

Modeling the Income Process

Part 2

Extract from “Earnings, Consumption and Lifecycle Choices”
by Costas Meghir and Luigi Pistaferri.
Handbook of Labor Economics, Vol. 4b, Ch. 9. (2011).

James J. Heckman
University of Chicago

AEA Continuing Education Program
ASSA Course: Microeconomics of Life Course Inequality
San Francisco, CA, January 5-7, 2016



Modeling the Income Process

- In this section we discuss the specification and estimation of the income process.
- Two main approaches will be discussed.
- The first looks at earnings as a whole, and interprets risk as the year-to-year volatility that cannot be explained by certain observables (with various degrees of sophistication).
- The second approach assumes that part of the variability in earnings is endogenous (induced by choices).
- In the first approach, researchers assume that consumers receive an uncertain but *exogenous* flow of earnings in each period.

- This literature has two objectives: (a) identification of the correct process for earnings, (b) identification of the information set - which defines the concept of an “innovation”.
- In the second approach, the concept of risk needs revisiting, because one first needs to identify the “primitive” risk factors.
- For example, if endogenous fluctuations in earnings were to come exclusively from people freely choosing their hours, the “primitive” risk factor would be the hourly wage.

- As for the issue of information set, the question that is being asked is whether the consumer knows more than the econometrician.
- This is sometimes known as the *superior information* issue.
- The individual may have advance information about events such as a promotion, that the econometrician may never hope to predict on the basis of observables (unless, of course, promotions are perfectly predictable on the basis of things like seniority within a firm, education, etc.).

- The correct DGP for income, earnings or wages will be affected by data availability.
- While the ideal data set is a long, large panel of individuals, this is somewhat a rare event and can be plagued by problems such as attrition (see Baker and Solon, 2003, for an exception).
- More frequently, researchers have available panel data on individuals, but the sample size is limited, especially if one restricts the attention to a balanced sample (for example, Baker, 1997; MaCurdy, 1982).
- Alternatively, one could use an unbalanced panel (as in Meghir and Pistaferri, 2004, and Heathcote, Storesletten and Violante, 2004).

- An important exception is the case where countries have available administrative data sources with reports on earnings or income from tax returns or social security records.
- The important advantage of such data sets is the accuracy of the information provided and the lack of attrition, other than what is due to migration and death.
- The important disadvantage is the lack of other information that is pertinent to modelling, such as hours of work and in some cases education or occupation, depending on the source of the data.

- Even less frequently, one may have available employer-employee matched data sets, with which it may be possible to identify the role of firm heterogeneity separately from that of individual heterogeneity, either in a descriptive way such as in Abowd, Kramarz and Margolis (1999), or allowing also for shocks, such as in Guiso, Pistaferri and Schivardi (2005), or in a more structural fashion as in Postel-Vinay and Robin (2002), Cahuc, Postel-Vinay and Robin (2006), Postel-Vinay and Turon (2009) and Lise, Meghir and Robin (2009).
- Less frequent and more limited in scope is the use of pseudo-panel data, which misses the variability induced by genuine idiosyncratic shocks, but at least allows for some results to be established where long panel data is not available (see Banks, Blundell and Brugiavini, 2001, and Moffitt, 1993).

Specifications

- Income processes found in the literature is implicitly or explicitly motivated by Friedman's permanent income hypothesis.
- We denote by $Y_{i,a,t}$ a measure of income (such as earnings) for individual i of age a in period t .
- This is typically taken to be annual earnings and individuals not working over a whole year are usually dropped.
- Issues having to do with selection and endogenous labour supply decisions will be dealt with in a separate section.
- Many of the specifications for the income process take the form

$$\ln Y_{i,a,t}^e = d_t^e + \beta^{e'} X_{i,a,t} + u_{i,a,t} \quad (1)$$

- In the above e denotes a particular group (such as education and sex) and $X_{i,a,t}$ will typically include a polynomial in age as well as other characteristics including region, race and sometimes marital status.
- d_t denote time effects.
- From now on we omit the superscript “e” to simplify notation. In (1) the error term $u_{i,a,t}$ is defined such that $E(u_{i,a,t}|X_{i,a,t}) = 0$.
- This allows us to work with residual log income $\widehat{y_{i,a,t}} = \ln Y_{i,a,t} - \widehat{d}_t - \widehat{\beta}'X_{i,a,t}$ where $\widehat{\beta}$ and the aggregate time effects \widehat{d}_t can be estimated using OLS.

- Henceforth we will ignore this first step and we will work directly with residual log income $y_{i,a,t}$, where the effect of observable characteristics and common aggregate time trends have been eliminated.
- The key element of the specification in (1) is the time series properties of $u_{i,a,t}$.

- A specification that encompasses many of the ideas in the literature is

$$u_{i,a,t} = a \times f_i + v_{i,a,t} + p_{i,a,t} + m_{i,a,t}$$

$$v_{i,a,t} = \Theta_q(L)\varepsilon_{i,a,t} \quad \text{Transitory process}$$

$$P_p(L)p_{i,a,t} = \zeta_{i,a,t} \quad \text{Permanent process}$$

(2)

- L is a lag operator such that $Lz_{i,a,t} = z_{i,a-1,t-1}$.

- In (2) the stochastic process consists of an individual specific lifecycle trend ($a \times f_i$)
- A transitory shock $v_{i,a,t}$, which is modelled as an MA process whose lag polynomial of order q is denoted $\Theta_q(L)$
- A permanent shock $P_p(L)p_{i,a,t} = \zeta_{i,a,t}$, which is an autoregressive process with high levels of persistence possibly including a unit root, also expressed in the lag polynomial of order p , $P_p(L)$
- and measurement error $m_{i,a,t}$ which may be taken as classical *iid* or not.

A Simple Model of Earnings Dynamics

- We start with the relatively simpler representation where the term $a \times f_i$ is excluded.
- Moreover we restrict the lag polynomials $\Theta(L)$ and $P(L)$: it is not generally possible to identify $\Theta(L)$ and $P(L)$ without any further restrictions.

- Thus we start with the typical specification used for example in MaCurdy (1982) and Abowd and Card (1989):

$$u_{i,a,t} = v_{i,a,t} + p_{i,a,t} + m_{i,a,t}$$

$$v_{i,a,t} = \varepsilon_{i,a,t} - \theta \varepsilon_{i,a-1,t-1} \quad \text{Transitory process} \quad (3)$$

$$p_{i,a,t} = p_{i,a-1,t-1} + \zeta_{i,a,t} \quad \text{Permanent process}$$

$$p_{i,0,t-a} = h_i$$

$m_{i,a,t}$ measurement error at age a and time t


with $m_{i,a,t}$, $\zeta_{i,a,t}$ and $\varepsilon_{i,a,t}$ all being independently and identically distributed and where h_i reflects initial heterogeneity, which here persists forever through the random walk ($a = 0$ is the age of entry in the labor market, which may differ across groups due to different school leaving ages).

- Generally, as we will show, the existence of classical measurement error causes problems in the identification of the transitory shock process.
- There are two principal motivations for the permanent/transitory decompositions: the first motivation draws from economics: the decomposition reflects well the original insights of Friedman (1957) by distinguishing how consumption can react to different types of income shock, while introducing uncertainty in the model.
- The second is statistical: At least for the US and for the UK the variance of income increases over the life-cycle (see Figure 1, which uses consumption data from the CEX and income data from the PSID).
- This, together with the increasing life cycle variance of consumption points to a unit root in income, as we shall see below.

- Moreover, income growth ($\Delta \ln y_{i,a,t}$) has limited serial correlation and behaves very much like an MA process of order 2 or three: this property is delivered by the fact that all shocks above are assumed *iid*. In our example growth in income has been restricted to an *MA(2)*.
- Even in such a tight specification identification is not straightforward: as we will illustrate we cannot separately identify the parameter θ , the variance of the measurement error and the variance of the transitory shock.
- But first consider the identification of the variance of the permanent shock.
- Define unexplained earnings growth as:

$$g_{i,a,t} \equiv \Delta y_{i,a,t} = \Delta m_{i,a,t} + (1 + \theta L)\Delta \varepsilon_{i,a,t} + \zeta_{i,a,t}. \quad (4)$$

Figure 1: The variance of log income (from the PSID, dashed line) and log consumption (from the CEX, continuous line) over the life cycle.



- Then the key moment condition for identifying the variance of the permanent shock is

$$E \left(\zeta_{i,a,t}^2 \right) = E \left[g_{i,a,t} \left(\sum_{j=-(1+q)}^{(1+q)} g_{i,a+j,t+j} \right) \right] \quad (5)$$

where q is the order of the moving average process in the original levels equation; in our example $q = 1$.

- Hence, if we know the order of serial correlation of the log income we can identify the variance of the permanent shock without any need to identify the variance of the measurement error or the parameters of the MA process.

- Indeed, in the absence of a permanent shock the moment in (5) will be zero, which offers a way of testing for the presence of a permanent component *conditional* on knowing the order of the MA process.
- If the order of the MA process is one in the levels, then to implement this we will need at least six individual-level observations to construct this moment.
- The moment is then averaged over individuals and the relevant asymptotic theory for inference is one that relies on a large number of individuals N .

- At this point we need to mention two potential complications with the econometrics.
- First, when carrying out inference we have to take into account that $y_{i,a,t}$ has been constructed using the pre-estimated parameters d_t and β in equation (1).
- Second, as said above to estimate such a model we may have to rely on panel data where individuals have been followed for the necessary minimum number of periods/years (6 in our example); this means that our results may be biased due to endogenous attrition.
- The order of the MA process for $v_{i,a,t}$ will not be known in practice and it has to be estimated.
- This can be done by estimating the autocovariance structure of $g_{i,a,t}$ and deciding *a priori* on the suitable criterion for judging whether they should be taken as zero.

Estimating and identifying the properties of the transitory shock.

- The next issue is the identification of the parameters of the moving average process of the transitory shock and those of measurement error.
- It turns out that the model is underidentified, which is not surprising: in our example we need to estimate three parameters, namely the variance of the transitory shock $\sigma_\varepsilon^2 = E(\varepsilon_{i,a,t}^2)$, the MA coefficient θ and the variance of the measurement error $\sigma_m^2 = E(m_{i,a,t}^2)$.
- To illustrate the under identification point suppose that $|\theta| < 1$ and assume that the measurement error is independently and identically distributed.

- We take as given that $q = 1$.
- Then the autocovariances of order higher than three will be zero, whatever the value of our unknown parameters, which is the root of the identification problem.
- The first and second order autocovariances imply

$$\sigma_{\varepsilon}^2 = \frac{E(g_{i,a,t}g_{i,a-2,t-2})}{\theta} \quad I \quad (6)$$

$$\sigma_m^2 = -E(g_{i,a,t}g_{i,a-1,t-1}) - \frac{(1+\theta)^2}{\theta} E(g_{i,a,t}g_{i,a-2,t-2}) \quad II$$

- The sign of $E(g_{i,a,t}g_{i,a-2,t-2})$ defines the sign of θ .

- Taking the two variances as functions of the MA coefficient we note two points.
- First, $\sigma_m^2(\theta)$ declines and $\sigma_\varepsilon^2(\theta)$ increases when θ declines in absolute value.
- Second, for sufficiently low values of $|\theta|$ the estimated variance of the measurement error $\sigma_m^2(\theta)$ may become negative.
- Given the sign of θ (defined by l in equation 6) this fact defines a bound for the MA coefficient.
- Suppose for example that $\theta < 0$, we have that $\theta \in [-1, \tilde{\theta}]$ where $\tilde{\theta}$ is the negative value of θ that sets σ_m^2 in (6) to zero.
- If θ was found to be positive the bounds would be in a positive range.
- The bounds on θ in turn define bounds on σ_ε^2 and σ_m^2 .

- An alternative empirical strategy is to rely on an external estimate of the variance of the measurement error, $\overline{\sigma_m^2}$.
- Define the moments, adjusted for measurement error as:

$$\begin{aligned}
 E \left[g_{i,a,t}^2 - 2\overline{\sigma_m^2} \right] &= \sigma_\zeta^2 + 2(1 + \theta + \theta^2) \sigma_\varepsilon^2 \\
 E \left(g_{i,a,t} g_{i,a-1,t-1} + \overline{\sigma_m^2} \right) &= -(1 + \theta)^2 \sigma_\varepsilon^2 \\
 E \left(g_{i,a,t} g_{i,a-2,t-2} \right) &= \theta \sigma_\varepsilon^2
 \end{aligned}$$

where $\overline{\sigma_m^2}$ is available externally.

- The three moments above depend only on θ , σ_ζ^2 and σ_m^2 .

- We can then estimate these parameters using a Minimum Distance procedure.
- Such external measures can sometimes be obtained through validation studies.
- For example, Bound and Krueger (1991) conduct a validation study of the CPS data on earnings and conclude that measurement error explains 28 percent of the overall variance of the rate of growth of earnings in the CPS.
- Bound et al. (1994) find a value of 22 percent using the PSID-Validation Study.

Estimating Alternative Income Processes

Time varying impacts

- An alternative specification with very different implications is one where

$$\ln Y_{i,a,t} = \rho \ln Y_{i,a-1,t-1} + d_t(X'_{i,a,t}\beta + h_i + v_{i,a,t}) + m_{i,a,t} \quad (7)$$

where h_i is a fixed effect while $v_{i,a,t}$ follows some MA process and $m_{i,a,t}$ is measurement error (see Holtz-Eakin, Newey and Rosen, 1988).

- This process can be estimated by method of moments following a suitable transformation of the model.

- Define $\theta_t = d_t/d_{t-1}$ and quasi-difference to obtain:

$$\ln Y_{i,a,t} = (\rho + \theta_t) \ln Y_{i,a-1,t-1} - \theta_t \rho \ln Y_{i,a-2,t-2} + d_t(\Delta X'_{i,a,t} \beta + \Delta v_{i,a,t}) + m_{i,a,t} - \theta_t m_{i,a-1,t-1} \quad (8)$$

- In this model the persistence of the shocks is captured by the autoregressive component of $\ln Y$ which means that the effects of time varying characteristics are persistent to an extent.
- Given estimates of the levels equation in (8) the autocovariance structure of the residuals can be used to identify the properties of the error term $d_t \Delta v_{i,a,t} + m_{i,a,t} - \theta_t m_{i,a-1,t-1}$.

- Alternatively, the fixed effect with the autoregressive component can be replaced by a random walk in a similar type of model.
- This could take the form

$$\ln Y_{i,a,t} = d_t(X'_{i,a,t}\beta + p_{i,a,t} + v_{i,a,t}) + m_{i,a,t} \quad (9)$$

- In this model $p_{i,a,t} = p_{i,a-1,t-1} + \zeta_{i,a,t}$ as before, but the shocks have a different effect depending on aggregate conditions.

- Given fixed T a linear regression in levels can provide estimates for d_t , which can now be treated as known.
- Now define $\theta_t = d_t/d_{t-1}$ and consider the following transformation

$$\ln Y_{i,a,t} - \theta_t \ln Y_{i,a-1,t-1} = d_t(\zeta_{i,a,t} + \Delta v_{i,a,t}) + m_{i,a,t} - \theta_t m_{i,a-1,t-1} \quad (10)$$

- The autocovariance structure of $\ln Y_{i,a,t} - \theta_t \ln Y_{i,a-1,t-1}$ can be used to estimate the variances of the shocks, very much like in the previous examples.
- In general again we will not be able to identify separately the variance of the transitory shock from that of measurement error, just like before.

- In general, one can construct a number of variants of the above model but we will move on to another important specification, keeping from now on any macroeconomic effects additive.
- It should be noted that (10) is a popular model among labor economists but not among macroeconomists.
- One reason is that it is hard to use in macro models – one needs to know the entire sequence of prices, address general equilibrium issues, etc.

Stochastic growth in Earnings

- Now consider generalizing in a different way the income process and allow the residual income growth (4) to become

$$g_{i,a,t} = f_i + \Delta m_{i,a,t} + (1 + \theta L)\Delta \varepsilon_{i,a,t} + \zeta_{i,a,t} \quad (11)$$

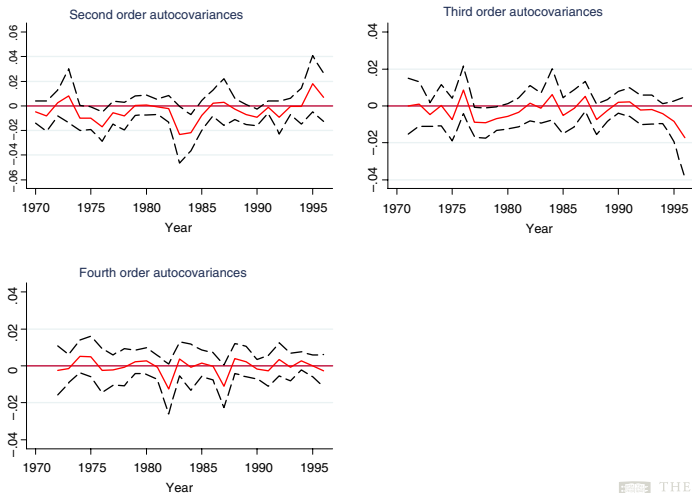
where the f_i is a fixed effect.

- The fundamental difference of this specification from the one presented before is that income growth of a particular individual will be correlated over time.
- In the particular specification above, all theoretical autocovariances of order three or above will be equal to the variance of the fixed effect f_i .
- Consider starting with the null hypothesis that the model is of the form presented in (3) but with an unknown order for the MA process governing the transitory shock $v_{i,a,t} = \Theta_q(L)\varepsilon_{i,a,t}$.

- In practice we will have a panel data set containing some finite number of time series observations but a large number of individuals, which defines the maximum order of autocovariance that can be estimated. In the PSID these can be about 30 (using annual data).
- The pattern of empirical autocovariances consistent with (4) is one where they decline abruptly and become all insignificantly different from zero beyond that point.
- The pattern consistent with (11) is one where the autocovariances are never zero but after a point become all equal to each other, which is an estimate of the variance of f_i .

- Evidence reported in MaCurdy (1982), Abowd and Card (1989), Topel and Ward (1992), Gottschalk and Moffitt (1994), Meghir and Pistaferri (2004) and others all find similar results: Autocovariances decline in absolute value, they are statistically insignificant after the 1st or 2nd order, and have no clear tendency to be positive.
- They interpret this as evidence that there is no random growth term.
- Figure 2 use PSID data and plot the second, third and fourth order autocovariances of earnings growth (with 95% confidence intervals) against calendar time.
- They confirm the findings in the literature: After the second lag no autocovariance is statistically significant for any of the years considered, and there are as many positive estimates as negative ones.
- In fact, there is no clear pattern in these estimates.

Figure 2: Second to fourth order autocovariances of earnings growth, PSID 1967-1997.



- With a long enough panel and a large number of cross sectional observations we should be able to detect the difference between the two patterns.
- However, there are a number of practical and theoretical difficulties.
- First, with the usual panel data, the higher order autocovariances are likely to be estimated based on a relatively low number of individuals.

- The other issue is that without a clearly articulated hypothesis we may not be able to distinguish among many possible alternatives, because we do not know the order of the MA process, q , or even if we should be using an MA or AR representation, or if the "permanent component" has a unit root or less.
- If we did, we could formulate a method of moments estimator and, subject to the constraints from the amount of years we observe, we could estimate our model and test our null hypothesis.

- The practical identification problem is well illustrated by an argument in Guvenen (2009).
- Consider the possibility that the component we have been referring to as permanent, $p_{i,a,t}$, does not follow a random walk, but follows some stationary autoregressive process.

- In this case the increase in the variance over the lifecycle will be captured by the term $a \times f_i$.
- The theoretical autocovariances of $g_{i,a,t}$ will never become exactly zero; they will start negative and gradually increase asymptoting to a positive number which will be the variance of f_i , say σ_f^2 .
- Specifically if $p_{i,a,t} = \rho p_{i,a-1,t-1} + \zeta_{i,a,t}$ with $|\rho| < 1$, there is no other transitory stochastic component, and the variance of the initial draw of the permanent component is zero, the autocovariances of order k have the form

$$E(g_{i,a,t}g_{i,a-k,t-k}) = \sigma_f^2 + \rho^{k-1} \left[\frac{\rho - 1}{\rho + 1} \right] \sigma_\zeta^2 \quad \text{for } k > 0 \quad (12)$$

- As ρ approaches one the autocovariances will approach σ_f^2 .
- However, the autocovariance in (12) is the sum of a positive and a negative component.
- Guvenen (2009) has shown based on simulations that it is almost impossible in practice with the usual sample sizes to distinguish the implied pattern of the autocovariances from (12) from the one estimated from PSID data.
- The key problem with this is that the usual panel data that is available either follows individuals for a limited number of time periods, or suffers from severe attrition, which is probably not random, introducing biases.
- Thus, in practice it is very difficult to identify the nature of the income process without some prior assumptions and without combining information with another process, such as consumption or labour supply.

- Haider and Solon (2006) provide a further illustration of how difficult is to distinguish one model from the other.
- They are interested in the association between current and lifetime income.
- They write current log earnings as

$$y_{i,a,t} = h_i + af_i$$

and lifetime earnings as (approximately)


$$\log V_i = r - \log r + h_i + r^{-1}f_i$$

- The slope of a regression of $y_{i,a,t}$ onto $\log V_i$ is:

$$\lambda_a = \frac{\sigma_h^2 + r^{-1}a\sigma_f^2}{\sigma_h^2 + r^{-1}\sigma_f^2}$$

- Hence, the model predicts that λ_a should increase linearly with age.
- In the absence of a random growth term ($\sigma_f^2 = 0$), $\lambda_a = 1$ at all ages.
- Figure 3, reproduced from Haider and Solon (2006) shows that there is evidence of a linear growth in λ_a only early in the life cycle (up until age 35); however, between age 35 and age 50 there is no evidence of a linear growth in λ_a (if anything, there is evidence that λ_a declines and one fails to reject the hypothesis $\lambda_a = 1$); finally, after age 50, there is evidence of a decline in λ_a that does not square well with any random growth term in earnings.

Figure 3: Estimates of λ_a from Haider and Solon (2006).



Other Enrichments/Issues

- The literature has addressed many other interesting issues having to do with wage dynamics, which here we only mention in passing.
- First, the importance of firm or match effects.
- Matched employer-employee data could be used to address these issues, and indeed some papers have taken important steps in this direction (see Abowd, Kramaz and Margolis, 1999; Postel Vinay and Robin, 2002; Guiso, Pistaferri and Schivardi, 2005).
- A number of papers have remarked that wages fall dramatically at job displacement, generating so-called "scarring" effects (Jacobson, Lalonde and Sullivan, 1993; von Wachter, Song and Manchester, 2007).

- The nature of these scarring effects is still not very well understood.
- On the one hand, people may be paid lower wages after a spell of unemployment due to fast depreciation of their skills (Ljunqvist and Sargent, 1998).
- Another explanation could be loss of specific human capital that may be hard to immediately replace at a random firm upon re-entry (see Low, Meghir and Pistaferri, 2010).

The conditional variance of earnings

- The typical empirical strategy followed in the precautionary savings literature in the attempt to understand the role of risk in shaping household asset accumulation choices typically proceeds in two steps.
- In the first step, risk is estimated from a univariate ARMA process for earnings (similar to one of those described earlier).
- Usually the variance of the residual is the assumed measure of risk.
- There are some variants of this typical strategy- for example, allowing for transitory and permanent income shocks.

- In the second step, the outcome of interest (assets, savings, or consumption growth) is regressed onto the measure of risk obtained in the first stage, or simulations are used to infer the importance of the precautionary motive for saving.
- Examples include Banks, Blundell and Brugiavini (2001) and Zeldes (1989).
- In one of the earlier attempts to quantify the importance of the precautionary motive for saving, Caballero (1990) concluded –using estimates of risk from MaCurdy (1982)- that precautionary savings could explain about 60% of asset accumulation in the US.
- A few recent papers have taken up the issue of risk measurement (i.e., modeling the conditional variance of earnings) in a more complex way.
- Here we comment primarily on Meghir and Pistaferri (2004).

Meghir and Pistaferri (2004)

- Returning to the model presented in section 13 we can extend this by allowing the variances of the shocks to follow a dynamic structure with heterogeneity.
- A relatively simple possibility is to use ARCH(1) structures of the form

$$E_{t-1}(\varepsilon_{i,a,t}^2) = \gamma_t + \gamma \varepsilon_{i,a-1,t-1}^2 + \nu_i \quad \textit{Transitory} \quad (13)$$

$$E_{t-1}(\zeta_{i,a,t}^2) = \varphi_t + \varphi \zeta_{i,a-1,t-1}^2 + \xi_i \quad \textit{Permanent}$$

where $E_{t-1}(\cdot)$ denotes an expectation conditional on information available at time $t - 1$.

- The parameters are all education-specific.
- Meghir and Pistaferri (2004) test whether they vary across education.
- The terms γ_t and φ_t are year effects which capture the way that the variance of the transitory and permanent shocks change over time, respectively.
- In the empirical analysis they also allow for life-cycle effects.
- In this specification we can interpret the lagged shocks $(\varepsilon_{i,a-1,t-1}, \zeta_{i,a-1,t-1})$ as reflecting the way current information is used to form revisions in expected risk.
- Hence it is a natural specification when thinking of consumption models which emphasize the role of the *conditional* variance in determining savings and consumption decisions.

- The terms ν_i and ξ_i are fixed effects that capture all those elements that are invariant over time and reflect long term occupational choices, etc.
- The latter reflects permanent variability of income due to factors unobserved by the econometrician.
- Such variability may in part have to do with the particular occupation or job that the individual has chosen.
- This variability will be known by the individuals when they make their occupational choices and hence it also reflects preferences.
- Whether this variability reflects permanent risk or not is of course another issue which is difficult to answer without explicitly modeling behavior.

- As far as estimating the mean and variance process of earnings is concerned, this model does not require the explicit specification of the distribution of the shocks; moreover the possibility that higher order moments are heterogeneous and/or follow some kind of dynamic process is not excluded.
- In this sense it is very well suited for investigating some key properties of the income process.
- Indeed this is important, because as we will see later on the properties of the variance of income will have implications for consumption and savings.

- However, this comes at a price: first, Meghir and Pistaferri (2004) need to impose linear separability of heterogeneity and dynamics in both the mean and the variance.
- This allows them to deal with the initial conditions problem without any instruments.
- Second, they do not have a complete model that would allow them to simulate consumption profiles.
- Hence the model must be completed by specifying the entire distribution.

Identification of the ARCH process

- If the shocks ε and ζ were observable it would be straightforward to estimate the parameters of the ARCH process in (13).
- However they are not.
- What we do observe (or can estimate) is $g_{i,a,t} = \Delta m_{i,a,t} + (1 + \theta L)\Delta \varepsilon_{i,a,t} + \zeta_{i,a,t}$. To add to the complication we have already argued that θ is not point identified.

- Nevertheless the following two key moment conditions identify the parameters of the ARCH process, conditional on the unobserved heterogeneity (ν and ξ):

$$E_{t-2} (g_{i,a+q+1,t+q+1} g_{i,a,t} - \theta \gamma_t - \gamma g_{it+q} g_{i,a-1,t-1} - \theta \nu_i) = 0 \quad \text{Transitory}$$

$$E_{t-q-3} \left[g_{i,a,t} \left(\sum_{j=-(1+q)}^{(1+q)} g_{i,a+j,t+j} \right) - \varphi_t - \varphi g_{i,a-1,t-1} \left(\sum_{j=-(1+q)}^{(1+q)} g_{i,a+j-1,t+j-1} \right) - \xi_i \right] = 0 \quad \text{Permanent}$$

(14)

- The important point here is that it is sufficient to know the order of the MA process q .
- We do not need to know the parameters themselves.

- The parameter θ that appears in (14) for the transitory shock is just absorbed by the time effects on the variance or the heterogeneity parameter.
- Hence measurement error, which prevents the identification of the MA process does not prevent identification of the properties of the variance, so long as such error is classical.

- The moments above are conditional on unobserved heterogeneity; to complete identification we need to control for that.
- As the moment conditions demonstrate, estimating the parameters of the variances is akin to estimating a dynamic panel data model with additive fixed effects.
- Typically we should be guided in estimation by asymptotic arguments that rely on the number of individuals tending to infinity and the number of time periods being fixed and relatively short.
- One consistent approach to estimation would be to use first differences to eliminate the heterogeneity and then use instruments dated $t - 3$ for the transitory shock and dated $t - q - 4$ for the permanent one.

- In this case the moment conditions become

$$E_{t-3} (\Delta g_{i,a+q+1,t+q+1} g_{i,a,t} - d_t^T - \gamma \Delta g_{it+q} g_{i,a-1,t-1}) = 0 \quad \text{Transitory}$$

$$E_{t-q-4} \left[\Delta g_{i,a,t} \left(\sum_{j=-(1+q)}^{(1+q)} g_{i,a+j,t+j} \right) - d_t^P - \varphi \Delta g_{i,a-1,t-1} \left(\sum_{j=-(1+q)}^{(1+q)} g_{ia+j-1t+j-1} \right) \right] = 0 \quad \text{Permanent}$$

(15)

where $\Delta x_t = x_t - x_{t-1}$. In practice, however, as Meghir and Pistaferri (2004) found out, lagged instruments suggested above may be only very weakly correlated with the entities in the expectations above.

- An alternative may be to use a likelihood approach, which will exploit all the moments implied by the specification and the distributional assumption; this however may be particularly complicated.
- A convenient approximation may be to use within groups on (14).
- This involves subtracting the individual mean off each expression on the right hand side, i.e. just replace all expressions in (14) by quantities where the individual mean has been removed.
- For example $g_{i,a+q+1,t+q+1}g_{i,a,t}$ is replaced by

$$g_{i,a+q+1,t+q+1}g_{i,a,t} - \frac{1}{T-q-1} \sum_{t=1}^{T-q-1} g_{i,a+q+1,t+q+1}g_{i,a,t}.$$

- Meghir and Pistaferri use individuals observed for at least 16 periods.
- Effectively, while ARCH effects are likely to be very important for understanding behavior, there is no doubt that they are difficult to identify.
- A likelihood based approach, although very complex may ultimately prove the best way forward.

Other Approaches

A summary of existing studies

Table 1: Income process studies

Table 2 Income process studies.

Authors	Year publ.	Data	Measure of income	Specification	Results
Lillard & Willis	1978	1967-73 PSID males	Annual earnings in levels	$u_{i,a,t} = h_i + p_{i,a,t}$ $p_{i,a,t} = \rho p_{i,a-1,t-1} + \zeta_{i,a,t}$	Individual fixed effects explain 73% of cross-sectional variance with no covariates (i.e., $\frac{\sigma_h^2}{\sigma_u^2} = 0.73$). Controls for standard wage equation covariates reduce this share to 60.6%; with additional controls for labor market conditions, the figure is 47.1%. AR shock has little persistence ($\rho = 0.35$ with full covariates, $\rho = 0.406$ with time effects only). ^a
Hause	1980	1964-69 Swedish males aged 21-26	Annual earnings in levels	$y_{i,a,t} = h_i + f_i a + u_{i,a,t}$ $u_{i,a,t} = \rho u_{i,a-1,t-1} + \varepsilon_{i,a,t}$ $\varepsilon_{i,a,t} \sim \text{iid}(0, \sigma_\varepsilon^2)$	Individual heterogeneity in slope and intercept of early-career earnings profile is substantial. Variance of AR innovations declines rapidly with time. In model with stationary process for $u_{i,a,t}$, $\sigma_{hf} < 0$, consistent with tradeoff between initial earnings and wage growth predicted by a human capital model. ^b

Table 2 (continued)

Authors	Year publ.	Data	Measure of income	Specification	Results
MaCurdy	1982	1967-76 PSID continuously married white males	Annual earnings in first-differences and levels, Average hourly wages in first-differences and levels	$u_{i,a,t} = h_i + e_{i,a,t}$ $e_{i,a,t} \sim \text{ARMA}(p, q)$	Estimated variance of individual fixed effect h_i is negative and insignificant, so individual heterogeneity is dropped in main specification. Both measures of income are stationary in first-differences and non-stationary in levels (i.e., the author finds a random walk component in levels). MA(2) or ARMA(1,1) is preferred for first-differences. ARMA(1,2) with a unit root ($\rho = 0.975$ for wages, $\rho = 0.974$ for earnings, not significantly different from 1) is preferred for levels.
Abowd & Card	1989	1969-79 PSID males 1969-79 PSID males excluding SEO 1966-75 NLS males 1971-75 SIME/DIME control group	Annual earnings in first-differences Annual hours in first-differences	$g_{i,a,t}^{\text{earnings}} = \mu v_{i,a,t} + \Delta m_{i,a,t}^{\text{earnings}} + e_{i,a,t}^{\text{earnings}}$ $g_{i,t}^{\text{hours}} = v_{i,a,t} + \Delta m_{i,a,t}^{\text{hours}} + e_{i,a,t}^{\text{hours}}$ <p>$v \sim \text{MA}(2)$, e, m serially uncorrelated. $m^{\text{earnings}} \perp m^{\text{hours}}$, e have unrestricted within period VCV. v, m, e mutually independent</p>	Extensive fitting procedure supports MA(2) for persistent shock v . Loading factor μ would capture behavioral responses to changes in the wage rate ($\mu = 1$ implies proportional changes in hours and earnings at a constant wage). However, changes in earnings do not seem to reflect behavioral responses to wage changes: $\mu = 1.09$ in PSID, 1.35 in PSID excluding SEO, 1.56 in NLS, 1.01 in SIME/DIME: $\mu = 1$ is not rejected in any sample.

(continued on next page)

Table 2 (continued)

Authors	Year publ.	Data	Measure of income	Specification	Results
Topel & Ward	1992	1957-72 LEED file, males only (matched firm-worker administrative records)	Quarterly SS earnings from a single employer, in annual first-differences ^c	$g_{i,a,t} = \Delta \eta_{i,a,t}$ where $\eta_{i,a,t} = p_{i,a,t} + e_{i,a,t}$ contains an AR(1) ($p_{i,a,t}$) + a white noise ($e_{i,a,t}$).	Raw autocovariance of earnings growth is strongly negative at one lag, then is small (insignificant) and negative at higher lags. AR coefficient $\rho = 0.970$ is insignificantly different from 1. Authors conclude on-the-job wage growth is a random walk, and so current wage is a sufficient statistic for the value of a job for early-career workers.
Gottschalk & Moffitt	1995	1969-87 PSID white males	Annual earnings in levels	$u_{iat} = \mu_t p_{i,a,t} + v_{i,a,t}$ $p_{ia} = p_{i,a-1,t-1} + \zeta_{i,a,t}$ $v_{ia} = \rho_t v_{i,a-1,t-1} + e_{i,a,t}$ $+ \lambda_t \theta e_{i,a-1,t-1}$ Legend: Loading of persistent shock (μ_t), AR coefficient (ρ_t), MA coefficient (λ_t), persistent earnings shock ($\zeta_{i,a,t}$), and transitory earnings shock ($e_{a,t}$).	Half the increase in cross-sectional variance is due to increase in the transitory innovation variance, and half is due to increase in the permanent innovation variance. Increase in transitory variance dominated in the second half of the 1980s.

Table 2 (continued)

Authors	Year publ.	Data	Measure of income	Specification	Results
Farber & Gibbons	1996	1979-91 NLSY males and females after 1st transition to work	Hourly wage rate in levels	$u_{i,a,t} = p_{i,a,t} + m_{i,a,t}$ $p_{i,a,t} = p_{i,a-1,t-1} + \zeta_{i,a,t}$	Authors reject hypothesis of martingale with classical measurement error or with AR(1) measurement error. Also run specification with stationary AR(1) in v_{it} and rejects it.
Baker	1997	1967-86 PSID males	Annual earnings in first-differences and levels	<p>Model 1 (HIP):</p> $u_{i,a,t} = h_i + f_i a + p_{i,a,t}$ $g_{i,a,t} = \tilde{f}_i + \Delta p_{i,a,t}$ <p>where $p_{i,a,t} = \rho p_{i,a-1,t-1} + \zeta_{i,a,t}$ (AR(1)).</p> <p>Model 2 (RIP with RW):</p> $u_{i,a,t} = h_i + e_{i,a,t}$ $g_{i,a,t} = \Delta e_{i,a,t}$ $e_{i,a,t} \sim \text{ARMA}(1,2) \text{ or } \text{ARMA}(1,1), \text{ time-varying variances for innovations to } e_{i,a,t} \text{ are estimated in both models.}$	Tests and rejects restrictions of no heterogeneity in growth rates and levels (in OLS estimates of HIP model). RIP specification does not reject RW. Nested model yields $\rho = 0.665$; first-differenced estimates of nested model yield much smaller AR coefficient. Monte Carlo evidence is presented suggesting that joint tests for zero higher-order autocovariances overreject with small samples or a large number of restrictions (as is the case here).

(continued on next page)

Table 2 (continued)

Authors	Year publ.	Data	Measure of income	Specification	Results
Chamberlain & Hirano	1999	1967-1991 PSID males aged 24-33	Annual earnings	$y_{i,a,t} = g_t(x(i, \beta)) + h_i$ $+ p_{i,a,t} + v_{i,a,t}$ $p_{i,a,t} = \rho p_{i,a-1,t-1} + \zeta_{i,a,t}$ Transitory shock $v_{i,a,t}$ heteroskedastic across individuals: $v_{i,a,t} \sim N(0, \frac{1}{h_i})$ $h_{i,a,t} \sim \text{Gamma.}$	Substantial heteroskedasticity in $v_{i,a,t}$ AR coefficient point estimate = 0.98.
Geweke & Keane	2000	1968-89 PSID males	Annual earnings	$y_{i,a,t} = \lambda y_{i,a-1,t-1} + (1 - \lambda)$ $[X_{i,a,t}\beta + h_i + \mu p_{i,0,t-a}] + p_{i,a,t}$ $p_{i,a,t} = \rho p_{i,a-1,t-1} + \zeta_{i,a,t}$ Initial conditions $y_{i,0,t-a} = X_{i,0,t-a}^0 \beta^0 + \zeta_{i,0,t-a}$ Innovations $\zeta_{i,a,t}$ and initial conditions draw $\zeta_{i,0,t-a}$ drawn from mixtures of 3 normals, allowing for non-normality of shocks. Initial conditions depend on different observables (X^0) than do current-period earnings (X). Marital status jointly modeled.	AR coefficient ρ on shock is 0.665, but not directly comparable to other AR coefficients because model includes lagged earnings. 60% to 70% of cross-section variance due to transitory shocks. Strong evidence of non-normality for initial conditions draw $\zeta_{i,0,t-a}$ and innovations $\zeta_{i,a,t}$: both shocks are left skewed and leptokurtic (density at mode about 3 times larger than predicted by normality). Non-normal shocks greatly improve fit to cross-sectional distribution and predictions of economic mobility. Non-normal model has less serial correlation.

Table 2 (continued)

Authors	Year publ.	Data	Measure of income	Specification	Results
Baker & Solon	2003	1975-83 Canadian males (administrative income tax records)	Annual earnings	$u_{i,a,t} = \mu_t[h_i + f_i a + p_{i,a,t}] + e_{i,a,t}$ $p_{i,a,t} = p_{i,a-1,t-1} + \zeta_{i,a,t}$ (random walk in permanent income) $e_{i,a,t} = \rho e_{i,a-1,t-1} + \lambda_t \epsilon_{i,a,t}$ (AR(1) with time-varying variance in transitory income) $\epsilon_{i,a,t} \sim \text{niid}(0, \sigma_{\text{age}}^2)$ (age-varying heteroskedasticity in transitory earnings innovation).	Estimated separately for two-year birth cohorts, both random walk component and profile heterogeneity (HIP and RIP) are important. Restricted specifications ($\sigma_\zeta = 0$, or $\sigma_f = 0$) inflate ρ and attribute more of the variance to transitory shocks (instability) than in the unrestricted model. Transitory innovation variance is U-shaped over the life cycle.
Meghir & Pistaferri	2004	PSID males 1968-1993	Annual earnings in first differences	Three education groups: High School Dropout (D), High School Graduate (H) and College (C). For each education group: $\ln y_{i,a,t} = f(a,t) + p_{i,a,t} + e_{i,a,t} + m_{i,a,t}$ $p_{i,a-1,t-1} + \epsilon_{i,a,t} e_{i,a,t} = \xi_{i,a,t} + \theta \xi_{i,a-1,t-1} m_{i,a,t}$ is i.i.d. measurement error $\epsilon_{i,a,t}$ and $\xi_{i,a,t}$ are serially uncorrelated model conditional variance of shocks as: $E_{t-1}(\epsilon_{i,a,t}) = d_{1t} + \zeta_{1i} + g_1(\text{age}) + \rho e \epsilon_{i,a-1,t-1}^2 E_{t-1}(\xi_{i,a,t}) = d_{2t} + \zeta_{2i} + g_2(\text{age}) + \rho_\xi \xi_{i,a-1,t-1}^2$	Tested for absence of unit root using autocovariance structure and reject. Error process set to random walk plus MA(1) transitory shock plus measurement error. Variances of shocks (permanent, transitory) D:(0.033, 0.055), H:(0.028, 0.027), C:(0.044, 0.005) pooled: (0.031, 0.030); ARCH effects (permanent, transitory): D:(0.33, 0.19), H:(0.89, 0.67), C:(0.028, 0.39), pooled: (0.56, 0.40)

(continued on next page)

Table 2 (continued)

Authors	Year publ.	Data	Measure of income	Specification	Results
Haider & Solon	2006	1951-91 HRS-SSA matched panel males ^d	Annual earnings (observe SS-taxable earnings)	Assume panel distribution of log yearly earnings $y_{i,a,t}$ is MVN, i.e., log earnings normal in each year, jointly distributed MVN. The authors can then impute censored earnings with a Tobit in each year. Pairwise ACVs across all years in panel are estimated with separate bivariate Tobits.	Measurement error and transitory shocks imply that annual earnings in any given year are a poor proxy for lifetime earnings in that it is subject to non-classical measurement error that varies over the life cycle. ^c
Browning, Alvarez, & Ejrnaes	2006	1968-93 PSID white males	Annual after-tax earnings	For each individual/age: $y_t = \delta(1 - \omega^t) + \alpha t + \beta^t y_0 + \sum_{s=0}^{t-1} \beta^s (\epsilon_{t-s} + \theta \epsilon_{t-s-1})^f$ $y_t^{obs} = y_t + m_t$ (classical measurement error), $\epsilon \sim \text{ARCH}(1)$ and m i.i.d. Individual heterogeneity allowed in $(v, \theta, \alpha, \beta, \delta, \omega)$. Distributions are parametrized as linear or logistic (for restricted parameters) functions of 6 independent normal latent factors.	The model is estimated under different assumptions regarding AR coefficient β : (1) β is a unit root for everyone, (2) $\beta < 1$ for everyone, and (3) β is a mixture of a unit root and a stable AR. Of these, a model where $\beta < 1$ for all agents is the only one not conclusively rejected by χ^2 tests. The median AR coefficient is 0.79.

Table 2 (continued)

Authors	Year publ.	Data	Measure of income	Specification	Results
Hryshko	2008	1968-97 PSID males excluding SEO	Annual earnings, first-differences and levels	$u_{i,a,t} = h_i + f_{i,a} + p_{i,a,t} + v_{i,a,t} + m_{i,a,t}$ $p_{i,a,t} = p_{i,a-1,t-1} + \zeta_{i,a,t}$ $v_{i,a,t} = \theta(L)\epsilon_{i,a,t}, \text{ i.e.,}$ heterogeneous intercept and slope, measurement error, RW in permanent income, and MA in transitory component.	Estimates in first-differences with σ_m^2 fixed at point estimate from another specification yield no heterogeneity in growth rates.
Altonji, Smith & Vidangos	2009	1978-96 PSID males	Annual earnings. Hours, wages, job transitions also used	$y_{i,a,t} = \gamma_0 + \gamma_X X_{i,a,t} + \gamma_w (w^{\text{latent}} - \gamma_0 - \gamma_X X_{i,a,t}) + \gamma_h (h_{i,a,t} - \gamma_0 - \gamma_X X_{i,a,t}) + e_{i,a,t}$ $e_{i,a,t} = \rho_e e_{i,a-1,t-1} + \epsilon_{i,a,t} y_{i,a,t}$ is log wages (not the residual); wage w and hours h are endogenous, with their own dynamic error structure. This is a joint statistical model of employment transitions, wages, hours worked, and earnings. ^h Discrete outcomes (employment transitions) are probit (usually with multiple error components): all shocks are independent normals. Wages, hours, and earnings are log-linear. Other important aspects: wage and hours include two individual fixed effects: μ_i ("ability") appears in all structural equations; η_i ("mobility") appears in all but the wage equation.	Authors present some simulated variance decompositions for lifetime and cross-sectional log earnings (not residuals) among white males. Earnings shocks and hours shocks contribute more than twice as much to cross-sectional variance than they do to lifetime variance (25% vs. 9% for both shocks combined). Search frictions (job-specific wage/hours shocks, job destruction, and job-to-job changes) generate 37% of variance in lifetime earnings, with job-specific wage shocks most important. Ability (μ) generates 11% of lifetime earnings variance, and education generates 31.4% of variance. ^g

(continued on next page) OF



Table 2 (continued)

Authors	Year publ.	Data	Measure of income	Specification	Results
Guvenen	2009	1968-93 PSID males	Annual earnings in levels	$u_{i,a,t} = h_i + f_i a + p_{i,a,t} + \mu_t v_{i,a,t}$ $p_{i,a,t} = \rho p_{i,a-1,t-1} + \lambda_t \zeta_{i,a,t}$ $v_{i,a,t} \sim \text{i.i.d.}$	Estimates of the process with slope heterogeneity yield estimates of AR coefficient ρ significantly below 1 (0.821 in the full sample), while estimates without heterogeneity ($\sigma_f = 0$) indicate a random walk in permanent income. MaCurdy's (1982) test for heterogeneity is criticized for low power regarding higher-order autocovariances.
Low, Meghir & Pistaferri	2010	SIPP	Hourly rate in first differences	$w_{i,j(t_0),a,t} = p_{i,a,t} + \epsilon_{i,a,t} + v_{i,j(t_0),a,t}$ $p_{i,a,t} = p_{i,a-1,t-1} + \zeta_{i,a,t}$ where $v_{i,j(t_0),a,t}$ is a match fixed effect. Allow for job mobility and participation. Estimates parameters using wage growth moments and allows for endogenous selection due to job mobility and employment.	Estimated standard deviation of permanent shocks is 0.10, of the match effect 0.23 and of the measurement error 0.09. Ignoring mobility increases st. dev of permanent shock to 0.15.

^a Authors cut sample by race (black/white).

^b No covariates, so profile heterogeneity captures differences across education groups (focus is on low education workers).

^c I.e., $s_{i,a,t} = y_{i,a,t} - y_{i,a-4,t-4}$ where t indexes quarters.

^d 1931-33 birth cohort only.

^e Sample average estimated ACVs pooled over full earnings history (from bivariate Tobit procedures) are very close to results from uncensored data in other studies (Baker and Solon, 2003, Böhlmark and Lindquist, 2006): ACV1 = 0.89, ACV2 = 0.82, ACV3 = 0.78, ACV4 = 0.75, ACV5 = 0.72, ACV6 = 0.69.

^f $[\delta]$ = "long-run" average earnings; $[\omega]$ = inverse speed of convergence to "long-run" average earnings; $[\alpha]$ = linear time trend; $[\beta]$ = AR(1) coefficient; $[\theta]$ = MA(1) coefficient; $[\epsilon]$ = ARCH WN, with constant v , ARCH coefficient $\frac{\exp(\phi)}{1+\exp(\phi)}$.

^g Parametrization of the model makes it difficult to compare point estimates to other results from the literature. Results for impulse-response to particular shocks are interesting results, but the less detailed models in the income-process literature reviewed here typically present unconditional dynamic behavior rather than distinguishing particular shocks.

^h "Joint" in the sense that it is more complex than the univariate earnings processes presented here, but still based only on labor market behavior; "statistical" in the sense that the model's structural equations are not derived from utility maximization.

- Guvenen (2009) compares what he calls a HIP (heterogeneous income profiles) income process and a RIP (restricted income profiles) income process and their empirical implications.
- The income process (in a simplified form) is as follows:

$$y_{i,a,t} = X'_{i,a,t}\beta_t + h_i + a \times f_i + p_{i,a,t} + d_t \varepsilon_{i,a,t}$$

$$p_{i,a,t} = \rho p_{i,a-1,t-1} + \varphi_t \zeta_{i,a,t}$$

with an initial condition equal to 0.

- Hryshko (2009) in an important paper sets out to resolve the random walk vs. stochastic growth process controversy by carrying out Monte Carlo simulations and empirical analysis on PSID data.
- First, he generates data based on a process with a random walk and persistent transitory shocks.
- He then fits a (misspecified) model assuming heterogeneous age profiles and an AR(1) component and finds that the estimated persistence of the AR component is biased downwards and that there is evidence for heterogeneous age profile.

- In the empirical data he finds that the model with the random walk cannot be rejected, while he finds little evidence in support of the model with heterogeneous growth rates.
- While these results are probably not going to be viewed as conclusive, what is clear is that the encompassing model of, say, Baker (1997) may not be a reliable way of testing the competing hypotheses.
- It also shows that the evidence for the random walk is indeed very strong and reinforces the results by Baker and Solon (2003), which support the presence of a unit root as well as heterogeneous income profiles.

- Browning, Ejrnaes and Alvarez (2006) extend this idea further by allowing the entire income process to be heterogeneous.
- Their model allows for all parameters of the income process to be different across individuals, including a heterogeneous income profile and a heterogeneous serial correlation coefficient restricted to be in the open interval $(0,1)$.
- This stable model is then mixed with a unit root model, with some mixing probability estimated from the data.
- This then implies that with some probability an individual faces an income process with a unit root; alternatively the process is stable with heterogeneous coefficients.

- They estimate their model using the same PSID data as Meghir and Pistaferri (2004) and find that the median AR(1) coefficient is 0.8, with a proportion of individuals (about 30%) having an AR(1) coefficient over 0.9.
- They attribute their result to the fact that they have decoupled the serial correlation properties of the shocks from the speed of convergence to some long run mean, which is governed by a different coefficient.