



Nonparametric Identification and Estimation of Earnings Dynamics using a Hidden Markov Model: Evidence from the PSID

Tong Zhou

EasyChair preprints are intended for rapid dissemination of research results and are integrated with the rest of EasyChair.

April 9, 2023

Nonparametric Identification and Estimation of Earnings Dynamics using a Hidden Markov Model: Evidence from the PSID

Tong Zhou

Department of Computer Science
Johns Hopkins University
Baltimore, United States
Email: tzhou11@jhu.edu

Abstract—This paper presents a hidden Markov model designed to investigate the complex nature of earnings persistence. The proposed model assumes that the residuals of log-earnings consist of a persistent component and a transitory component, both following general Markov processes. Nonparametric identification is achieved through spectral decomposition of linear operators, and a modified stochastic EM algorithm is introduced for model estimation. Applying the framework to the Panel Study of Income Dynamics (PSID) dataset, we find that the earnings process displays nonlinear persistence, conditional skewness, and conditional kurtosis. Additionally, the transitory component is found to possess non-Gaussian properties, resulting in a significantly asymmetric distributional impact when high-earning households face negative shocks or low-earning households encounter positive shocks. Our empirical findings also reveal the presence of ARCH effects in earnings at horizons ranging from 2 to 8 years, further highlighting the complex dynamics of earnings persistence.

Index Terms—Hidden Markov Model, Panel Data, Nonparametric Identification, Modified Stochastic EM, PSID

I. INTRODUCTION

Earnings dynamics is a fascinating and important area in economics, with significant implications for understanding economic agents' consumption decisions. Macroeconomists employ life-cycle models and profiles of agents' earnings dynamics to examine their various responses within the economy, laying the foundation for the creation of sensible policies to manage business cycles. In a broader context, the nature of earnings dynamics is crucial in addressing a wide range of economic issues, including income inequality, optimal design of fiscal policies and insurance programs, economic mobility, and human capital development. As such, accurately characterizing earnings dynamics enables more effective management and a deeper understanding of a country's economy.

We utilize a parsimonious specification of the earnings process, where log-earnings consist of an unobserved persistent shock and an unobserved transitory shock. The literature on earnings process specifications varies in its focus on the distinction between these two types of shocks, a concept that can be traced back to Nobel laureate Milton Friedman's renowned permanent income hypothesis (PIH). Although there are numerous models of earnings dynamics, most tend to

concentrate on linear specifications for these two hidden components, inherently excluding the possibility of nonlinear transmission of earnings shocks.

In this paper, we introduce a new nonparametric framework to explore earnings dynamics. Both the persistent and transitory components are modeled as two generic first-order Markov processes. Apart from the first-order restriction, no further assumptions are imposed on the model. In essence, our specification establishes a hidden Markov model (HMM) with two latent state variables. Our focus is on identifying the two Markov kernels, specifically, the conditional distributions of the persistent component given its past and the conditional distribution of the transitory component given its past.

We propose a two-step stochastic EM algorithm for estimating the model. In the E-step, we use an MCMC procedure to obtain draws for the two hidden components through a likelihood-based approach. In the M-step, we perform a maximization procedure on a series of quantile regressions with imputed values for hidden covariates. The iteration continues until the expected likelihood is maximized.

II. MATERIALS AND METHODS

A. Model

1) *Setup*: In line with the conventions of earnings dynamic literature, we use $\log Y$ to represent the real (log) earnings, and it can be decomposed into the explanatory part, a persistent component U and a transitory component V . The earnings process for each household i at time t is as follows:

$$\log(Y_{it}) = \mathbf{z}'_{it}\boldsymbol{\beta} + U_{it} + V_{it}, \quad (1)$$

where \mathbf{z}_{it} is a set of observed demographics and known by agents at t . We let $y_{it} = \log(Y_{it}) - \mathbf{z}'_{it}\boldsymbol{\beta}$ denote the log of real income net of predictable individual components.

We assume both U_{it} and V_{it} follow some general unknown functions $H_t(U_{i,t-1}, \eta_{it})$ and $Q_t(V_{i,t-1}, \varepsilon_{it})$, where η_{it} and ε_{it} are assumed to follow conditional standard uniform distributions, i.e.

$$\eta_{it}|(U_{i,t-1}, U_{i,t-2}, \dots) = \text{Unif}(0, 1), \quad t = 2, \dots, T \quad (2)$$

$$\varepsilon_{it}|(V_{i,t-1}, V_{i,t-2}, \dots) = \text{Unif}(0, 1), \quad t = 2, \dots, T. \quad (3)$$

This general nonparametric setting offers greater flexibility for studying the persistence of earnings dynamics and encompasses many earnings dynamic models as special cases including the *canonical earnings dynamics models*, where the persistent component follows a unit-root process. Given that both processes are unobserved, a Bernoulli instrumental variable $\omega(C_{it})$ is required to differentiate them, where $\omega(\cdot)$ is a known transformation of consumption data C_{it} for agent i at t . Since the purpose of $\omega(C_{it})$ is merely to distinguish the two Markov kernels, it suffices for our purposes to use a logistic function, i.e. $\mathbb{P}(\omega(C_{it}) = 1|U_{it}) = 1/(1 + \exp(-\beta_0 - \beta_1 U_{it}))$. The rationale behind this setup can be found in the works of [3], [4], [5].

Putting above discussions together, we have the complete model setup

$$\begin{aligned} y_{it} &= U_{it} + V_{it} \\ U_{it} &= H_t(U_{i,t-1}, \eta_{it}) \\ V_{it} &= Q_t(V_{i,t-1}, \varepsilon_{it}) \\ \mathbb{P}(\omega(C_{it}) = 1) &= 1/(1 + \exp(-\beta_0 - \beta_1 U_{it})) \end{aligned}$$

2) *Assumptions*: We will outline the assumptions needed to identify the model. Our identification strategy relies on the powerful spectral decomposition of linear operators. A thorough overview of this approach can be found in the work of [11].

- Assumption 1.** 1) (*First-order Markov*) Both U_{it} and V_{it} follow a generic first-order Markov process;
- 2) (*Conditional uniform distribution*) Both η_{it} and ε_{it} follow conditional standard uniform distributions
- 3) (*Monotonicity*) The unknown condition quantile function $\tau \mapsto H_t(U_{i,t-1}, \tau)$ and $\tau \mapsto Q_t(V_{i,t-1}, \tau)$ are strictly increasing for $\tau \in (0, 1)$.
- 4) (*Invertibility*) The conditional distribution functions $F(U_{it}|U_{i,t-1})$ and $F(V_{it}|V_{i,t-1})$ are both invertible w.r.t. their respective arguments U_{it} and V_{it} for each i and t .

Assumption 1) states that U_{it} and V_{it} have only one-period memory of their past. This condition imposes dynamic exclusion restrictions that aid in obtaining nonparametric identifications. This assumption is also commonly made in structural economic models. Although it can be relaxed to allow for higher-order Markov process, we maintain the first-order Markovian assumption in this paper for simplicity. Assumption 2) normalizes the error terms η_{it} and ε_{it} to follow standard uniform distributions. This setup enables us to discuss consequences of shocks along the rank of $U_{i,t-1}$ and $V_{i,t-1}$. This representation also nests the canonical model of earnings dynamics as a special case where U_{it} is assumed to follow a unit-root process, i.e.

$$U_{i,t+1} = U_{it} + \nu_{i,t+1}$$

where $\nu_{i,t+1} = F^{-1}(\eta_{i,t+1})$ is the inverse function of the CDF of $\eta_{i,t+1}$. Assumption 3) guarantees that U_{it} and V_{it} have absolutely continuous distributions. Assumption 1) - 3)

combined imply that for all $\tau \in (0, 1)$, $H_t(U_{i,t-1}, \tau)$ happens to be the τ -conditional quantile of U_{it} given $U_{i,t-1}$. This relationship also holds for $Q_t(V_{i,t-1}, \tau)$. Assumption 4) is furnished to facilitate identification of the nonlinear functions H_t and Q_t . The monotonicity restriction on H_t and Q_t are necessary for the existence of their marginal densities $f(U_{it}|U_{i,t-1})$ and $f(V_{it}|V_{i,t-1})$. It is not a sufficient condition because a stronger condition of absolute continuity on the distribution function $F_{V_{it}|V_{i,t-1}}$ cannot be weakened. However, since it is rare that a distribution function is continuous but not absolutely continuous, assumption 4) can be almost equivalent to the existence of the two marginal densities.

Assumption 2 (independence). *Two random vectors $(\eta_{i2}, \dots, \eta_{iT}, U_{i1})$ and $(\varepsilon_{i2}, \dots, \varepsilon_{iT}, V_{i1})$ are statistically independent. The Bernoulli random variable $\omega(C_t)$ is independent of V_t for all t .*

This assumption suggests that the persistent process $\{U_{it}\}_{t=1}^T$ and the transitory process $\{V_{it}\}_{t=1}^T$ are statistically independent. This restriction allows for the common deconvolution technique of separating two unknown probability densities. For instance, once one of the marginal densities $f(U_{it})$ or $f(V_{it})$ is identified, the other one will also be automatically identified through the deconvolution argument.

Since our identification strategy relies on the technique of manipulating linear operators, we provide the definition of linear operator here to facilitate our later discussions. Let $\mathcal{L}^p(F_U)$ denote the collection of functions of variable U for which its p -th moment is finite, i.e. $g \in \mathcal{L}^p(F_U)$ implies

$$\|g\|_{\mathcal{L}^p} = \left(\int_{\mathcal{U}} g(u) dF_U(u) \right)^{\frac{1}{p}} < \infty, \quad (4)$$

where \mathcal{U} denotes the support of U . The definition for the space $\mathcal{L}^q(F_V)$ is similar.

Now we define a linear operator

$$\mathcal{L}_{V|U} : \mathcal{L}^p(F_U) \rightarrow \mathcal{L}^q(F_V), \quad (5)$$

where $p, q \geq 1$. Specifically, for any $g \in \mathcal{L}^p(F_U)$, we have

$$\mathcal{L}_{V|U} g = \int_{\mathcal{U}} f_{V|U}(v|u) g(u) du \in \mathcal{L}^q(F_V), \quad (6)$$

where the function $f_{V|U}$ is called the kernel of the linear operator $\mathcal{L}_{V|U}$. This expression is particularly useful when multiple linear operators are present, since we do not need to introduce new notations for each involved operator.

Assumption 3. *There exist variables Y_{it} such that*

- 1) *For any y_t can \tilde{c}_t , there exists a y_{t-1} and \tilde{c}_{t-2} and a neighborhood \mathcal{N}^r around $(y_t, \tilde{c}_{t-1}, y_{t-1}, \tilde{c}_{t-2})$ such that, for any $(y'_t, \tilde{c}'_{t-1}, y'_{t-1}, \tilde{c}'_{t-2}) \in \mathcal{N}^r$, the linear operator $\mathcal{L}_{Y_{t-2}, y'_{t-1}, \tilde{c}'_{t-2}, y'_t, \tilde{c}'_{t-1}, Y_{t+1}}$ is one-to-one.*
- 2) *For any y_t and \tilde{c}_{t-1} , the linear operator $\mathcal{L}_{Y_{t=1}|y_t, \tilde{c}_{t-1}, U_{t-1}, V_t}$ is one-to-one.*
- 3) *For any y_{t-1} and \tilde{c}_{t-2} , the linear operator $\mathcal{L}_{Y_{t-2}, y_{t-1}, \tilde{c}_{t-2}, Y_t}$ is one-to-one.*

Assumption 4. 1) The characteristic function of (U_{i1}, \dots, U_{iT}) and (V_{i1}, \dots, V_{iT}) do not vanish on the real line.

2) The characteristic function of (U_{i1}, \dots, U_{iT}) and (V_{i1}, \dots, V_{iT}) are absolutely continuous.

Assumption 4.1) is commonly made to achieve nonparametric identification (see [6]). For univariate distributions, this assumption rules out certain families of distributions, e.g., truncated normal, symmetric uniform and many discrete distributions. Assumption 4.2) is made to facilitate the deconvolution argument and also implies that the joint distributions of (U_{i1}, \dots, U_{iT}) and (V_{i1}, \dots, V_{iT}) exist.

To avoid cluttered notations, we simplify the notations by omitting the subscript i without causing confusions. In the following derivations, we define $\tilde{C}_{t-1} := \omega(C_{t-1})$.

Assumption 5 (Uniqueness of spectral decomposition). For any (Y_t, \tilde{C}_{t-1}) and any $(u_{t-1}, v_t) \neq (u'_{t-1}, v'_t)$, there exists a $(y_{t-1}, \tilde{c}_{t-2})$ and corresponding neighborhood \mathcal{N}^r satisfying Assumption 3.1), such that for some $(y'_t, \tilde{c}'_{t-1}, y'_{t-1}, \tilde{c}_{t-2}) \in \mathcal{N}^r$ with $(y'_t, \tilde{c}_{t-1}) \neq (y_t, \tilde{c}_{t-1})$ and $(y'_{t-1}, \tilde{c}_{t-2}) \neq (y_{t-1}, \tilde{c}_{t-2})$:

$$0 < k(y_t, \tilde{c}_{t-1}, y'_{t-1}, \tilde{c}_{t-2}, y_{t-1}, \tilde{c}_{t-2}, u_{t-1}, v_t) < C < \infty$$

and

$$k(y_t, \tilde{c}_{t-1}, y'_{t-1}, \tilde{c}_{t-2}, y_{t-1}, \tilde{c}_{t-2}, u_{t-1}, v_t) \neq k(y_t, \tilde{c}_{t-1}, y'_{t-1}, \tilde{c}_{t-2}, y_{t-1}, \tilde{c}_{t-2}, u'_{t-1}, v'_t)$$

where

$$k(y_t, \tilde{c}_{t-1}, y'_{t-1}, \tilde{c}_{t-2}, y_{t-1}, \tilde{c}_{t-2}, u_{t-1}, v_t) = \frac{f(y_t, \tilde{c}_{t-1} | y_{t-1}, \tilde{c}_{t-2}, u_{t-1}, v_t) f(y'_t, \tilde{c}'_{t-1} | y'_{t-1}, \tilde{c}'_{t-2}, u_{t-1}, v_t)}{f(y'_t, \tilde{c}'_{t-1} | y_{t-1}, \tilde{c}_{t-2}, u_{t-1}, v_t) f(u_{t-1}, v_t | y'_{t-1}, \tilde{c}'_{t-2}, u_{t-1}, v_t)}$$

Assumption 6 (normalization). The Markov kernels are normalized by $\mathbb{E}[U_{t+1}|U_t] = U_t$ and $\mathbb{E}[V_{t+1}|V_t] = 0$.

In the eigenfunctions $f(y_{t+1}|y_t, \tilde{c}_{t-1}, u_{t-1}, v_t)$, both u_{t-1} and v_t are unobserved and continuously distributed. Assumption 6 is made to differentiate and identify the two components.

Assumption 7 (Stationarity). For any $2 \leq t \leq T$, the Markov kernels is time-invariant, i.e.,

$$f(Y_t, \tilde{C}_{t-1}, U_{t-1}, V_t | Y_{t-1}, \tilde{C}_{t-2}, U_{t-2}, V_{t-1}) = f(Y_3, \tilde{C}_2, U_2, V_3 | Y_2, \tilde{C}_1, U_1, V_2)$$

Assumption 7 is not necessary for identification of the Markov density. It eases our derivations. From the next section, we can see that only five periods of data is sufficient for achieving nonparametric identification.

B. Identification

Based on the assumptions made in the previous section, identification can be accomplished by applying Theorem 9 from [11] in dynamic settings. This strategy can be better understood in Figure 1, where the dependence structures and dynamic exclusion restrictions can be easily visualized.




Fig. 1. Graphical illustration of earnings dynamics

We state the main identification theorem

Theorem 1 (Identification). Under Assumption 1 - Assumption 7, the density

$$f(Y_{t+1}, \tilde{C}_t, Y_t, \tilde{C}_{t-1}, Y_{t-1}, \tilde{C}_{t-2}, Y_{t-2}, \tilde{C}_{t-3})$$

for any $t \in \{4, \dots, T-1\}$ uniquely determines the densities $f(Y_t, \tilde{C}_{t-1}, U_{t-1}, V_t | Y_{t-1}, \tilde{C}_{t-2}, U_{t-2}, V_{t-1})$.

Theorem 1 implies that our interests of Markov kernels can be identified by basic probability rules, the Bayes rules and the deconvolution technique.

Corollary 1. Under Assumption 1 - Assumption 7, the Markov kernels $f_{V_t|V_{t-1}}$, $f_{U_t|U_{t-1}}$ and marginal distributions f_{U_t} and f_{V_t} are uniquely identified, for $t = 4, \dots, T-1$.

C. Estimation

We introduce a *modified stochastic EM algorithm* (MSEM) to estimate this HMM, while the stochastic EM was originally proposed by [10]. The MSEM provides a much more faster implementation of the estimation by replacing the likelihood with the objective functions of quantile regression models. The MSEM is similar to the one presented in [1]. The difference lies in the fact that their paper involves only one state variables. Specifically, for any $\tau \in (0, 1)$ we employ the following estimating equations

$$U_{it} = \sum_{k=0}^{K_1} a_k^H(\tau) \phi_k(U_{i,t-1}, \text{age}_{it})$$

$$V_{it} = \sum_{k=0}^{K_2} a_k^Q(\tau) \phi_k(V_{i,t-1}, \text{age}_{it})$$

$$U_{i1} = \sum_{k=0}^{K_3} a_k^{H_1}(\tau) \phi_k(\text{age}_{i1})$$

$$V_{i1} = \sum_{k=0}^{K_2} a_k^{Q_1}(\tau) \phi_k(\text{age}_{i1}),$$

where ϕ_k is the Hermite polynomials. We selected different orders of polynomials for the four equations to maximize the likelihood. The quantile-based estimation strategy provides a flexible specification of the Markov kernels. [2] applies this estimation strategy to estimate the smoking effects of women during pregnancy on children's birthweights.

Another advantage of using quantile-based estimation is that the original nonparametric estimation problem is reduced to estimating a finite number of parameters, i.e., the coefficients of the Hermite polynomials. We discuss U_{it} as an example: the functions $a_k^H(\tau)$ are modeled as piecewise-polynomial interpolating splines on equi-length intervals $[\tau_1, \tau_2], [\tau_2, \tau_3], \dots, [\tau_{T-1}, \tau_T]$ that partition the unit interval $(0, 1)$. In other words, we need to estimate $a_k^H(\tau)$ for each interval of τ and k . Additionally, the objective function of quantile regressions can be used as a surrogate likelihood. Since it is a convex function, the implementation can be fast. Once those a_τ^H 's are obtained, we are finished with the estimation of U_{it} .

We still take U_{it} as an example to illustrate the MSEM algorithm. We start with an initial value for the parameter vector $\hat{\theta}^{(0)}$. Each iteration follows the following two steps until convergence of the $\hat{\theta}^{(s)}$ in the s -th iteration:

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

$$\hat{\theta}^{(s+1)} = \arg \min_{\theta} \sum_{i=1}^N \sum_{m=1}^M R(y_i, U_i^{(m)}; \theta),$$

where $R(\cdot)$ is the surrogate likelihood, i.e. the objective function of the piece-wise quantile regressions. In the E-step, we use a random-walk Metropolis-Hastings algorithm for drawing $U_i^{(m)}$ in the E-step. The M-step consists of a number of quantile regressions. For instance, for each ℓ , the parameters $a_k^H(\tau_\ell)$ are updated as

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

where $\rho_\tau(u) = u(\tau - \mathbf{1}(u \leq 0))$ is the check function in standard quantile regressions and $\mathbf{1}(\cdot)$ denotes the indicator function, first introduced by [8].

[10] examined the statistical properties of the stochastic EM algorithm within a likelihood case. He provided certain conditions under which the Markov chain $\hat{\theta}^{(s)}$ is ergodic. He also outlined the asymptotic distribution of $\hat{\theta}$. [1] characterized the asymptotic distribution of $\hat{\theta}$ in a manner that aligns with our model, specifically when utilizing the surrogate likelihood during the M-step.

The MSEM algorithm can be summarized as follows

Algorithm 1: Modified Stochastic EM Algorithm

Stochastic E-Step:

for $s \leftarrow 1$ to S do

for $m \leftarrow 1$ to M do

Draw M values $u_i^{(m)} = (u_{i1}^{(m)}, \dots, u_{iT-1}^{(m)})$ from

$$f(u_i^T | y_i^T, \tilde{C}_i^T; \hat{\theta}^{(s)}) \propto \prod_{t=2}^T f_{V_{it}|V_{i,t-1}}(y_{it} - u_{it} | y_{i,t-1} - u_{i,t-1}; \hat{\theta}^{(s)}) \cdot \prod_{t=2}^T f(u_{it} | u_{i,t-1}; \hat{\theta}^{(s)}) \cdot f(u_{i1}; \hat{\theta}^{(s)}) \cdot f(v_{i1}; \hat{\theta}^{(s)}) \cdot \prod_{t=1}^T f(\tilde{C}_{it} | U_{it}, V_{it})$$

M-step:

for $\ell \leftarrow 1$ do

$$\begin{aligned} \hat{a}_{\ell}^{H,(s+1)} &= \arg \min_{(a_{0\ell}^H, \dots, a_{K\ell}^H)} \sum_{i=1}^N \sum_{t=2}^T \sum_{m=1}^M \rho_{\tau_\ell} \left(u_{it}^{(m)} - \sum_{k=0}^K a_{k\ell}^H \varphi_k(u_{i,t-1}^{(m)}, \mathbf{age}_{it}) \right), \\ \hat{a}_{\ell}^{Q,(s+1)} &= \arg \min_{(a_{0\ell}^Q, \dots, a_{K\ell}^Q)} \sum_{i=1}^N \sum_{t=2}^T \sum_{m=1}^M \rho_{\tau_\ell} \left(y_{it} - u_{it}^{(m)} - \sum_{k=0}^K a_{k\ell}^Q \varphi_k(v_{i,t-1}^{(m)}, \mathbf{age}_{it}) \right), \\ \hat{a}_{\ell}^{H_1,(s+1)} &= \arg \min_{(a_{0\ell}^{H_1}, \dots, a_{K_1\ell}^{H_1})} \sum_{i=1}^N \sum_{m=1}^M \rho_{\tau_\ell} \left(u_{i1}^{(m)} - \sum_{k=0}^{K_1} a_{k\ell}^{H_1} \varphi_k(\mathbf{age}_{i1}) \right), \\ \hat{a}_{\ell}^{Q_1,(s+1)} &= \arg \min_{(a_{0\ell}^{Q_1}, \dots, a_{K_1\ell}^{Q_1})} \sum_{i=1}^N \sum_{m=1}^M \rho_{\tau_\ell} \left(y_{i1} - u_{i1}^{(m)} - \sum_{k=0}^{K_1} a_{k\ell}^{Q_1} \varphi_k(\mathbf{age}_{i1}) \right), \end{aligned}$$

Update tail parameters; Impose normalization restrictions $\implies \hat{\theta}^{(s+1)}$.

III. RESULTS AND DISCUSSIONS

A. Data

As the longest running household panel survey in the world, the PSID dataset contains a large amount of data in the US over 50 years. The dataset includes a wide range of variables on income, employment, education, economic, social, health-related factors and many other aspects of life for each individual and their family members. However, the size, variety and complexity of the dataset make it challenging to analyze using traditional statistical methods.

The second challenge comes from the longitudinal nature, which means that it follows the same individuals and families over time. This presents unique challenges for analysis, such as handling missing data, attrition, and changes in the variables of interest over time.

The third one is data quality. The quality of PSID can vary over time, as changes in survey methodology or sample composition can affect the accuracy and reliability of the data. This requires careful attention to data cleaning and quality control procedures.

Our study draws on the PSID as our primary data source. To ensure high quality data and better compare our empirical findings with other literature, our sample selection procedure and data preprocessing mainly follow the works of [3], [4].

B. Empirical Findings

1) *Densities and Moments*: Figure 2 illustrates the marginal distributions of the persistent and transitory earnings components at the mean age. The persistent component U_{it} displays small deviations from Gaussianity. However, the marginal distribution of V_{it} provides strong evidence to reject Gaussianity

owing to its high kurtosis and fat tails. It is worth noting that the density of V_{it} in our model is less spiky than that in [1]. A possible explanation for this difference could be the mutual dependence structure of V_{it} , which is governed by its first-order Markovian property, whereas they are assumed to be mutually independent across t in their paper. In Figure 3, we report the conditional skewness for $\tau = 11/12$, for both the U component and the V component. These two panels show a similar pattern: $U_{it}(V_{it})$ is positively skewed for low values of $U_{i,t-1}(V_{i,t-1})$, and negatively skewed for high values of $U_{i,t-1}(V_{i,t-1})$.




Fig. 2. Marginal distributions of persistent and transitory earnings components




Fig. 3. Conditional skewness of U component and V component.

2) *Nonlinear persistence*: This experiment examines the marginal effects of persistent and transitory shocks. For linear specifications of earnings dynamics, their marginal effects would be constant by design. In the canonical model of earnings dynamics, for example, where the innovation is a random walk, then the marginal effect is 1, regardless of $U_{i,t-1}$ and τ . In contrast, our model allows the persistence of $U_{i,t-1}$ to depend on the magnitude and direction of the shock. As a result, the persistence of a shock to $U_{i,t-1}$ depends on the size and sign of current and future shocks. In particular, our model enables specific shocks to erase the memory of past shocks. Furthermore, the interaction between the shock η_{it} and the lagged persistent component $U_{i,t-1}$ is a central feature of our nonlinear approach. We then estimate the earnings model, and, given the estimated parameters, we simulate the model. Figure 5 shows that our nonlinear model reproduces the patterns of nonlinear persistence well.

Figure 4 indicates the presence of nonlinear persistence, which depends on both the percentile of past earnings (τ_{init}) and the percentile of the quantile innovation (τ_{shock}). Figure 6 then shows the estimated persistence of the earnings component U_{it} . Specifically, the graph shows the marginal effects, evaluated at percentiles τ_{init} and τ_{shock} and at the mean age

in the sample. Persistence in U 's is higher than persistence in log-earnings residuals, consistently with the fact that Figure 6 is net of transitory shocks. One observation that sets this study apart from [1] is that the persistence in Figure 6 is higher than 1. For high-earnings households hit by good shocks and low-earnings households hit by bad shocks, persistence is even above 1.5, with the persistence in the latter being higher than that in the former.

Figures 7 and 8 demonstrate that the persistence in V 's is generally lower in magnitude than that in U . One notable feature is that when high-earnings households are hit by bad shocks and low-earnings households are hit by good shock, the persistence can be negative. The high degree of nonlinearity displayed here strongly rejects that V_t follows an independent process across t . If it did, its nonlinear persistence measures would remain constant for any t .




Fig. 4. 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




Fig. 5. Estimates of the average derivative of the conditional quantile function of simulated model.

3) *ARCH effects*: Figure 9 presents estimates of log-earnings residuals growth at various horizons, from 2 to 8 years. All of them suggest the presence of ARCH effects, which is consistent with findings in the existing literature, such as [9]. The data also reveal that log-earnings growth is non-Gaussian and displays negative skewness and high kurtosis. [7] finds similar features in U.S. administrative data. [1] further highlights the skewness and excess kurtosis of log-earnings growth at long horizons are primarily due to the non-Gaussianity of the transitory component.




Fig. 6. Estimates of the average derivative of the conditional quantile function of the persistent component U .




Fig. 7. Estimates of the average derivative of the conditional quantile function of the transitory component V .




Fig. 8. Estimates of the average derivative of the conditional quantile function of the persistent component V .

IV. CONCLUSION

We develop a nonparametric identification strategy for modeling earnings dynamics, differentiating the two unobserved components based on their distinct impact on household consumption. We also propose an modified stochastic EM algorithm for estimating this model.

In analyzing PSID, the empirical results reveal notable nonlinearities in both persistence component and transitory component. Specifically, substantial nonlinear persistence and conditional skewness are observed in both components. These findings suggest that the earnings shocks to a household depend on both the history of past shocks and the household's past relative wealth. In particular, persistence is higher for

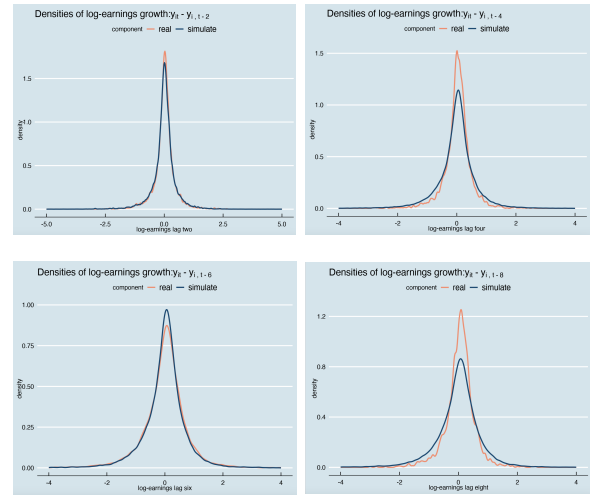


Fig. 9. Densities of log-earnings growth at various horizons.

high-earnings households hit by good shocks and low-earnings households hit by bad shocks, while it is lower for high-earnings household hit by bad shocks and low-earnings households hit by good shocks. These features align with similar observations in the PSID that previous earnings dynamic models cannot capture. We also find some other features such as ARCH effects that have been documented in other literature.

ACKNOWLEDGMENT

The authors would like to thank the University of Michigan for providing the Panel Study of Income Dynamics (PSID) data, which was essential to the success of our research. We are also deeply grateful to Professor Roger Koenker for his valuable suggestions and clarifications on issues about quantile regression models.

REFERENCES

- [1] Arellano, Manuel, Richard Blundell, and Stéphane Bonhomme. "Earnings and consumption dynamics: a nonlinear panel data framework." *Econometrica* 85.3 2017: 693-734.
- [2] Arellano, Manuel, and Stéphane Bonhomme. "Nonlinear panel data estimation via quantile regressions." *Econometric Theory*. 2016: C61-C94.
- [3] Blundell, Richard, Luigi Pistaferri, and Ian Preston. "Consumption inequality and partial insurance." *American Economic Review* 98.5 2008: 1887-1921.
- [4] Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. "Consumption inequality and family labor supply." *American Economic Review* 106.2 2016: 387-435.
- [5] Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. "Children, time allocation, and consumption insurance." *Journal of Political Economy* 126.S1 2018: S73-S115.
- [6] D'Haultfoeulle, Xavier. "On the completeness condition in nonparametric instrumental problems." *Econometric Theory* 27.3 2011: 460-471.
- [7] Guvenen F, Karahan F, Ozkan S, Song J. "What do data on millions of US workers reveal about life-cycle earnings risk?". National Bureau of Economic Research; 2015 Feb 2.
- [8] Koenker R, Bassett Jr G. "Regression quantiles". *Econometrica: journal of the Econometric Society*. 1978 Jan 1:33-50.
- [9] Meghir C, Pistaferri L. "Income variance dynamics and heterogeneity". *Econometrica*. 2004 Jan;72(1):1-32.
- [10] Nielsen SF. The stochastic EM algorithm: estimation and asymptotic results. *Bernoulli*. 2000 Jun 1:457-89.
- [11] Schennach S. "Measurement systems. *Journal of Economic Literature*". 2022 Dec;60(4):1223-63.