Earnings, Consumption and Lifecycle Choices

by Costas Meghir and Luigi Pistaferri

from Chapter 9 of *Handbook of Labor Economics, Volume 4b*, 2011

James J. Heckman
University of Chicago

Econ 350, Winter 2021
This draft, February 23, 2021 7:27pm

Introduction

- Distinction is between *ex-ante* and *ex-post* household responses to risk.
- *Ex-ante* responses answer the question:
  "What do people do in the anticipation of shocks to their economic resources?".
- *Ex-post* responses answer the question:
  "What do people do when they are actually hit by shocks to their economic resources?".

The Life Cycle-Permanent Income Hypothesis

- To see how the degree of persistence of income shocks and the nature of income changes affects consumption
- Consider a simple example in which income is the only source of uncertainty of the model.
- Preferences are quadratic, consumers discount the future at rate $\frac{1-\beta}{\beta}$ and save on a single risk-free asset with deterministic real return r, $\beta(1+r)=1$ (this precludes saving due to returns outweighing impatience), the horizon is finite (the consumer dies with certainty at age A and has no bequest motive for saving), and credit markets are perfect.

The Life Cycle-Permanent Income Hypothesis

- $c_{i,a,t}$ is consumption at age a and time t; $y_{i,a,t}$ is income at age a and time t
- $\Omega_{i,a,t}$: information set of agent i, at age a in year t.
- The change in household consumption can be written as

$$\Delta c_{i,a,t} = \pi_a \sum_{j=0}^{A} \frac{E\left(y_{i,a+j,t+j}|\Omega_{i,a,t}\right) - E\left(y_{i,a+j,t+j}|\Omega_{i,a-1,t-1}\right)}{(1+r)^j} \quad (1)$$

  a indexes age and t time,
- $\pi_a = \frac{r}{1+r}\left[1 - \frac{1}{(1+r)^{A-a+1}}\right]^{-1}$: "annuity" parameter that increases with age and $\Omega_{i,a,t}$ is the consumer's information set at age a.
- Three key issues regarding the response of consumption to changes in the economic resources of the household.

- First, consumption responds to news in the income process, but not to expected changes (heterogeneity vs uncertainty).
- The second key issue emerging from equation (1) is that the life cycle horizon also plays an important role (the term π_a).

- The last key feature of equation (1) is the persistence of innovations (how they affect the information sets).
- More persistent innovations have a larger impact than short-lived innovations.
- Suppose that income follows an ARMA(1,1) process:

$$y_{i,a,t} = \underbrace{\rho y_{i,a-1,t-1}}_{\text{AR}} + \underbrace{\varepsilon_{i,a,t} + \theta\varepsilon_{i,a-1,t-1}}_{\text{MA1}} \tag{2}$$

- In this case, substituting (2) in (1), the consumption response is given by

$$\begin{aligned}\Delta c_{i,a,t} =& \left(\frac{r}{1+r}\right)\left[1-\frac{1}{(1+r)^{A-a+1}}\right]^{-1} \\ & \left[1+\frac{\rho+\theta}{1+r-\rho}\left(1-\left(\frac{\rho}{1+r}\right)^{A-a}\right)\right]\varepsilon_{i,a,t} \\ =& \kappa\left(r,\rho,\theta,A-a\right)\varepsilon_{i,a,t}\end{aligned}$$

- Table 1 shows the value of the marginal propensity to consume κ for various combinations of ρ, θ, and $A - a$ (setting $r = 0.02$).

Table 1: The response of consumption to income shocks under quadratic preferences

ρ	θ	$A-a$	κ
1	-0.2	40	0.81
1	0	10	1
0.99	-0.2	40	0.68
0.95	-0.2	40	0.39
0.8	-0.2	40	0.13
0.95	-0.2	30	0.45
0.95	-0.2	20	0.53
0.95	-0.2	10	0.65
0.95	-0.1	40	0.44
0.95	-0.01	40	0.48
1	0	∞	1
0	-0.2	40	0.03

- A number of facts emerge.
- If the income shock represents an innovation to a random walk process ($\rho = 1, \theta = 0$), consumption responds one-to-one to it regardless of the horizon (the response is attenuated only if shocks end after some period, say $L < A$).
- A decrease in the persistence of the shock lowers the value of κ. When $\rho = 0.8$ (and $\theta = -0.2$) for example, the value of κ is a modest 0.13.

- A decrease in the persistence of the MA component acts in the same direction (but the magnitude of the response is much attenuated).
- In this case as well, the presence of liquidity constraints may invalidate the sharp prediction of the model.
- For example, more and less persistent shocks may have a similar effect on consumption.
- When the consumer is hit by a short-lived negative shock, she can smooth the consumption response over the entire horizon by borrowing today (and repaying in the future when income reverts to the mean).
- If borrowing is precluded, a short-lived or long-lived shock have similar impacts on consumption.

- The income process (2) considered above is restrictive, because there is a single error component which follows an ARMA(1,1) process.
- A very popular characterization in calibrated macroeconomic models is to assume that income is the sum of a random walk process and a transitory i.i.d. component:

$$y_{i,a,t} = p_{i,a,t} + \varepsilon_{i,a,t} \tag{3}$$

$$p_{i,a,t} = p_{i,a-1,t-1} + \zeta_{i,a,t} \tag{4}$$

- The appeal of this income process is that it is close to the process used in a Friedman's permanent income hypothesis.

- In this case, the response of consumption to the two types of shocks is:

$$\Delta c_{i,a,t} = \pi_a \varepsilon_{i,a,t} + \zeta_{i,a,t} \tag{5}$$

- Consumption responds one-to-one to permanent shocks but the response of consumption to a transitory shock depends on the time horizon.
- For young consumers (with a long time horizon), the response should be small.
- The response should increase as consumers age.

- Figure 1 plots the value of the response for a consumer who lives until age 75.

Figure 1: The response of consumption to a transitory income shock

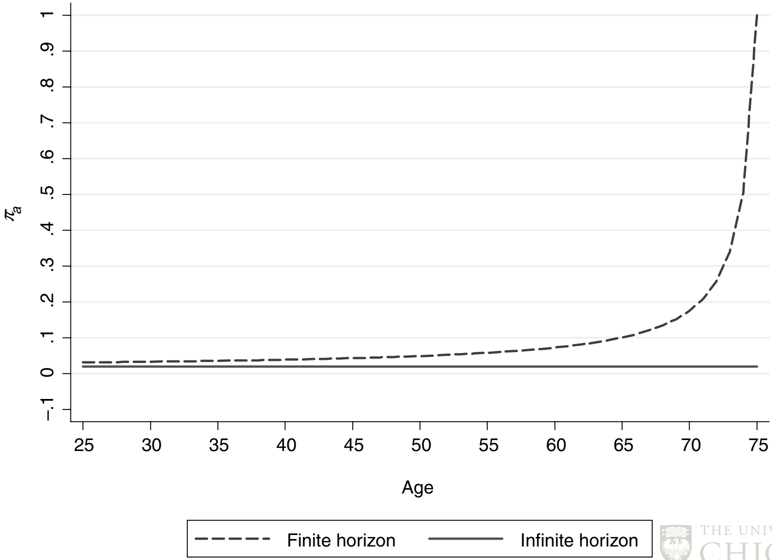

- Clearly, it is only in the last 10 years of life or so that there is a substantial response of consumption to a transitory shock.
- The graph also plots for the purpose of comparison the expected response in the infinite horizon case.
- An interesting implication of this graph is that a transitory unanticipated stabilization policy is likely to affect substantially only the behavior of older consumers (unless liquidity constraints are important—which may well be the case for younger consumers).

- Note finally that if the permanent component were literally permanent ($p_{i,a,t} = p_i$), it would affect the level of consumption but not its change.
- In the classical version of the LC-PIH the *size* of income changes does not matter.
- One reason why the size of income changes may matter is because of adjustment costs: Consumers tend to smooth consumption and follow the theory when expected income changes are large, but are less likely to do so when the changes are small and the cost of adjusting consumption are not trivial.

- Suppose for example that consumers who want to adjust their consumption upwards in response to an expected income increase need to face the cost of negotiating a loan with a bank.
- This "magnitude hypothesis" has been formally tested by Scholnick (2010), who use a large data set provided by a Canadian bank that includes information on both credit cards spending as well as mortgage payment records.
- As in Stephens (2008) he argues that the final mortgage payment represent an expected shock to disposable income (that is, income net of pre-committed debt service payments).

- Outside the quadratic preference world, uncertainty about future income realizations will also impact consumption.

Beyond the PIH

- The model with quadratic preferences gives very sharp predictions regarding the impact on consumption of various types of income shocks.
- For example, there is the sharp prediction that permanent shocks are entirely consumed (an MPC of 1).
- Unfortunately, quadratic preferences have well known undesirable features, such as increasing risk aversion with wealth and lack of a precautionary motive for saving.
- Do the prediction of this model survive under more realistic assumptions about preferences?
- The answer is: only qualitatively.

- The problem with more realistic preferences, such as CRRA, is that they deliver no closed form solution for consumption — that is, there is no analytical expression for the "consumption function" and hence the value of the propensity to consume in response to risk (income shocks) is not easily derivable.
- This is also the reason why the literature moved on to estimating Euler equations after Hall (1978).
- The advantage of the Euler equation approach is that one can be silent about the sources of uncertainty faced by the consumer (including crucially the stochastic structure of the income process).

- However, in the Euler equation approach only a limited set of parameters (preference parameters such as the elasticity of intertemporal substitution or the intertemporal discount rate) can be estimated.
- Some dissatisfaction in the literature regarding the evidence coming from Euler equation estimates (see Browning and Lusardi, 1996; Attanasio and Weber, 2010).

- Recently there has been an attempt to go back to the concept of a "consumption function".
- Two approaches have been followed.
- First, the Euler equation that describe the expected dynamics of the growth in the marginal utility can be approximated to describe the dynamics of consumption growth.
- Blundell, Pistaferri and Preston (2008), extending Blundell and Preston (1998) (see also Blundell and Stoker, 1994), derive an approximation of the mapping between the expectation error of the Euler equation and the income shock.
- Second, there is work by Carroll (2001) and Kaplan and Violante (2010) discuss numerical simulations in the buffer-stock and Bewley model, respectively.
- We discuss the results of these two approaches in turn.

Approximation of the Euler equation

- Blundell, Pistaferri and Preston (2008) consider the consumption problem faced by household i of age a in period t.
- Assuming that preferences are of the CRRA form, the objective is to choose a path for consumption C so as to:

$$\max_{C} E_a \sum_{j=0}^{A-a} \beta^j \frac{C_{i,a+j,t+j}^{1-\gamma} - 1}{1-\gamma} e^{Z'_{i,a+j,t+j}\vartheta_{a+j}}. \qquad (6)$$

where $Z_{i,a+j,t+j}$ incorporates taste shifters (such as age, household composition, etc.), and we denote with $E_a(.) = E(.|\Omega_{i,a,t})$.

- Maximization of (6) is subject to the budget constraint which in the self-insurance model assumes individuals have access to a risk free bond with real return r

$$A_{ia+j+1} = (1+r)\left(A_{i,a+j,t+j} + Y_{i,a+j,t+j} - C_{i,a+j,t+j}\right) \quad (7)$$
$$A_{i,A} = 0 \quad (8)$$

with $A_{i,a,t}$ given.

- Blundell, Pistaferri and Preston (2008) set the retirement age after which labor income falls to zero at L, assumed known and certain, and the end of the life-cycle at age A.

- They assume that there is no uncertainty about the date of death.
- With budget constraint (7), optimal consumption choices can be described by the Euler equation (assuming for simplicity that there is no preference heterogeneity, or $\vartheta_a = 0$):

$$C_{i,a-1,t-1}^{-\gamma} = \beta \left(1 + r\right) E_{a-1} C_{i,a,t}^{-\gamma}. \tag{9}$$

- As it is, equation (9) is not useful for empirical purposes.

- Blundell, Pistaferri and Preston (2008) show that the Euler equation can be approximated as follows:

$$\Delta \log C_{i,a,t} \simeq \eta_{i,a,t} + f_{i,a,t}^{C}$$

where $\eta_{i,a,t}$ is a consumption shock with $E_{a-1}(\eta_{i,a,t}) = 0$, $f_{i,a,t}^{c}$ captures any slope in the consumption path due to interest rates, impatience or precautionary savings and the error in the approximation is $O(E_a \eta_{i,a,t}^2)$.

- Suppose that any idiosyncratic component to this gradient to the consumption path can be adequately picked up by a vector of deterministic characteristics $\Gamma_{i,a,t}^{c}$ and a stochastic individual element $\xi_{i,a,t}$

$$\Delta \log C_{i,a,t} - \Gamma_{i,a,t}^{c} = \Delta c_{i,a,t} \simeq \eta_{i,a,t} + \xi_{i,a,t}.$$

- Assume log income is

$$\log Y_{i,a,t} = p_{i,a,t} + \varepsilon_{i,a,t} \quad (10)$$
$$p_{i,a,t} = \Gamma^{y}_{i,a,t} + p_{i,a-1,t-1} + \zeta_{i,a,t} \quad (11)$$

where $\Gamma^{y}_{i,a,t}$ represent observable characteristic influencing the growth of income.

- Income growth can be written as:

$$\Delta \log Y_{i,a,t} - \Gamma^{y}_{i,a,t} = \Delta y_{i,a,t} = \zeta_{i,a,t} + \Delta \varepsilon_{i,a,t}.$$

- The (ex-post) intertemporal budget constraint is

$$\sum_{j=0}^{A-a} \frac{C_{i,a+j,t+j}}{(1+r)^j} = \sum_{j=0}^{L-a} \frac{Y_{i,a+j,t+j}}{(1+r)^j} + A_{i,a,t}$$

where A is the age of death and L is the retirement age.

- Applying the approximation above and taking differences in expectations gives

$$\eta_{i,a,t} \simeq \Xi_{i,a,t} \left[\zeta_{i,a,t} + \pi_a \varepsilon_{i,a,t}\right]$$

where π_a is an annuitization factor defined below in technical notes, $\Xi_{i,a,t} = \frac{\sum_{j=0}^{A-a} \frac{Y_{i,a+j,t+j}}{(1+r)^j}}{\sum_{j=0}^{A-a} \frac{Y_{i,a+j,t+j}}{(1+r)^j} + A_{i,a,t}}$ is the share of future labor income in current human and financial wealth, and the error of the approximation is $O(\left[\zeta_{i,a,t} + \pi_a \varepsilon_{i,a,t}\right]^2 + E_{a-1}\left[\zeta_{i,a,t} + \pi_a \varepsilon_{i,a,t}\right]^2)$.

Link to Appendix

- Then

$$\Delta \log C_{i,a,t} \simeq \xi_{i,a,t} + \Xi_{i,a,t}\zeta_{i,a,t} + \pi_a \Xi_{i,a,t}\varepsilon_{i,a,t} \qquad (12)$$

with a similar order of approximation error.

- The random term $\xi_{i,a,t}$ can be interpreted as the innovation to higher moments of the income process.

- The interpretation of the impact of income shocks on consumption growth in the PIH model with CRRA preferences is straightforward.
- For individuals a long time from the end of their life with the value of current financial assets small relative to remaining future labor income, $\Xi_{i,a,t} \simeq 1$, and permanent shocks pass through more or less completely into consumption whereas transitory shocks are (almost) completely insured against through saving.
- Precautionary saving can provide effective self-insurance against permanent shocks only if the stock of assets built up is large relative to future labor income, which is to say $\Xi_{i,a,t}$ is appreciably smaller than unity, in which case there will also be some smoothing of permanent shocks through self insurance.

- The most important feature of the approximation approach is to show that the effect of an income shock on consumption depends not only on the persistence of the shock and the planning horizon (as in the LC-PIH case with quadratic preferences), but also on preference parameters.
- *Ceteris paribus*, the consumption of more prudent households will respond less to income shocks.
- The reason is that they can use their accumulated stock of precautionary wealth to smooth the impact of the shocks (for which they were saving precautiously against in the first place).
- Simulation results (below) confirm this basic intuition.

Kaplan and Violante

- Kaplan and Violante (2010) investigate the amount of consumption insurance present in a life-cycle version of the standard incomplete markets model with heterogenous agents (e.g., Rios-Rull, 1995; Huggett, 1996).
- Kaplan and Violante's setup differs from that in Blundell, Pistaferri and Preston by adding the uncertainty component μ_a to life expectancy, and by omitting the taste shifters from the utility function.
- μ_a is the probability of dying at age a.
- It is set to 0 for all $a < L$ (the known retirement age) and it is greater than 0 for $L \leq a \leq A$.

- The KV model also differs from BPP by specifying a realistic social security system.
- Two baseline setups are investigated - a natural borrowing constraint setup (henceforth NBC), in which consumers are only constrained by their budget constraint, and a zero borrowing constraint setup (henceforth ZBC), in which consumers have to maintain non-negative assets at all ages.
- The income process is similar to BPP.
- Part of KV's analysis is designed to check whether the amount of insurance predicted by the Bewley model can be consistently estimated using the identification strategy proposed by BPP and whether BPP's estimates using PSID and CEX data conform to values obtained from calibrating their theoretical model.

- Kaplan and Violante (2010) model is calibrated to match the US data.
- Survival rates are obtained from the NCHS, the intertemporal discount rate is calibrated to match a wealth-income ratio of 2.5, the permanent shock parameters (σ_ζ^2 and the variance of the initial draw of the process) are calibrated to match PSID data and the variance of the transitory shock (σ_ε^2) is set to the 1990-1992 BPP point estimate (0.05).
- The KV model is solved numerically.
- This allows for the calculation of both the "true" and the BPP estimators of the "partial insurance parameters" (the response of consumption to permanent and transitory income shocks).

- Figure 2 is reproduced from Kaplan and Violante (2010).

Figure 2: Age profile of MPC coefficients for transitory and permanent income shocks. (Source: Kaplan and Violante (2010)).

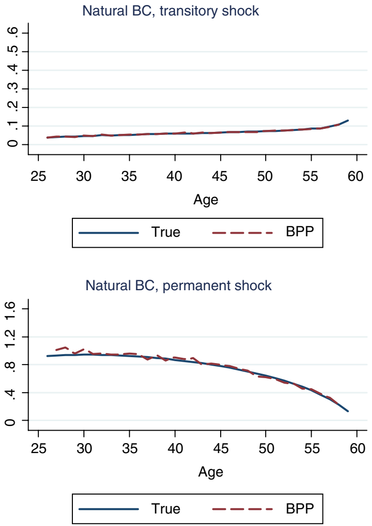

- First, in the NBC environment the MPC with respect to transitory shocks is fairly low throughout the life cycle, and similarly to what is shown in Figure 1, increases over the life cycle due to reduced planning horizon effect. The life cycle average MPC is 0.06.
- Second, there is considerable insurance also against the permanent shock, which increases over the life cycle due to the ability to use the accumulated wealth to smooth these shocks. The life cycle average MPC is 0.77, well below the MPC of 1 predicted by the infinite horizon PIH model.

- Third, the ZBC environment affects only the ability to insure transitory shocks (which depend on having access to loans), but not the ability to insure permanent shocks (which depend on having access to a storage technology, and hence it is not affected by credit restrictions).
- Fourth, the performance of the BPP estimators is remarkably good. Only in the case of the ZBC environment and the permanent shock does the BPP estimator display an upward bias, and even in that case only very early in the life cycle.
- According to KV the source of the bias is the failure of the orthogonality condition used by BPP for agents close to the borrowing constraint.
- It is worth noting that the ZBC environment is somewhat extreme as it assumes no unsecured borrowing.

- Finally, KV compare the average MPCs obtained in their model (0.06 and 0.77) with the actual estimates obtained by BPP using actual data.
- As we shall see, BPP find an estimate of the MPC with respect to permanent shocks of 0.64 (s.e. 0.09) and an estimate of the MPC with respect to transitory shocks of 0.05 (s.e. 0.04).
- Clearly, the "theoretical" MPCs found by KV lie well in the confidence interval of BPP's estimates.

- One thing that seems not to be borne out in the data is that theoretically the degree of smoothing of permanent shocks should be strictly increasing and convex with age, while BPP report increasing amount of insurance with age as a non-significant finding.
- As discussed by Kaplan and Violante (2010), the theoretical pattern of the smoothing coefficients is the result of two forces: a wealth composition effect and a horizon effect.
- The increase in wealth over the life cycle due to precautionary and retirement motives means that agents are better insured against shocks.
- As the horizon shortens, the effect of permanent shock resembles increasingly that of a transitory shock.
- Given that the response of consumption to shocks of various nature is so different (and so relevant for policy in theory and practice), it is natural to turn to studies that analyze the nature and persistence of the income process.

- In this section we discuss papers that use consumption and income data jointly.
- More recently, a number of papers have argued that consumption and income data jointly can be used to measure the extent of risk faced by households and understand its nature.
- This approach starts from the consideration that the use of income data alone is unlikely to be conclusive about the extent of risk that people face.
- The idea is to use actual individual choices (such as consumption, labor supply, human capital investment decisions) to infer the amount of risk that people face.

- Among the papers pursuing this idea, Blundell and Preston (1995), and Heckman et al. (2006) deserve a special mention.
- As correctly put by Cunha and Heckman (2007), "purely statistical decompositions cannot distinguish uncertainty from other sources of variability. Transitory components as measured by a statistical decomposition may be perfectly predictable by agents, partially predictable or totally unpredictable."
- Another reason why using forward looking " choices" allows us to learn about features of the earnings process is that consumption choices should reflect the nature of income changes.
- For example, if we were to observe a large consumption response to a given income change, we could infer that the income change is unanticipated and persistent (Blundell and Preston, 1995; Guvenen and Smith, 2009).
- We discuss these two approaches, together with notable contributions, in turn.

Approach 1: Identifying insurance for a given information set

- Most papers propose a statistical representation of the following type:

$$
\begin{aligned}
Y &= Y^P + Y^T \\
C &= C^P + C^T \\
Y^P &= X^P \beta^P + \zeta \\
Y^T &= X^T \beta^T + \varepsilon \\
C^P &= \kappa^P Y^P \\
C^T &= \kappa^T Y^T + \eta
\end{aligned}
$$

  in which Y (C) is current income (consumption), divided in permanent Y^P (C^P) and transitory Y^T (C^T).
- The main objective of most papers is to estimate κ_P, test whether $\kappa_P > \kappa_T$, and or/test whether $\kappa_P = 1$ (the income proportionality hypothesis).

Hall and Mishkin (1982)

- The authors in the papers above do not write explicitly the stochastic structure of income.
- For example, in the statistical characterization above permanent income is literally permanent (a fixed effect).
- The first paper to use micro panel data to decompose income shocks into permanent and transitory components writing an explicit stochastic income process is Hall and Mishkin (1982), who investigate whether households follow the rational expectations formulation of the permanent income hypothesis using PSID data on income and food consumption.

- Their setup assumes quadratic preferences (and hence looks at consumption and income changes), imposes that the marginal propensity to consume with respect to permanent shocks is 1, and leaves only the MPC with respect to transitory shocks free for estimation.
- The income process is described by equations (3) and (4) (enriched to allow for some serial correlation of the MA type in the transitory component), so that the change in consumption is given by equation (5):

$$\Delta c_{i,a,t} = \zeta_{i,a,t} + \pi_a \varepsilon_{i,a,t}$$

- Since the PSID has information only on food consumption, this equation is recast in terms of food spending (implicitly assuming separability between food and other non-durable goods):

$$\Delta c^F_{i,a,t} = \alpha(\zeta_{i,a,t} + \pi_a \varepsilon_{i,a,t}) + \Delta m^F_{i,a,t}$$

where α is the proportion of income spent of food, and m^F is a stochastic element added to food consumption (measurement error), not correlated with the random elements of income ($\zeta_{i,a,t}$ and $\varepsilon_{i,a,t}$).
- The model is estimated using maximum likelihood assuming that all the random elements are normally distributed.

- Hall and Mishkin (1982) also allow for the possibility that the consumer has some "advance information" (relative to the econometrician) about the income process.
- Calling Υ the degree of advance information, they rewrite their model as:

$$\Delta c^F_{i,a,t} = \alpha\Upsilon(\zeta_{i,a+1,t+1} + \pi_{a+1}\varepsilon_{i,a+1,t+1}) + \alpha\left(1-\Upsilon\right)\left(\zeta_{i,a,t} + \pi_a\varepsilon_{i,a,t}\right) + \Delta m^F_{i,a,t} \qquad (13)$$

- Their estimates of (13) only partly confirm the PIH.
- Their estimates of Υ is 0.25 and their estimate of π (which they assume to be constant over the life cycle) is 0.29, too high to be consistent with plausible interest rates.

Approach 2: Identifying information set for given insurance configuration

- Why can consumption and income data be useful in identifying information set or learn more about the nature of the income process?
- To see very clearly this point, consider a simple extension of an example used by Browning, Hansen and Heckman (1999).
- Assume that the income process is given by the sum of a random walk ($p_{i,a,t}$), a transitory shock ($\varepsilon_{i,a,t}$) and a measurement error ($m_{i,a,t}$, which may even reflect " superior information" , i.e., information that is observed by the individual but not by an econometrician):

$$
\begin{aligned}
y_{i,a,t} &= p_{i,a,t} + \varepsilon_{i,a,t} + m_{i,a,t} \\
p_{i,a,t} &= p_{i,a-1,t-1} + \zeta_{i,a,t}
\end{aligned}
$$

- Written in first differences, this becomes

$$\Delta y_{i,a,t} = \zeta_{i,a,t} + \Delta\varepsilon_{i,a,t} + \Delta m_{i,a,t}$$

- One cannot separately identify transitory shocks and measurement error (unless access to validation data gives us an estimate of the amount of variability explained by measurement error, as in Meghir and Pistaferri, 2004; or higher order restrictions are invoked, as in Cragg, 1997; or assumptions about separate serial correlation of the two components are imposed).
- Assume as usual that preferences are quadratic, $\beta(1+r)=1$ and that the consumer's horizon is infinite for simplicity.

- The change in consumption is given by equation (5) adapted to the infinite horizon case:

$$\Delta C_{i,a,t} = \zeta_{i,a,t} + \frac{r}{1+r}\varepsilon_{i,a,t} \tag{14}$$

- The component $m_{i,a,t}$ does not enter (14) because consumption does not respond to measurement error in income.
- However, if $m_{i,a,t}$ represented "superior information", then this assumption would have behavioral content: it would be violated if liquidity constraints were binding - and hence $m_{i,a,t}$ would belong to (14).

- Suppose a researcher has access to panel data on consumption and income (a very stringent requirement, as it turns out).
- Then one can use the following covariance restrictions:

$$\begin{aligned} var\left(\Delta y_{i,a,t}\right) &= \sigma_{\zeta}^{2} + 2\left(\sigma_{\varepsilon}^{2} + \sigma_{m}^{2}\right) \\ cov\left(\Delta y_{i,a,t}, \Delta y_{i,a-1,t-1}\right) &= -\left(\sigma_{\varepsilon}^{2} + \sigma_{m}^{2}\right) \\ var\left(\Delta c_{i,a,t}\right) &= \sigma_{\zeta}^{2} + \left(\frac{r}{1+r}\right)^{2} \sigma_{\varepsilon}^{2} \end{aligned}$$

- As is clear from the first two moments, σ_{ε}^{2} and σ_{m}^{2} cannot be told apart from income data alone (although the variance of permanent shocks can actually be identified - e.g., using $\sigma_{\zeta}^{2} = var\left(\Delta y_{i,a,t}\right) + 2cov\left(\Delta y_{i,a,t}, \Delta y_{i,a-1,t-1}\right)$, the stationary version of equation above).

- However, the availability of consumption data solves the identification problem. In particular, one could identify the variance of transitory shocks from, e.g.

$$\sigma_{\varepsilon}^{2} = \left(\frac{r}{1+r}\right)^{-2} \left[var\left(\Delta c_{i,a,t}\right) - var\left(\Delta y_{i,a,t}\right) - 2cov\left(\Delta y_{i,a,t}, \Delta y_{i,a-1,t-1}\right)\right] \tag{15}$$

- Note also that if one is willing to use the covariance between changes in consumption and changes in income ($cov\left(\Delta c_{i,a,t}, \Delta y_{i,a,t}\right) = \sigma_{\zeta}^{2} + \left(\frac{r}{1+r}\right)\sigma_{\varepsilon}^{2}$), then there is even an overidentifying restriction that can be used to test the model.
- It is useful at this point to separate the literature into two sub-branches — those devoted to learning features of the income process, and those devoted to identifying information set.

Is the increase in income inequality permanent or transitory?

- Blundell and Preston (1998) use the link between the income process and consumption inequality to understand the nature and causes of the increase in inequality of consumption and the relative importance of changes in the variance of transitory and permanent shocks.
- Their motivation is that they have only repeated cross-section data, and the variances of income shocks are changing over time due to, for example, rising inequality.

- Hence for a given cohort, say, and even ignoring measurement error one has:

$$var\left(y_{i,a,t}\right) = var\left(p_{i,0,t-a}\right) + \sum_{j=0}^{a} var\left(\zeta_{i,j,t-a+j}\right) + var\left(\varepsilon_{i,a,t}\right)$$

  where $j = 0$ corresponds to the age of entry of this cohort in the labor market.

- With repeated cross-sections one can write the change in the variance of income for a given cohort as

$$\Delta var\left(y_{i,a,t}\right) = var\left(\zeta_{i,a,t}\right) + \Delta var\left(\varepsilon_{i,a,t}\right)$$

- Hence, a rise in inequality (the left-hand side of this equation) may be due to a rise in "volatility" $\Delta var\left(\varepsilon_{i,a,t}\right) > 0$ or the presence of a persistent income shock, $var\left(\zeta_{i,a,t}\right)$.
- In repeated cross-sections the problem of distinguishing between the two sources is unsolvable if one focuses just on income data.
- Suppose instead one has access to repeated cross-section data on consumption (which, conveniently, may or may not come from the same data set - the use of multiple data set is possible as long as samples are drawn randomly from the same underlying population).

- Then we have that the change in consumption inequality for a given cohort is):

$$\Delta var\left(c_{i,a,t}\right) = var\left(\zeta_{i,a,t}\right) + \left(\frac{r}{1+r}\right)^2 var\left(\varepsilon_{i,a,t}\right)$$

assuming one can approximate the variance of the change with the change of the variances (see Deaton and Paxson, 1994, for a discussion of the conditions under which this approximation is acceptable).
- Here one can see that the growth in consumption inequality is dominated by the permanent component (for small r the second term on the right hand side vanishes).
- Indeed, assuming $r \approx 0$, we have that the change in consumption inequality identifies the variance of the permanent component and that the difference between the change in income inequality and the change in consumption inequality identifies the change in the variance of the transitory shock.

- Related to Blundell and Preston (1998) is a paper by Hryshko (2008).
- He estimates jointly a consumption function (based on the CRRA specification) and an income process.
- Based on the evidence from Hryshko (2009) and the literature, as well as the need to match the increasing inequality of consumption over the lifecycle, he assumes that the income process is the sum of a random walk and a transitory shock.
- In addition, just like Blundell, Pistaferri and Preston (2008) he estimates the proportion of the permanent and the transitory shock that are insured finding that 37% of permanent shocks are insured via channels other than savings; transitory shocks are only insured via savings.

Identifying Information Set

Cunha, Heckman and Navarro (2005)

- The framework of their paper has been extended in Cunha and Heckman (2007), where the authors show that a large fraction of the increase in inequality in recent years is due to the increase in the variance of the unforecastable components.
- In particular, they estimate the fraction of future earnings that is forecastable and how this fraction has changed over time using college decision choices.
- For less skilled workers, roughly 60% of the increase in wage variability is due to uncertainty.
- For more skilled workers, only 8% of the increase in wage variability is due to uncertainty.

- Suppose as usual that preferences are quadratic, $\beta(1+r)=1$, initial assets are zero, the horizon is infinite, but the consumer receives income only in two periods, t and $t+1$.
- Consumption is therefore

$$C_{i,a,t}=\frac{r}{1+r}y_{i,a,t}+\frac{r}{(1+r)^2}E\left(y_{i,a+1,t+1}|\Omega_{i,a,t}\right)$$

- Write income in $t+1$ as

$$y_{i,a+1,t+1}=X'_{i,a+1,t+1}\beta+\zeta^{A}_{i,a+1,t+1}+\zeta^{U}_{i,a+1,t+1}$$

where $X'_{i,a+1,t+1}\beta$ is observed by both the individual and the econometrician, $\zeta^{A}_{i,a+1,t+1}$ is potentially observed only by the individual, and $\zeta^{U}_{i,a+1,t+1}$ is unobserved to both.

- The idea is that one can form the following "deviation" variables

$$
\begin{aligned}
z^{C}_{i,a,t} &= C_{i,a,t} - \frac{r}{1+r} y_{i,a,t} - \frac{r}{(1+r)^2} X'_{i,a+1,t+1}\beta \\
z^{Y}_{i,a+1,t+1} &= y_{i,a+1,t+1} - X'_{i,a+1,t+1}\beta
\end{aligned}
$$

- If $cov\left(z^{C}_{i,a,t}, z^{Y}_{i,a+1,t+1}\right) \neq 0$, there is evidence of "superior information", i.e., the consumer used more than just $X'_{i,a+1,t+1}\beta$ to decide how much to consume in period t.

Primiceri and Van Rens (2009)

- Using CEX data, they find that all of the increase in income inequality over the 1980-2000 period can be attributed to an increase in the variance of permanent shocks but that most of the permanent income shocks are anticipated by individuals; hence consumption inequality remains flat even though income inequality increases.
- While their results challenge the common view that permanent shocks were important only in the early 1980s (see Card and Di Nardo, 2002; Moffitt and Gootschalk, 2000), they could be explained by the poor quality of income data in the CEX (see Heathcote, Storesletten and Violante, 2009).

- The log income process is specified as follows

$$y_{i,a,t} = p_{i,a,t} + \varepsilon_{i,a,t} \quad (16)$$
$$p_{i,a,t} = p_{i,a-1,t-1} + \zeta^U_{i,a,t} + \zeta^A_{i,a,t} \quad (17)$$

where $\varepsilon_{i,a,t}$ and $\zeta^U_{i,a,t}$ are unpredictable to the individual and $\zeta^A_{i,a,t}$ is predictable to the individual but unobservable to the econometrician.

- Using CRRA utility with incomplete markets (there is only a risk free bond) log consumption can be shown to follow (approximately):

$$c_{i,a,t} = c_{i,a-1,t-1} + \zeta^U_{i,a,t} \quad (18)$$

- From equations (16), (17) and (18), the following cohort-specific moment conditions are implied:

$$\begin{aligned}\Delta var_t(y) &= var_t\left(\zeta^U\right) + var_t\left(\zeta^A\right) + \Delta var_t(\varepsilon)\\ \Delta var_t(c) &= var_t\left(\zeta^U\right)\\ \Delta cov_t(y,c) &= var_t\left(\zeta^U\right)\\ cov_t(\Delta y, y_{-1}) &= -var_{t-1}(\varepsilon)\end{aligned}$$

- Using these moment conditions, it is possible to (over)identify $var_t\left(\zeta^U\right)$ and $var_t\left(\zeta^A\right)$ for $t = 1, ..., T$ and $var_t(\varepsilon)$ for $t = 0, ..., T$.

Guvenen (2006) and Guvenen and Smith (2009)

- In Guvenen's (2007) model, income data are generated by the heterogeneous income profile specification.
- However, individuals do not know the parameters of their own profile (in particular, they ignore the slope of life-cycle profile f_i and the value of the persistent component) and need to learn about them, using Bayesian updating, by observing successive income realizations, which are noisy because of the mean reverting transitory shock.
- He shows that this model can be made to fit the consumption data very well (both in terms of levels and variance over the lifecycle) and in some ways better than the process that includes a unit root.

- By introducing learning, Guveven relaxes the restriction linking the income process to consumption and as a result weakens the identifying information implied by this link.
- This allows the income process to be stationary and consumption to behave as if income is not stationary.
- Thus, from a welfare point of view the individual is facing essentially as much uncertainty as they would under the random walk model, which is why the model can fit the increasing inequality over the lifecycle. In Guvenen's model it is just the interpretation of the nature of uncertainty that has changed.
- The fact that the income process conditional on the individual is basically deterministic (except for the small transitory shock) has lost its key welfare implications.
- Thus whether the income is highly uncertain or deterministic becomes irrelevant for issues that have to do with insurance and precautionary savings: individuals perceive it as highly uncertain and this is all that matters.

- Guvenen and Smith (2009) use consumption data jointly with income data to estimate the structural parameters of the model.
- They extend the consumption imputation procedure of Blundell, Pistaferri and Preston (2008) to create a panel data of income and consumption data in the PSID.

- The authors estimate the structural parameters of their model applying an indirect inference approach - a simulation based approach suitable for models in which it is very difficult to specify the criterion function.
- As in Guvenen (2009) the authors find that income shocks are less persistent in the HIP case ($\rho = 0.76$) than in the RIP case (ρ close to one), and that there is a significant evidence for heterogeneity in income growth.
- In addition, they find that prior uncertainty is quite small ($\Lambda = 0.19$, meaning that about 80 percent of the uncertainty about the random trend component is resolved in the first period of life).

- They therefore argue that the amount of uninsurable lifetime income risk that households perceive is smaller than what is typically assumed in calibrated macroeconomic models.
- Statistically speaking, the estimate is very imprecise and one could conclude that everything about the random trend term is known early on in the life cycle.

Information or Insurance?

- In the three examples above it is possible to solve the identification problem by making the following assumptions.
- First, consumption responds to signal, not to noise.
- In a related way, consumption responds to unanticipated, not to forecastable changes in income.
- While the orthogonality of consumption to measurement error in income is not implausible, the orthogonality to anticipated changes in income has behavioral content.

- Households will respond to anticipated changes in income, causing the theory to fail, if there are intertemporal distortions induced by, e.g. liquidity constraints.
- Second, the structure of markets is such that the econometrician can predict response of consumption to income shocks on the basis of a model of individual behavior.
- For example, in the strict version of the PIH with infinite horizon, the marginal propensity to consume out of permanent shock is 1 and the marginal propensity to consume out of transitory shock is equal to the annuity value $\frac{r}{1+r}$.
- That is, one identifies the variances of interest only under the assumption that the chosen model of behavior describes the data accurately.

- But what if there is more insurance than predicted by, for example, the simple PIH version of the theory?
- There are alternative theories that predict that consumers may insure their income shocks to a larger extent than predicted by a simple model with just self-insurance through a risk-free bond. One example is the full insurance model.
- Clearly, it is hard to believe full insurance is literally true.
- The model has obvious theoretical problems: private information, limited enforcement, etc. And there are of course also empirical problems.

- But outside the extreme case of the full insurance model, there is perhaps more insurance than predicted by the strict PIH version with just a risk-free bond.
- Previously we have seen that standard Bewley-type models can generate some insurance even of permanent shocks as long as people accumulate some precautionary wealth.
- To achieve this result, one does not require sophisticated contingent Arrow-Debreu markets.
- All is needed is a simple storage technology (such as a saving account).

- A recent macroeconomic literature has explored a number of theoretical alternatives to the insurance configurations described above.
- These alternative models fall under two broad groups: those that assume public information but limited enforcement of contracts, and those that assume full commitment but private information.
- These models prove that the self-insurance case is Pareto-inefficient even conditioning on limited enforcement and private information issues.
- In both types of models, agents typically achieve more insurance than under a model with a single non-contingent bond, but less than under a complete markets environment.

- These models show that the relationship between income shocks and consumption depends on the degree of persistence of income shocks.
- Alvarez and Jermann (2000), for example, explore the nature of income insurance schemes in economies where agents cannot be prevented from withdrawing participation if the loss from the accumulated future income gains they are asked to forgo becomes greater than the gains from continuing participation.
- Such schemes, if feasible, allow individuals to keep some of the positive shocks to their income and therefore offer only partial income insurance.
- If income shocks are persistent enough and agents are infinitely lived, then participation constraints become so severe that no insurance scheme is feasible.
- With finite lived agents, the future benefits from a positive permanent shock exceed those from a comparable transitory shock.

- This suggests that the degree of insurance should be allowed to differ between transitory and permanent shocks and should also be allowed to change over time and across different groups.
- Krueger and Perri (2006) provide an empirical review of income and consumption inequality in the 80's and 90's.
- They then suggest a theoretical macro model based on self insurance with limited commitment trying to explain the moderate expansion in consumption inequality compared to income inequality.
- Their hypothesis is that an increase in the volatility of idiosyncratic labour income has not only been an important factor in the increase in income inequality, but has also caused a change in the development of financial markets, allowing individual households to better insure against the bigger idiosyncratic income fluctuations.

- Another reason for partial insurance is moral hazard.
- This is the direction taken in Attanasio and Pavoni (2012).

- We now want to provide a simple example of the identification issue: is the attenuated response of consumption to income shocks reflecting " insurance/smoothing" or " information" ?
- Assume that log income and log consumption changes are given by the following equations:

$$\begin{aligned}\Delta y_{i,a,t} &= \Delta \varepsilon_{i,a,t} + \zeta^{A}_{i,a,t} + \zeta^{U}_{i,a,t} \\ \Delta c_{i,a,t} &= \zeta^{U}_{i,a,t} + \pi_a \varepsilon_{i,a,t}\end{aligned}$$

- In this case, income shifts because of anticipated permanent changes in income (a pre-announced promotion) and unanticipated permanent changes in income.

- In theory, consumption changes only in response to the unanticipated component.
- Suppose that our objective is to estimate the extent of "information", i.e., how large are permanent changes in income that are unanticipated:

$$\Upsilon = \frac{\sigma^2_{\zeta^U}}{\sigma^2_{\zeta^U} + \sigma^2_{\zeta^A}}$$

- A possible way of identifying this parameter is to run a simple IV regression of
  $\Delta c_{i,a,t}$ onto $\Delta y_{i,a,t}$ using $(\Delta y_{i,a-1,t-1} + \Delta y_{i,a,t} + \Delta y_{i,a+1,t+1})$ as an instrument (see Guiso, Pistaferri and Schivardi, 2005).
- This yields indeed:

$$\frac{cov\left(\Delta c_{i,a,t}, \Delta y_{i,a-1,t-1} + \Delta y_{i,a,t} + \Delta y_{i,a+1,t+1}\right)}{cov\left(\Delta y_{i,a,t}, \Delta y_{i,a-1,t-1} + \Delta y_{i,a,t} + \Delta y_{i,a+1,t+1}\right)} = \frac{\sigma^2_{\zeta^U}}{\sigma^2_{\zeta^U} + \sigma^2_{\zeta^A}} = \Upsilon$$

- Suppose now that $\sigma^2_{\zeta^A} = 0$ (no advance or superior information), but there is some insurance against permanent and transitory shocks, measured by the partial insurance parameters Φ and Ψ.

- What is the IV regression above identifying? The model now is

$$\Delta y_{i,a,t} = \zeta^{U}_{i,a,t} + \Delta\varepsilon_{i,a,t} \quad (19)$$

$$\Delta c_{i,a,t} = \Phi\zeta^{U}_{i,a,t} + \Psi\varepsilon_{i,a,t} \quad (20)$$

and the IV parameter takes the form

$$\frac{cov\left(\Delta c_{i,a,t}, \Delta y_{i,a-1,t-1} + \Delta y_{i,a,t} + \Delta y_{i,a+1,t+1}\right)}{cov\left(\Delta y_{i,a,t}, \Delta y_{i,a-1,t-1} + \Delta y_{i,a,t} + \Delta y_{i,a+1,t+1}\right)} = \frac{\phi\sigma^2_{\zeta^U}}{\sigma^2_{\zeta^U}} = \Phi,$$

which is what Blundell, Pistaferri and Preston (2008) assume.

- Hence, the same moment $\frac{cov\left(\Delta c_{i,a,t}, \Delta y_{i,a-1,t-1}+\Delta y_{i,a,t}+\Delta y_{i,a+1,t+1}\right)}{cov\left(\Delta y_{i,a,t}, \Delta y_{i,a-1,t-1}+\Delta y_{i,a,t}+\Delta y_{i,a+1,t+1}\right)}$ has two entirely different interpretations depending on what assumptions one makes about information and insurance.
- What if we have both an anticipated component and partial insurance?
- It's easy to show that in this case

$$\frac{cov\left(\Delta c_{i,a,t}, \Delta y_{i,a-1,t-1} + \Delta y_{i,a,t} + \Delta y_{i,a+1,t+1}\right)}{cov\left(\Delta y_{i,a,t}, \Delta y_{i,a-1,t-1} + \Delta y_{i,a,t} + \Delta y_{i,a+1,t+1}\right)} = \Phi\Upsilon$$

  a combination of information and insurance.
- In sum, suppose that a researcher finds that consumption responds very little to what the econometrician defines to be a shock to economic resources (for the moment, neglect the distinction between transitory and permanent shocks).

- There are at least two economically interesting reasons why this might be the case.
- First, it is possible that what the econometrician defines to be a shock is not, in fact, a shock at all when seen from the point of view of the individual.
- In other words, the change in economic resources identified by the econometrician as an innovation might be predicted in advance (at least partly) by the consumer.
- Hence if the consumer is rational and not subject to borrowing constraints, her consumption will not respond to changes in income that are anticipated.
- It follows that the "extent of attenuation" of consumption in response to income shocks measures the extent of "superior information" that the consumers possess.

- The other possibility is that what the econometrician defines to be a shock is correctly a shock when seen from the point of view of the individual.
- However, suppose that the consumer has access to insurance mechanisms over and above self-insurance (for example, government insurance, intergenerational transfers. etc.).
- Hence, consumption will react little to the shock (or less than predicted by a model with just self-insurance).
- In this case, the "extent of attenuation" of consumption in response to income shocks measures the extent of "partial insurance" that the consumers has available against income shocks.

- More broadly, identification of information set requires taking a stand on the structure of (formal and informal) credit and insurance markets.
- What looks like lack of information may be liquidity constraints in disguise (consumer responds too much to negative transitory shock, say).
- What looks like superior information may be insurance in disguise (consumer responds too little to permanent shocks).

Approaching the information/insurance conundrum

- The literature has considered two approaches to solve the information/insurance identification issue.
- A first method attempts to identify episodes in which income changes unexpectedly, and to evaluate in a quasi-experimental setting how consumption reacts to such changes.
- A second approach estimates the impact of shocks combining realizations and expectations of income or consumption in surveys where data on subjective expectations are available (see Hayashi, 1985, and Pistaferri, 2001, for means, and Kaufmann and Pistaferri, 2009, for covariance restrictions).
- First consider Blundell, Pistaferri and Preston (2008) (BPP).
- Imposes assumptions about the information set(s) of the agents and estimates insurance, but provides a test of "superior information".

Blundell, Pistaferri and Preston (2008)

- The consumption model considered in Blundell, Pistaferri and Preston (2008) is given by equation (12), while their income process is given by (10) and (11).
- In their study they create panel data on a comprehensive consumption measure for the PSID using an imputation procedure based on food demand estimates from the CEX.
- Table 2 reproduces their main results.

Table 2: Partial Insurance Estimates from Blundell, Pistaferri and Preston (2008)

		Whole sample	Born 1940s	Born 1930s	No College	Low wealth
Φ (Partial insurance perm. shock)		0.6423 (0.0945)	0.7928 (0.1848)	0.6889 (0.2393)	0.9439 (0.1783)	0.8489 (0.2848)
Ψ (Partial insurance trans. shock)	0.0533 (0.0435)	0.0675 (0.0705)	−0.0381 (0.0737)	0.0768 (0.0602)	0.2877 (0.1143)	

- They find that consumption is nearly insensitive to transitory shocks (the estimated coefficient is around 5 percent, but higher among poor households), while their estimate of the response of consumption to permanent shocks is significantly lower than 1 (around 0.65, but lower for the college educated and those near retirement and higher for poor or less educated households), suggesting that households are able to insure at least part of the permanent shocks.

- These results show
  1. (a) that the estimates of the insurance coefficients in the baseline case are statistically consistent with the values predicted by the calibrated Kaplan-Violante model;
  2. (b) that younger cohorts have harder time smoothing their shocks, presumably because of the lack of sufficient wealth;
  3. (c) groups with actual or presumed low wealth are not able to insure permanent shocks (as expected from the model) and have even difficulties smoothing transitory shocks (credit markets can be not available for people with little or no collateral).

- While the setting of Blundell, Pistaferri and Preston (2008) cannot be used to distinguish between insurance and information, their paper provides a test of their assumption about richness of the information set.
- In particular, they follow Cunha, Heckman and Navarro (2005) and test whether unexpected consumption growth (defined as the residual of a regression of consumption growth on observable household characteristics) is correlated with future income changes (defined also as the residual of a regression of income growth on observable household characteristics).
- If it was the case, then consumption contains more information than used by the econometrician.

- Their test of superior information reported in Table 3 shows that consumption is not correlated with future income changes.
- BPP find little evidence of anticipation.
- This suggests the persistent labour income shocks that were experienced in the 1980s were not anticipated.
- These were largely changes in the returns to skills, shifts in government transfers and the shift of insurance from firms to workers.

Table 3: Test of Superior Information, from Blundell, Pistaferri andPreston (2008)

Test $cov\left(\Delta y_{a+1}, \Delta c_a\right)$ for all a	p-value 0.25
Test $cov\left(\Delta y_{a+2}, \Delta c_a\right)$ for all a	p-value 0.27
Test $cov\left(\Delta y_{a+3}, \Delta c_a\right)$ for all a	p-value 0.74
Test $cov\left(\Delta y_{a+4}, \Delta c_a\right)$ for all a	p-value 0.68

- Finally, the results of Blundell, Pistaferri and Preston (2008) can be used to understand why consumption inequality in the US has grown less than income inequality during the past two decades.
- Their findings suggest that the widening gap between consumption and income inequality is due to the change in the durability of income shocks.
- In particular, a growth in the variance of permanent shocks in the early eighties was replaced by a continued growth in the variance of transitory income shocks in the late eighties.
- Since they find little evidence that the degree of insurance with respect to shocks of different durability changes over this period, it is the relative increase in the variability of more insurable shocks rather than greater insurance opportunities that explains the disjuncture between income and consumption inequality.

Solution 1: The quasi-experimental approach

- The approach we discuss in this section does not require estimation of an income process, or even observing the individual shocks.
- Rather, it compares households that are exposed to shocks with households that are not (or the same households before and after the shock), and assumes that the difference in consumption arise from the realization of the shocks.
- The idea here is to identify episodes in which changes in income are unanticipated, easy to characterize (i.e., persistent or transient), and (possibly) large.

- The first of such attempts dates back to a study by Bodkin (1959), who laid down fifty years ago all the ingredients of the quasi-experimental approach.
- In this pioneering study the experiment consists of looking at the consumption behavior of WWII veterans after the receipt of unexpected dividend payments from the National Service Life Insurance.
- Bodkin assumes that the dividend payments are unanticipated and represent a windfall source of income, and finds a point estimate of the marginal propensity to consume non-durables out of this windfall income is as high as 0.72, a strong violation of the permanent income model.

- Recent papers in the quasi-experimental framework look at the effect of unemployment shocks on consumption, and the smoothing benefits provided by unemployment insurance (UI) schemes.
- As pointed out by Browning and Crossley (2001) unemployment insurance provides two benefits to consumers.
- First, it provides "consumption smoothing benefits" for consumers that are liquidity constrained.
- In the absence of credit constraints, individuals who faced a negative transitory shock such as unemployment would borrow to smooth their consumption.

- If they are unable to borrow they would need to adjust their consumption downward considerably.
- Unemployment insurance provides some liquidity and hence it has positive welfare effects.
- Second, unemployment insurance reduces the conditional variance of consumption growth and hence the need to accumulate precautionary savings.

- One of the earlier attempts to estimate the welfare effects of unemployment insurance is Gruber (1997).
- Gruber finds a large smoothing effect of UI, in particular that a 10 percentage point rise in the replacement rate reduces the fall in consumption upon unemployment by about 3 percent.
- Browning and Crossley (2001) extend Gruber's idea to a different country (Canada instead of the US), using a more comprehensive measure of consumption (instead of just food) and legislated changes in UI (instead of state-time variation).
- But this small effect masks substantial heterogeneity, with low-assets households at time of job loss exhibiting elasticities as high as 20 percent.

- First, some of these shocks may not come as a surprise, and individuals may have saved in their anticipation.
- Second, the theory predicts that consumers smooth marginal utility, not consumption per se.
- If an unemployment shock brings more leisure and if consumption is a substitute for leisure, an excess response of consumption to the transitory shock induced by losing one's job does not necessarily represent a violation of the theory.
- Finally, even if unemployment shocks are truly fully unanticipated, they may be partially insured through government programs such as unemployment insurance (and disability insurance in case of disability shocks).

- Wolpin (1982) and Paxson (1992) study the effect of weather shocks in India and Thailand, respectively.
- In agricultural economies, weather shocks affect income directly through the production function and deviations from normal weather conditions are truly unanticipated events.
- Wolpin (1982) uses Indian regional time series data on rainfall to construct long run moments as instruments for current income (which is assumed to measure permanent income with error).

- A second limitation of the approach is that some of the income shocks (in particular, unemployment and disability shocks), cannot be considered as truly exogenous events.

Solution 2: Subjective expectations

- As pointed out in Sections 4.1. and 4.2, identifying income shocks is difficult because people may have information that is not observed by the econometrician.
- When the news is realized, the econometrician will measure as a shock what is in fact an expected event.
- The literature based on subjective expectations attempts to circumvent the problem by asking people to report quantitative information on their expectations, an approach forcefully endorsed by Manski (2004).
- Hayashi (1985) is the first study to adopt this approach. He uses a four-quarter panel of Japanese households containing respondents' expectations about expenditure and income in the following quarter.
- His results are in line with Hall and Mishkin (1982), suggesting a relatively high sensitivity of consumption to income shocks.

- Pistaferri (2001) combines income realizations and quantitative subjective income expectations contained in the 1989-93 Italian Survey of Household Income and Wealth (SHIW) to identify separately the transitory and the permanent income shocks.
- To see how subjective income expectations allow estimating transitory and income shocks for each household, consider the income process of equations (3) and (4).
- Define $E(x_{i,a,t}|\Omega_{i,a-1,t-1})$ as the subjective expectation of $x_{i,a,t}$ given the individual's information set at age $a-1$.
- It is worth pointing out that $\Omega_{i,a-1,t-1}$ is the set of information possessed at individual level; the econometrician's information set is generally less rich.
- The assumption of rational expectations implies that the transitory shock at time t can be point identified by:

$$\varepsilon_{i,a,t} = -E\left(\Delta y_{i,a,t}|\Omega_{i,a-1,t-1}\right) \tag{21}$$

- Using equations (3), (4). and (21), the permanent shock at time t is identified by the expression:

$$\zeta_{i,a,t} = \Delta y_{i,a,t} - E\left(\Delta y_{i,a,t} | \Omega_{i,a-1,t-1}\right) + E\left(\Delta y_{i,a+1,t+1} | \Omega_{i,a,t}\right)$$

  e.g., the income innovation at age a adjusted by a factor that takes into account the arrival of new information concerning the change in income between a and $a+1$.
- Thus, the transitory and permanent shocks can be identified if one observes, for at least two consecutive time periods, the conditional expectation and the realization of income, a requirement satisfied by the 1989-93 SHIW.
- Pistaferri estimates the saving for a rainy day equation of Campbell (1987) and finds that consumers save most of the transitory shocks and very little of the permanent shocks, supporting the saving for a rainy day model.

- Kaufmann and Pistaferri (2009) use the same Italian survey used by Pistaferri (2000).
- Their results are reproduced in Table 4.
- Their most general model separates transitory changes in income in anticipated (with variance $\sigma^2_{\varepsilon^A}$),
- unanticipated ($\sigma^2_{\varepsilon^U}$), and measurement error (σ^2_y); permanent changes in income in anticipated ($\sigma^2_{\zeta^A}$) and unanticipated ($\sigma^2_{\zeta^U}$);
- allows for measurement error in consumption and subjective income expectations (σ^2_c and σ^2_e, respectively),
- allows for partial insurance with respect to transitory shocks (Ψ) and permanent shocks (Φ).

Table 4: EWMD Results, from Kaufmann and Pistaferri (2009)

Parameter	(1)	(2)	(3)
$\sigma^2_{\varepsilon^U}$	0.1056 (0.0191)	0.1172 (0.0175)	0.0197 (0.0208)
$\sigma^2_{\varepsilon^A}$	0	0	0.0541 (0.0163)
σ^2_{y}	0	0	0.0342 (0.0215)
$\sigma^2_{\zeta^U}$	0.0301 (0.0131)	0.0253 (0.0113)	0.0208 (0.0133)
$\sigma^2_{\zeta^A}$	0	0	0.0127 (0.0251)
σ^2_{c}		0.0537 (0.0062)	0.0474 (0.0097)
σ^2_{e}			0.1699 (0.0225)
Ψ		0.1442 (0.0535)	0.3120 (0.4274)
Φ		0.6890 (0.2699)	0.9341 (0.5103)
χ^2 (df; p-value)	3.2440 (1; 7%)	16.4171 (5; 0.6%)	36.4001 (12; 0.03%)

- A number of interesting facts emerge.
- First, the transitory variation in income is split between anticipated component (about 50%), the unanticipated component (20%) and measurement error (30%).
- This lowers the estimated degree of insurance with respect to transitory shocks.
- Similarly, a good fraction of the permanent variation (about 1/3) appears anticipated, and this now pushes the estimated insurance coefficient towards 1 - i.e., these results show evidence that there is no insurance whatsoever with respect to permanent shocks.

Subjective expectations: data problems

- The Italian SHIW offers the opportunity to test some simple hypotheses regarding the validity of subjective data.

- The type of income processes do not distinguish between fluctuations in income caused by exogenous shocks and those caused by endogenous responses to shocks.
- One reason is that different type of shocks may be differently insurable, raising important policy implications.
- Moreover, it may allow us to better characterize behavior.

- In a key contribution in this direction Abowd and Card (1989) extended the earlier literature to consider joint movements of hours and wages.
- Having established that both hours and earnings growth can be represented by an MA(2) process, they then link the two based on the lifecycle model.
- Their approach can reveal how much of the variation in earnings comes from genuine shocks to wages and how much is due to responses of these shocks through hours of work.

- Their conclusion was that the common components in the variation of earnings and hours could not be explained by variation in productivity.
- With their approach they opened up the idea of considering the stochastic properties of different related quantities jointly and to use this framework to assess how much of the fluctuations can be attributed to risk, as opposed to endogenous responses, such as changing hours.
- Of course, to the extent that hours may be driven by short term demand for labour in the workplace, rather than voluntary adjustments, such fluctuations may also represent risk.

Extending the income process to allow for endogenous fluctuations

- The key issue highlighted by the Abowd and Card approach is to disentangle the effect of shocks from the responses to shocks.
- While Abowd and Card do not go all the way in that direction, they do relate the fluctuations in earnings and hours.
- The first important modification is that they are now explicit about modelling wages per unit of time.
- In the specific application the unit of time is a quarter and the individual may either be working over this period or not.

- Extending the framework to a richer labour supply framework (the intensive margin) is relatively straightforward.
- The second modification is allowing for match effects; this implies that a source of fluctuations is obtaining a different job; what job one samples is a separate source of risk, to the extent that match effects are important.
- However, individuals can accept or reject job offers, a fact that needs to be recognized when combining such a process with a model of lifecycle consumption and labour supply.

- In what follows we use the notation w for (hourly) wages.
- Hence we specify

$$\ln w_{i,a,t} = d_t + x'_{i,a,t}\psi + u_{i,a,t} + e_{i,a,t} + a_{ij(t_0)} \tag{22}$$

where $w_{i,a,t}$ is the real hourly wage, d_t represents the log price of human capital at time t, $x_{i,a,t}$ a vector of regressors including age, $u_{i,a,t}$ the permanent component of wages, and $e_{i,a,t}$ the transitory error component.

- All parameters of the wage process are education specific (subscripts omitted for simplicity).
- In principle, the term $e_{i,a,t}$ might be thought of as representing a mix between a transitory shock and measurement error.
- In the usual decomposition of shocks into transitory and permanent components, researchers work with annual earnings data where transitory shocks may well be important because of unemployment spells.

- The term $a_{ij(t_0)}$ denotes a firm-worker match-specific component where $j\left(t_0\right)$ indexes the firm that the worker joined in period $t_0 \leq t$.
- It is drawn from a normal distribution with mean zero and variance σ_a^2.
- Low, Meghir and Pistaferri model the match effect as constant over the life of the worker-employer relationship.
- If the worker switches to a different employer between t and $t+1$, however, there will be some resulting wage growth which we can term a mobility premium denoted as $\xi_{i,a+1,t+1} = a_{ij(t+1)} - a_{ij(t_0)}$.

- The match effect is assumed normally distributed and successive draws of $a_{ij(t)}$ are assumed independent; however, because of the endogenous mobility decisions successive *realizations* of the match effect will be correlated.
- Since offers can be rejected when received, only a censored distribution of $\xi_{i,a+1,t+1}$ is observed. The match effect $a_{ij(.)}$ is complementary to individual productivity.

- Assume that the permanent component of wages follows a random walk process:

$$u_{i,a,t} = u_{i,a-1,t-1} + \zeta_{i,a,t} \tag{23}$$

- The random shock to the permanent process, $\zeta_{i,a,t}$ is normally distributed with mean zero and variance σ_{ζ}^2 and is independent over time.
- Assume this shock reflects uncertainty.

- In order to make sense of such a process, we need to make further assumptions relating to firm behavior.
- Thus it is simpler to assume that there are constant returns to scale in labor implying that the firm is willing to hire anyone who can produce non-negative rents.

- First, we have the shocks to productivity $\zeta_{i,a,t}$; second, there are shocks to job opportunities: these are reflected in the job arrival rate when employed (λ^e) and when unemployed (λ^n), as well as by the possibility of a lay off (job destruction, δ).
- Finally, there is the draw of a match specific effect.
- Individuals can respond to these by quitting into unemployment and accepting or rejecting a job offer.
- This model clarifies what aspect of earnings fluctuations reflects risk and what reflects an endogenous reaction to risk.
- The discussion also highlights the distinction between just describing the fluctuations of income vis-à-vis estimating a model of income fluctuations whose intention is to understand the welfare implications of risk.

Other Approaches on Endogenizing Volatility

- Here we discuss other approaches endogenizing wage or earnings volatility.
- Postel-Vinay and Turon (2009) test whether the observed covariance structure of earnings in the UK may be generated by a structural job search model with on-the-job search.

- Low and Pistaferri (2010) use data on subjective reports of work limitations available from the PSID to identify health shocks separately from other shocks to productivity.
- Their framework is similar to LMP.
- Huggett, Ventura and Yaron (2007) study human capital accumulation.
- In their model individuals may choose to divert some of their working time to the production of human capital.
- People differ in initial human capital (schooling, parents' teachings, etc.), initial financial wealth, and the innate ability to learn.
- Among other things, their framework generalizes Ben-Porath (1967) to allow for risk, i.e., shocks to the existing stock of human capital.

- Their questions of interest are:
  a. How much of lifetime inequality is determined before entry in the labor market (initial conditions)? and
  b. How much is due to episodes of good or bad luck over the life cycle (shocks)? The answers to these two questions have clear policy relevance.
- If the answer to (a) is " a lot" , one would want early intervention policies (e.g., public education).
- If the answer to (b) is " a lot" , one would want to expand income maintenance programs (UI, means-tested welfare, etc.).
- In HVY wages grow because of shocks to existing human capital, or systematic fanning out due to differences in learning abilities. Old people do not invest, hence only the first force is present.

- This provides an important idea for identification: Data on old workers can be used to identify the distribution of shocks to human capital.
- They next construct an age profile for the first, second, and third moment of earnings.
- Age, time, cohort effects are not separately identifiable, so need to impose some restrictions, such as: (a) No time effects (b) No cohort effects.
- Finally, they calibrate the distribution of initial conditions (initial human capital and learning ability) and the shape of the human capital production function to match the age profile of the first three moments of earnings, while fixing the remaining parameters to realistic values (from the literature).
- HVY use their model to do two things: (1) compute how much lifetime inequality is due to initial conditions and how much to shocks, and (2) run counterfactual experiments (shutting down risk to human capital or learning ability differences).

- Their results are that between 60% and 70% of the variability in lifetime utility (or earnings) is due to variability in initial conditions.
- Among initial conditions, the lion's share is taken by heterogeneity in initial human capital (rather than initial wealth or innate ability).
- Eliminating learning ability heterogeneity makes the age profile of inequality flat (even declining over a good fraction of the working life, 35-55).
- Eliminating shocks to human capital generates a more moderate U-shape age profile of inequality.

- For our purposes, one of the main points of the paper is that the standard incomplete markets model (for example, Heathcote, Storesletten and Violante, 2008) – which assumes an exogenous income process – may exaggerate the weight played by shocks as opposed to initial conditions in determining lifetime inequality.
- Hence, it may overestimate the welfare gain of government insurance programs and underestimate the welfare gain of providing insurance against "bad initial conditions" (bad schools, bad parents, bad friends, etc.).
- Note however that the "exaggeration" effect of IM models only holds under the assumption that initial conditions are fully known to the agents at the beginning of the life cycle.

- If people have to "learn" their initial conditions, then they will face unpredictable innovations to these processes.
- Recent work by Guvenen (2007) estimates that people can forecast only about 60% of their "learning ability" – the remaining 40% is uncertainty revealed (quite slowly) over the life cycle.
- Similar conclusions are reached in work by Cunha, Heckman, and Navarro (2005).

Shocks and labour market equilibrium

- We have moved from the standard reduced form models of income fluctuations to the more structural approach of Low, Meghir and Pistaferri (2010).
- However, there is further to go. What is missing from this framework is an explicit treatment of equilibrium pay policies.
- More specifically, in LMP the wage shocks are specified as shocks to the match specific effect, without specifying how these shocks arise.

- If we think about the match specific effect as being produced by a combination of the qualities of the worker and of the firm, then as in Postel-Vinay and Robin (2002), we can work out the pay policy of the firm under different assumptions on the strategies that individuals and firms follow.
- In that framework income/earnings, but only because individuals either receive alternative job offers, to which the incumbent firm responds, or because they move to an alternative firm.

- Lise, Meghir and Robin (2009) generalize this framework to allow for shocks to the firm's productivity.
- In this context, the observed wage shocks are further decomposed as originating from shocks to the productivity of the firm, responses to alternative offers or to moving to new jobs, either via unemployment or directly by firm to firm transition.
- In this context, the shocks are specified as changes in basic underlying characteristics of the firm as well as due to search frictions and comes closest to providing a full structural interpretation of income shocks, allowing also for the behavior of firms and strategies that lead to wages not being always responsive to the underlying shocks.

- While this offers a way forward in understanding the source of fluctuations, the approach is not complete because it assumes that both individuals and firms are risk neutral.
- In this sense individuals have no interest in insurance and do not save for precautionary reasons.
- Extending such models to allow for risk aversion, wage contracts that partially insure the worker and for savings, is the natural direction for obtaining an integrated approach of earnings fluctuations and an analysis of the effects of risk.

Appendix

Technical Discussion Drawn from Blundell, Low, and Preston (2008)
Approximating the Euler Equation

- The household plan at age t is to maximize the expected remaining lifetime utility:

$$E_t \sum_{\tau=0}^{T-t} \frac{U(c_{i,t+\tau})}{(1+\delta)^{\tau}}$$

- We begin by calculating the error in approximating the Euler equation.

$$E_t U'\left(c_{it+1}\right) = U'\left(c_{it}\right)\left(\frac{1+\delta}{1+r}\right) = U'\left(c_{it} e^{\kappa_{it+1}}\right) \tag{24}$$

for some κ_{it+1}

- By exact Taylor expansion of period $t+1$ marginal utility in $\ln c_{it+1}$ and around $\ln c_{it} + \kappa_{it+1}$, there exists a $\tilde{c}$ between $c_{it} e^{\kappa_{it+1}}$ and c_{it+1} such that

$$\begin{aligned} U'(c_{it+1}) =& U'\left(c_{it} e^{\kappa_{it+1}}\right) \Bigg[1 + \frac{1}{\kappa\left(c_{it} e^{\kappa_{it+1}}\right)} \left[\Delta \ln c_{it+1} - \kappa_{it+1}\right] \\ & + \frac{1}{2} \beta\left(\tilde{c}, c_{it} e^{\kappa_{it+1}}\right) \left[\Delta \ln c_{it+1} - \kappa_{it+1}\right]^2 \Bigg] \end{aligned} \tag{25}$$

where $\kappa(c) \equiv U'(c)/cU''(c) < 0$ and $\beta(\tilde{c}, c) \equiv [\tilde{c}^2 U'''(\tilde{c}) + \tilde{c} U''(\tilde{c})]/U'(c)$.

- Taking expectations

$$E_t U'(c_{it+1}) = U'\left(c_{it} e^{\kappa_{it+1}}\right)\left[1 + \frac{1}{\kappa\left(c_{it} e^{\kappa_{it+1}}\right)} E_t \left[\Delta \ln c_{it+1} - \kappa_{it+1}\right] + \frac{1}{2} E_t \left\{\beta\left(\tilde{c}, c_{it} e^{\kappa_{it+1}}\right)\left[\Delta \ln c_{it+1} - \kappa_{it+1}\right]^2\right\}\right] \quad (26)$$

THE UNIVERSITY OF CHICAGO

- Substituting for $E_t U'(c_{it+1})$ from (24)

$$\frac{1}{\kappa(c_{it}e^{\kappa_{it+1}})} E_t \left[\Delta \ln c_{it+1} - \kappa_{it+1}\right] + \frac{1}{2} E_t \left\{ \beta(\tilde{c}, c_{it}e^{\kappa_{it+1}}) \left[\Delta \ln c_{it+1} - \kappa_{it+1}\right]^2 \right\} = 0 \qquad (27)$$

and thus

$$\Delta \ln C_{it+1} = \kappa_{it+1} - \frac{\kappa(c_{it}e^{\kappa_{it+1}})}{2} E_t \left\{ \beta\left(\tilde{c}, c_i te^{\kappa_{it+1}}\right) \left[\Delta \ln c_{it} + 1e^{\kappa_{it+1}}\right]^2 \right\} + \varepsilon_{it+1} \qquad (28)$$

where the consumption innovation ε_{it+1} satisfies $E_t\varepsilon_{it+1} = 0$. As $E_t\varepsilon^2_{it+1} \to 0$, $\beta(\tilde{c}, c_{it}e^{\kappa_{it+1}})$ tends to a constant and therefore by Slutsky's theorem

$$\Delta \ln c_{it+1} = \varepsilon_{it+1} + \kappa_{it+1} + \mathcal{O}\left(E_t \left|\varepsilon_{it+1}\right|^2\right). \qquad (29)$$

- If preferences are CRRA then κ_{it+1} does not depend on c_it and is common to all households, say κ_{t+1}.
- The log of consumption therefore follows a martingale process with common drift

$$\Delta \ln c_{it+1} = \varepsilon_{it+1} + \kappa_{t+1} + \mathcal{O}\left(E_t \left|\varepsilon_{it+1}\right|^2\right). \qquad (30)$$

Approximating the Lifetime Budget Constraint

- The second step in the approximation is relating income risk to consumption variability.
- In order to make this link between the consumption innovation ε_{it+1} and the permanent and transitory shocks to the income process, we loglinearise the intertemporal budget constraint using a general Taylor series approximation (extending the idea in Campbell 1993).

- Define a function F: $\mathbb{R}^{N+1} \to \mathbb{R}$ by $F(\boldsymbol{\xi}) = \ln \sum_{j=0}^{N} \exp \xi_j$.
- By exact Taylor expansion around an arbitrary point $\boldsymbol{\xi}^0 \in \mathbb{R}^{N+1}$

$$F(\boldsymbol{\xi}) = \ln \sum_{j=0}^{N} \exp \xi_j^0 + \sum_{j=0}^{N} \frac{\exp \xi_j^0}{\sum_{k=0}^{0} \exp \xi_k^0} (\xi_j - \xi_j^0) + \frac{1}{2} \sum_{j=0}^{N} \sum_{k=0}^{N} \frac{\partial^2 F(\tilde{\boldsymbol{\xi}})}{\partial \xi_j \partial \xi_k} \left(\xi_j - \xi_j^0\right) \left(\xi_k - \xi_k^0\right) \tag{31}$$

where $\tilde{\boldsymbol{\xi}}$ lies between $\boldsymbol{\xi}$ and $\boldsymbol{\xi}^0$ and is used to make the expansion exact.

- The coefficients in the remainder term are given by

$$\frac{\partial^2 F(\tilde{\boldsymbol{\xi}})}{\partial \xi_j \partial \xi_k} = \frac{\exp \tilde{\xi}_j}{\sum_k \exp \tilde{\xi}_k} \left(\delta_{jk} - \frac{\exp \tilde{\xi}_j}{\sum_k \exp \tilde{\xi}_k} \right), \tag{32}$$

where δ_{jk} denotes the Kronecker delta.

- These coefficients are bounded because $0 < \exp \tilde{\xi}_j / \sum_k \exp \tilde{\xi}_k < 1$.

- Hence, taking expectations of (32) subject to information set $\mathcal{I}$

$$E_{\mathcal{I}}\left[F(\boldsymbol{\xi})\right] = \ln \sum_{j=0}^{N} \exp \xi_j^0 + \sum_{j=0}^{N} \frac{\exp \xi_j^0}{\sum_{k=0}^{N} \exp \xi_k^0} (E_{\mathcal{I}} \xi_j - \xi_j^0)$$
$$+ \frac{1}{2} \sum_{j=0}^{N} \sum_{k=0}^{N} E_{\mathcal{I}} \left(\frac{\partial^2 F(\tilde{\boldsymbol{\xi}})}{\partial \xi_j \partial \xi_k} (\xi_j = \xi_j^0)(\xi_k - \xi_k^0) \right). \tag{33}$$

- We apply this expansion firstly to the expected present value of consumption, $\sum_{j=0}^{T-t} c_{it+j}(1+r)^{-j}$.
- Let $N = T - t$ and let

$$\begin{aligned} \xi_j &= \ln c_{it+j} - j \ln(1+r) \\ \xi_j^0 &= E_{t-1} \ln c_{it+j} - j \ln(1+r), \quad i = 0, \ldots, T-t. \end{aligned} \tag{34}$$

- Then, substituting equation (34) into equation (33) and noting only the order of magnitude for the remainder term,

$$E_{\mathcal{I}}\left[\ln \sum_{j+0}^{T-t} \frac{c_{it+j}}{(1+r)^j}\right] = \ln \sum_{j+0}^{T-t} \exp\left[E_{t-1} \ln c_{it+j} - j\ln(1+r)\right] + \sum_{j=0}^{T-t} \theta_{it+j}\left[E_{\mathcal{I}} \ln c_{it+j} - E_{t-1} \ln c_{it+j}\right] + \mathcal{O}(E_{\mathcal{I}} \|\boldsymbol{\varepsilon}_{it}^{T}\|^2) \qquad (35)$$

where

$$\theta_{it+j} = \frac{\exp \xi_j^0}{\sum_{k=0}^{N} \exp \xi_k^0} = \frac{\exp\left[E_{t-1} \ln c_{it+j} - j\ln(1+r)\right]}{\sum_{k=0}^{T-t} \exp\left[E_{t-1} \ln c_{it+k} - kl(1-r)\right]}$$

and $\boldsymbol{\varepsilon}_{it}^{T}$ denotes the vector of future consumption innovations $(\varepsilon_{it}, \varepsilon_{it+1}, \ldots, \varepsilon_{iT})'$.

- The term θ_{it+j} can be seen as an annuitisation factor for consumption.
- We now apply the expansion (33) to the expected present vale of resources, $\sum_{j=0}^{R-t-1}(1+r)^{-j}y_{it+j}+A_{iT+1}(1+r)^{-(T-t)}$.
- Let $N=R-T$ and let

$$\begin{aligned}
\xi_j &= \ln y_{it+j} - j\ln(1-r) \\
\xi_j^0 &= E_{t-1}\ln y_{it+j} - j\ln(1+r) \quad j=0,\ldots,R-t-1 \\
\xi_N &= \ln\left[A_{it}-A_{iT+1}(1+r)^{-(T-t)}\right] \\
\xi_N^0 &= E_{t-1}\ln[A_{it}-A_{iT-1}(a+r)^{-(T-t)}].
\end{aligned} \tag{36}$$

- Then, substituting equation (36) into equation (33), and again noting only the order of magnitude for the remainder term,

$$
\begin{aligned}
& E_{\mathcal{I}} \ln \left(\sum_{j=0}^{R-r-1} \frac{y_{it+j}}{(1+r)^j} + A_{it} - \frac{A_{iT+1}}{(1+r)^{T-t}} \right) = \\
& \ln \left[\sum_{j=0}^{R-t-1} \exp \left[E_{t-1} \ln y_{it+j} - j \ln(1+r) \right] + \exp E_{t-1} \ln \left[A_{it} - \frac{A_{iT+1}}{(1+r)^{T-t}} \right] \right] \\
& + \pi_{it} \sum_{j=0}^{R-t-1} \alpha_{t+j} \left[E_{\mathcal{I}} \ln y_{it+j} - E_{t-1} \ln y_{it+j} \right] \\
& + (1-\pi_{it}) \left[E_{\mathcal{I}} \ln \left[A_{it} - \frac{A_{iT+1}}{(1+r)^{T-1}} \right] - E_{t-1} \ln \left[A_{it} - \frac{A_{iT+1}}{(a+r)^{T-t}} \right] \right] \\
& + \mathcal{O} \left(E_{t-1} \| (\boldsymbol{\nu}_{it}^{R-1}) \|^2 \right) \qquad (37)
\end{aligned}
$$

THE UNIVERSITY OF CHICAGO

where

$$\pi_{t+j} =$$

$$= \frac{\exp[E_{t-1} \ln y_{it+j} - j \ln(1+r)]}{\sum_{k=0}^{R-t-1} \exp[E_{t-1} \ln y_{it+k} - k \ln(1+r)]}$$

$$= \frac{\exp\left[\sum_{k=0}^{j}(\eta_{t+k} + E_{t-1}\bar{u}_{t+k}) + E_{t-1}\bar{u}_{t+j} - j \ln(1+r)\right]}{\sum_{k=0}^{R-t-1} \exp\left[\sum_{t=0}^{k}(\eta_{t+l} + E_{t-1}\bar{u}_{t+l}) + E_{t-1}\bar{u}_{t+k} - k \ln(1+r)\right]}$$

(This is the same as $\pi_{i,a,t}$ in Blundell et al., 2008.)

…can be seen as an annuitisation factor for income (common within a cohort because of the assumption of common income trends) and

$$\begin{aligned}\Xi_{i,a,t} =& 1 - \frac{\exp \xi_N^0}{\sum_{k=0}^{N} \exp \xi_k^0} \\ =& \frac{\sum_{j=0}^{R-t-1} \exp[E_{t-1} \ln y_{t+j} - j \ln(1+r)]}{\sum_{j=0}^{R-t-1} \exp[E_{t-1} \ln y_{it+j} - j \ln(1+r)] + \exp E_{t-1} \ln[A_{it} - A_{iT+1}/(1+R)^{T-t}]}\end{aligned}$$

is (roughly) the share of expected future labor income in current human and financial wealth (net of terminal assets) and $\boldsymbol{\nu}_{it}^{R-1}$ denotes the vector of future income shocks $(\nu_{it}', \nu_{it+1}', \ldots, \nu_{iR-1}')'$.

This corresponds to the $\Xi_{i,a,t}$.

- We are able to equate the subjects of equations (35) and (37) because the realised budget must balance and $\sum_{j=0}^{R-t} \frac{c_{it+j}}{(1+r)^j}$ and $\sum_{j=0}^{R-t-1} \frac{y_{it+j}}{(1+r)^j} + A_{it} - \frac{A_i T+1}{(1+r)^{T-t}}$ therefore have the same distribution.

- We use (35) and (37), taking differences between expectations at the start of the period, before the shocks are realised, and at the end of the period, after the shocks are realised.
- This gives

$$\begin{aligned}\varepsilon_{it}+&\mathcal{O}(E_t \parallel \boldsymbol{\varepsilon}_{it}^{T} \parallel^2 + E_{t-1} \parallel \boldsymbol{\varepsilon}_{it}^{T} \parallel^2) \\ &=\pi_{it}(v_{it}+\alpha_t u_{it})+\pi_{it}\Omega_t \\ &\quad + \mathcal{O}(E_t \parallel \boldsymbol{\nu}_{it}^{R-1} \parallel^2 + E_{t-1} \parallel \boldsymbol{\nu}_{it}^{R-1} \parallel^2),\end{aligned}$$

where the left hand side is the innovation to the expected present value of consumption and the right hand side is the innovation to the expected present value of income and

$$\Omega_t = \sum_{j=0}^{R-t-1} \pi_{t+j} \sum_{k=0}^{j} (E_t - E_{t-1})\omega_{t+k}$$

captures the revision to expectations of current and future common shocks.

- Squaring the two sides, taking expectations and inspecting terms reveals that the terms which are $\mathcal{O}(E_t \parallel \varepsilon_{it}^{T} \parallel^2 + E_{t-1} \parallel \varepsilon_{it}^{T} \parallel^2)$ are $\mathcal{O}(E_t \parallel \boldsymbol{\nu}_{it}^{R-1} \parallel^2 + E_{t-1} \parallel \boldsymbol{\nu}_{it}^{R-1} \parallel^2)$.
- Furthermore, since, for all $j \geq 0, \parallel \nu_{it+j} \parallel^2 = \mathcal{O}_p(E_t \parallel \nu_{it+j} \parallel^2)$ by Chebyshev's inequality, $E_t \parallel \boldsymbol{\nu}_{it}^{R-1} \parallel^2 = \mathcal{O}_p(E_{t-1} \parallel \boldsymbol{\nu}_{it}^{R-1} \parallel^2)$.
- Thus

$$\varepsilon_{it} = \Xi_{i,a,t}(v_{it} + \pi_{i,t} u_{it}) + \Xi_{i,a,t}\Omega_t + \mathcal{O}_p(E_{t-1} \parallel \boldsymbol{\nu}_{it}^{R-1} \parallel^2)$$

and therefore

$$\Delta \ln c_{it} = \kappa_t + \Xi_{i,a,t}(v_{it} + \alpha_t u_{it}) + \Xi_{i,a,t}\Omega_t + \mathcal{O}_p(E_{t-1} \parallel \boldsymbol{\nu}_{it}^{R-1} \parallel^2).$$

End of Digression.

Return to main text