Nonlinear Persistence and Partial Insurance:

Innovations in the panel data dynamics of income and consumption

Richard Blundell (UCL & IFS),
joint work with Manuel Arellano (CEMFI), Stéphane Bonhomme (Chicago) and Jack Light (Chicago)

Special session organised by:
INTERNATIONAL ASSOCIATION OF APPLIED ECONOMETRICS

ASSA, January 4, 2019

The Panel Data Dynamics of Income and Consumption

The aim of this research is to examine the transmission of income "shocks" through to consumption:

> The 'larger' objective is to model the links between earnings, income, and consumption inequality - the *distributional dynamics* of inequality

- early papers: Deaton and Paxson (1994), Blundell and Preston (1998), Krueger and Perri (2005), Blundell, Pistaferri and Preston (BPP, 2008),..

> Here the focus is on *nonlinear persistence and partial insurance:*

The Panel Data Dynamics of Income and Consumption

The aim of this research is to examine the transmission of income "shocks" through to consumption:

> The 'larger' objective is to model the links between earnings, income, and consumption inequality - the *distributional dynamics* of inequality

- early papers: Deaton and Paxson (1994), Blundell and Preston (1998), Krueger and Perri (2005), Blundell, Pistaferri and Preston (BPP, 2008),..

> Here the focus is on *nonlinear persistence and partial insurance:*

1. We consider *alternative ways of modelling persistence*. Explore the nonlinear nature of income shocks with the implications for the insurance of income shocks for consumption,

2. We examine *improved sequential computational methods for the nonlinear latent/hidden quantile Markov model*, with *extensions for heterogeneity and selection at the extensive margin*.

The Panel Data Dynamics of Income and Consumption

The aim of this research is to examine the transmission of income "shocks" through to consumption:

$>$ The 'larger' objective is to model the links between earnings, income, and consumption inequality - the *distributional dynamics* of inequality

- early papers: Deaton and Paxson (1994), Blundell and Preston (1998), Krueger and Perri (2005), Blundell, Pistaferri and Preston (BPP, 2008),..

$>$ Here the focus is on *nonlinear persistence and partial insurance:*

1. We consider *alternative ways of modelling persistence*. Explore the nonlinear nature of income shocks with the implications for the insurance of income shocks for consumption,

2. We examine *improved sequential computational methods for the nonlinear latent/hidden quantile Markov model*, with *extensions for heterogeneity and selection at the extensive margin*.

$*$ Exploiting new US Household Panel data and Norwegian Register data.

New data on consumption, assets and income

I. Newly designed panel surveys: e.g. PSID 1999 - 2015.

- Collection of consumption and assets had a major revision in 1999
  - Around 90% of consumption from 2005. We focus on this series.
  - Food at home, food away from home, gasoline, health, transportation, utilities, clothing, leisure activities, etc - choice of purchase frequency.
  - Earnings and hours for all earners; Assets/debts measured in each wave.
- for background see Blundell, Pistaferri & Saporta-Eksten (BPS, 2016), and Arellano, Blundell and Bonhomme (ABB, 2017).

II. Administrative linked data: e.g. Norwegian population register.

- Linked registry databases with unique individual identifiers.
  - Containing records for **every Norwegian from 1999 to 2014.**
  - Detailed socioeconomic information (market income, cash transfers). Recent links to financial transactions data on real estate and assets; and to hours of work$\Rightarrow$ new consumption measurements.
- for background see Blundell, Graber & Mogstad (BGM, 2015) and Eika, Mogstad and Vestad (2018).

- In the prototypical "canonical" panel data model, (log) family (earned) income y_{it} is:

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

where y_{it} is net of a *systematic component*, η_{it} is a *random walk* with innovation v_{it},

$$\eta_{it} = \eta_{it-1} + v_{it}, \quad i = 1, ..., N, \quad t = 1, ..., T.$$

and ε_{it} is a *transitory shock*.

- In the prototypical "canonical" panel data model, (log) family (earned) income y_{it} is:

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

where y_{it} is net of a *systematic component*, η_{it} is a *random walk* with innovation v_{it},

$$\eta_{it} = \eta_{it-1} + v_{it}, \quad i = 1, ..., N, \quad t = 1, ..., T.$$

and ε_{it} is a *transitory shock*.

- Consumption growth is then related to the two latent income shocks:

$$\triangle c_{it} = \phi_t v_{it} + \psi_t \varepsilon_{it} + \nu_{it}, \quad i = 1, ..., N, \quad t = 1, ..., T.$$

where c_{it} is log total consumption net of a systematic component,
$>$ ϕ_t measures the *transmission* of persistence shocks v_{it}, and
$>$ ψ_t measures the *transmission* of transitory shocks;
- the ν_{it} are taste shocks.

- In the prototypical "canonical" panel data model, (log) family (earned) income y_{it} is:

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

where y_{it} is net of a *systematic component*, η_{it} is a *random walk* with innovation v_{it},

$$\eta_{it} = \eta_{it-1} + v_{it}, \quad i = 1, ..., N, \quad t = 1, ..., T.$$

and ε_{it} is a *transitory shock*.

- Consumption growth is then related to the two latent income shocks:

$$\triangle c_{it} = \phi_t v_{it} + \psi_t \varepsilon_{it} + \nu_{it}, \quad i = 1, ..., N, \quad t = 1, ..., T.$$

where c_{it} is log total consumption net of a systematic component,
> ϕ_t measures the *transmission* of persistence shocks v_{it}, and
> ψ_t measures the *transmission* of transitory shocks;
- the ν_{it} are taste shocks.

- In the prototypical "canonical" panel data model, (log) family (earned) income y_{it} is:

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

where y_{it} is net of a *systematic component*, η_{it} is a *random walk* with innovation v_{it},

$$\eta_{it} = \eta_{it-1} + v_{it}, \quad i = 1, ..., N, \quad t = 1, ..., T.$$

and ε_{it} is a *transitory shock*.

- Consumption growth is then related to the two latent income shocks:

$$\triangle c_{it} = \phi_t v_{it} + \psi_t \varepsilon_{it} + \nu_{it}, \quad i = 1, ..., N, \quad t = 1, ..., T.$$

where c_{it} is log total consumption net of a systematic component,

> ϕ_t measures the *transmission* of persistence shocks v_{it}, and

> ψ_t measures the *transmission* of transitory shocks;

- the ν_{it} are taste shocks.

$\Rightarrow$ ϕ_t and ψ_t are transmission or "*partial insurance*" parameters that depend on age t.

- In the prototypical "canonical" panel data model, (log) family (earned) income y_{it} is:

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

where y_{it} is net of a *systematic component*, η_{it} is a *random walk* with innovation v_{it},

$$\eta_{it} = \eta_{it-1} + v_{it}, \quad i = 1, ..., N, \quad t = 1, ..., T.$$

and ε_{it} is a *transitory shock*.

- Consumption growth is then related to the two latent income shocks:

$$\triangle c_{it} = \phi_t v_{it} + \psi_t \varepsilon_{it} + \nu_{it}, \quad i = 1, ..., N, \quad t = 1, ..., T.$$

where c_{it} is log total consumption net of a systematic component,
> ϕ_t measures the *transmission* of persistence shocks v_{it}, and
> ψ_t measures the *transmission* of transitory shocks;
- the ν_{it} are taste shocks.

$\Rightarrow$ ϕ_t and ψ_t are transmission or "*partial insurance*" parameters that depend on age t. Note that Blundell, Low & Preston (2014) extend to ARMA processes and to allow partial insurance to depend on assets.

Assessing the simple partial insurance framework

• This "*standard*" partial insurance framework implies a set of extended covariance restrictions for panel data on consumption and income,

$\triangleright$ allowing the transmission (insurance) parameters (ϕ and ψ), and variances (σ_v^2 and σ_ε^2) to depend on age and education is key, e.g. results for PSID US panel data and Norwegian population register, in extra slides.

$\Rightarrow$ can show (over-)identification and efficient estimation via GMM, see Blundell, Preston and Pistaferri (AER, 2008).

Assessing the simple partial insurance framework

• This "*standard*" partial insurance framework implies a set of extended covariance restrictions for panel data on consumption and income,

▷ allowing the transmission (insurance) parameters (ϕ and ψ), and variances (σ_v^2 and σ_ε^2) to depend on age and education is key, e.g. results for PSID US panel data and Norwegian population register, in extra slides.

⇒ can show (over-)identification and efficient estimation via GMM, see Blundell, Preston and Pistaferri (AER, 2008).

▷ However, linearity *rules out the nonlinear transmission of shocks*.

• The aim in recent work, e.g. ABB (Ecta, 2017), is to develop a framework that allows for:

⇒ "*unusual*" shocks to wipe out the memory of past shocks, and

⇒ "*future persistence*" of a current shock to depend on the future shocks

• Introduce an *alternative approach to modeling persistence*:

Nonlinear persistence

• Using the permanent/transitory model, consider a cohort of households, $i = 1, ..., N$, of age t. Let y_{it} denote log-labor income, net of age dummies

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

▷ η_{it} follows a general first-order Markov process (can be generalised).

Nonlinear persistence

• Using the permanent/transitory model, consider a cohort of households, $i = 1, ..., N$, of age t. Let y_{it} denote log-labor income, net of age dummies

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

▷ η_{it} follows a general first-order Markov process (can be generalised).

• Denoting the τth conditional quantile of η_{it} given $\eta_{i,t-1}$ as $Q_t(\eta_{i,t-1}, \tau)$, we specify

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

▷ ε_{it} zero mean, independent over time. Q_t and F_{ε_t} are *age (t) specific*.

Nonlinear persistence

• Using the permanent/transitory model, consider a cohort of households, $i = 1, ..., N$, of age t. Let y_{it} denote log-labor income, net of age dummies

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

▷ η_{it} follows a general first-order Markov process (can be generalised).

• Denoting the τth conditional quantile of η_{it} given $\eta_{i,t-1}$ as $Q_t(\eta_{i,t-1}, \tau)$, we specify

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

▷ ε_{it} zero mean, independent over time. Q_t and F_{ε_t} are *age (t) specific*.

Consider the following measure of persistence:

$$\rho_t(\eta_{i,t-1}, \tau) = \frac{\partial Q_t(\eta_{i,t-1}, \tau)}{\partial \eta}.$$

$\Rightarrow \rho_t(\eta_{i,t-1}, \tau)$ measures the persistence of $\eta_{i,t-1}$ when, at age t, it is hit by a shock u_{it} that has rank τ. Measures the *persistence of histories*.

Nonlinear persistence

• Using the permanent/transitory model, consider a cohort of households, $i = 1, ..., N$, of age t. Let y_{it} denote log-labor income, net of age dummies

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

▷ η_{it} follows a general first-order Markov process (can be generalised).

• Denoting the τth conditional quantile of η_{it} given $\eta_{i,t-1}$ as $Q_t(\eta_{i,t-1}, \tau)$, we specify

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

▷ ε_{it} zero mean, independent over time. Q_t and F_{ε_t} are *age (t) specific*.

Consider the following measure of persistence:

$$\rho_t(\eta_{i,t-1}, \tau) = \frac{\partial Q_t(\eta_{i,t-1}, \tau)}{\partial \eta}.$$

$\Rightarrow \rho_t(\eta_{i,t-1}, \tau)$ measures the persistence of $\eta_{i,t-1}$ when, at age t, it is hit by a shock u_{it} that has rank τ. Measures the *persistence of histories*.

▷ Conditional heteroscedasticity, skewness & kurtosis. Condition on X_{it}.

Motivating evidence from quantile autoregressions of log family earnings

$$\frac{\partial Q_{y_t|y_{t-1}}(y_{i,t-1},\tau)}{\partial y}$$

PSID data

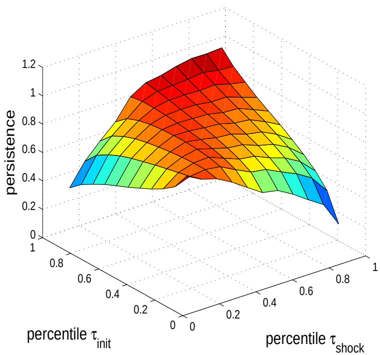

Norwegian administrative data

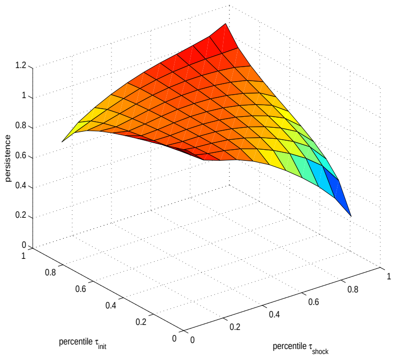

Notes: ABB (2017). Household earnings, Age 25-60, 1999-2009 (US) and 2005-2006 (Norway). Estimates of average derivative of conditional quantile of y_{it} given $y_{i,t-1}$ with respect to $y_{i,t-1}$, on grid of 11-quantiles and 3rd degree Hermite polynomial.

Modelling consumption and partial insurance

- Let c_{it} and a_{it} denote log-consumption and assets (beginning of period) net of age dummies.

- Our empirical specification is based on

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

where ν_{it} are independent across periods, and g_t is a nonlinear, age-dependent function, monotone in taste shifter ν_{it}. Can allow for individual heterogeneity, advance information and habits.

Modelling consumption and partial insurance

- Let c_{it} and a_{it} denote log-consumption and assets (beginning of period) net of age dummies.

- Our empirical specification is based on

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

where ν_{it} are independent across periods, and g_t is a nonlinear, age-dependent function, monotone in taste shifter ν_{it}. Can allow for individual heterogeneity, advance information and habits.

- The consumption responses to η and ε are given by

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

▷ $\phi_t(a, \eta, \varepsilon)$ and $\psi_t(a, \eta, \varepsilon)$ reflect the transmission of the persistent and transitory earnings components, respectively.

- These are the partial insurance coefficients.

Identification:

Income:

- For $T = 3$, Wilhelm (2012) gives conditions for distribution of ε_{i2}. In particular, completeness of the *pdf*s of $(y_{i2}|y_{i1})$ and $(\eta_{i2}|y_{i1})$.
- Apply the result to each of the three-year sub-panels:

$\Rightarrow$ The marginal distribution of ε_{it} are identified for $t \in \{2, 3, ..., T-1\}$.

$\Rightarrow$ By independence the distribution of $(\varepsilon_{i2}, \varepsilon_{i3}, ..., \varepsilon_{i,T-1})$ is identified.

$\Rightarrow$ By deconvolution the distribution of $(\eta_{i2}, \eta_{i3}, ..., \eta_{i,T-1})$ is identified.

- The distribution of ε_{i1}, η_{i1}, and ε_{iT}, η_{iT} are not identified in general.

Identification:

Income:

- For $T = 3$, Wilhelm (2012) gives conditions for distribution of ε_{i2}. In particular, completeness of the *pdf*s of $(y_{i2}|y_{i1})$ and $(\eta_{i2}|y_{i1})$.
- Apply the result to each of the three-year sub-panels:

$\Rightarrow$ The marginal distribution of ε_{it} are identified for $t \in \{2, 3, ..., T-1\}$.
$\Rightarrow$ By independence the distribution of $(\varepsilon_{i2}, \varepsilon_{i3}, ..., \varepsilon_{i,T-1})$ is identified.
$\Rightarrow$ By deconvolution the distribution of $(\eta_{i2}, \eta_{i3}, ..., \eta_{i,T-1})$ is identified.

- The distribution of ε_{i1}, η_{i1}, and ε_{iT}, η_{iT} are not identified in general.

Consumption:

- Assume u_{it} and ε_{it} are independent of past earnings shocks and past asset holding, for $t \geq 1$, where $\eta_{it} = Q_t(\eta_{i,t-1}, u_{it})$.
- Let η_{i1} and a_{i1} be arbitrarily dependent;
- Denoting $\eta_i^t = (\eta_{it}, \eta_{i,t-1}, ..., \eta_{i1})$, assume that a_{it} is independent of $(\eta_i^{t-1}, a_i^{t-2}, \varepsilon_i^{t-2})$ given $(a_{i,t-1}, c_{i,t-1}, y_{i,t-1})$; -> consistent with the accumulation rule in the standard life-cycle model.

Initial Assets: Let $y = (y_1, ..., y_T)$. We have

$$
\begin{aligned}
f(a_1|y) &= \int f(a_1|\eta_1, y) f(\eta_1|y) d\eta_1 \\
&= \int f(a_1|\eta_1) f(\eta_1|y) d\eta_1,
\end{aligned}
$$

where we have used that u_{it} and ε_{it} are independent of a_{i1}.

- Note that $f(\eta_1|y)$ is identified from the earnings process alone.

- If $f(\eta_1|y)$ is complete, then $f(a_1|\eta_1)$ is identified – structure is as in the ill-posed NPIV problem where η_1 endogenous and y is the instrument.

Initial Assets: Let $y = (y_1, ..., y_T)$. We have

$$\begin{aligned} f(a_1|y) &= \int f(a_1|\eta_1, y) f(\eta_1|y) d\eta_1 \\ &= \int f(a_1|\eta_1) f(\eta_1|y) d\eta_1, \end{aligned}$$

where we have used that u_{it} and ε_{it} are independent of a_{i1}.

• Note that $f(\eta_1|y)$ is identified from the earnings process alone.

• If $f(\eta_1|y)$ is complete, then $f(a_1|\eta_1)$ is identified – structure is as in the ill-posed NPIV problem where η_1 endogenous and y is the instrument.

First period consumption: We have

$$f(c_1, a_1|y) \equiv \int f(c_1, a_1|\eta_1, y) f(\eta_1|y) d\eta_1$$

and given our assumptions (and f($a_1|\eta_1$) can be treated as known)

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

• If completeness in $(y_2, ..., y_T)$ of $f(\eta_1|y_1, y_2, ..., y_T)$, then $f(c_1|a_1, \eta_1, y_1)$, is identified. Subsequent periods follow similar arguments.

Empirical specification:

- **For income,** the quantile function of η_{it} given $\eta_{i,t-1}$ is specified as

$$\begin{aligned} Q_t(\eta_{t-1},\tau) &= Q(\eta_{t-1}, age_t, \tau) \\ &= \sum_{k=0}^{K} a_k^Q(\tau)\varphi_k(\eta_{t-1}, age_t), \end{aligned}$$

where φ_k, $k = 0, 1, ..., K$, are polynomials (Hermite). Similar specification for quantile functions of ε_{it} and η_{i1}. Can condition on covariates X_{it}.

Empirical specification:

- **For income,** the quantile function of η_{it} given $\eta_{i,t-1}$ is specified as

$$\begin{aligned} Q_t(\eta_{t-1}, \tau) &= Q(\eta_{t-1}, age_t, \tau) \\ &= \sum_{k=0}^{K} a_k^Q(\tau)\varphi_k(\eta_{t-1}, age_t), \end{aligned}$$

where φ_k, $k = 0, 1, ..., K$, are polynomials (Hermite). Similar specification for quantile functions of ε_{it} and η_{i1}. Can condition on covariates X_{it}.

- **The consumption function** (log) is specified as:

$$\begin{aligned} g_t(a_t, \eta_t, \varepsilon_t, \tau) &= g(a_t, \eta_t, \varepsilon_t, age_t, \tau) \\ &= \sum_{k=1}^{K} b_k^g \widetilde{\varphi}_k(a_t, \eta_t, \varepsilon_t, age_t) + b_0^g(\tau) \end{aligned}$$

– additivity in the taste heterogeneity, not essential. Similar specification for conditional quantiles of a_{i1} given η_{i1} and age.

Implementation and Estimation

• Model quantile coefficients $a_k^Q(\tau)$ as piecewise-linear interpolating splines (Wei and Carroll, 2009) on a grid $0 < \tau_1 < \tau_2 < ... < \tau_L < 1$, convenient as the likelihood function is available in closed form.

• Note extend the specification of the intercept coefficient $a_0^Q(\tau)$ on $(0, \tau_1]$ and $[\tau_L, 1)$ using a parametric model: exponential (λ).

• In practice, for the PSID data, we take $L = 11$ and $\tau_\ell = \ell/L + 1$. φ_k and $\widetilde{\varphi}_k$ are low-dimensional tensor products of Hermite polynomials.

Implementation and Estimation

• Model quantile coefficients $a_k^Q(\tau)$ as piecewise-linear interpolating splines (Wei and Carroll, 2009) on a grid $0 < \tau_1 < \tau_2 < ... < \tau_L < 1$, convenient as the likelihood function is available in closed form.

• Note extend the specification of the intercept coefficient $a_0^Q(\tau)$ on $(0, \tau_1]$ and $[\tau_L, 1)$ using a parametric model: exponential (λ).

• In practice, for the PSID data, we take $L = 11$ and $\tau_\ell = \ell / L + 1$. φ_k and $\widetilde{\varphi}_k$ are low-dimensional tensor products of Hermite polynomials.

• Use stochastic EM algorithm in estimation: a simulated version of the classical EM algorithm of Dempster et al (1977), where new draws from latent Markov process η are computed in every iteration of the algorithm.

• 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.

Model restrictions: income

- Let θ be the income-related parameters with true values $\overline{\theta}$.

- Let $\rho_\tau(u) = u(\tau - \mathbf{1}\{u \leq 0\})$ denote the "check" function of quantile regression, and let $\overline{a}_{k\ell}^Q$ denote the value of $a_{k\ell}^Q = a_k^Q(\tau_\ell)$ evaluated at the true $\overline{\theta}$. The model implies

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

with additional restrictions involving the other parameters in θ.

- Note that the objective function is smooth (due to the presence of the integrals) and convex (because of the check function).

- f_i denotes the posterior density of $(\eta_{i1}, ..., \eta_{iT})$ given the income data

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

Overview of estimation

- A compact notation for the restrictions implied by the income model is

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

- We use a "stochastic EM" algorithm (in a non-likelihood setup).

Starting with $\widehat{\theta}^{(0)}$ we iterate on s=0,1,... the following two steps until convergence of the Markov Chain:

1. Stochastic E-step: draw $\eta_i^{(m)} = (\eta_{i1}^{(m)}, ..., \eta_{iT}^{(m)})$ for $m = 1, ..., M$ from $f_i(\cdot; \widehat{\theta}^{(s)})$. ABB orginally used a random-walk Metropolis-Hastings (MCMC) sampler. Here we make use of particle filter methods that can be more numerically stable in complex models

2. M-step: update

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

Improved numerical methods - SMC

Consider the stochastic E step of the algorithm. We first note that the Markovian structure of the latent earnings components allows use of SMC methods, see Doucet and Johansen (2009)

- At time $t = 1$, draw N 'particles' $\{\eta_1^k\}$ from a suitable proposal distribution $q(\eta_1|y_1)$,
- Re-sampling with weights $\{W_1^k\} \propto \frac{p(\eta_1|y_1)}{q(\eta_1|y_1)}$ gives N particles approximately distributed according to $p(\eta_1|y_1)$,
- At $t = 2$, we now aim to approximate: $p(\eta_1, \eta_2|y_1, y_2) \propto p(\eta_1|y_1)f(\eta_2|\eta_1)g(y_2|\eta_2)$,
- ...but since we already have N samples approximately distributed according to $p(\eta_1|y_1)$, we can extend each of these using a second proposal distribution $q(\eta_2|y_2, \eta_1)$,
- Subsequent re-sampling of the particles $\{\eta_{1:2}^k\}$ with appropriately chosen weights yields approximate samples from $p(\eta_1, \eta_2|y_1, y_2)$.

Improved numerical methods - P(Particle)MCMC

SMC methods can also be used to generate highly-efficient proposals within an MCMC algorithm, see Andrieu, Doucet and Hollenstein (2010)

- Within an Metropolis-Hastings algorithm they lead to a very simple implementation in our current latent Markov settings
- Initialize by running an SMC algorithm to generate approximate samples $\eta_1^{(1)}, ..., \eta_T^{(1)}$ from $f(\eta_1, ..., \eta_T | y_1, ..., y_T)$ and store the marginal likelihood estimate $\hat{f}^{(1)}(y_1, ..., y_T)$
- At subsequent iterations run SMC algorithms to propose new samples $\eta_1^{(*)}, ..., \eta_T^{(*)}$ and generate new marginal likelihood estimates $\hat{f}^*(y_1, ..., y_T)$, which are accepted with probability $\frac{\hat{f}^*(y_1,...,y_T)}{\hat{f}^{(i-1)}(y_1,...,y_T)}$
- Inferior to standard SMC in simple settings, but allows for flexibility when used in conjunction with other MCMC transitions.

=> Particularly useful for extensions of the latent Markov framework with selection at the extensive margin and unobserved heterogeneity.

(A related set of) Model restrictions for consumption

- Letting μ (true value $\overline{\mu}$) be the consumption-related parameters, the model implies

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

and

$$\overline{\sigma}^2 = \mathbb{E}\left[\int \left(c_{it} - \sum_{k=1}^{K} b_k^g \widetilde{\varphi}_k(a_{it}, \eta_{it}, y_{it} - \eta_{it}, age_{it})\right)^2 g_i(\eta_i^T; \overline{\theta}, \overline{\mu}) d\eta_i^T\right],$$

with additional restrictions involving the other parameters in μ.

- Here g_i denotes the posterior density of $(\eta_{i1}, ..., \eta_{iT})$ given the earnings, consumption, asset data, and income process parameters $(\overline{\theta})$

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

Empirical results:

PSID 2005 - 2015 with full consumption items, family earnings, after tax earnings and assets.

Using improved SMC and PMCMC methods.

Nonlinear persistence of y_{it}: PSID

$$\frac{\partial Q_{y_t|y_{t-1}}(y_{i,t-1},\tau)}{\partial y}$$

(a) PSID panel data (b) Nonlinear model

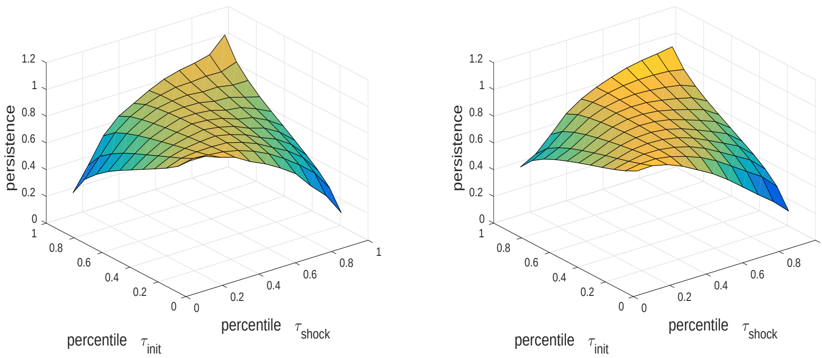

Notes: PSID 2005-2015. (b) presents PMCMC estimates of the average derivative of conditional quantile function of y_{it} given $y_{i,t-1}$ with respect to $y_{i,t-1}$, evaluated at τ_{shock} and $y_{i,t-1}$ corresponding to τ_{init} percentile of dist. of $y_{i,t-1}$.

Nonlinear persistence of η_{it} (new PSID):

$$\rho_t(\eta_{i,t-1}, \tau) = \frac{\partial Q_{\eta_t|\eta_{t-1}}(\eta_{i,t-1}, \tau)}{\partial \eta}$$

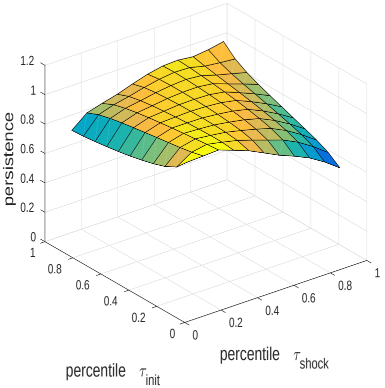

Notes: Updated PSID 2005-2015: PMCMC estimates of the average derivative of the conditional quantile function of η_{it} on $\eta_{i,t-1}$ with respect to $\eta_{i,t-1}$, evaluated at percentile τ_{shock} and at a value of $\eta_{i,t-1}$ that corresponds to the τ_{init} percentile of the distribution of $\eta_{i,t-1}$. Evaluated at mean age in the sample.

Nonlinear persistence of y_{it}: Norway

$$\frac{\partial Q_{y_t|y_{t-1}}(y_{i,t-1},\tau)}{\partial y}$$

(a) Norwegian register data

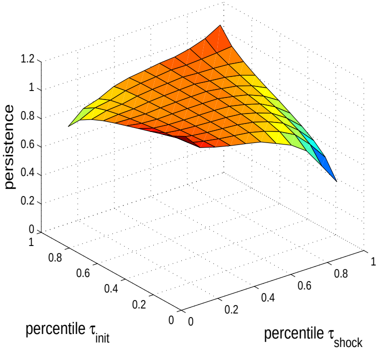

(b) Nonlinear model

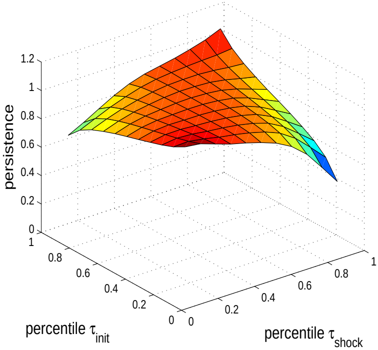

Note: Estimates of the average derivative of the conditional quantile function of y_{it} given $y_{i,t-1}$ with respect to $y_{i,t-1}$, evaluated at percentile τ_{shock} and at a value of $y_{i,t-1}$ that corresponds to the τ_{init} percentile of the dist. of $y_{i,t-1}$.

Nonlinear persistence for y_{it}, 95% confidence bands

(a) Earnings, PSID data

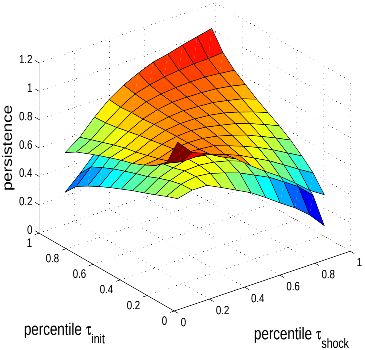

(b) Earnings, nonlinear model

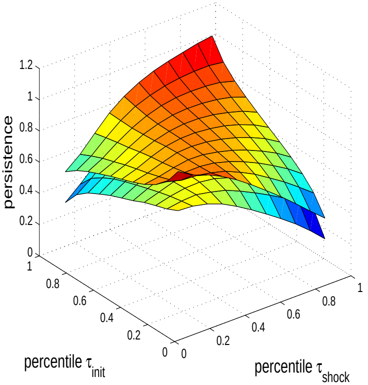

Notes: From original estimation in ABB (2017). Pointwise 95% confidence bands. Parametric bootstrap, 500 replications.

Estimates of persistence with different samplers

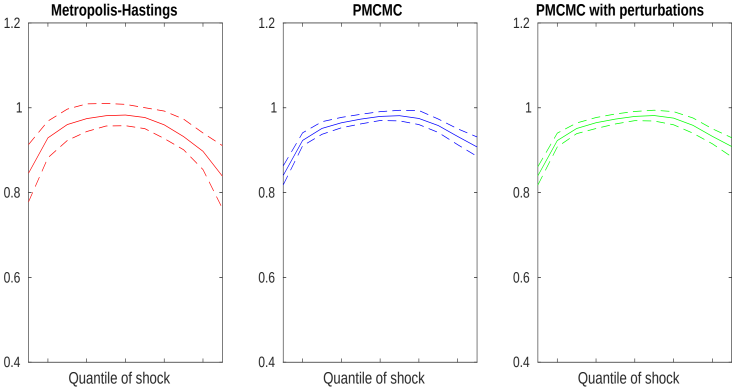

Figure: Stability of average persistence estimates across 200 repeated implementations.

Notes: Perturbations in the final panel are to the initial conditions within the Stochastic EM algorithm which add random noise $\sim N(0, 0.01)$ to all initial coefficient values at all quantiles.

Robustness of original ABB results (1999 - 2007, PSID)

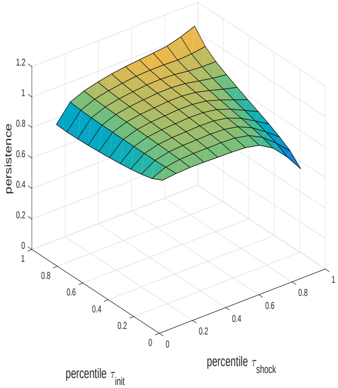

Figure: Baseline result in *Arellano et al. 2017*

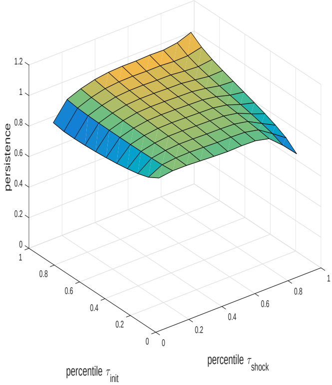

Figure: Replicated result with updated SMC methods

Consumption response to η_{it} by τ_{assets} and τ_{age} $\mathbb{E}\left[\frac{\partial g_t(a,\eta_{it},\varepsilon_{it},\nu_{it})}{\partial \eta}\right]$

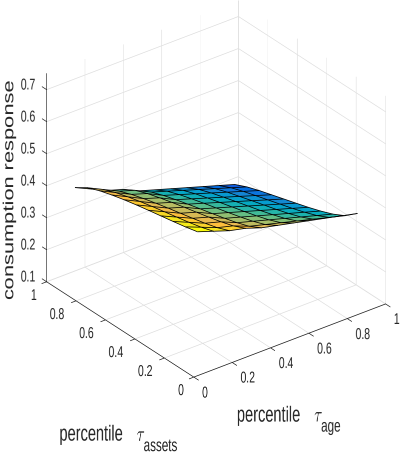

Figure: Household labor income

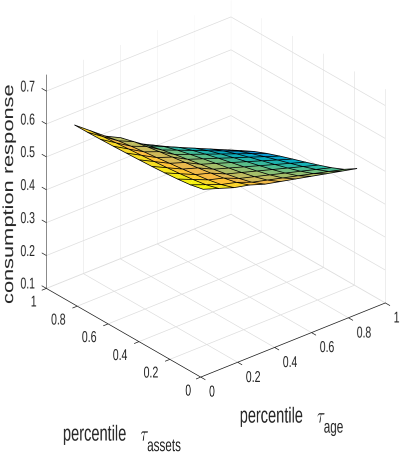

Figure: Household disposable income

Notes: Partial Insurance 'coefficients', New PSID 2010-2015, PMCMC Estimates.

Extensions: covariate specification

- Consider a specification for wages as follows:

$$W_{it} = X_{it}\beta + \eta_{it} + \epsilon_{it}, \text{where } E[\epsilon_{it}|X_{it}, \eta_{it}] = 0 \quad (1)$$

$$\eta_{it} = Q_t(X^*_{it}, \eta_{i,t-1}, V_{it}) \quad (2)$$

- ϵ_{it} and V_{it} independent at all lags and leads; normalize V to be standard uniform; $X^*_{it} \subset X_{it}$
- Our approach allows us to easily include X^*_{it} inside the whole estimation, not in a prior step as is almost always done in the literature, at a very small cost.
- Note we would like $X^*_{it} = X_{it}$, but the curse of dimensionality presents an issue - we want to estimate $Q_t(.)$ as flexibly as possible but X is usually large. An attractive as an alternative we may set $X^*_{it} = X'_{it}\beta$ alternative is to include index functions directly, e.g. a linear index we would simply include (functions of) $X'_{it}\beta$ directly as controls in equation (2).

Extensions: Selection at Extensive Margin

- The idea is that adverse persistent shocks increase probability of selection out of the labor market, i.e. non-random selection on η_{it}.
- Consider the specification for wages:

$$W_{it}^{*} = X_{it}\beta + \eta_{it} + \epsilon_{it}, \text{where } E[\epsilon_{it}|X_{it}] = 0 \tag{3}$$

$$\eta_{it} = Q_t(X_{it}^{*}, \eta_{i,t-1}, V_{it}) \tag{4}$$

$$D_{it} = 1\{U_{it} \leq q_t(X_{it}, \eta_{it}, \eta_{i,t-1}, Z_{it})\} \tag{5}$$

- Wages are observed at W_{it}^{*} when $D_{it} = 1$; ϵ_{it}, V_{it} and U_{it} independent at all lags and leads; normalize U and V to be standard uniform.
- Nonparametric identification shown under large support of the propensity score (e.g. large support of for Z), examples include simulated tax instruments, other income, etc.
- The integrated moment condition corresponding to the τth conditional quantile in (4)

$$\sum_{t=1}^{T} E\left[\int \rho_{\tau_l}(\eta_{it} - Q_t(X_{it}, \eta_{i,t-1}, \tau) f_i(\eta_i^T | X_i^T, W_i^T) d\eta_i^T\right]$$

continues to hold, although the posterior density takes a different form

Extensive margin selection and stochastic EM

- Estimation simply requires we change the form of the likelihood in the stochastic E-step, whilst estimates of (4) now include both participants and non-participants.
- Posterior density in the stochastic E-step takes the from:

$$f_i(\eta_i^T | W_i^T, D_i^T, Z_i^T) \propto f_i(\eta_i^T, W_i^T, D_i^T | Z_i^T) \tag{6}$$

$$= \prod_{t=1}^{T} f(W_{it} | \eta_{it}, D_{it}, X_{it})^{D_{it}} f(D_{it} | \eta_{it}, Z_{it}) \tag{7}$$

$$\prod_{t=2}^{T} f(\eta_{it} | \eta_{i,t-1}, X_{it}) f(\eta_{i1} | X_{it}) \tag{8}$$

Summary

- Use the framework developed in Arellano, Blundell and Bonhomme (2017) to shed new light on the *nonlinear transmission of income shocks to consumption and the nature of insurance to income shocks*.
- Use *new data on consumption, assets, earnings and family income* from the the new PSID and from the Norwegian population register.
- The new framework involves a *Markovian permanent-transitory model of income, that reveals significant asymmetric persistence of shocks.*
- We develop a flexible age-dependent nonlinear consumption rule that is a function of assets, permanent income and transitory income.
- We provide *conditions for nonparametric identification and show how a simulation-based sequential Quantile Regression method is feasible*.
- The Markovian structure for latent earnings components allows us to make use of *particle SMC methods to improve the MCMC algorithm*.
- The new framework provides *robust measures of nonlinear persistence for family earnings and income, and new estimates of the degree of partial insurance of income shocks for consumption*.

Extra Slides

1. Consumption in the PSID and Norwegian Register
2. Income components in the Norwegian Register Data
3. The Metropolis-Hastings method.
4. Identification and extensions in consumption model.
5. Summary statistical properties.

Consumption in the (New) PSID Data

• PSID From 1999.

• C_{it}: Information on food expenditures, rents, health expenditures, utilities, car-related expenditures, ...

• A_{it}: Assets holdings are the sum of financial assets, real estate value, pension funds, and car value, net of mortgages and other debt. (Net worth).

• y_{it} are residuals of log total pre-tax household labor earnings on a set of demographics. Note, c_{it} and a_{it} are residuals, using the same set of demographics as for earnings.

▷ cohort and calendar time dummies, family size and composition, education, race, and state dummies.

• As in BPS, we select married male heads aged between 25 and 59.

Consumption in the Norwegian Register Data

The analysis combines several data sources for the period 1994-2014

- Tax records on income and wealth
- Real estate transactions from Norwegian Land Register
- Transactions in listed and unlisted stocks from Norwegian Registry of Securities.

The initial sample covers all households where the household's oldest is at least 18 years old, everyone above 17 years has filed a tax return

- The number of household-year observations in the initial panel is 44,302,000.
- In each year, we keep only households with a male head, age 30 - 60, cohort 1945 - 1975, with non-missing information on schooling and location.

Detailed description of the dataset and consumption measurement in Eika, Mogstad and Vestad (2018).

Measuring Consumption

Let $W_{ikt} = p_{kt}A_{ikt}$, total household consumption expenditure:

$$C_{it} = \underbrace{\left(E_{it} - \tau_{it} + \sum_k r_{kt} A_{ikt-1} \right)}_{\textit{disposable income}}$$

$$\underbrace{- \underbrace{\sum_k (W_{ikt} - W_{ikt-1})}_{\textit{changes in wealth}} + \underbrace{\sum_k (p_{kt} - p_{kt-1}) A_{ikt-1}}_{\textit{capital gains}}}_{-\textit{net savings}}$$

where

- Y_{it} : labour income and cash transfers
- τ_{it} : taxes
- A_{it-1} : assets held at the end of period $t-1$.

Combining the last two terms using financial and real estate transactions data has been a key insight.

The Income Process in the Norwegian population register

Variance of permanent shocks to income

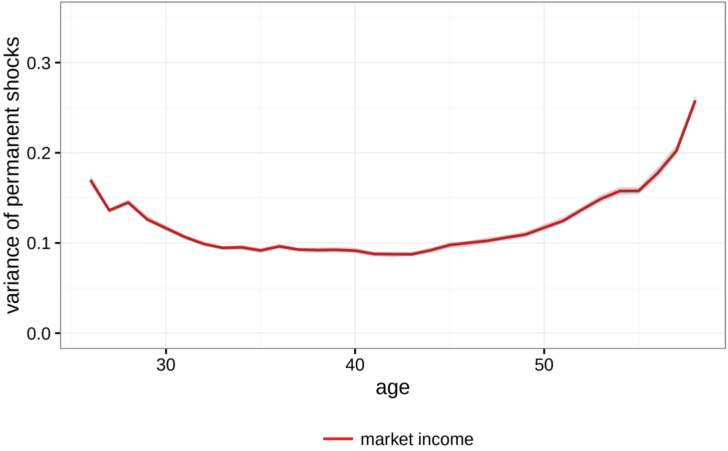

Source: Blundell, Graber and Mogstad (2015).

Norwegian population register data

Variance of permanent shocks to income

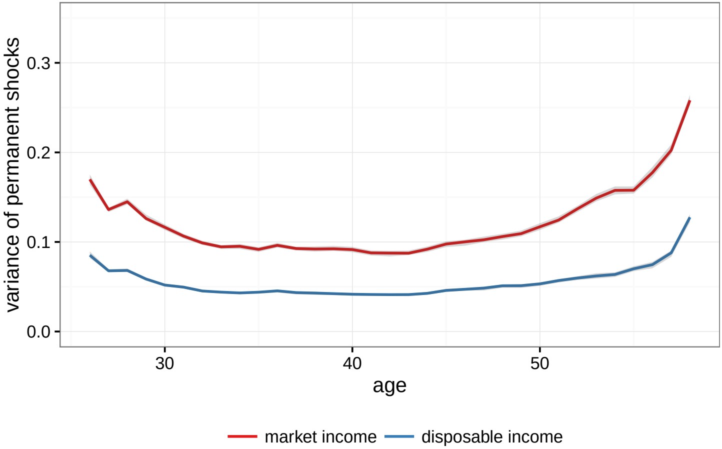

Source: Blundell, Graber and Mogstad (2015).

Variance of permanent shocks to income

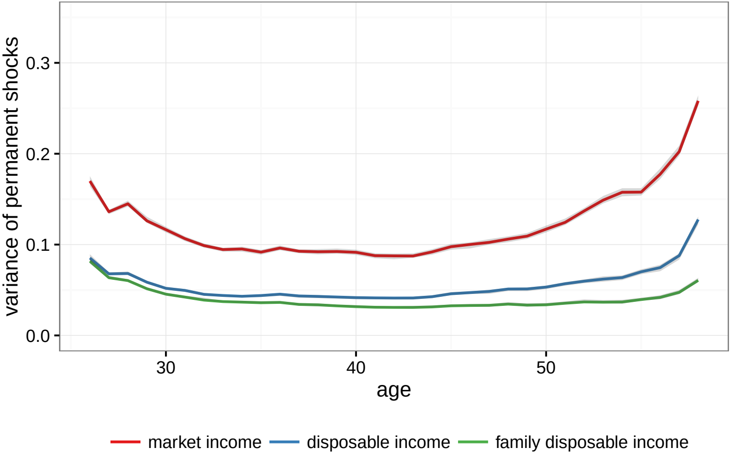

Source: Blundell, Graber and Mogstad (2015).

Norwegian population register data

Variance of permanent shocks to income (low skilled)

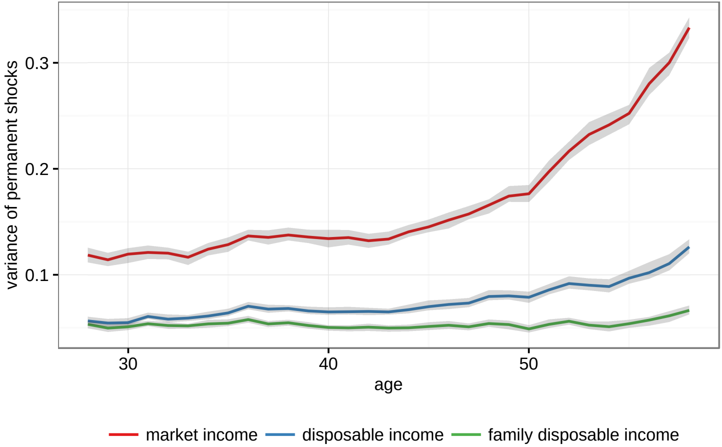

Source: Blundell, Graber and Mogstad (2015).

Metropolis-Hastings method

The stochastic EM algorithm requires samples from the posterior density $f(\eta_1, ..., \eta_T | y_1, ..., y_T)$

Metropolis-Hastings: the 'standard'MCMC technique used in ABB. Letting $q(.)$ denote proposal distributions and $\pi(.)$ the full joint density, the basic algorithm is outlined as,

- Initialize $\eta_1^{(0)}, ..., \eta_T^{(0)} \sim q(\eta_1, ..., \eta_T | y_1, ..., y_T)$
- For subsequent iterations propose updates to each element of $\eta_1, ..., \eta_T$ sequentially through time as $\eta_t^* \sim q(\eta_t^{(i)} | \eta_t^{(i-1)})$
- Accept each subsequent proposal with probability $\min\{1, \frac{q(\eta_t^{(i-1)} | \eta_t^*)\pi(\eta_t^*, .)}{q(\eta_t^* | \eta_t^{(i-1)})\pi(\eta_t^{(i-1)}, .)}\}$
- If a proposal is accepted, set $\eta_t^{(i)} = \eta_t^*$, otherwise set $\eta_t^{(i)} = \eta_t^{(i-1)}$
- Resulting Markov chain yields correlated samples.

Consumption: identification in subsequent periods

- We have

$$
\begin{aligned}
f(a_2|c_1, a_1, y) &= \int f(a_2|c_1, a_1, \eta_1, y_1) f(\eta_1|c_1, a_1, y) d\eta_1 \\
f(c_2|a_2, c_1, a_1, y) &= \int f(c_2|a_2, \eta_2, y_2) f(\eta_2|a_2, c_1, a_1, y) d\eta_2.
\end{aligned}
$$

- By induction it can be shown that 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 $(c_i^{t-1}, a_i^{t-1}, y_i^{t-1}, y_{i,t+1}, ..., y_{iT})$.

- Intuition: lagged consumption and assets, as well as lags and leads of earnings, are used as instruments for η_{it}.

Extensions to the Consumption Model

• Consumption rule with *unobserved heterogeneity:*

$$c_{it} = g_t \left(a_{it}, \eta_{it}, \varepsilon_{it}, \xi_i, \nu_{it} \right).$$

• We assume that u_{it} and ε_{it}, for $t \geq 1$, are independent of (a_{i1}, ξ_i).

• The distribution of $(a_{i1}, \xi_i, \eta_{i1})$ is unrestricted.

• A combination of the above identification arguments and the main result of Hu and Schennach (08) identifies:

– the period-t consumption distribution $f(c_t | a_t, \eta_t, y_t, \xi)$, and

– the distribution of initial conditions $f(\eta_1, \xi, a_1)$.

Statistical properties, notes.

• Nielsen (2000) studies the properties of this algorithm in a likelihood case. He provides conditions for the Markov Chain $\widehat{\theta}^{(s)}$ to be ergodic (for a fixed sample size).

• He also shows that $\sqrt{N}\left(\widehat{\theta}^{(s)} - \overline{\theta}\right)$ converges to a Gaussian autoregressive process as N tends to infinity.

• Arellano and Bonhomme [AB] (2015) adapt Nielsen's arguments to derive the form of the asymptotic variance in a non-likelihood case.

• AB also study consistency as K (number of polynomial terms) and L (number of knots) tend to infinity with N.