THE
ECONOMIC
JOURNAL

The Economic Journal, **124** (*May*), 289–318. Doi: 10.1111/ecoj.12133 © 2014 The Author(s). The Economic Journal published by John Wiley & Sons Ltd on behalf of Royal Economic Society. Published by John Wiley & Sons, 9600 Garsington Road, Oxford OX4 2DQ, UK and 350 Main Street, Malden, MA 02148, USA.

INCOME DYNAMICS AND LIFE-CYCLE INEQUALITY: MECHANISMS AND CONTROVERSIES*

Richard Blundell

This study focuses on the transmission of inequality over the working life. A model of constrained intertemporal choice is used to provide structure to the distributional dynamics of wages, earnings, income and consumption. The mechanisms used to insure labour market shocks are examined in a partial-insurance setting where the manner and scope for insurance depends on the access to credit, the information available to consumers and the durability of income shocks. Drawing on recent research, family labour supply, the credit market and the tax system are all shown to play a key role. These mechanisms vary in importance across different points of the life cycle and the business cycle.

Economic inequality has many linked dimensions. Labour economists typically focus on inequality in hourly wages or earnings, public economists on disposable income and wealth and household economists on consumption. These different dimensions capture different aspects of inequality, and analysed together they can considerably enhance our understanding of inequality dynamics. The link between the various measures of inequality is mediated by multiple 'insurance' mechanisms. These mechanisms include credit markets, labour supply, taxation, welfare benefits, formal insurance, informal gifts, transfers etc.

The objective of the research reported in this article is to use the framework of constrained intertemporal choice over consumption, saving and family labour supply to provide a structure for the distributional dynamics of wages, earnings, income and consumption. Here, we focus on the evolution of inequality over the working life. At the heart of this analysis is the study of labour income dynamics. The dynamics of labour income and wages are the foundation for thinking about the transmission of inequality over the life cycle. It is the key to exploring the mechanisms used by families to 'insure' against labour market shocks.

The definition of insurance adopted in this work is very broad. It covers formal and informal mechanisms that are used to attenuate the impact of shocks to earned income. These mechanisms will vary in importance across different types of households at different points of their life cycle and at different points in the business cycle. The manner and scope for insurance depends on the access to credit, the

* Corresponding author: Richard Blundell, Department of Economics, University College London, London WC1E6BT, UK. Email: r.blundell@ucl.ac.uk.

This study is based on my 2013 Royal Economic Society Presidential Lecture. I would like to thank Cormac O'Dea, Ian Preston, Luigi Pistaferri, Itay Sapota-Eksten, the Editor and referee for helpful comments. The research is part of the programme of research of the ESRC Centre for the Microeconomic Analysis of Public Policy at IFS. Funding from the ESRC, grant number R000239865, is gratefully acknowledged. Material from the FES made available by the ONS through the ESRC Data Archive has been used by permission of the controller of HMSO. Neither the ONS nor the ESRC Data Archive bears responsibility for the analysis or the interpretation of the data reported here. The usual disclaimer applies.

information available to consumers and the durability of income shocks. To contrast empirical observation with the standard incomplete markets self-insurance model, a general partial-insurance framework following Blundell *et al.* (2008) is used in which transmission parameters between income shocks and consumption growth indicate the degree of insurance.

In one of the foundation studies in this area, Deaton and Paxson (1994) noted that an implication of the permanent income hypothesis was that, for any birth cohort, inequality of consumption and income should grow with age. Examining survey data on income and consumption from a wide range of countries, they found this to be true. The income variance will increase with permanent income shocks and the variance of consumption will also cumulate permanent income shocks. The degree to which these move in line will depend on the degree of precautionary savings and access to credit. This was a key insight for the subsequent work in this field. Recent evidence on the growth in consumption inequality over the life cycle for different birth cohorts in the UK and the US shows a strong increase in inequality across cohorts. Younger birth cohorts face higher overall consumption inequality during their working life than similarly aged older cohorts. Figures 1 and 2 show the striking evidence from the UK[1] and from the US[2] respectively. Income inequality growth displays some similarities, but a clearly different pattern; see Figure 3 for the UK, for example.

The aim of this study is to look behind these Figures and investigate the linkages that underlie these inequality measures. Understanding the importance of labour market

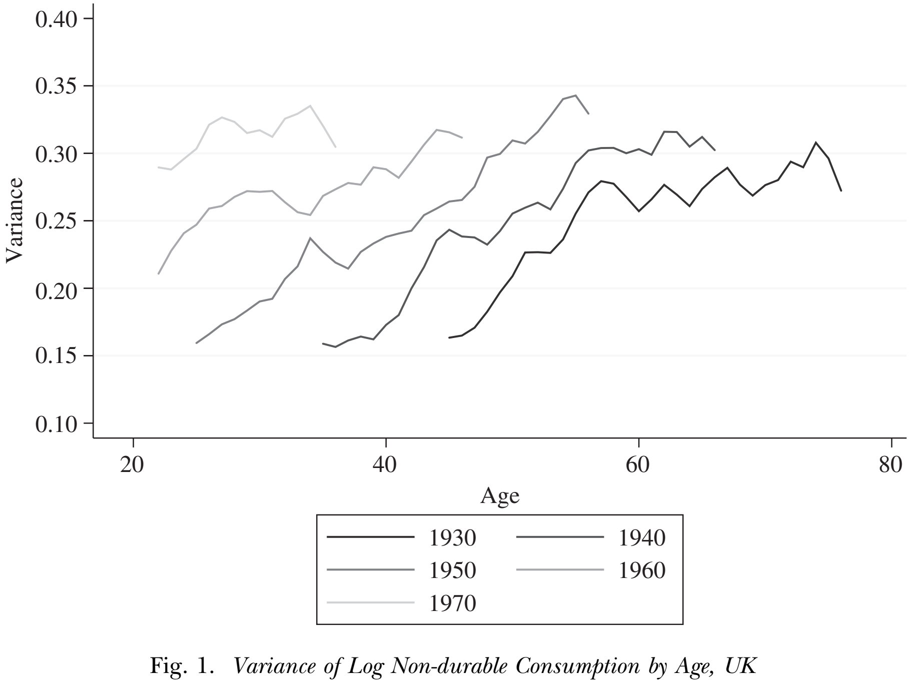

Fig. 1. *Variance of Log Non-durable Consumption by Age, UK*

[1] I am grateful to Cormac O'Dea for these Figures from the expenditure data (FES and EFS) for Britain. See Brewer and O'Dea (2012) for further description and data sources.
[2] I am grateful to Luigi Pistaferri for these Figures from the expenditure data (PSID and CEX) for the US.

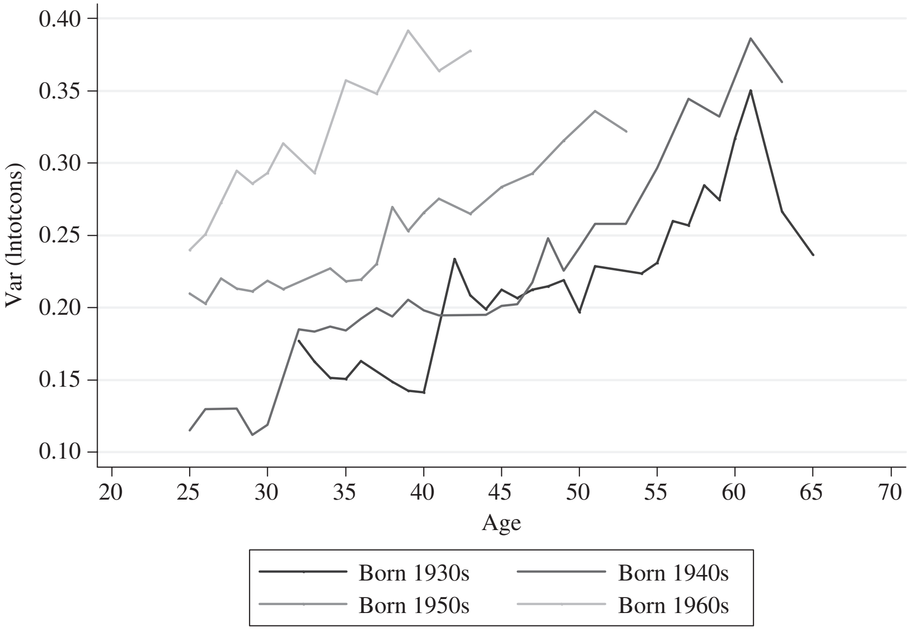

Fig. 2. *Variance of Log Non-durable Consumption by Age, US*

shocks over the working life and uncovering the mechanisms that link these measures of inequality across the lifetime of each birth cohort is a key motivation for this research.

This study attempts to bridge three literatures. The first concerns the examination of the evolution in inequality over time for consumption and income.[3] This leads us naturally to a second set of literature which concerns the panel data dynamics of income. Typically, this literature has focused on the variance decomposition of male earnings.[4] Finally, there is the literature on intertemporal consumption decisions under uncertainty, especially those which examine excess sensitivity and partial insurance.[5] Recently, there has also been growth in an important and directly related literature on information, family labour supply, learning and human capital.[6]

[3] In particular, studies from the BLS in the US, comprehensively summarised in Johnson *et al.* (2005). This work emanated from the important early work by Cutler and Katz (1992). In the UK, Blundell and Preston (1995) provided a comparison of income and consumption inequality, focusing on the strong inequality growth episode of the 1980s highlighted in the work of Atkinson (1999) and earlier references therein.

[4] Foremost among these are the studies by Lillard and Willis (1978), Lillard and Weiss (1979), MaCurdy (1982), Abowd and Card (1989), Moffitt and Gottschalk (1995, 2002), Baker (1997), Haider (2001), Meghir and Pistaferri (2004), Haider and Solon (2006), Guvenen (2009), Browning *et al.* (2010), Altonji *et al.* (2013) and DeBacker *et al.* (2013). In terms of the durability of income shocks, the path-breaking work is the Moffitt and Gottschalk (1994) study of US male earnings.

[5] Among the key works are Hall and Mishkin (1982), Campbell and Deaton (1989), Cochrane (1991), Deaton and Paxson (1994), Townsend (1994), Attanasio and Davis (1996), Blundell and Preston (1998), Krueger and Perri (2004), Heathcote *et al.* (2004), Storesletten *et al.* (2004), Krueger and Perri (2006), Heathcote *et al.* (2007), Attanasio and Pavoni (2011), Blundell *et al.* (2008), Primiceri and Van Rens (2009), Low *et al.* (2010), Guvenen and Smith (2012) etc.

[6] Including Pistaferri (2003), Cunha *et al.* (2005), Cunha *et al.* (2007), Guvenen (2007), Kaufmann and Pistaferri (2009), Huggett *et al.* (2011) and Blundell *et al.* (2012).

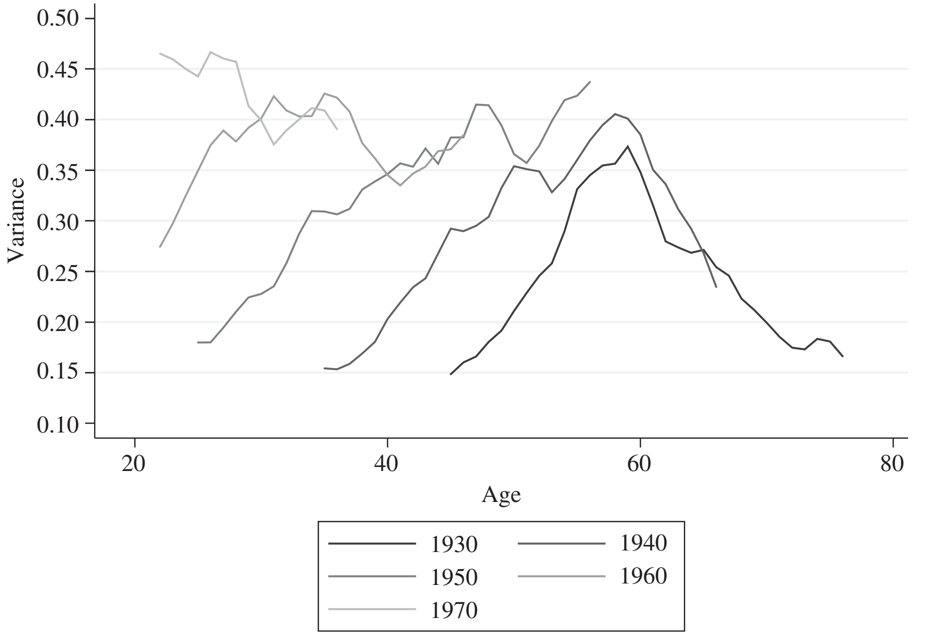

Fig. 3. *Variance of Log Disposable Income by Age, UK*

The layout of the remainder of the article is as follows. The next Section examines the key ingredients of labour market dynamics. Section 2 develops the linkage between these models and intertemporal models of consumer behaviour. Models of partial insurance are considered in Section 3 and some summary empirical results are presented. Section 4 looks at the role of family labour supply as an additional insurance mechanism. Robustness issues are discussed in Section 5, and Section 6 concludes with a summary of what has been learned so far, drawing out some directions for future research.

1. Some Panel Data Income Dynamics

There is an extensive applied econometrics literature modelling income dynamics using the extensive panel data now available for most modern economies. The focus in this article is on non-stationarity and on the persistence of shocks. These are some of the key components of labour income dynamics as they impact on consumption and saving decisions.

A convenient general specification of log income $y_{i,a,t}(\equiv \ln Y_{i,a,t})$ for consumer i of age a in time period t with observable characteristics $Z_{i,a,t}$ is given by

$$y_{i,a,t} = B'_{i,a,t}f_i + Z'_{i,a,t}\varphi + y^P_{i,a,t} + y^T_{i,a,t}, \tag{1}$$

where y^P_{it} is a persistent process of income shocks which adds to the individual-specific trend (by age and time) $B'_{i,a,t}f_i$ and where y^T_{it} is a transitory shock represented by some low-order MA process. A key consideration is to allow variances (or factor loadings) of y^P and y^T to vary with age and time for each birth cohort. Of course, not all three effects can be identified without restrictions. A simple way to extract these effects is to assume common life-cycle age effects and time effects, across cohorts.

For any cohort, a reasonably general specification for the idiosyncratic effects $B'_{i,t}f_i$, which we explore further below, is given by

$$B'_{i,t}f_i = f_{0i} + p_t f_{1i}, \tag{2}$$

where f_{0i} is an individual effect, and p_t is some time (or age) trend so that $p_t f_{1i}$ represents an idiosyncratic trend.

Suppose $y^T_{i,t}$ is represented by a low-order MA(q)

$$v_{it} = \sum_{j=0}^{q} \theta_j \varepsilon_{i,t-j} \text{ with } \theta_0 \equiv 1, \tag{3}$$

and y^P_{it} by

$$y^P_{it} = \rho y^P_{it-1} + \zeta_{it}. \tag{4}$$

Suppose also that we assume the deterministic term $Z'_{i,t}\varphi$ is already removed from $y_{i,t}$, the dynamic panel data income model becomes

$$y_{i,t} = p_t f_{1i} + f_{0i} + y^P_{i,t} + \sum_{j=0}^{q} \theta_j \varepsilon_{i,t-j}, \tag{5}$$

as implemented in Blundell *et al.* (2014), for example.

If $q = 1$, then this implies a key quasi-difference moment restriction

$$\text{cov}(\Delta^\rho y_t, \Delta^\rho y_{t-2}) = \text{var}(f_0)(1-\rho)^2 + \text{var}(f_1)\Delta^\rho p_t \Delta^\rho p_{t-2} - \rho\theta_1 \text{var}(\varepsilon_{t-2}), \tag{6}$$

where $\Delta^\rho = (1 - \rho L)$ is the quasi-difference operator. For large $\rho = 1$ and small θ_1, (6) implies

$$\text{cov}(\Delta y_t, \Delta y_{t-2}) \simeq \text{var}(f_1)\Delta p_t \Delta p_{t-2}. \tag{7}$$

Consequently, for near unit root permanent shocks and innovation transitory shocks, if we set the individual trends to zero ($\text{var}(f_1) = 0$), there are no autocovariances of order two or above remaining in the growth rates of the income variable y. Allowing for a higher MA process relaxes this; but at some point, the autocovariance structure for income growth drops to zero. This observation is a key source of identification in 'permanent–transitory' panel data models of income dynamics (MaCurdy 1982; Meghir and Pistaferri, 2004, 2011).

1.1. *Idiosyncratic Trends*

But what of idiosyncratic trends? The trend term $p_t f_{1i}$ in (5) could take a number of forms. Two alternatives worth highlighting are as follows:

(*a*) deterministic idiosyncratic trend:

$$p_t f_{1i} = r(t) f_{1i},$$

where r is a known function of t, e.g. $r(t) = t$, and (*b*) stochastic trend in 'ability prices':

$$p_t = p_{t-1} + \xi_t.$$

with $E_{t-1}\xi_t = 0$.

Evidence points to cases where each of these could be of key importance. Deterministic trends, as in (*a*), appear most prominently early in the working life (Haider and Solon, 2006) and for the higher educated (Blundell *et al.*, 2014). Formally, this is a life-cycle effect, an age effect for any birth cohort. The earlier in the career we select individuals in a panel and the higher their education the more likely this is to be a dominant effect.[7] Alternatively, stochastic trends, as in (*b*), are most likely to occur during periods of technical change when skill prices are changing across the unobserved ability distribution. Formally, this is a calender time effect within any labour market.

A key question in relation to (*b*) is how many skill price factors do we need? Heckman and Scheinkman (1987) show that maybe many such terms are required. In any panel data study sampled from a large disparate economy like the US or the UK, it may prove difficult to identify skill prices in local labour markets. For such panel data studies, a stochastic trend, represented in the permanent idiosyncratic shock term (4), may provide a good approximation.[8] As we will see these considerations have important implications for the distribution of consumption growth rates.

1.2. *The Permanent–transitory Model of Income Dynamics*

As a simple representation of non-stationary income dynamics for each household i, the permanent–transitory decomposition provides a useful baseline. We rewrite (5) as

$$y_{it} = y_{it}^P + y_{it}^T \tag{8}$$

with

$$y_{it}^P = y_{it-1}^P + \zeta_{it} \tag{9}$$

and transitory or mean reverting component, $y_{it}^T = v_{i,t}$

$$v_{it} = \sum_{j=0}^{q} \theta_j \varepsilon_{i,t-j} \text{ with } \theta_0 \equiv 1. \tag{10}$$

This formulation implies a restrictive structure for the autocovariances of $\Delta y_{it}(=\zeta_{it} + \Delta v_{it})$

$$\text{cov}(\Delta y_t, \Delta y_{t+s}) = \begin{cases} \text{var}(\zeta_t) + \text{var}(\Delta v_t) & \text{for } s = 0 \\ \text{cov}(\Delta v_t, \Delta v_{t+s}) & \text{for } s \neq 0. \end{cases} \tag{11}$$

Allowing for an MA(q) process, for example, adds $q-1$ extra parameter (the $q-1$ MA coefficients) but also $q-1$ extra moments, so that identification is unaffected.

[7] For example, Blundell *et al.* (2014) find idiosyncratic age trends to be important only for those with college education.

[8] This would contrast to the study of a sample from a specific high-skill profession data (scientists, for example) as used in the innovative early study by Lillard and Weiss (1979). See also Lillard and Willis (1978).

It seems reasonable that the variance of shocks to income, especially more permanent shocks, should be largest at either end of the working life. At the beginning of the working life, individuals are sorting themselves into occupations and firms that best suit their human capital resulting in positive and negative longer term 'surprises'. At the end of the working life, health shocks are likely to become more relevant.

We can see this process taking place most clearly in studies that use detailed lifetime profiles from population register panel data. Figures 4 and 5 are derived from a study of income dynamics from the Norwegian Population Register Panel as analysed in Blundell *et al.* (2014). The Figures plot the variances of the permanent shocks to labour market income and disposable income for men during their working life. The second Figure separates out the low-educated group and shows the strong increase in the variance of permanent shocks at older working ages for this group. Indeed, this study suggests that the overall U shape for variances over the life cycle may reflect an aggregation over high-educated workers, whose shocks are largest earlier their lifetime, and low-educated workers, who face larger variances to persistent income shocks later in their working life.

1.3. *Some (Simple) Empirics of Income Dynamics*

To examine the ability of the permanent–transitory income model to provide a good representation of income dynamics in the UK and the US, Tables 1 and 2 present the autocovariance structure of the Panel Study of Income Dynamics in the US, PSID and the British Household Panel Study in the UK (BHPS) respectively. These tables

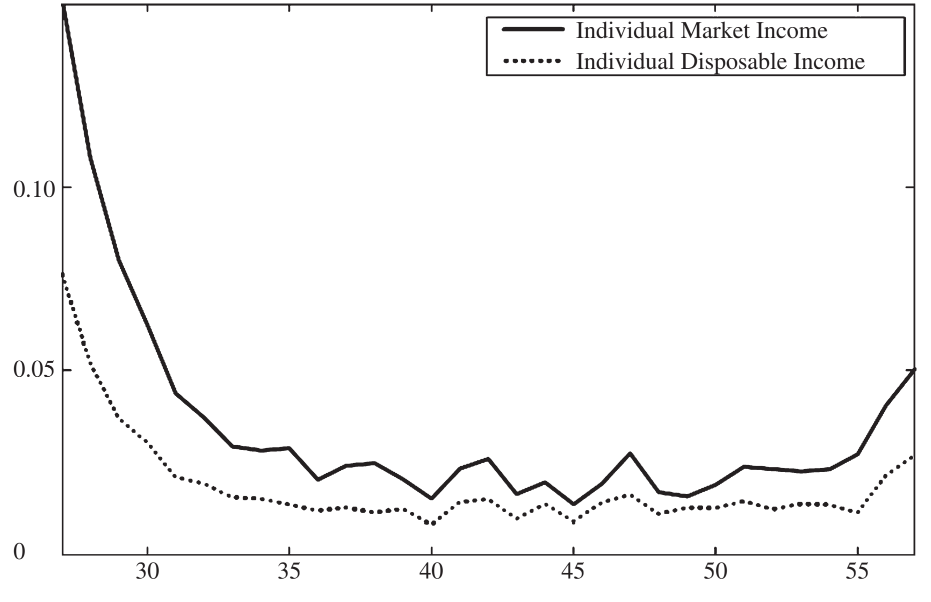

Fig. 4. *Variance of Permanent Shocks by Age, Norway*
Source. Blundell *et al.* (2014).

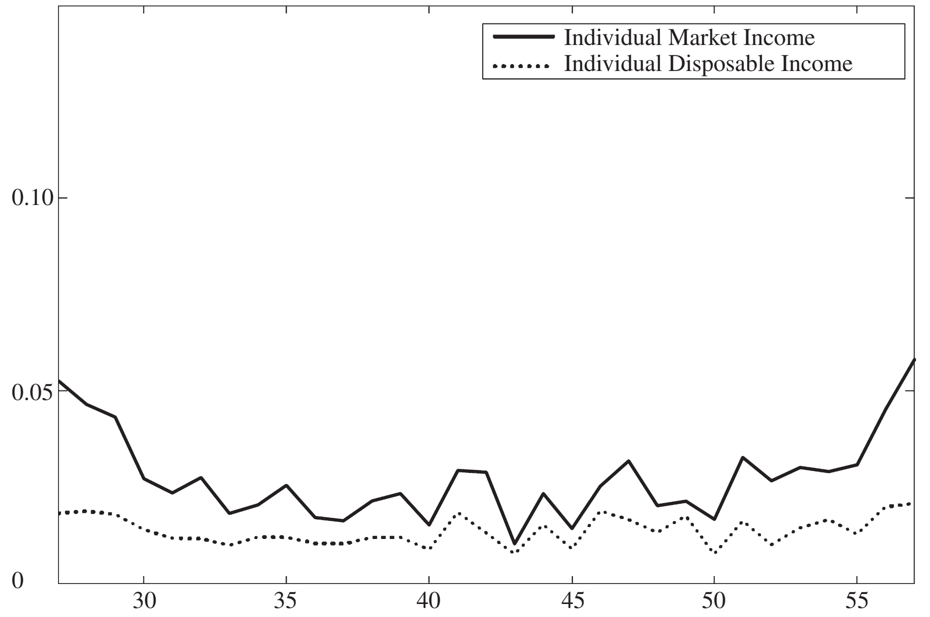

Fig. 5. *Variance of Permanent Shocks by Age (Low Educated), Norway*
Source. Blundell *et al.* (2014).

suggest that the latent factor structure (11) aligns 'well' with the autocovariance structure of income in both of these panel data sets. In addition, there is evidence of non-stationary behaviour with the variances changing over time.[9] However, there is little evidence of autocovariances extending much beyond the first order. For the PSID data, the test that $\text{cov}(\Delta y_{t+1}, \Delta y_t) = 0$ for all t, has a p-value of 0.0048; for the $t + 2$ term, it is 0.0125 and for the $t + 3$ and $t + 4$ terms, the p-values are 0.6507 and 0.9875 respectively.

Why may it be hard to detect idiosyncratic trends? To a large extent it is probably a matter of the age selection used.[10] As noted above, forecastable components and differential trends are most important early in the life cycle. Tables 1 and 2 use samples where the head is male, lives in a couple and prime aged (aged between 30 and 60 years). Such a selection largely removes the early career trends and the later career health effects. Moreover, the baseline specifications (8)–(10) allow for general fixed effects and initial conditions. Indeed, this structure is easily extended to account for higher moments, and a regular deconvolution argument can be used to show identification of variances and complete distributions of the factors, see Bonhomme and Robin (2010).

A key point is to allow for non-stationarity. In this identification/estimation approach, the variances (or factor loadings) of the permanent and transitory factors

[9] The UK BHPS data start in 1992 and therefore misses the large increase in inequality during the 1980s, see Blundell and Preston (1998).
[10] Indeed, as noted above, the Blundell *et al.* (2014) population register study finds that idiosyncratic trends are only significant for the young and high-educated group.

Table 1

The Autocovariance Matrix of Income Growth PSID

Year	var(Δy_t)	cov(Δy_{t+1},Δy_t)	cov(Δy_{t+2},Δy_t)
1980	0.0832	−0.0196	−0.0018
	(0.0089)	(0.0035)	(0.0032)
1981	0.0717	−0.0220	−0.0074
	(0.0075)	(0.0034)	(0.0037)
1982	0.0718	−0.0226	−0.0081
	(0.0051)	(0.0035)	(0.0026)
1983	0.0783	−0.0209	−0.0094
	(0.0066)	(0.0034)	(0.0042)
1984	0.0805	−0.0288	−0.0034
	(0.0055)	(0.0036)	(0.0032)
1985	0.1090	−0.0379	−0.0019
	(0.0180)	(0.0074)	(0.0038)
1986	0.1023	−0.0354	−0.0115
	(0.0077)	(0.0054)	(0.0038)
1987	0.1116	−0.0375	0.0016
	(0.0097)	(0.0051)	(0.0046)
1988	0.0925	−0.0313	−0.0021
	(0.0080)	(0.0042)	(0.0032)
1989	0.0883	−0.0280	−0.0035
	(0.0067)	(0.0059)	(0.0034)
1990	0.0924	−0.0296	−0.0067
	(0.0095)	(0.0049)	(0.0050)
1991	0.0818	−0.0299	NA
	(0.0059)	(0.0040)	
1992	0.1177	NA	NA
	(0.0079)		

Source. Blundell *et al.* (2008).

are permitted to vary non-parametrically with cohort, education and time. Indeed, the relative variance of these two factors is a measure of persistence or durability of labour income shocks. This is what drives the changing relationship between the dispersion of income and consumption.

2. Intertemporal Choice and the Evolution of the Consumption Distribution

2.1. *Self-insurance*

Recall that the main objective of the research described in this article is to uncover the transmission of inequality from earnings to consumption over the working life. To provide a underlying economic framework for thinking about this, a model of constrained intertemporal consumer choice with partial insurance is developed in this section.

As a baseline specification of the household consumption decision, suppose that at time t each household i maximises the conditional expectation of the discounted stream of time separable, differentiable utility:

$$\max_C \mathrm{E}_t \sum_{j=0}^{T-t} u_t(C_{i,t+j}, Z_{i,t+j}),$$

Table 2
The Autocovariance Matrix of Income Growth BHPS

Year	var(Δy_t)	cov(Δy_{t+1},Δy_t)	cov(Δy_{t+2},Δy_t)
1980	0.1429	−0.0504	−0.0044
	(0.0071)	(0.0042)	(0.0039)
1981	0.0717	−0.0220	−0.0074
	(0.0075)	(0.0034)	(0.0037)
1982	0.0718	−0.0226	−0.0081
	(0.0051)	(0.0035)	(0.0026)
1983	0.0783	−0.0209	−0.0094
	(0.0066)	(0.0034)	(0.0042)
1984	0.0805	−0.0288	−0.0034
	(0.0055)	(0.0036)	(0.0032)
1985	0.1090	−0.0379	−0.0019
	(0.0180)	(0.0074)	(0.0038)
1986	0.1023	−0.0354	−0.0115
	(0.0077)	(0.0054)	(0.0038)
1987	0.1116	−0.0375	0.0016
	(0.0097)	(0.0051)	(0.0046)
1988	0.0925	−0.0313	−0.0021
	(0.0080)	(0.0042)	(0.0032)
1989	0.0883	−0.0280	−0.0035
	(0.0067)	(0.0059)	(0.0034)
1990	0.0924	−0.0296	−0.0067
	(0.0095)	(0.0049)	(0.0050)
1991	0.0818	−0.0299	NA
	(0.0059)	(0.0040)	
1992	0.1177	NA	NA
	(0.0079)		

Source. Blundell and Etheridge (2008).

where $Z_{i,t+j}$ represents taste shifters, demographic non-separabilities and discount rate heterogeneity. The retirement age, set at L, is assumed known and certain, as is the end of the life cycle at $T(> L)$.

Individuals can self-insure using a credit market with access to a risk-free bond with real return r_{t+j}. Consequently, in this simple credit market framework, consumption and income are linked through the intertemporal budget constraint

$$A_{i,t+j+1} = (1 + r_{t+j})(A_{i,t+j} + Y_{i,t+j} - C_{i,t+j}) \text{ with } A_{i,T} = 0. \tag{12}$$

With self-insurance and constant relative risk averse (CRRA) preferences

$$u_t(C_{i,t+j}, Z_{i,t+j}) \equiv \frac{1}{(1+\delta)^j} \frac{C^{\beta}_{i,t+j} - 1}{\beta} \mathrm{e}^{Z'_{i,t+j}\vartheta}, \tag{13}$$

the first-order conditions become

$$C^{\beta-1}_{i,t-1} = \frac{1 + r_{t-1}}{1 + \delta} \mathrm{e}^{\Delta Z'_{i,t}\vartheta_t} \mathrm{E}_{t-1} C^{\beta-1}_{i,t}.$$

Building on the studies of Campbell (1993) and Blundell and Stoker (1999), Blundell *et al.* (2013) derive a general approximation for consumption growth for this self-insurance model given by

$$\Delta c_{i,t} \simeq \Delta Z'_{i,t}\vartheta'_t + \eta_{i,t} + \Gamma_{i,t}, \tag{14}$$

where $c_{i,t} \equiv \Delta \log C_{i,t}$, $\vartheta'_t = (1-\beta)^{-1}\vartheta_t$, $\eta_{i,t}$ is a consumption growth shock with $E_{t-1}\eta_{i,t} = 0$, $\Gamma_{i,t}$ captures any slope in the consumption path due to interest rates, impatience or precautionary savings and the error in the approximation is $\mathcal{O}(E_{t-1}\eta^2_{i,t})$. Conveniently, with CRRA preferences, $\Gamma_{i,t}$ is independent of $C_{i,t}$.

2.2. *Linking the Evolution of the Consumption and Income Distributions*

For log income growth in the permanent–transitory model (9, 10) we have

$$\Delta y_{i,t+k} = \zeta_{i,t+k} + \sum_{j=0}^{q} \theta_j \varepsilon_{i,t+k-j}. \tag{15}$$

The intertemporal budget constraint (12) can be written as

$$\sum_{k=0}^{T-t} Q_{t+k} C_{i,t+k} = \sum_{k=0}^{L-t} Q_{t+k} Y_{i,t+k} + A_{i,t},$$

where Y is the level of income, T is death, L is retirement and Q_{t+k} is appropriate discount factor $\prod_{i=1}^{k}(1+r_{t+i})$, $k = 1,\ldots,T-t$ (and $Q_t = 1$).

Define

$$\pi_{i,t} = \sum_{k=0}^{L-t} Q_{t+k} Y_{i,t-k} \Big/ \sum_{k=0}^{L-t} Q_{t+k} Y_{i,t-k} + A_{i,t}, \tag{16}$$

the share of future labour income in current human and financial wealth, and

$$\gamma_t \simeq \frac{r}{1+r}\left[1 + \sum_{j=1}^{q} \theta_j \Big/ (1+r)^j\right] \tag{17}$$

the annuity factor (for $r_t = r$).

Blundell *et al.* (2013) show that the stochastic individual element $\eta_{i,t}$ in consumption growth (14) is approximated by

$$\eta_{i,t} \simeq \pi_{i,t}\left(\zeta_{i,t} + \gamma_{Lt}\varepsilon_{i,t}\right),$$

where

$$\pi_{it} \approx \frac{\text{Human wealth}_{it}}{\text{Assets}_{it} + \text{Human wealth}_{it}}$$

and γ_{Lt} is the annuity value of a transitory shock for an individual aged t retiring at age L. A link between consumption growth and the income shocks can be expressed, to order $O(\|v_t\|^2)$, where $v_t = (\zeta_t, \varepsilon_t)'$, as

$$\Delta \ln C_{it} \cong \Gamma_{it} + \Delta Z'_{it}\varphi^c + \pi_{it}\zeta_{it} + \pi_{it}\gamma_{Lt}\varepsilon_{it} + \xi_{it}. \tag{18}$$

As the expression for π suggests, when assets are close to zero, permanent shocks flow completely into consumption in this self-insurance model. Outside the credit market, there are no additional mechanisms in this simple framework for insuring labour market shocks.

Each of the terms on the rhs of (18) captures key factors in the determination of intertemporal consumption growth. In particular, Γ_{it} reflects impatience, precautionary savings, intertemporal substitution; $\Delta Z_{it}^{\prime}\varphi^{c}$ captures deterministic preference shifts and labour supply non-separabilities; $\pi_{it}\zeta_{it}$ measures the impact of permanent income shocks. The $\pi_{it}\gamma_{Lt}\varepsilon_{it}$ term measures the impact of transitory income shocks, where $\gamma_{Lt} < 1$ is the annuitisation factor. Finally, ξ_{it} allows for the impact of shocks to higher income moments, etc.[11]

As noted above, the degree of self-insurance is reflected in the π parameter (16), which corresponds to the ratio of human capital wealth to total wealth (financial plus human capital wealth). The term $(1 - \pi_{it})$ reflects the extent to which 'permanent' shocks are insurable in a finite horizon model with incomplete markets. For given level of human capital wealth, past savings imply higher financial wealth today, and hence a lower value of π_{it}: consumption responds less to income shocks (precautionary saving). Individuals approaching retirement will have a lower value of π. In the certainty equivalence version of the PIH, $\pi_{it} \simeq 1$ and $\gamma_{Lt} \simeq 0$. This shows the importance of measuring assets. Without data on the level of net assets it is difficult to accurately measure π_{it}, and consequently difficult to examine deviations from the simple self-insurance model.

Blundell *et al.* (2013) report a number of simulation experiments for this specification. For reasonable values of parameters, π_{it} declines with age as future labour income diminishes and assets are built up. A high discount rate discourages saving as it is more costly in terms of utility for individuals to self-insure. A high elasticity of intertemporal substitution also discourages saving. All these cases therefore involve diminished self-insurance and raise π_{it}. High-income growth also reduces the need for saving as individuals do not want to accumulate savings and move resources into the future when income is high. Eliminating transitory variance growth raises π_{it} but not by very much. Introducing a social security pension (equal to half of final income) raises π and also makes it flatter over the life cycle. Liquidity constraints are found to have a similar impact.

2.3. *When Does Consumption Inequality Measure Welfare Inequality?*

We pause at this point and ask whether consumption better reflects household welfare than does some measure of current income. To do this we define $\tilde{Y}_i$ as that *certain* present discounted value of lifetime income which would allow the individual to achieve the same expected utility. The consumption stream $\tilde{\mathbf{C}}_i = \tilde{\mathbf{C}}(EU_i)$ that would be chosen given $\tilde{Y}_i$ satisfies

$$\sum_t u_t(\tilde{C}_{it}) \equiv \mathrm{E}\left[\sum_t u_t(C_{it})\right] = EU_i.$$

[11] Introducing higher moments and non-linearities in shocks to the income process will be a key area of development in future research.

Suppose we now look at comparisons across individuals facing different income risk. First, consider the case of Constant Absolute Risk Aversion (CARA) preferences,

$$u_t(C_{it}) = -\alpha_t \exp(-\gamma_t C_{it}) \quad \alpha_t, \gamma_t > 0, t > 0.$$

Blundell and Preston (1998) show that $C_{it} \geq C_{jt}$ implies $EU_i \geq EU_j$ whenever individuals i and j share the same year of birth if and only if $\mathbf{C}_i = \tilde{\mathbf{C}}(EU_i)$, whatever the distribution of future income. What about more reasonable risk preferences? This result is a special case of a more general result for decreasing absolute risk aversion preferences that include the more familiar CRRA case (13). These preferences imply $C_{i0} < \tilde{C}_{i0}$, i.e. that there is 'excess' precautionary saving if higher incomes decrease risk aversion. Consumption overestimates the welfare cost of income risk.

Neither income nor consumption accurately reflects consumer welfare. In some circumstances, consumption can be considered an improvement on income. The view taken here is that consumption and income are both valuable in understanding the dynamics of inequality and the underlying mechanisms behind changes in inequality. This is a positive rather than a normative analysis.

3. Partial Insurance

In the partial-insurance approach (Blundell *et al.*, 2008), transmission parameters are specified that link the shocks to income with consumption growth at the decision unit level, generalising the self-insurance formulation (18). These transmission parameters can change across time and may differ across individuals according to their birth cohorts. They reflect the degree of 'insurance' available to individuals experiencing the income shocks and encompass self-insurance through simple credit markets as well as other mechanisms used to smooth the impact of labour market shocks on consumption.

The specification of the transmission parameters will allow the degree of insurance to differ depending on whether the shock to income is short lived or is persistent. To quote Deaton and Paxson (1994), 'one of the main reasons for measuring consumption inequality and its evolution is to help understand and calibrate the way in which the economy handles risk'.

In theory it is possible to construct economic environments that insure consumption fully against idiosyncratic income shocks. This is the perfect insurance case in which $\pi_{it} = 0$. Such economic environments would struggle to be achievable in reality due to moral hazard and limited enforcement, among other compelling reasons for incomplete credit markets. Indeed, the assumptions required for complete insurance have been soundly rejected in careful empirical studies, Attanasio and Davis (1996) being a leading example. To allow for more general 'insurance' mechanisms, Blundell *et al.* (2008) introduce 'partial insurance'. This is designed to capture the possibility of 'excess insurance' and also 'excess sensitivity'.

To understand this approach, first note that the stochastic Euler equation for consumption growth is consistent with many stochastic processes for consumption. It does not say anything about the variance of consumption. In the full information perfect market model with separable preferences, the variance of consumption growth is zero. In comparison with the self-insurance model, the intertemporal budget constraint based on a single asset is violated. Partial insurance allows some, but not full, additional insurance. For example, Attanasio and Pavoni (2011) consider an economy with moral hazard and hidden asset accumulation – individuals now have hidden access to a simple credit market. They show that, depending on the cost of shirking and the persistence of the income shock, some partial insurance is possible. A linear insurance rule can be obtained as an 'exact' solution in a dynamic Mirrlees model with CRRA utility.

3.1. *Consumption Dynamics with Partial Insurance*

To allow for partial insurance, we need to account for additional 'insurance' mechanisms and excess sensitivity. For this, Blundell *et al.* (2008) introduce transmission parameters ϕ_t and ψ_t. For any birth cohort, the consumption growth relationship (18) is now written as

$$\Delta \ln C_{it} \cong \Gamma_{it} + \Delta Z_{it}' \varphi^c + \xi_{it} + \phi_t \zeta_{it} + \psi_t \varepsilon_{it}, \tag{19}$$

in which partial insurance w.r.t. permanent shocks implies $0 \leq 1 - \phi_t \leq 1$, whereas partial insurance w.r.t. transitory shocks implies $0 \leq 1 - \psi_t \leq 1$. The expressions $1 - \phi_t$ and $1 - \psi_t$ then measure the fractions insured and subsume π_t and γ_t from the self-insurance model. With the consumption growth relationship (19), we now have a (latent) factor structure that provides the panel data moments linking the evolution of distribution of consumption to the evolution of labour income distribution. This structure describes how consumption updates in response to income shocks.

3.2. *The Key Panel Data Moments*

Taking the models for income dynamics (15) and consumption dynamics (19) together we can now derive the second-order panel data variances and autocovariances that serve to identify the unknown transmission parameters, ϕ_t and ψ_t, of the partial-insurance specification. The autocovariance structure for log-adjusted income growth ($\Delta y_t \equiv \Delta \ln Y_t - \Delta Z_t' \phi^y$) is given in (11). For log consumption ($\Delta c_t \equiv \Delta \ln\ C_{it} \cong \Gamma_{it} + \Delta Z_{it}' \varphi^c$)

$$\mathrm{cov}(\Delta c_t, \Delta c_{t+s}) = \phi_t^2 \mathrm{var}(\zeta_t) + \psi_t^2 \mathrm{var}(\varepsilon_t) + \mathrm{var}(\xi_t) \tag{20}$$

for $s = 0$ and zero otherwise. For the cross-moments between income and consumption growth:

$$\mathrm{cov}(\Delta c_t, \Delta y_{t+s}) = \begin{cases} \phi_t \mathrm{var}(\zeta_t) + \psi_t \mathrm{var}(\varepsilon_t) \\ \psi_t \mathrm{cov}(\varepsilon_t, \Delta v_{t+s}) \end{cases} \tag{21}$$

for $s = 0$ and $s > 0$ respectively. A simple summary of the key panel data moments is given by

$$
\begin{aligned}
\mathrm{var}(\Delta y_t) &= \mathrm{var}(\zeta_t) + \mathrm{var}(\Delta\varepsilon_t) \\
\mathrm{cov}(\Delta y_{t+1}, \Delta y_t) &= -\mathrm{var}(\varepsilon_t) \\
\mathrm{var}(\Delta c_t) &= \phi_t^2 \mathrm{var}(\zeta_t) + \psi_t^2 \mathrm{var}(\varepsilon_t) + \mathrm{var}(\xi_{it}) + \mathrm{var}(u_{it}^c) \\
\mathrm{var}(\Delta c_t, \Delta c_{it+1}) &= -\mathrm{var}(u_{it}^c) \\
\mathrm{cov}(\Delta c_t, \Delta y_t) &= \phi_t \mathrm{var}(\zeta_t) + \psi_t \mathrm{var}(\varepsilon_t) \\
\mathrm{cov}(\Delta c_t, \Delta y_{t+1}) &= -\psi_t \mathrm{var}(\varepsilon_t)
\end{aligned} \tag{22}
$$

where u_{it}^c is a term reflecting measurement error in the consumption series. There are six time-varying autocovariances on the right-hand side of (19) and six unknown time-varying parameters – the two transmission parameters ϕ_t and ψ_t, and the four unknown factor variances $\mathrm{var}(\zeta_t), \mathrm{var}(\varepsilon_t), \mathrm{var}(u_{it}^c)$ and $\mathrm{var}(\xi_{it})$. This leads naturally to a discussion of identification.[12]

3.3. *Identification*

To assess the identification of the transmission parameters in this partial-insurance specification, consider the simplest model with no measurement error, serially uncorrelated transitory component and stationarity. The model can be identified with four years of data ($t+1$, t, $t-1$, $t-2$). The central parameters to identify are as follows: $\phi, \psi, \sigma_\xi^2, \sigma_\zeta^2$ and σ_ε^2. Standard results imply $\mathrm{E}[\Delta y_t(\Delta y_{t-1} + \Delta y_t + \Delta y_{t+1})] = \sigma_\zeta^2$ and also that $\mathrm{E}(\Delta y_t \Delta y_{t-1}) = \mathrm{E}(\Delta y_{t+1} \Delta y_t) = -\sigma_\varepsilon^2$. Identification of σ_ε^2 rests on the idea that income growth rates are autocorrelated due to mean reversion caused by the transitory component. Identification of σ_ζ^2 rests on the idea that the variance of income growth $[\mathrm{E}(\Delta y_t\, \Delta y_t)]$, less the contribution of the mean reverting component $[\mathrm{E}(\Delta y_t \Delta y_{t-1}) + \mathrm{E}(\Delta y_t\, \Delta y_{t+1})]$, coincides with the permanent innovations. In general, if one has T years of data, only $T-3$ variances of the permanent shock can be identified, and only $T-2$ variances of the i.i.d. transitory shock can be identified.

Blundell *et al.* (2008) show that identification of the transmission parameter for transitory shocks (ψ) uses the fact that income and lagged consumption may be correlated through the transitory component $[\mathrm{E}(\Delta c_t \Delta y_{t+1}) = \psi\sigma_\varepsilon^2]$. Scaling this by $\mathrm{E}(\Delta y_t \Delta y_{t+1}) = \sigma_\varepsilon^2$ identifies ψ. There is also a simple IV interpretation: ψ is identified by a regression of Δc_t on Δy_t using Δy_{t+1} as an instrument. A similar reasoning applies to the permanent shock transmission term ϕ, where the current covariance between consumption and income growth $[\mathrm{E}(\Delta c_t \Delta y_t)]$, stripped of the contribution of the transitory component, reflects the arrival of permanent income shocks

$$\mathrm{E}[\Delta c_t(\Delta y_{t-1} + \Delta y_t + \Delta y_t)] = \phi\sigma_\zeta^2.$$

Scaling this by the variance of permanent income shock, identified by using income moments alone, identifies the loading factor ϕ. Note again the IV interpretation: ϕ is identified by a regression of Δc_t on Δy_t using $(\Delta y_{t-1} + \Delta y_t + \Delta y_{t+1})$ as an instrument.

The variance of the component σ_ξ^2 is identified using a residual variability idea: the variance of consumption growth, stripped of the contribution of permanent and

[12] Below we show that under additional assumptions these can also be turned into identifying moments for repeated cross-section data.

transitory income shocks, reflects heterogeneity in the consumption gradient. Blundell *et al.* (2008) also show identification in the presence of measurement error. Measurement error in consumption induces serial correlation in consumption growth. Because consumption is a martingale with drift in the absence of measurement error, the variance of measurement error can still be recovered and the other parameters of interest remain identified. One obvious reason for the presence of measurement error in consumption occurs if an imputation procedure is used to measure consumption in the income panel. Note that we would expect such measurement error to be non-stationary. For measurement error in income, we can show that ϕ and $\sigma^2_{u^c}$ remain identified.[13]

3.4. *Non-stationarity*

We have already seen that allowing for non-stationarity is likely to be empirically important. Suppose we have T years of data, then

$$\mathrm{E}[\Delta y_s(\Delta y_{s-1} + \Delta y_s + \Delta y_{s+1})] = \sigma^2_{\zeta,s}$$

for $s = 3, 4, \ldots, T-1$. The variance of the transitory shock can be identified using:

$$-\mathrm{E}(\Delta y_s \Delta y_{s+1}) = \sigma^2_{\varepsilon,s}$$

for $s = 2, 3, \ldots, T-1$. In the case of time-varying partial-insurance parameters, identification is achieved through moment conditions

$$\frac{\mathrm{E}(\Delta c_s \Delta y_{s+1})}{\mathrm{E}(\Delta y_s \Delta y_{s+1})} = \psi_s$$

and

$$\frac{\mathrm{E}[\Delta c_s(\Delta y_{s-1} + \Delta y_s + \Delta y_{s+1})]}{\mathrm{E}[\Delta y_s(\Delta y_{s-1} + \Delta y_s + \Delta y_{s+1})]} = \phi_s$$

for all $s = 2, 3, \ldots, T-1$ and $s = 3, 4, \ldots, T-2$ respectively.

3.5. *What can we Learn from Repeated Cross Sections?*

Historically many countries have collected income and expenditure surveys that have no (or very little) longitudinal structure.[14] This is true for the Family Expenditure Survey (FES) in the UK, the Consumers Expenditure Survey (CEX) in the US, Family Expenditure Survey (FAMEX) in Canada, Budget de Famille in France, Household Expenditure Survey (HES) in Australia and the National Survey of Family Income and Expenditure (NSFIE) in Japan. Even without panel data, is it still possible to learn about the importance of permanent shocks?

With repeated cross-section measurements of income alone we cannot distinguish permanent from transitory income shocks, let alone identify the evolution of those

[13] Blundell *et al.* (2008) note that for the PSID, a back-of-the-envelope calculation shows that the variance of measurement error in earnings accounts for approximately 30% of the variance of the overall transitory component of earnings.

[14] The primary motivation for the data collection in these cases stemmed from the requirement to construct consumer price indices.

variances and the insurance parameters. However, adding measurements on consumption, and strengthening our assumptions somewhat, Blundell and Preston (1998) show that we can identify both. Indeed, with measurements of consumption and income in the same survey, we can identify the insurance parameter on permanent shocks as well as the evolution of permanent and transitory income shock variances.

To see how this insight works, assume the cross-sectional covariances of the shocks with previous periods' incomes are zero and the annuitisation value of transitory shocks is negligible. The variance–covariance structure now (18) has the form

$$\begin{aligned} \Delta\text{Var}(y_t) &= \text{Var}(\zeta_t) + \Delta\text{Var}(\varepsilon_t) \\ \Delta\text{Var}(c_t) &= \phi_t^2\text{Var}(\zeta_t) + \mathcal{O}(\text{E}_{t-1}\|v_{it}\|^3) \\ \Delta\text{Cov}(c_t, y_t) &= \phi_t\text{Var}(\zeta_t) + \mathcal{O}(\text{E}_{t-1}\|v_{it}\|^3). \end{aligned} \tag{23}$$

These moments can be used to identify $\text{Var}(\zeta_t)$, $\Delta\text{Var}(\varepsilon_t)$ and ϕ_t. This is the approach taken in Blundell *et al.* (2013) for the repeated cross sections of income and consumer surveys in the UK. They report a transmission parameter for permanent shocks ϕ_t that is stable over this period and takes on values around 0.8. They also use this approach to show for birth cohorts, who are in the labour market in the UK during the 1980s, that there is a strong spike in the variance of the permanent shocks at the depth of the recession.

3.6. *Partial-insurance Parameters for the US*

The PSID was the first household panel data study to systematically collect longitudinal data on family incomes, extensive demographic variables and some items of consumption at frequent intervals. It became a key data source for examining various measures of inequality up until the early 1990s. Taking on a central place in the studying of income especially following the huge run-up in income inequality in the US during the 1980s, documented the ground-breaking study of Moffitt and Gottschalk (1994), for example. In the late 1990s, the PSID was extensively revamped, moving from an annual to a biennial panel and increasing its collection of asset and consumption data. We return to the exploration of this remarkable new data below. But for the study of the 1980s inequality boom, the original panel survey provides compelling evidence.

The results summarised here refer to the PSID sample used in Blundell *et al.* (2008) covering the period 1978–1992 and consisting of all the possible subpanels of $5 \leq \text{length} \leq 15$ years for households with a male head aged 30–60. For the consumption data, a series from the CEX was drawn using the 1980–1992 surveys. The CEX collects information on a variety of socio-demographic and economic variables that are also collected in the PSID. Expenditure in the CEX is reported in each quarter and refers to the previous quarter; income is reported in the second and fifth interview (with some exceptions), and refers to the previous 12 months. For consistency with the timing of consumption, fifth-quarter income data are used.

Although the CEX is a quarterly panel, it consists of four quarters only and the income data are only freshly surveyed in the first and fourth quarters. Rather than using this data directly, a structural inverse demand equation with time-varying parameters is used to link the two data sources. The PSID contains a measure of total food expenditures. To line up the measures as best as is possible, five-quarter respondents only (annual expenditure measures) from the CEX are utilised. Otherwise, the sample selection is similar to that for the PSID. Further detail of this approach and a comparison of both data sources are in Blundell *et al.* (2004), which builds on the earlier work of Skinner (1987). Table 3 presents the implied autocovariance structure between consumption and income growth.

Table 4 provides the estimates of the partial-insurance parameters ϕ and ψ from (19) for the baseline specification. It also shows results for specifications which allow the

Table 3

The Autocovariance Matrix of Consumption Growth in the US

Year	var(Δc_t)	cov(Δc_{t+1},Δc_t)	cov(Δc_{t+2},Δc_t)
1980	0.1275	−0.0526	0.0022
	(0.0097)	(0.0076)	(0.0056)
1981	0.1197	−0.0573	0.0025
	(0.0116)	(0.0084)	(0.0043)
1982	0.1322	−0.0641	0.0006
	(0.0110)	(0.0087)	(0.0060)
1983	0.1532	−0.0691	−0.0056
	(0.0159)	(0.0100)	(0.0067)
1984	0.1869	−0.1003	−0.0131
	(0.0173)	(0.0163)	(0.0089)
1985	0.2019	−0.0872	NA
	(0.0244)	(0.0194)	
1986	0.1628	NA	NA
	(0.0184)		
1987	NA	NA	NA
1988	NA	NA	NA
1989	NA	NA	NA
1990	0.1751	−0.0602	−0.0057
	(0.0221)	(0.0062)	(0.0067)
1991	0.1646	−0.0696	NA
	(0.0142)	(0.0100)	
1992	0.1467	NA	NA
	(0.0130)		

Source. Blundell *et al.* (2008).

Table 4

Partial-insurance Parameter Estimates

Transmission parameters	Whole sample	No college	College	Born 1940s	Born 1930s
ϕ(Partial-insurance permanent shock)	0.6423	0.9439	0.4194	0.7928	0.6889
	(0.0945)	(0.1783)	(0.0924)	(0.1848)	(0.2393)
ψ(Partial-insurance transitory shock)	0.0533	0.0768	0.0675	0.0273	−0.0381
	(0.0435)	(0.0602)	(0.0550)	(0.0705)	(0.0737)
p-value test of equal ϕ	0.23	0.99	0.08	0.81	0.18
p-value test of equal ψ	0.75	0.33	0.29	0.76	0.04

Notes. Standard errors in parenthesis. This Table reports DWMD results of the parameters of interest. See Blundell *et al.* (2008) for results allowing for time-varying variances of measurement error in consumption.
Source. Blundell *et al.* (2008).

transmission parameters to differ by birth cohort and by education level. Overall, the results suggest around 65% of permanent income shocks are not insured; that is, they find their way directly into consumption. Virtually, all transitory shocks appear to be insured. As we will see below, using direct measures of household assets and also allowing for family labour supply responses provide clear new insights into these baseline partial-insurance results.

3.7. *The Importance of Measuring Assets*

The PSID over the 1978–92 period contained some broad measures of financial and housing wealth. Combining these and selecting households with (initial) wealth in the lowest 30% of the wealth distribution is found to imply a very different pattern of transmission parameters. The second column of results in Table 5 shows the impact on this selection by initial wealth. Comparing with the baseline specification reproduced in column 1, these results show a larger point estimate for the transmission of permanent shocks and, more importantly perhaps, a strongly significant transmission parameter for transitory shocks. Unsurprisingly, with limited access to financial or housing wealth, even transitory shocks to income impact on consumption. Column 3 shows this is not the case for the higher wealth group.

If borrowing at, or close to, the risk-free rate is difficult, families can experience welfare loss even for short-run falls in income reflected in the consequent reduction in consumption. However, as Browning and Crossley (2009) so elegantly note, the service flow from durable consumption can be maintained by running down durables and holding back on replacement or maintenance. Any fall in non-durable consumption can thereby be attenuated, at least for a short while. The upshot of this is that we would expect to find even greater sensitivity to transitory income shocks among low-wealth households for consumption measures that include durable purchases.

To assess the importance of this mechanism for low-wealth families, we can examine the same selection of low-wealth households but now include durable expenditures in our consumption measure. In the final column in Table 5, the transmission parameter for transitory shocks is now even larger than column 2 and the permanent shock parameter has a point estimate of unity. Once durable expenditures are included, consumption growth is even more sensitive to transitory

Table 5

Wealth and Durables

Consumption: Income: Sample:	Non-durable Net income Baseline	Non-durable Net income Low wealth	Non-durable Net income High wealth	Total Net income Low wealth
ϕ (Partial-insurance permanent shock)	0.6423	0.8489	0.6248	1.0342
	(0.0945)	(0.2848)	(0.0999)	(0.3517)
ψ (Partial-insurance transitory shock)	0.0533	0.2877	0.0106	0.3683
	(0.0435)	(0.1143)	(0.0414)	(0.1465)

Note. See Table 4.

shocks for low-wealth families. As noted in the discussion of the partial-insurance model (19), the transmission parameters subsume self-insurance and do not allow us to separate the various insurance mechanisms. If we could accurately measure π for each family in the survey, this would allow this distinction. The enhanced asset data in the recent PSID data allow us to investigate these points in more detail in Section 4 below.

3.8. *Excess Insurance?*

The 65% insurance of permanent income shocks result is a somewhat puzzling and controversial result. It implies that around 35% of permanent shocks do not find their way into consumption. This is substantially in excess of what would be reasonably accommodated by the rough measures of savings for families in the PSID over this period, suggesting additional insurance mechanisms over and above simple self-insurance.

We have already seen that this 'average' result is very sensitive to the inclusion of low-wealth households even when using the rough measures of wealth in the PSID surveys from the early 1980s. We might also want to compare across education groups. For example, in the college/no-college comparison, the estimates from Blundell *et al.* (2008) suggest that low-education groups failed to insure almost any permanent income shocks. As also might be expected, in a comparison of younger *versus* older birth cohorts, the older cohort was able to insure more.

The 'excess insurance' result also provided a strong motivation for the Blundell *et al.* (2012) to which we turn to in the next section. This study uses the recently enhanced PSID data on assets, consumption and labour supply, to dig further into possible 'insurance' mechanisms.

3.9. *Inequality During the Recession*

The Blundell *et al.* (2008) results also suggest a sharp rise in the variance of permanent shocks in the early 1980s, coincident with the recession in the US. The estimated variance of transitory shocks continued to rise through the middle of the decade but by that time the importance of permanent shocks declined. There is little evidence of changes in the ϕ and ψ transmission parameters over time, suggesting that the degree of insurance was fairly stable over this period.[15]

These results on inequality during the recession coincide with those for the UK over the 1980s. For this period there is no household panel data available in the UK, so Blundell *et al.* (2013) follow the approach for the repeated cross sections of income and consumer surveys described earlier. As noted above, they report a transmission parameter for permanent shocks ϕ that is stable over this period and takes on values around 0.8. For birth cohorts who are in their early and prime-age working lives during the 1980s, they document a strong spike in the variance of the permanent shocks at the depth of the recession.

[15] They do find a strong rejection of constancy when food in PSID is used, driven by food becoming more of a necessity over this period.

4. Additional 'Insurance' Mechanisms

Overall we can think of four types of insurance we might wish to consider: Self-insurance through the credit market; social insurance through taxes and welfare benefits; family labour supply through the labour supply responses of other family members and informal contracts, gifts, etc. The transmission parameter approach taken in the partial-insurance model specifications above combines all these mechanisms into single transmission parameters. It does not therefore allow the separation of the different mechanisms involved.

4.1. *Taxes and Transfers*

Understanding the degree to which taxes and transfer programmes act as insurance for permanent and transitory income shocks among the poor is important in assessing their effectiveness and their design. There exists an important empirical literature on redistributive mechanisms provided by social insurance, transfers, progressive taxation which include Gruber and Yelowitz (1999), Gruber (2000), Kniesner and Ziliak (2002) and Blundell and Pistaferri (2003). There are also studies on informal transfers within extended families and across interpersonal networks; for example, Kotlikoff and Spivak (1981) and Attanasio and Rios-Rull (2000).

A simple way to assess the 'insurance' value of the tax and transfer system in the context of the partial-insurance approach is to examine the impact on the insurance parameters of changing the income definition to be gross of taxes and transfers. A reduction in the transmission parameters would indicate the degree of additional insurance. The second column of Table 6 shows the results of such an experiment using the partial-insurance modelling framework and PSID–CEX data source analysis above. The reduction in the estimated transmission parameter for permanent shocks ϕ from 0.64 to 0.37 indicates the important role of taxes and transfers in insuring family incomes. The final column points to the importance of family labour supply to which we now turn.

4.2. *Family Labour Supply*

Perhaps one of the most interesting avenues of 'insurance' within families is the use of family labour supply. Family labour supply acts as a natural mechanism for smoothing income to permanent wage across different family members, see Stephens (2009),

Table 6

Taxation and Other Earnings

Consumption: Income:	Non-durable Net income	Non-durable Earnings only	Non-durable Male earnings
ϕ (Partial-insurance permanent shock)	0.6423 (0.0945)	0.3700 (0.0574)	0.2245 (0.0493)
ψ (Partial-insurance transitory shock)	0.0533 (0.0435)	0.0633 (0.0309)	0.0502 (0.0294)

Note. See Table 4.

Attanasio *et al.* (2005, 2008), Heathcote *et al.* (2007), Lise and Seitz (2011). Moreover, family labour supply may be used as an insurance to transitory shocks in the absence of access to credit. It is inefficient for one family member to use their labour supply to insure transitory shocks to another family members wages when the family can access credit. But this is not so when there are borrowing constraints or frictions in the credit market.

To introduce family labour supply behaviour consider total income Y_t as the sum of two sources, Y_{1t} and $Y_{2t} \equiv W_{2t}h_{2t}$. Assume that the labour supplied by the primary earner to be fixed, so that it is the secondary worker who has flexibility in labour supply h_2 in response to shocks in wages W_2 and the primary worker's income Y_1.

To keep things manageable, assume that the primary income Y_1 and the hourly wage rate of the secondary worker W_{2t} follow the simple permanent–transitory structure

$$\Delta \ln Y_{1t} = \gamma_{1t} + \Delta u_{1t} + v_{1t},$$

and

$$\Delta \ln W_{2t} = \gamma_{2t} + \Delta u_{2t} + v_{2t}.$$

respectively. Household decisions are taken to maximise a household utility function

$$\sum_k (1+\delta)^{-k}[U(C_{t+k}) - V(h_{2,t+k})].$$

Using the same approach as in the pure consumption model above the consumption growth and intertemporal labour supply equations

$$\begin{aligned} \Delta \ln C_{t+k} &\simeq \sigma_{t+k} \Delta \ln \lambda_{t+k} \\ \Delta \ln h_{2,t+k} &\simeq -\rho_{t+k}(\Delta \ln \lambda_{t+k} + \Delta \ln W_{2,t+k}), \end{aligned}$$

where λ is the marginal utility of wealth and where $\sigma_t \equiv U_t'/C_t U_t'' < 0$ and $\rho_t \equiv -V_t'/h_{2t}V_t'' > 0$. In this formulation, the interest rate and discount rates have been set equal to focus on the wage and income effects.[16]

These stochastic growth relationships for $\ln C$, $\ln Y_1$, $\ln Y_2$ and $\ln W_2$ imply a specific variance–covariance structure. This structure generates the key panel data moments:

$$\begin{aligned} \mathrm{Var}(\Delta c_t) &\simeq \beta^2\sigma^2 s^2 \mathrm{Var}(v_{1t}) + \beta^2\sigma^2(1-\rho)^2(1-s)^2\mathrm{Var}(v_{2t}) \\ &\quad + 2\beta^2\sigma^2(1-\rho)s(1-s)\mathrm{Cov}(v_{1t}, v_{2t}) \\ \mathrm{Var}(\Delta y_{1t}) &\simeq \mathrm{Var}(v_{1t}) + \Delta \mathrm{Var}(u_{1t}) \\ \mathrm{Var}(\Delta y_{2t}) &\simeq (1-\psi)^2\mathrm{Var}(u_{2t}) - \beta^2\rho^2 s^2\mathrm{Var}(v_{1t}) \\ &\quad + \beta^2\sigma^2(1-\rho)^2\mathrm{Var}(v_{2t}) - 2\beta^2\sigma(1-\rho)s\mathrm{Cov}(v_{1t}, v_{2t}) \\ \mathrm{Var}(\Delta w_{2t}) &\simeq \mathrm{Var}(v_{2t}) + \Delta \mathrm{Var}(u_{2t}), \end{aligned}$$

where $\beta = 1/[\sigma + \rho(1-s)]$ and s_t is the ratio of the mean value of the primary earner's earnings to that of the household $\bar{Y}_{1t}/\bar{Y}_t$. These moments are sufficient to identify permanent and transitory shock distribution for ln Y_1 and W_2, and their the

[16] See Attanasio *et al.* (2002).

evolution over time. Note that when there is a positive labour supply elasticity $-\rho > 0$, then the secondary worker provides insurance for permanent shocks to Y_1.

The last two columns of Table 6 show the impact of moving from a total family earnings measure to a male (Head) earnings measure for their sample of PSID couples. The fall in the transmission parameter for permanent shocks ϕ from 0.37 to 0.225 shows the importance of family labour supply in insuring permanent shocks to male earnings in this period.

In the pure self-insurance framework where consumption smoothing is provided through borrowing and saving at the risk-free rate, although the secondary worker's labour supply responds positively to her own transitory wage shocks, she (or he) does not respond to transitory shocks to primary worker income. It is always preferable for the family to use the credit market to smooth such transitory shocks. However, when credit at the risk-free rate is not available, transitory shocks to primary incomes will be insured by secondary worker labour supply.

To investigate the family labour supply story further, Blundell *et al.* (2012) use the enhanced data from the post-1996 PSID to estimate a model of consumption inequality and family labour supply for couples. The new asset data allow a *direct* measure of π_{it} and the more comprehensive consumption data avoid the need for imputation. Their analysis extends previous work and expresses the distributional dynamics of consumption and earnings growth as functions of Frisch elasticities, 'insurance parameters' and wage shocks.

For example, the impact of a permanent shock to male wages w_m is shown to generalise the transmission parameter $\pi_{i,t}$ in expression (18) to take the form:

$$\pi_{i,t} s_{i,m,t} \frac{\eta_{c,p}\left(1+\eta_{h_m,w_m}\right)}{\eta_{c,p}+\left(1-\pi_{i,t}\right)\bar{\eta}_{h,w}}, \tag{24}$$

where $s_{i,m,t}$ is the share of the male earnings in future human capital wealth, and the $\eta_{c,p}$ and $\eta_{h,w}$ parameters are the Frisch consumption and hours of work elasticities respectively.[17] Consequently, this new transmission parameter (24) captures the impact of a permanent wage shock on consumption allowing for family labour supply responses, savings and risk aversion.

The expression (24) assumes additive separability between consumption and labour supply, something strongly rejected in the Blundell *et al.* (2012) study. Once this assumption is relaxed, the generalised partial-insurance specification is shown to fit the dynamics of wages, earnings and consumption well and implies some remarkably clear predictions for the different insurance mechanisms over the life-cycle emerge.

For example, one key contribution of this approach is summarised in Figure 6 which describes the average estimated response of consumption to a 10% permanent decrease in the male's wage rate at different ages (of the household head). As the male share of earnings is around 70%, the impact without accounting for family labour supply or self-insurance is around −7%. This impact on consumption is reduced considerably once family labour supply and self-insurance are accounted for.

[17] Note that the definition of π used in Blundell *et al.* (2012) is the ratio of assets to assets plus human capital wealth. Consequently, the measure is $1-$ the measure of π here.

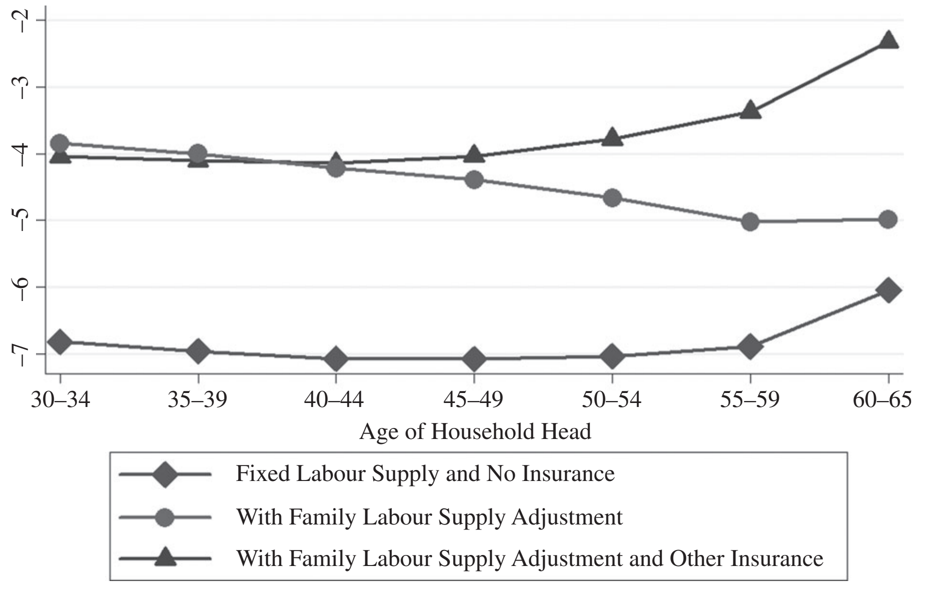

Fig. 6. *The Average Impact on Consumption of a Permanent Shock to Male Wage, US*
Source. Blundell *et al.* (2012).

At younger working ages, say, those in their 30s, all the impact is through family labour supply responses. Few households at this point in their life cycle have built up enough savings to be of use in insuring permanent labour income shocks. But the labour supply of their family members can respond and does so to cut the average impact on consumption to −4% on average. It is not until older working ages that assets play a role. Figure 6 also shows that at older working ages, say in the late 50s, the average impact is further reduced to below 3%. The adjustment is shared evenly between the use of savings and family labour supply.

Of course, and as documented in the article, there is much heterogeneity in the response to permanent shocks across household types according to their assets, education, demographics, labour market attachment etc. This analysis suggests that family labour supply can be a key mechanism for insuring idiosyncratic labour market shocks for couples, especially for those with limited access to assets.

5. Robustness Issues

There are many issues surrounding the robustness of the results presented above. There are information and anticipation issues. Among these are the consumer's ability to distinguish between permanent and transitory shocks and the degree of anticipation or advance information of income 'shocks'. Then, there are issues around the robustness to assumptions about the nature of the economy; in particular, to assumptions about the non-stationarity of the income distribution and also to assumptions about the credit market and borrowing constraints. I briefly consider each of these in turn.

5.1. *Distinguishing Permanent from Transitory Shocks*

Suppose we return to the partial-insurance model which ignores labour supply (19) and suppose that the consumer cannot separately identify transitory ε_{it} from permanent ζ_{it} income shocks. For a consumer who simply observed the income innovation ε_{it} in $y_{it} = y_{i,t-1} + \varepsilon_{it} - \theta_t \varepsilon_{i,t-1}$, we have consumption innovation:

$$\eta_{it} = \rho_t(1-\theta_{t+1})\varepsilon_{it} + \frac{r}{1+r}\theta_{t+1}\varepsilon_{it}, \tag{25}$$

where $\rho_t = 1 - (1+r)^{-(R-t+1)}$. The evolution of θ_t is directly related to the evolution of the variances of the transitory and permanent innovations to income; see Blundell and Preston (1998).

The permanent effects component in this decomposition can be thought of as capturing news about both current and past permanent effects as

$$\mathrm{E}\left(\sum_{j=0} \zeta_{i,t-j} | \varepsilon_{it}, \varepsilon_{i,t-1}, \ldots\right) - \mathrm{E}\left(\sum_{j=0} \zeta_{i,t-j} | \varepsilon_{i,t-1}, \ldots\right) = (1-\theta_{t+1})\varepsilon_{it}.$$

This represents the best prediction of the permanent/transitory split.

Suppose instead that the researcher decides to ignore the split between the two shocks. The partial-insurance coefficient is now a weighted average of the coefficients of partial insurance ϕ and ψ, with weights given by the importance of the variance of permanent (transitory) shocks. In a period where the importance of permanent shocks is declining such as the mid to late 1980s, one would have the impression that insurance is growing. Instead, it is the relative importance of more insurable shocks that is growing.

5.2. *Anticipation and Information*

The overlap between the econometrician's information set and that of the consumer is likely to be far from exact, see Cunha *et al.* (2005), for example. The 'shocks' to income may well already be known to the consumer. In this case, the transmission parameter will subsume the advance information available to the consumer. What appears to be insurance may just reflect advance information.

One way to assess the extent of advance information is to examine the covariances between current consumption growth and future income growth. If information arrives and consumers act on this, it should be immediately included in consumption. For the 1978–92 US panel data in Blundell *et al.* (2008), we find little evidence of significant anticipation. The p-value for the test $\mathrm{cov}(\Delta y_{t+1}, \Delta c_t) = 0$ for all time periods is 0.3305, and for the next three higher order covariances, the p-values are even larger, 0.6058, 0.8247 and 0.7752 respectively. This 'suggests' that the shocks that were experienced in the US in the 1980s for this PSID sample of prime-age headed families can be interpreted as largely unanticipated. In the US and in the UK, it seems sensible to conclude that the key dominant changes in earnings inequality over this period reflected three changes – shifts in the returns to skills, shifts in government transfers and the shift of insurance from firms to workers; see Machin and Van Reenen (2008). It is difficult to argue that these were easily anticipated by consumers.

5.3. *Alternative 'Economies'*

The simulation studies reported in Kaplan and Violante (2010) and Blundell *et al.* (2013) examine various aspects of the economic assumptions underlying the partial-insurance approach. Both studies simulate a baseline life-cycle version of the incomplete markets model in which households have access to a single risk-free bond and have time-separable expected CRRA utility; see Huggett (1993).

Blundell *et al.* (2013) focus on a non-stationary economy where the permanent variance follows a two-state, first-order Markov process with the transition probability between alternative variances. The aim is to see how well the evolution of the variances and covariances, especially in the case of repeated cross-section data, identifies the changes in the permanent and transitory variances in this non-stationarity Markov switching model. They show that under a wide range of assumptions on income processes, discount rates and substitution elasticities, reasonably accurate estimates of the variances and their changes over time can be inferred from the evolution of consumption and income inequality.

The important Kaplan and Violante (2010) study focuses more on the nature of the insurance available ranging between an incomplete markets framework with borrowing constraints, the Permanent Income model and an autarky framework, all in a stationary setting. They also consider advance information, where a proportion of the shocks are known in advance to the consumer and also where returns from human capital are known and correlated with initial conditions. They find that the transmission parameter approach does well in most of the situations considered, although the interpretation of the transmission parameters as measures of insurance depends on assumptions about information and the persistence of shocks. Perhaps most interestingly, their study highlights the importance of specifying the income dynamics correctly when assessing whether the standard self-insurance model accords well with the data. As noted above, excess insurance can simply reflect a lower persistence in the shocks to income. Guvenen (2007) investigates this further in a useful and creative generalisation of the consumption growth model with persistent income dynamics, as in (5), but where individuals learn about their idiosyncratic trends.

6. Summary and Future Directions

The research reported in this study seeks to understand the transmission of inequality over the working life. Taking labour market shocks as the primary source of uncertainty, the aim is to examine the linkages between the distribution of wages, earnings, joint labour supply, savings and consumption. A large variety of 'insurance' mechanisms that link these various measures are considered, including credit markets, labour supply, taxation, welfare benefits, formal insurance, informal gifts and transfers. The key idea is to use the framework of constrained intertemporal choice to provide a structure for the distributional dynamics of inequality over the working life.

The dynamics of labour income and wages are a central part of this analysis and much of the discussion in the article is on understanding the nature of labour market income dynamics. These dynamics have been shown to be the key to exploring the mechanisms used by families to 'insure' against labour market shocks.

The definition of insurance that we have adopted is very broad. It covers formal and informal mechanisms that are used to attenuate the impact of shocks to income. These mechanisms have been found to vary in importance across different types of households at different points of their life cycle and at different points in time: The manner and scope for insurance depending on the access to credit, the durability of income shocks and the ability of other family members to adjust their labour supply. The workhorse framework used has been the standard incomplete markets self-insurance model generalised to a partial-insurance framework developed in Blundell *et al.* (2008), in which transmission parameters are used to indicate the degree of insurance to labour market shocks.

Gathering up the results, I have uncovered clear evidence for non-stationarity over the life cycle and over time. Variances (of persistent shocks) tend to be larger at the beginning and at the end of the working life. Recessions also often result in spikes in the variances of shocks. For most families, the partial-insurance framework points to a large pass through of longer term 'persistent' shocks to consumption, this is the key link between the income and consumption distributions. For low-wealth households, however, shorter term 'transitory' shocks provide an important additional linkage. Recent work in this area has also found that family labour supply is an important mechanism for smoothing consumption, especially for those with limited net assets. Indeed, Blundell *et al.* (2012) argue that once family labour supply, assets and taxes (and benefits) are properly accounted for, there is little evidence for additional insurance. Together, these results resolve many of the puzzles and controversies in the empirical literature as to exactly how best to develop the microfoundations of the simple permanent income model that allows for a more realistic setting for household behaviour.

What of future research?: There remains much to be done to dig deeper into the underlying stochastic processes and the different insurance mechanisms. In particular, exploring the role of idiosyncratic trends, non-stationarity, higher order moments and nonlinearities in the income process. Relaxing the inherent linearity assumptions, underlying most dynamic panel data models of labour income will be a key development in this research. These have the potential to impact strongly on consumption decisions. Heterogeneous trends can lead to an exaggeration of the importance of persistent shocks; see Guvenen and Smith (2012).[18] Higher order moments and non-linearities in persistent shocks increasing in importance the larger the degree of risk aversion.

There is also the need to understand further the mechanism and market incentives for insurance, here the key articles to date have been Krueger and Perri (2006) and Attanasio and Pavoni (2011). Related to this is the importance of advance information and learning along the life cycle. The study by Cunha *et al.* (2005) is central in this respect.

Finally, there is much more to be done on the interaction between credit, durable expenditures and family labour supply. How does the irreversibility of durable purchase decisions interact with uncertainty? Which goods are complementary and

[18] Even so, the evolution of variance of consumption can still be used to identify the variance of permanent shocks; see Blundell *et al.* (2013).

which are substitutes for labour supply? How important is the degree of liquidity of housing and financial wealth in insuring against shocks? What happens when shocks in the labour market and asset markets are correlated, much as was the case in the recent financial recession? The models developed here suggest a magnified transmission effect as lower asset values provide less ability for self-insurance and correlated labour market shocks providing less opportunity for insurance through family labour supply. One thing is for sure, the results of the research presented here provide a strong motivation for collecting consumption data, along with asset and earnings data, in new longitudinal household surveys and linked administrative register data.

University College London and Institute for Fiscal Studies

References

Abowd, J. and Card, D. (1989). 'On the covariance structure of earnings and hours changes', *Econometrica*, vol. 57(2), pp. 411–45.

Altonji, J., Smith, A. and Vidangos, I. (2013). 'Modeling earnings dynamics', *Econometrica*, vol. 81(4), pp. 1395–454.

Atkinson, A.B. (1997). 'Bringing income distribution in from the cold', ECONOMIC JOURNAL, vol. 107(441), pp. 297–321.

Atkinson, A.B. (1999). 'The distribution of income in the UK and OECD countries in the twentieth century', *Oxford Review of Economic Policy*, vol. 15(4), pp. 56–75.

Attanasio, O., Berloffa, G., Blundell, R. and Preston, I. (2002). 'From earnings inequality to consumption inequality', ECONOMIC JOURNAL, vol. 112(478), pp. C52–9.

Attanasio, O. and Davis, S. (1996). 'Relative wage movements and the distribution of consumption', *Journal of Political Economy*, vol. 104(6), pp. 1227–62.

Attanasio, O., Low, H. and Sanchez-Marcos, V. (2005). 'Female labor supply as an insurance against idiosyncratic risk', *Journal of the European Economic Association*, vol. 3(2–3), pp. 755–64.

Attanasio, O., Low, H. and Sanchez-Marcos, V. (2008). 'Explaining changes in female labor supply in a life-cycle model,' *American Economic Review*, vol. 98(4), pp. 1517–52.

Attanasio, O. and Pavoni, N. (2011). 'Risk sharing in private information models with asset accumulation: explaining the excess smoothness of consumption', *Econometrica*, vol. 79(4), pp. 1027–68.

Attanasio, O. and Rios-Rull, V. (2000). 'Consumption smoothing in island economies: can public insurance reduce welfare?', *European Economic Review*, vol. 44(7), pp. 1225–58.

Baker, M. (1997). 'Growth-rate heterogeneity and the covariance structure of life-cycle earnings', *Journal of Labor Economics*, vol. 15(2), pp. 338–75.

Baker, M. and Solon, G. (2003). 'Earnings dynamics and inequality among Canadian men, 1976–1992: evidence from longitudinal income tax records', *Journal of Labor Economics*, vol. 21(2), pp. 267–88.

Blundell, R. and Etheridge, B. (2010). 'Consumption, income and earnings inequality in Britain', *Review of Economic Dynamics*, vol. 13(1), pp. 76–102.

Blundell, R., Graber, M. and Mogstad, M. (2014), 'Labor income dynamics and the insurance from taxes, transfers, and the family', IFS Working Paper 01/14, January.

Blundell, R., Low, H. and Preston, I. (2013). 'Decomposing changes in income risk using consumption data', *Quantitative Economics*, vol. 4(1), pp. 1–37.

Blundell, R. and Pistaferri, L. (2003). 'Income volatility and household consumption: the impact of food assistance programs', *Journal of Human Resources*, vol. 38(Supplement), pp. 1032–50.

Blundell, R., Pistaferri, L. and Preston, I. (2004). 'Imputing consumption in the PSID using food demand estimates from the CEX', Working Paper, Institute for Fiscal Studies, WP04/27.

Blundell, R., Pistaferri, L. and Preston, I. (2008). 'Consumption inequality and partial insurance', *American Economic Review*, vol. 98(5), pp. 1887–921.

Blundell, R., Pistaferri, L. and Saporta-Eksten, I. (2012). 'Consumption inequality and family labor supply', NBER Discussion Paper, November.

Blundell, R. and Preston, I. (1996). 'Income, expenditure and the living standards of UK households', *Fiscal Studies*, vol. 16(3), pp. 40–54.

Blundell, R. and Preston, I. (1998). 'Consumption inequality and income uncertainty', *Quarterly Journal of Economics*, vol. 113(2), pp. 603–40.

Blundell, R. and Stoker, T. (1999). 'Consumption and the timing of income risk', *European Economic Review*, vol. 43(3), pp. 475–507.

Bonhomme, S. and Robin, J.M. (2010). 'Generalized non-parametric deconvolution with an application to earnings dynamics', *Review of Economic Studies*, vol. 77(2), pp. 491–533.

Brewer, M. and O'Dea, C. (2012). 'Measuring living standards with income and consumption: evidence from the UK', IFS Working Papers W12/12, Institute for Fiscal Studies.

Browning, M. and Crossley, T.F. (2009).'Shocks, stocks and socks: smoothing consumption over a temporary income loss', *Journal of the European Economic Association*, vol. 7(6), pp. 1169–92.

Browning, M., Ejrnes, M. and Alvarez, J. (2010). 'Modelling income processes with lots of heterogeneity', *Review of Economic Studies*, vol. 77(4), pp. 1353–81.

Campbell, J.Y. (1993). 'Intertemporal asset pricing without consumption data', *American Economic Review*, vol. 83, pp. 487–512.

Campbell, J. and Deaton, A. (1989). 'Why is consumption so smooth?', *Review of Economic Studies*, vol. 56(6), pp. 357–73.

Carroll, C. (2001). 'Precautionary saving and the marginal propensity to consume out of permanent income'. NBER Working Paper 8233.

Cochrane, J.H. (1991). 'A simple test of consumption insurance', *Journal of Political Economy*, vol. 99(5), pp. 957–76.

Cunha, F., Heckman, J. and Navarro, S. (2005). 'Separating uncertainty from heterogeneity in life cycle earnings', *Oxford Economic Papers*, vol. 57(2), pp. 191–261.

Cutler, D. and Katz, L. (1992). 'Rising inequality? Changes in the distribution of income and consumption in the 1980s'. *American Economic Review*, vol. 82(2), pp. 546–51.

Deaton, A. and Paxson, C. (1994). 'Intertemporal choice and inequality', *Journal of Political Economy*, vol. 102 (3), pp. 384–94.

DeBacker, J., Heim, B., Panousi, V., Ramnath, S. and Vidangos, I. (2013). 'Rising inequality: transitory or permanent? New evidence from a panel of tax returns', Brookings Spring Panel on Economic Activity, March.

Dynarski, S. and Gruber, J. (1997). 'Can families smooth variable earnings?' *Brooking Papers on Economic Activity*, vol. 28(1), pp. 229–305.

Gruber, J. (2000). 'Cash welfare as a consumption smoothing mechanism for single mothers', *Journal of Public Economics*, vol. 75(2), pp. 157–82.

Gruber, J. and Yelowitz, A. (1999). 'Public health insurance and private savings', *Journal of Political Economy*, vol. 107(6), pp. 1249–74.

Guvenen, F. (2007). 'Learning your earning: are labor income shocks really very persistent?', *American Economic Review*, vol. 97(3), pp. 687–712.

Guvenen, F. (2009). 'An empirical investigation of labor income processes', *Review of Economic Dynamics*, vol. 12(1), pp. 58–79.

Guvenen, F. and Smith, A. (2012). 'Inferring labor income risk from economic choices: an indirect inference approach', mimeo, Yale University.

Haider, S.J. (2001). 'Earnings instability and earnings inequality in the United States, 1967–1991', *Journal of Labor Economics*, vol. 19(4), pp. 799–836.

Haider, S.J. and Solon, G. (2006). 'Life-cycle variation in the association between current and lifetime earnings', *American Economic Review*, vol. 96(4), pp. 1308–20.

Hall, R. and Mishkin, F. (1982). 'The sensitivity of consumption to transitory income: estimates from panel data of households', *Econometrica*, vol. 50(2), pp. 261–81.

Hayashi, F., Altonji, A. and Kotlikoff, L. (1996). 'Risk sharing between and within families', *Econometrica*, vol. 64(2), pp. 261–94.

Heathcote, J., Storesletten, K. and Violante, G.L. (2007). 'Consumption and labour supply with partial insurance: an analytical framework,' CEPR Discussion Paper 6280.

Heathcote, J., Storesletten, K. and Violante, G. (2008). 'Insurance and opportunities: a welfare analysis of labor market risk', *Journal of Monetary Economics*, vol. 55(3), pp. 501–25.

Heathcote, J., Storesletten, K. and Violante, G.L. (2010). 'The macroeconomic implications of rising wage inequality in the United States', *Journal of Political Economy*, vol. 118(4), pp. 681–722.

Heckman, J. and Cunha, F. (2007). 'The Technology of Skill Formation', *American Economic Review*, vol. 97(2), pp. 31–47.

Heckman, J.J. and Scheinkman, J.A. (1987). 'The importance of bundling in a Gorman-Lancaster model of earnings', *Review of Economic Studies*, vol. 54(2), pp. 243–55.

Huggett, M. (1993), 'The risk free rate in heterogeneous-agent, incomplete-insurance economies', *Journal of Economic Dynamics and Control*, vol. 17(5–6), pp. 953–69.

Huggett, M., Ventura, G. and Yaron, A. (2011). 'Sources of lifetime inequality', *American Economic Review*, vol. 101(7), pp. 2923–54.

Johnson, D., Smeeding, T. and Torrey, B.B. (2005). 'United States inequality through the prisms of income and consumption', *Monthly Labor Review*, vol. 128(4), pp. 11–24.

Kaplan, G. and Violante, G. (2010). 'How much consumption insurance beyond self-insurance', *American Economic Journal*, vol. 2(4), pp. 53–87.

Kaufmann, K. and Pistaferri, L. (2009). 'Disentangling insurance and information in intertemporal consumption choices', *American Economic Review: Papers & Proceedings*, vol. 99(2), pp. 387–92.
Kniesner, T.J. and Ziliak, J.P. (2002). 'Tax reform and automatic stabilization', *American Economic Review*, vol. 92(3), pp. 590–612.
Kotlikoff, L. and Spivak, A. (1981). 'The family as an incomplete annuities market', *Journal of Political Economy*, vol. 89(2), pp. 372–91.
Krueger, D. and Perri, F. (2006). 'Does income inequality lead to consumption inequality? Evidence and theory', *Review of Economic Studies*, vol. 73(1), pp. 163–93.
Lillard, L. and Weiss, Y. (1979). 'Components of variation in panel earnings data: American Scientists, 1960–70', *Econometrica*, vol. 47(2), pp. 437–54.
Lillard, L.A. and Willis, R.J. (1978). 'Dynamic aspects of earning mobility', *Econometrica*, vol. 46(5), pp. 985–1012.
Lise, J. and Seitz, S. (2011). 'Consumption inequality and intra-household allocations', *Review of Economic Studies*, vol. 78(1), pp. 328–55.
Low, H. (2005). 'Self-insurance in a life-cycle model of labour supply and savings', *Review of Economic Dynamics*, vol. 8(4), pp. 945–75.
Low, H., Meghir, C. and Pistaferri, L. (2010). 'Wage risk and employment risk over the life cycle', *American Economic Review*, vol. 100(4), pp. 1432–67.
Machin, S. and Van Reenen, J. (2008). 'Changes in wage inequality', in (S.N. Durlauf and L.E. Blume, eds.), *The New Palgrave Dictionary of Economics*, 2nd ed.
MaCurdy, T. (1982). 'The use of time series processes to model the error structure of earnings in a longitudinal data analysis', *Journal of Econometrics*, vol. 18(1), pp. 82–114.
Meghir, C. and Pistaferri, L. (2004), 'Income variance dynamics and heterogeneity', *Econometrica*, vol. 72(1), pp. 1–32.
Meghir, C. and Pistaferri, L. (2011). 'Earnings, consumption and life cycle choices', *Handbook of Labor Economics*, Amsterdam: Elsevier.
Moffitt, R. and Gottschalk, P. (1994). 'Trends in the autocovariance structure of earnings in the US: 1969–1987,' Unpublished.
Moffitt, R. and Gottschalk, P. (1995). 'Trends in the autocovariance structure of earnings in the US: 1969–1987,' Brookings Papers.
Moffitt, R. and Gottschalk, P. (2002). 'Trends in the transitory variance of earnings in the United States', Economic Journal, vol. 112(468), pp. C68–73.
Pistaferri, L. (2003). 'Anticipated and unanticipated wage changes, wage risk, and intertemporal labor supply', *Journal of Labor Economics*, vol. 21(3), pp. 729–54.
Primiceri, G.E. and van Rens, T. (2009). 'Heterogeneous life-cycle profiles, income risk and consumption inequality', *Journal of Monetary Economics*, vol. 56(1), pp. 20–39.
Skinner, J. (1987). 'A superior measure of consumption from the panel study of income dynamics', *Economic Letters*, vol. 23(2), pp. 213–16.
Stephens, M. (2002). 'Worker displacement and the added worker effect', *Journal of Labor Economics*, vol. 20(3), pp. 504–37.
Storesletten, K., Telmer, C. and Yaron, A. (2004). 'Consumption and risk sharing over the life cycle', *Journal of Monetary Economics*, vol. 51(3), pp. 609–33.
Townsend, R.M. (1994), 'Risk and insurance in village India', *Econometrica*, vol. 62(3), pp. 539–91.