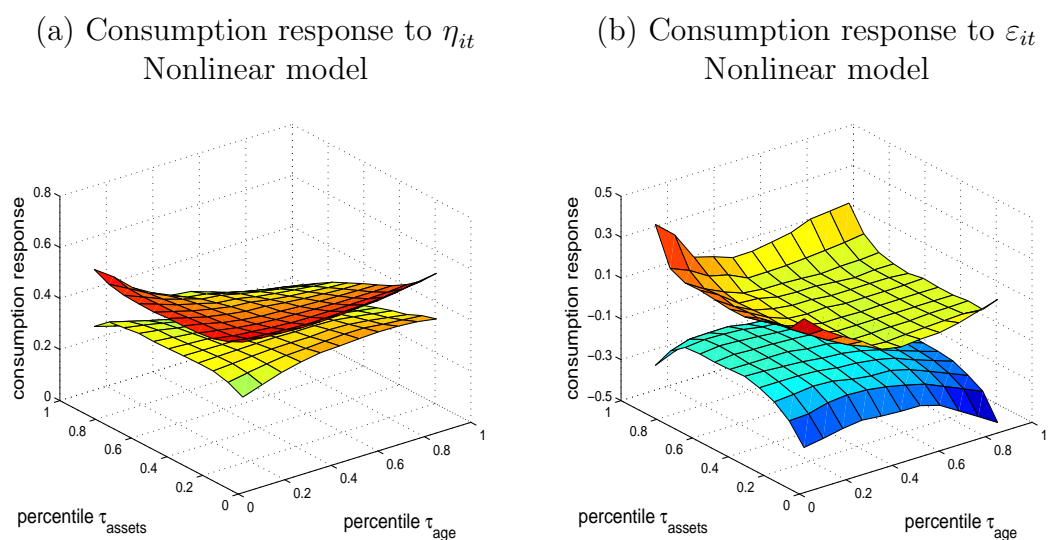
Note: See notes to Figure 3. Pointwise 95% confidence bands. Parametric bootstrap, 500 replications.

Figure D4: Conditional skewness of log-earnings residuals and η component, 95% confidence bands



Note: See notes to Figure 5. Pointwise 95% confidence bands. Parametric bootstrap, 500 replications.


Figure D5: Consumption responses to earnings shocks, by assets and age, 95% confidence bands




Note: See notes to Figure 6. Pointwise 95% confidence bands. Parametric bootstrap, 200 replications.

Figure D6: Consumption responses to earnings shocks, by assets and age, model with household-specific unobserved heterogeneity


(a) Consumption response to earnings
PSID data




(b) Consumption response to earnings
Nonlinear model



(c) Consumption response to η_{it}
Nonlinear model




(d) Consumption response to ε_{it}
Nonlinear model



Note: See the notes to Figure 6. Consumption and assets model with household-specific unobserved heterogeneity.

Figure D7: Impulse responses by age and initial assets, model with household-specific unobserved heterogeneity



Note: See notes to Figures 9. Consumption and assets model with household-specific unobserved heterogeneity. In the simulations ξ_i is set to zero.