

IZA DP No. 6125

Decomposing Changes in Income Risk Using Consumption Data

Richard Blundell
Hamish Low
Ian Preston

November 2011

Decomposing Changes in Income Risk Using Consumption Data

Richard Blundell

*University College London,
IFS and IZA*

Hamish Low

*University of Cambridge
and IFS*

Ian Preston

*University College London
and IFS*

Discussion Paper No. 6125
November 2011

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit organization supported by Deutsche Post Foundation. The center is associated with the University of Bonn and offers a stimulating research environment through its international network, workshops and conferences, data service, project support, research visits and doctoral program. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ABSTRACT

Decomposing Changes in Income Risk Using Consumption Data

We develop a new approach to the decomposition of income risk within a nonstationary model of intertemporal choice. The approach allows for changes in income risk over the life-cycle and with the business cycle. It requires only repeated cross-section data and can allow for mixtures of persistent and transitory components in the dynamic process for income. Evidence from a stochastic simulation of consumption choices in a nonstationarity environment is used to show the robustness of the method for decomposing income risk. The approach is used to investigate the changes in income risk in Britain across the inequality growth period from the late 1970s to the late 1990s. Peaks in the variance of permanent shocks are shown to occur in the middle of the 1980s and the early 1990s.

JEL Classification: C30, D52, D91

Keywords: income risk, consumption, nonstationarity, inequality

Corresponding author:

Richard Blundell
Department of Economics
University College London
Gower Street
London WC1E 6BT
United Kingdom
E-mail: r.blundell@ucl.ac.uk

1 Introduction

Over time, consumers face changing profiles of income risk. They have to make their decisions about how much to spend and how much to save in a stochastic environment where the amount of risk is evolving over time. Some periods will be characterised by uncertainty over family income where shocks will permanently shift expected incomes whereas, in other periods, uncertainty will reflect shorter-term less persistent variation. Either kind of income risk will result in a rise in the inequality of income across consumers but the different types of shocks will have very different implications for consumption behaviour and these differences can be used to identify the evolution of different sorts of risk. The main objective of this paper is to use the joint distribution of income and consumption together with a model of intertemporal consumption choice and stochastic volatility to estimate the changing nature of income risk.

Over the last three decades many developed economies have experienced episodes of strong growth in income inequality, most notably the UK and the US (see Atkinson (1999) and Saez and Piketty (2003), respectively). Our analysis is motivated by these episodes but seeks to explore the underlying nature of income risk and the way it relates to changes in income inequality. We draw on two separate literatures. First, is the analysis of the distribution of permanent and transitory shocks to income and how this has changed over time; see Moffitt and Gottschalk (2002). Second, is the analysis of consumption inequality within an intertemporal choice setting, see Deaton and Paxson (1994). To understand the changes in uncertainty over time within an intertemporal choice setting, we extend this literature to model individuals as facing a non-stationary stochastic environment. Our method allows for the estimation of income risk and its evolution over time within such an environment.

Self-insurance and insurance contracts in general are harder to write when shocks are more persistent. Consumption therefore reacts more strongly to income shocks when the shocks are persistent. By studying intertemporal

choice in an environment where the amount of uncertainty is stochastic, we show that the evolution of the joint distribution of consumption and income can identify reliably changes in income risk over the life-cycle and with the business cycle.

We use our method on data from Britain to uncover different patterns of income persistence and consumption inequality over the business cycles of the 1980s and the 1990s, across different birth cohorts. We find that there was a distinct spike in the variance of permanent shocks in the early 1980s coinciding with the UK economy emerging from a deep recession. This spike is found across all birth cohorts. The early 1990s recession was less severe than the 1980s recession but it still resulted in a strong drop in GDP and a strong rise in unemployment in the UK. Again we find evidence of a spike in the variance of permanent shocks in the early 1990s but only for the younger birth cohorts.

Our method identifies components of income risk using income and consumption data from expenditure surveys. This generalizes the approach of Blundell and Preston (1998) by allowing for a general autoregressive income processes which includes as a special case, the permanent-transitory shock decomposition; and in being explicit about the approximation error. Our approach is based on a new approximation to the optimal consumption growth rule in an environment where the variance of shocks is itself stochastic. Allowing the variance to be stochastic means that the observation that the variance changes over time is consistent with the model of consumer behaviour. We use evidence from simulations of a dynamic stochastic model of intertemporal consumption choices in which the nonstationarity in income risk is modeled by a two state Markov switching model is used to show that this approximation can provide a robust method for decomposing income risk. The approach is found to work well for dynamic income models with a mixture of persistent (not necessarily unit root) and transitory processes. We also examine robustness to unobserved heterogeneity in consumption growth and to idiosyncratic trends in income growth.

In addition to identifying the variances of the persistent and transitory components of income shocks, our method identifies the transmission parameter from income shocks into consumption. If the persistent shock were to follow a unit root process and if there were no insurance against these shocks, the transmission parameter would be 1. The presence of different forms of insurance reduces this parameter below 1. Blundell, Preston and Pistaferri (2008) estimate values for the US which differ for households of different type but typically lie between 0.6 and 0.8. The difference from 1 is an estimate of the total amount of insurance provided against permanent shocks and is a measure of “partial insurance”: this is a measure that comprises both self-insurance, through saving and borrowing, and all other mechanisms such as family insurance. It is worth stressing that these estimates of partial insurance are not measures of the amount of insurance over and above self-insurance, but rather a combination of all forms of insurance. The advantage of using consumption choices to identify this transmission parameter is that we can remain agnostic about the source of partial insurance.

In the case where the persistent process has an AR(1) coefficient less than unity, the transmission parameter will be less than unity because of the lack of persistence in addition to the presence of any partial insurance, as discussed in Kaplan and Violante (2010). In such a case we are able to show that the availability of time paths for the joint moments of income and consumption allow both the partial insurance parameter and the AR(1) coefficient to be identified from repeated cross-section data. The results also show that the variances of the persistent and transitory components of income shocks remain accurately measured.

An important by-product of this paper is to show the value of using repeated cross-section data on income and consumption when longitudinal information is unavailable. Panel data surveys on consumption and income are limited. This means panel data is not able to identify the insurance against those risks. In Blundell, Pistaferri and Preston (2008), this was achieved for the US by imputing consumption data into the PSID, and we compare

our estimates of the transmission parameter and of the permanent variance to their estimates.¹ However, repeated cross-section household expenditure surveys that contain measurements on both consumption and income in the same survey are commonly available in many economies and over long periods of time.² For example, the data we use in our empirical application is from the Family Expenditure Survey (FES) in Britain which has been available on a consistent annual basis since the late 1960s (see Blundell and Preston, 1995). In the US, the Consumer Expenditure Survey has been available since 1980 (see Cutler and Katz, 1992, and Johnson, Smeeding and Torrey, 2005) and there are many other examples from other countries. Of course, the availability of panel data on consumption and income would allow both the identification of richer income processes, as well as the identification of additional transmission or ‘insurance’ parameters.

The layout of the rest of the paper is as follows. In section 2 we derive the approximations which link consumption inequality to income risk and which underly our methodology. The usefulness of having consumption and income data in the same survey is explored in detail. The methodology allows for mixtures of persistent and transitory income processes and does not require the persistent process to follow a unit root. The presence of a unit root can be tested using the method. On the other hand, if a unit root component is assumed when the correct specification is an AR(1), the model is likely to over-predict the level of insurance to income shocks. The income process considered in this section is more general than that considered in papers such as Blundell and Preston (1998) and Blundell, Pistaferri and Preston (2008) but the section is also novel in deriving meticulously the order of approximation error involved. Section 3 develops an approach for idiosyncratic trends in consumption and income and also discuss the robustness of our approxi-

¹Since 1997, the PSID has collected more consumption information but the survey is only every second year.

²There are, of course, limitations associated with the use of cross-section evidence alone. These are laid out clearly in what follows. For example, we assume that the cross-section covariance of shocks in any period with income in the previous period is zero.

mation to liquidity constraints and heterogeneity in discount rates. Section 4 describes the non-stationary environment we simulate and reports the results of our Monte Carlo experiments.

In Section 5 we present new estimates of the decomposition of income risk for Britain from the recession and inequality growth episodes of the 1980s and early 1990s. For the British data a mixture of a transitory and permanent (unit root) component with changing variances for each of these components is found to provide a good representation of the data. The results suggest a spike in the variance of permanent risk in each recession with a continuously growing variance for transitory risk over the period. An estimated transmission parameter for permanent income shocks on consumption of around 0.8 accords well with the standard self-insurance model. Section 6 concludes.

2 The Evolution of Income and Consumption Variances

2.1 The income process

Consider an individual i living for T periods. Until retirement at age R they work fixed hours to earn an income which evolves stochastically according to a process with a permanent-transitory decomposition. Specifically suppose log income in period t can be written

$$\ln y_{it} = \ln Y_{it} + \omega_t + u_{it} + V_{it} \quad t = 1, \dots, R-1$$

where Y_{it} represents a nonstochastic component of income, ω_t is a stochastic term common to the members of the cohort, u_{it} is an idiosyncratic transitory shock in period t and V_{it} evolves according to a process

$$V_{it} = \rho V_{it-1} + v_{it}$$

where v_{it} is an idiosyncratic persistent shock and ρ captures the persistence of the shock. The nonstochastic part of income contains a common³ deterministic trend η_t

$$\Delta \ln Y_{it} = \eta_t.$$

The final $T - R + 1$ periods of life are spent in mandatory retirement with no labour income.

The process for income can therefore be written

$$\begin{aligned} \Delta \ln y_{it} = & \rho \Delta \ln y_{it-1} + [\eta_t - \rho \eta_{t-1}] + \Delta[\omega_t - \rho \omega_{t-1}] \\ & + \Delta v_{it} + \Delta[u_{it} - \rho u_{it-1}], \end{aligned} \quad (1)$$

which simplifies to

$$\Delta \ln y_{it} = \eta_t + \Delta \omega_t + v_{it} + \Delta u_{it}, \quad (2)$$

if $\rho = 1$.

We let $\nu_{it} = (v_{it}, u_{it}, \omega_t - E_{t-1}\omega_t)'$ denote the vector of shocks in period t and $\boldsymbol{\nu}_{it}^s = (\nu'_{it}, \nu'_{it+1}, \dots, \nu'_{is})'$ denotes the stacked vector of idiosyncratic income shocks from period t to s .

We assume the idiosyncratic shocks u_{it} and v_{it} are orthogonal and unpredictable given prior information so that

$$E(u_{it} | v_{it}, \boldsymbol{\nu}_{i1}^{t-1}, Y_{i0}) = E(v_{it} | u_{it}, \boldsymbol{\nu}_{i1}^{t-1}, Y_{i0}) = 0.$$

We make no assumptions about the time series properties⁴ of the common shocks ω_t .

Setting ρ to 1 gives the popular specification with a transitory and permanent shock, compatible with an MA(1) process for idiosyncratic changes

³In section 3, we consider extending analysis to the case where this income trend is individual specific and we show how the approximation can be used in the presence of this heterogeneity in income growth.

⁴The lack of specificity about the time series properties of ω_t means we should refrain from thinking of it as specifically permanent or transitory in nature.

in log income⁵.

We assume that the variances of the shocks v_{it} and u_{it} are the same in any period for all individuals in any cohort but allow that these variances are not constant over time and indeed can evolve stochastically. Define $\text{Var}(u_t)$ to be the cross-section variance of transitory shocks in period t for a particular cohort and $\text{Var}(v_t)$ to be the corresponding variance of persistent shocks. These are the idiosyncratic components of persistent and transitory risk facing individuals.

Assuming the cross-sectional covariances of the shocks with previous periods' incomes to be zero, then the variance of income follows a second order difference process

$$\Delta [\text{Var}(\ln y_t) - \rho^2 \text{Var}(\ln y_{t-1})] = \Delta \text{Var}(v_t) + \Delta [\text{Var}(u_t) - \rho^2 \text{Var}(u_{t-1})]$$

In the case that $\rho = 1$ this expression is a first order process

$$\Delta \text{Var}(\ln y_t) = \text{Var}(v_t) + \Delta \text{Var}(u_t).$$

Permanent risk ($\text{Var}(v_t)$) or growth in transitory uncertainty ($\Delta \text{Var}(u_t)$) both result in growth of income inequality. Observing the cross-section distribution of income cannot, on its own, distinguish these.

2.2 Consumption choice

Consumption and income are linked through the intertemporal budget constraint

$$\sum_{s=0}^{T-t} \frac{c_{it+s}}{(1+r)^s} + \frac{A_{iT+1}}{(1+r)^{T-t}} = \sum_{s=0}^{R-t-1} \frac{y_{it+s}}{(1+r)^s} + A_{it} \quad (3)$$

where c_{it} denotes consumption in period t , A_{it} is assets at beginning of period t and r is a real interest rate, assumed for simplicity to be constant. The

⁵See Macurdy 1982, Abowd and Card 1989, Moffitt and Gottschalk 2002, Meghir and Pistaferri 2004 for examples of papers modelling the time series properties of individual earnings using longitudinal data.

terminal condition that $A_{iT+1} = 0$ implies that individuals will not borrow more than the discounted sum of the greatest lower bounds on income that they will receive in each remaining period.

Suppose the household plans at age t to maximise expected remaining lifetime utility

$$E_t \left[\sum_{\tau=0}^{T-t} \frac{U(c_{it+\tau})}{\delta_{t+\tau}} \right] \quad (4)$$

where $\delta_{t+\tau}$, $\tau = 0, 1, \dots$ is a sequence of subjective discount factors, assumed for the moment to be common, and $U : \mathbb{R} \rightarrow \mathbb{R}$ is a concave, three times continuously differentiable utility function.

The solution to the consumer problem requires expected constancy of discounted marginal utility $\lambda_{it+\tau}$ across all future periods

$$\begin{aligned} U'(c_{it+\tau}) &= \lambda_{it+\tau} \\ E_t \lambda_{it+\tau} &= \left(\frac{\delta_{t+\tau}}{\delta_t(1+r)} \right)^\tau \lambda_{it}, \quad \tau = 0, 1, \dots, T-t \end{aligned} \quad (5)$$

This is the familiar Euler condition for consumption over the life-cycle (see Hall 1978, Attanasio and Weber 1993, for example).

We show in Appendix A.1 that

$$\Delta \ln c_{it} = \varepsilon_{it} + \Gamma_{it} + \mathcal{O}(E_{t-1} \varepsilon_{it}^2) \quad (6)$$

where ε_{it} is an innovation term; Γ_{it} is an anticipated gradient to the consumption path, reflecting precautionary saving, impatience and intertemporal substitution. $\mathcal{O}(x)$ denotes a term with the property that there exists a $K < \infty$ such that

$$|\mathcal{O}(x)| < K |x|.$$

If preferences are CRRA and there is a common discount rate, then the gradient term does not depend on c_{it-1} and is common to all households, see Appendix A1. In section 3, we consider allowing Γ_{it} to vary within a cohort. The anticipated gradient to the consumption path could vary across individuals because of heterogeneity in the discount rate or in the coefficient

of relative risk aversion. We return in section 3 to the issue of how well the approximation would deal with this heterogeneity.

Thus, considering cross-sectional variation in consumption,

$$\Delta \text{Var}(\ln c_t) = \text{Var}(\varepsilon_t) + \mathcal{O}(E_{t-1} |\varepsilon_{it}|^3) \quad (7)$$

This has the implication that, up to a term which is $\mathcal{O}(E_{t-1} |\varepsilon_{it}|^3)$, the growth of the consumption variance should always be positive, as noted, for example, by Deaton and Paxson (1994).

2.3 Linking income and consumption shocks

The innovation ε_{it} is tied to the income shocks ω_t , u_{it} and v_{it} through the lifetime budget constraint (3). We show in the Appendix that we can approximate the relation between these innovations through a formula

$$\varepsilon_{it} = \phi_{it} v_{it} + \psi_{it} u_{it} + \Omega_t + \mathcal{O}_p \left(E_{t-1} \|\boldsymbol{\nu}_{i0}^{R-1}\|^2 \right) \quad (8)$$

where $\mathcal{O}_p(x)$ denotes a term with the property (see Mann and Wald 1943) that for each $\kappa > 0$ there exists a $K < \infty$ such that

$$P(|\mathcal{O}_p(x)| > K|x|) < \kappa,$$

where Ω_t is a common shock, defined in the Appendix, and ϕ_{it} and ψ_{it} are transmission parameters for the persistent and transitory shocks.

Specifically,

$$\begin{aligned} \phi_{it} &= \pi_{it} \sum_{j=0}^{R-t-1} \alpha_{t+j} \rho^j, \\ \psi_{it} &= \pi_{it} \alpha_t, \end{aligned} \quad (9)$$

where two additional parameters are introduced

- α_t : an annuitisation factor, common within a cohort.

- π_{it} : a self-insurance factor capturing the significance of lifetime earned income as a component of total human and financial wealth.⁶

To quantify the annuitisation factor, we need information on the time horizon, the interest rate and expected wage growth. To quantify the self-insurance factor we need to add to this information on current asset holdings and income levels. Precise definitions of these terms are given in the Appendix. Typically, the transmission of persistent shocks into consumption, ϕ_{it} , will be large relative to the transmission of transitory shocks, ψ_{it} . In particular, if we consider, for expositional purposes, the special case in which R and T go to infinity and in which $e^{\eta t} - 1 = \eta < r$ is constant then

$$\begin{aligned}\phi_{it} &= \frac{r - \eta}{1 + r - \rho(1 + \eta)} \pi_{it}, \\ \psi_{it} &= \frac{r - \eta}{1 + r} \pi_{it}.\end{aligned}$$

If we take the case where the persistent shocks are permanent, $\rho = 1$, then, in this infinite horizon case, $\phi_{it} = \pi_{it}$. If shocks are not permanent and $\rho < 1$, the value of ϕ_{it} will be less than π_{it} . Estimates of ϕ_{it} will therefore overestimate the amount of actual insurance, π_{it} , if $\rho < 1$. This is similar to the point in Kaplan and Violante (2009).

Let $\bar{\phi}_t$, $\bar{\psi}_t$, $\text{Var}_t(\phi_t)$ and $\text{Var}_t(\psi_t)$ be the cross section means and variances of ϕ_{it} and ψ_{it} . Since r , η and ρ are common within cohorts, variation in ϕ_{it} or in ψ_{it} comes only from variation in π across individuals. Any such variation arises due to differences in the expected amount of partial-insurance across individuals. Differences in the expected amount of insurance will arise because of differences across individuals in initial asset holdings, or more generally, if different individuals have differential access to insurance mechanisms. This interpretation of the variation in π arises because the approximation is being taken around the path of consumption that would be realised if, in each period, the individual received the mean shock to income. If individuals all

⁶If there were other mechanisms for insurance against permanent shocks, such as through the family, these would need to be included in the definition of π_{it} .

face the same income process and have the same level of initial assets, there will be no variation in this path and so no variation in π .

The growth in the cross-section variance and covariances of income and consumption take the form indicated in the following theorem:

Theorem 1 *Assuming an income process*

$$\begin{aligned}\Delta \ln y_{it} &= \rho \Delta \ln y_{it-1} + [\eta_t - \rho \eta_{t-1}] + \Delta[\omega_t - \rho \omega_{t-1}] \\ &\quad + \Delta v_{it} + \Delta[u_{it} - \rho u_{it-1}],\end{aligned}$$

then

$$\begin{aligned}\Delta \text{Var}(\ln y_t) &= \rho^2 \Delta \text{Var}(\ln y_{t-1}) + \Delta \text{Var}(u_t) - \rho^2 \Delta \text{Var}(u_{t-1}) + \Delta \text{Var}(v_t) \\ \Delta \text{Var}(\ln c_t) &= (\bar{\phi}_t^2 + \text{Var}(\phi_t)) \text{Var}(v_t) + (\bar{\psi}_t^2 + \text{Var}(\psi_t)) \text{Var}(u_t) \\ &\quad + \text{Var}(\pi_t) \Omega_t^2 + 2 \text{Cov}(\pi_t, \ln c_0) \Omega_t \\ &\quad + \mathcal{O}(E_{t-1} \|\boldsymbol{\nu}_{i0}^{R-1}\|^3) \\ \Delta \text{Cov}(\ln c_t, \ln y_t) &= (\rho - 1) \text{Cov}(\ln c_{t-1}, \ln y_{t-1}) + \bar{\phi}_t \text{Var}(v_t) \\ &\quad + \bar{\psi}_t \text{Var}(u_t) - \rho \bar{\psi}_{t-1} \text{Var}(u_{t-1}) \\ &\quad + \text{Cov}(\pi_t, \ln y_0) \Omega_t - \rho \text{Cov}(\pi_{t-1}, \ln y_0) \Omega_{t-1} \\ &\quad + \mathcal{O}(E_{t-1} \|\boldsymbol{\nu}_{i0}^{R-1}\|^3)\end{aligned}\tag{10}$$

Proof: See Appendix A1.

Corollary 1 *Assuming an income process $\Delta \ln y_{it} = \eta_t + \Delta \omega_t + \Delta u_{it} + v_{it}$, then*

$$\begin{aligned}\Delta \text{Var}(\ln y_t) &= \text{Var}(v_t) + \Delta \text{Var}(u_t) \\ \Delta \text{Var}(\ln c_t) &= (\bar{\phi}_t^2 + \text{Var}(\phi_t)) \text{Var}(v_t) + (\bar{\phi}_t^2 + \text{Var}(\phi_t)) \alpha_t^2 \text{Var}(u_t) \\ &\quad + \text{Var}(\pi_t) \Omega_t^2 + 2 \text{Cov}(\pi_t, c_{t-1}) \Omega_t + \mathcal{O}(E_{t-1} \|\boldsymbol{\nu}_{i0}^{R-1}\|^3) \\ \Delta \text{Cov}(\ln c_t, \ln y_t) &= \bar{\phi}_t \text{Var}(v_t) + \Delta [\bar{\pi}_t \alpha_t \text{Var}(u_t)] \\ &\quad + \Delta [\text{Cov}(\phi_t, \ln y_0) \Omega_t] + \mathcal{O}(E_{t-1} \|\boldsymbol{\nu}_{i0}^{R-1}\|^3).\end{aligned}\tag{11}$$

From these expressions derived from the life-cycle model of consumption, we can identify approximately the growth in the transitory variance and the level of the permanent variances from the growth in consumption and income variances. The approximation used can take differing degrees of accuracy depending on the information available and assumptions made about ρ , π_{it} , ϕ_{it} and α_t . The value of this theorem is that it clarifies assumptions required to justify simple approaches and how extra information could be used to improve on approximations as that extra information becomes available.

1. Particularly simple forms follow by allowing $\rho \simeq 1$, $\bar{\phi}_t \simeq 1$, $\text{Var}(\pi_t) \simeq 0$ and $\alpha_t \simeq 0$. More precisely, the assumption that $\alpha_t \simeq 0$ is that α_t in any particular period is negligible relative to its sum over the remaining periods. This implies there is a long horizon, and this in turn means $\psi_{it} = 0$ and transitory shocks will be smoothed completely. The assumption $\bar{\phi}_t \simeq 1$ implies that there is on average no insurance against permanent shocks and that such shocks get transmitted into consumption one-for-one. The assumption that $\text{Var}(\pi_t) \simeq 0$ means there is no heterogeneity in the extent of self-insurance and so common shocks do not generate any variability in consumption. Together these assumptions lead equation (10) to simplify to:

$$\begin{aligned}\Delta \text{Var}(\ln y_t) &= \text{Var}(v_t) + \Delta \text{Var}(u_t) \\ \Delta \text{Var}(\ln c_t) &\simeq \text{Var}(v_t) \\ \Delta \text{Cov}(\ln c_t, \ln y_t) &\simeq \text{Var}(v_t)\end{aligned}\tag{12}$$

so that the within cohort growth in the variance of consumption identifies the variance of permanent shocks. The difference between the growth in the within cohort variances of income and consumption then identifies the growth in the variance of transitory shocks through the first equation in (12). The evolution of the covariance should follow that of the consumption variance and this provides one testable overidentifying restriction per period of the data. This approximation is analogous

to Blundell and Preston (1998). The difference is that Blundell and Preston (1998) derive an exact relationship in levels using quadratic utility as opposed to the the current paper which is an approximation in logs and assumes CRRA utility.

Violations in the assumptions made to generate this simple approximation will result in approximation error. For example, if there is some insurance against permanent shocks, or if there is heterogeneity across individuals in the degree of this insurance, then this will generate an approximation error, and we assess how large this might be in section 4.2.2.

2. We can generalise this simplest approximation by relaxing the assumption that $\bar{\phi}_t \simeq 1$. This means that there will be partial insurance against permanent shocks, and so the transmission of permanent shocks into consumption will be less than one-for-one. However, maintaining the assumption that $\text{Var}(\pi_t) \simeq 0$ means that there is no heterogeneity in initial assets and the expected amount of partial insurance is common across individuals of the same age. Keeping the other assumptions that $\rho \simeq 1$, and $\alpha_t \simeq 0$ implies

$$\begin{aligned}\Delta \text{Var}(\ln y_t) &= \text{Var}(v_t) + \Delta \text{Var}(u_t) \\ \Delta \text{Var}(\ln c_t) &\simeq \bar{\phi}_t^2 \text{Var}(v_t) \\ \Delta \text{Cov}(\ln c_t, \ln y_t) &\simeq \bar{\phi}_t \text{Var}(v_t).\end{aligned}\tag{13}$$

These formulae are likely to provide a significant improvement to the approximation if reasonable values for $\bar{\phi}_t$ can be used. Two possible sources could be considered:

- With extraneous information on assets and incomes and assumptions about income growth, estimates of $\bar{\phi}_t$ could be calculated directly as the estimated fraction of human capital in total human and financial wealth

- Given the overidentification implied by availability of variance and covariance information on consumption and income, $\bar{\phi}_t$ could be estimated simultaneously with the variances of the shocks by, say, minimum distance methods. In principle, sufficient degrees of freedom exist to estimate $\bar{\phi}_t$ separately for each period; in practice, it would make sense to impose some degree of smoothness on the path of $\bar{\phi}_t$ over time, for example by estimating a suitable parametric time path, thereby retaining some degrees of freedom for testing.

Using the last two expressions in (13), the identification of $\bar{\phi}_t$ can be seen to come from the ratio of the evolution of the variance of consumption to the evolution of the covariance:

$$\bar{\phi}_t = \frac{\Delta \text{Var}(\ln c_t)}{\Delta \text{Cov}(\ln c_t, \ln y_t)} \quad (14)$$

3. If in addition to allowing for self-insurance, we allow $\rho < 1$, then, as in equation (10), the income variance and consumption-income covariance no longer obey simple difference equations. Maintaining the assumptions of no heterogeneity in self-insurance ($\text{Var}(\pi_t) \simeq 0$) and a long horizon ($\alpha_t \simeq 0$), the approximation becomes

$$\begin{aligned} \Delta \text{Var}(\ln y_t) &= \rho^2 \Delta \text{Var}(\ln y_{t-1}) + \Delta \text{Var}(u_t) - \rho^2 \Delta \text{Var}(u_{t-1}) + \Delta \text{Var}(v_t) \\ \Delta \text{Var}(\ln c_t) &= \bar{\phi}_t^2 \text{Var}(v_t) \\ \Delta \text{Cov}(\ln c_t, \ln y_t) &= (\rho - 1) \text{Cov}(\ln c_{t-1}, \ln y_{t-1}) + \bar{\phi}_t \text{Var}(v_t). \end{aligned} \quad (15)$$

The unknowns in this approximation are the variances of the persistent and transitory shocks, the value of the transmission parameter over time, $\bar{\phi}_t$, and the value of ρ . With enough periods, all these parameters can be estimated. In particular, we can use this approximation to test for whether the persistent shock follows a unit root or not.

As before, $\bar{\phi}_t$ is identified from a comparison of the evolutions of the variance of consumption and its covariance with income

$$\bar{\phi}_t = \frac{\Delta \text{Var}(\ln c_t)}{\text{Cov}(\ln c_t, \ln y_t) - \rho \text{Cov}(\ln c_{t-1}, \ln y_{t-1})}. \quad (16)$$

while information on ρ is available from the autoregressive properties of the income variance process.

4. Information allowing us to estimate sensible values for higher moments of ϕ_t and values of $\alpha_t \neq 0$ would in principle allow full use to be made of all terms in (10). Estimates of common shocks Ω_t could, for instance, be recovered since differences across individuals in the extent of self-insurance and hence in the transmission parameters would mean that common income shocks would create heterogeneous consumption shocks. In practice, such information is unlikely to be available and any such identification would be tenuous.

As we have stressed throughout, the main data requirement is cross section variances and covariances of log income and consumption, and panel data is not required. These variances and covariances can be estimated by corresponding sample moments with precision given by standard formulae. The underlying variances of the shocks can then be inferred by minimum distance estimation using (10), alongside estimation of values for $\bar{\phi}_t$, $\text{Var}(\pi_t)$, ρ and α_t , depending on what sophistication of approximation is used. The minimised distance providing a χ^2 test of the overidentifying restrictions.

3 Idiosyncratic Trends

In our discussion of the approximation in section 2, we assumed that there were no idiosyncratic trends in consumption or income. In this section, we show the extent to which heterogeneity in the income and consumption trends affects the approximations.

3.1 Consumption Trends

Heterogeneity in consumption trends may arise because of differences in impatience, or differences in the timing of needs over the life-cycle, or because of differences in the elasticity of intertemporal substitution. We can allow for such heterogeneity by introducing heterogeneous consumption trends Γ_{it} into equation (6):

$$\Delta \ln c_{it} = \varepsilon_{it} + \Gamma_{it} + \mathcal{O}(E_{t-1}\varepsilon_{it}^2)$$

Keeping to the assumption that $\rho \simeq 1$, $\text{Var}(\pi_t) \simeq 0$ and $\alpha_t \simeq 0$ leads the equations for the evolution of variances to be modified to give:

$$\begin{aligned} \Delta \text{Var}(\ln y_t) &\simeq \text{Var}(v_t) + \Delta \text{Var}(u_t) \\ \Delta \text{Var}(\ln c_t) &\simeq \bar{\phi}_t^2 \text{Var}(v_t) + 2\text{Cov}(c_{t-1}, \Gamma_t) \\ \Delta \text{Cov}(\ln c_t, \ln y_t) &\simeq \bar{\phi}_t \text{Var}(v_t) + \text{Cov}(y_{t-1}, \Gamma_t) \end{aligned} \quad (17)$$

The evolution of $\text{Var}(\ln c_t)$ is no longer usable because consumption trends must be correlated with levels of consumption at some points in the lifecycle so that $\text{Cov}(c_{t-1}, \Gamma_t) \neq 0$ for some t . In other words, the evolution of the cross-section variability in log consumption no longer reflects only the permanent component and so it cannot be used for identifying the variance of the permanent shock. By contrast, the evolution of $\text{Var}(\ln y_t)$ is unaffected and the evolution of $\text{Cov}(\ln c_t, \ln y_t)$ will also be unaffected if there is no reason for income paths to be associated with consumption trends (so that we assume that $\text{Cov}(y_{t-1}, \Gamma_t) = 0$). We can therefore still recover the permanent variance and the evolution of the transitory variance, but without any over-identifying conditions. The lack of over-identifying restrictions means that either we need an external estimate of $\bar{\phi}_t$ or we can only use our simplest approximation assuming $\bar{\phi}_t = 1$. Alternatively, we can impose that $\bar{\phi}_t$ is constant over time, or follows a parametric path, and this generates over-identifying restrictions when we have multiple periods of data.

The assumption that $\text{Cov}(y_{t-1}, \Gamma_t) = 0$ may be violated in particular cases: for example, if liquidity constraints are binding, then the path of

consumption will be affected by the timing of income. We explore in the simulations below the implications of liquidity constraints for our method. A second example is if heterogeneity in consumption paths is driven by a similar factor that drives income paths. Differential skill acquisition might be such a factor: a lower discount rate will lead to greater investment in skills which is likely to be associated with faster income growth.

3.2 Income Trends

Individuals also differ in their expectations about income growth, particularly across occupations and across education groups. For example, Baker (1997), Haider (2001) and Guvenen (2007) argue strongly for the importance of heterogeneity in income trends. Haider and Solon (2006) suggest that such heterogeneity in trends may be most important early in the life-cycle and late in the life-cycle. Where these differences are driven by observable characteristics (education, for example), the original approximation can be implemented after conditioning appropriately on group membership. To the extent, however, that these differences are unobservable, they will contaminate the evolution of the cross-section variance in income.

Letting

$$\Delta \ln Y_{it} = \eta_{it}$$

but maintaining the assumptions that $\rho = 1$, $\text{Var}(\pi_t) \simeq 0$ and $\alpha_t \simeq 0$, the equations for the evolution of the variances become

$$\begin{aligned} \Delta \text{Var}(\ln y_t) &\simeq \text{Var}(v_t) + \Delta \text{Var}(u_t) + 2\text{Cov}(y_{t-1}, \eta_t) \\ \Delta \text{Var}(\ln c_t) &\simeq \bar{\phi}_t^2 \text{Var}(v_t) \\ \Delta \text{Cov}(\ln c_t, \ln y_t) &\simeq \bar{\phi}_t \text{Var}(v_t) + \text{Cov}(c_{t-1}, \eta_t) \end{aligned} \tag{18}$$

The evolution of the cross-section variance of income is no longer informative about uncertainty. This implies that the link between the cross-section variability of income and uncertainty (as exploited by Meghir and Pistaferri, 2004, and Blundell, Pistaferri and Preston, 2005) is broken. The evolution

of $\text{Var}(\ln y_t)$ is no longer usable because income trends must be correlated with levels of income (differently at different dates but not always zero). However, the evolution of $\text{Var}(\ln c_t)$ is unaffected and can be used to identify the variance of permanent shocks given a value for $\bar{\phi}_t$. The evolution of the transitory variance cannot be identified and the role of the covariance term is useful as an overidentifying restriction only if the levels of consumption are uncorrelated with the income trend. This is unlikely to hold in practice because incomplete markets mean that the timing of income matters for consumption. However, the strength of this approach for identifying the permanent variance is that the consumption information identifies the unexpected component in income growth (for a given value of $\bar{\phi}_t$) and the permanent variance can be distinguished from expected variability.

4 Simulating Consumption Choices in a Non-stationary Environment

In the approach we have developed in this paper, moments are used to estimate variances of shocks by ignoring terms which are $\mathcal{O}(E_{t-1}\|\nu_{it}\|^3)$ and by ignoring heterogeneity in self-insurance by setting $\text{Var}(\pi_t) = 0$. The aim of the simulation analysis is to examine the accuracy with which changes to the underlying structural variances can be recovered in a nonstationary environment. To do this, we simulate the consumption behaviour of individuals in a life-cycle model under a range of assumptions about discounting, risk aversion, liquidity constraints and the income process.

The specific simulation designs are motivated by the sorts of numbers found in recent studies which have looked at the changing pattern of permanent and transitory shocks to income (see, for example, Moffitt and Gottschalk, 2002; Meghir and Pistaferri, 2004; Blundell, Pistaferri and Preston, 2008). From the simulations we construct cross-sections of income and consumption which we then use to assess our approach to decompose changes in income risk into permanent and transitory components.

4.1 Intertemporal Preferences and the Income Process

Consumers' within period utility is given by the constant elasticity of substitution form,

$$U(c_{it}) = \frac{\gamma_i}{1 + \gamma_i} c_{it}^{1+1/\gamma_i}. \quad (19)$$

When we allow for preference heterogeneity, this can enter through γ_i and also through the discount rate, δ_i , in equation (4).

The income process is outlined above in section 2.1, and transitory and permanent shocks to income are assumed to be log-normally distributed.⁷ When we allow for heterogeneity in the deterministic rate of income growth, this enters through η_{it} . Transitory shocks are assumed to be *i.i.d.* within period with variance growing at a deterministic rate. The permanent shocks are subject to stochastic volatility. We model the permanent variance as following a two-state, first-order Markov process with the transition probability between alternative variances, $\sigma_{v,L}^2$ and $\sigma_{v,H}^2$, given by β .

$$\begin{array}{c} \sigma_{v,L}^2 \quad \sigma_{v,H}^2 \\ \begin{array}{cc} \sigma_{v,L}^2 & \begin{array}{cc} 1 - \beta & \beta \end{array} \\ \sigma_{v,H}^2 & \begin{array}{cc} \beta & 1 - \beta \end{array} \end{array} \end{array} \quad (20)$$

This process means that consumers believe that the permanent variance has an ex-ante probability β of changing in each t . In the simulations, the variance actually switches only once and this happens in period S , which we assume is common across all individuals.⁸

The common stochastic terms ω_t are set at values which ensure that the uncertainty in log income is associated with no growth in the expected level of income and therefore ω_t also follows a two-state first-order Markov process.

⁷In the numerical implementation, we truncate the distribution at four standard deviations below the mean. The extent of truncation can affect the consumption function because individuals are able to borrow up to the amount they can repay with certainty.

⁸In solving the model for a particular individual, it is irrelevant whether a particular shock is idiosyncratic or common because the model is partial equilibrium.

While individuals therefore encounter a particularly large common shock in period S , there are smaller non-zero common shocks in all periods in the sense that $\omega_t \neq E_{t-1}\omega_t$ for all t .

Individuals begin their working lives with no assets. As discussed above, the terminal condition that $A_{iT+1} = 0$ restricts borrowing to the discounted sum of greatest lower bounds on incomes. In addition we consider the effect of introducing an explicit liquidity constraint:

$$A_{it} \geq 0 \tag{21}$$

We simulate individuals from age 20 to age 70 ($T = 70$), with the last 10 years of life spent in mandatory retirement. Individuals can use asset holdings to increase consumption in retirement. Parameters used in the baseline are summarised in Table 1.

We consider 14 experiments where we vary the parameters of the model. For each experiment, we simulate consumption, earnings and asset paths for 50,000 individuals. To obtain estimates of the variance for each period, we draw random cross sectional samples of 2000 individuals for each period from age 30 to 50. We repeat this process 1000 times to provide information on the properties of the estimators.

The way in which parameters are varied across experiments is described in Table 2. A first block of experiments considers the effect of changing the preference parameters, allowing for higher and lower values for the discount rate and the EIS. We also consider heterogeneity across individuals in the preference parameters which leads to heterogeneous consumption trends. A second block of experiments considers changing the income process, allowing for heterogeneity in the expected income paths that individuals face across their whole lives and heterogeneity in paths only for the first 10 years. We also simulate an AR(1) process, first with persistence equal to 0.95 and also with persistence equal to 0.90. Finally three further experiments consider further modifications: first, setting the growth in transitory variance to zero; second, reducing the number of retirement years to discourage asset accumulation

Table 1: **Baseline Parameter Values**

Description		Value
Discount Rate	δ	0.02
EIS	$-\gamma$	0.67
Income Growth Rate	η_t	0.0
Interest Rate	r	0.015
Change in Transitory Var.	$\Delta\sigma_{u_t}^2$	0.01
Permanent Variance	$\sigma_{v_t}^2 \quad t < S$	0.015
	$\sigma_{v_t}^2 \quad t \geq S$	0.005
Transition Probability	β	0.05
AR(1) Parameter	ρ	1
Age at Variance Switch	S	40
Retirement Age	R	60
Terminal Age	T	70

and introducing liquidity constraints; and finally allowing for social security pensions linked to final salary.

As discussed above, we calculate several estimates of differing subtlety. The simplest approximation, based on equation (12), would be accurate if it were not possible to insure at all against permanent shocks and if there were complete insurance against transitory shocks. In practice, individuals can use savings to insure partially against permanent shocks because individuals have finite horizons, and in the data, there may exist other mechanisms to smooth shocks, such as family transfers. We might therefore expect the accuracy of this simple approximation to depend on the utility cost of saving and the presence of other insurance mechanisms. We label such estimates $\phi = 1$.

We can improve on this simplest approximation by allowing for these insurance mechanisms. We do this by estimating $\bar{\phi}_t$, and hence the amount of insurance, jointly with the variances of the shocks by minimum distance assuming a linear path for ϕ over time.⁹ This is estimation based on equation (13). We label such estimates *MDE* ϕ . When we allow for an AR(1) process rather than a unit root, estimation is based on (15). When we allow for heterogeneity, estimation uses equation (17) for consumption heterogeneity, and equation (18) for income heterogeneity.

In each case the moments (11) are fitted by minimum distance using asymptotically optimum weights based on the estimated sampling precision of the sample moments. Estimated variances are smoothed by applying a third order moving average.

⁹This estimate of $\bar{\phi}_t$ should in principle capture any type of insurance although there is, of course, only self-insurance in the actual simulations.

Table 2: **Experiment Parameter Values**

Description	ρ	δ	EIS	η	$\Delta\sigma_{u_t}^2$	T
Baseline	1.0	0.02	0.67	0.0	0.1	70
Varying Preferences						
High discount rate	1.0	0.04	0.67	0.0	0.1	70
Low discount rate	1.0	0.01	0.67	0.0	0.1	70
High EIS	1.0	0.02	2.00	0.0	0.1	70
Low EIS	1.0	0.02	0.20	0.0	0.1	70
Hetero discount rate	1.0	$\left\{ \begin{array}{c} 0.01 \\ 0.02 \\ 0.04 \end{array} \right\}$	0.67	0.0	0.1	70
Hetero EIS	1.0	0.02	$\left\{ \begin{array}{c} 0.20 \\ 0.67 \\ 2.00 \end{array} \right\}$	0.0	0.1	70
Varying Income Process						
Hetero income growth	1.0	0.02	0.67	$\left\{ \begin{array}{c} -0.01 \\ 0.0 \\ 0.02 \end{array} \right\}$	0.1	70
Early hetero income growth	1.0	0.02	0.67	$\left\{ \begin{array}{c} -0.01 \\ 0.0 \\ 0.02 \end{array} \right\}$	0.1	70
AR(1) Process	0.95	0.02	0.67	0.0	0.1	70
AR(1) Process	0.90	0.02	0.67	0.0	0.1	70
Other Variations						
No transitory variance growth	1.0	0.02	0.67	0.0	0.0	70
Liquidity constrained	1.0	0.02	0.67	0.0	0.1	62
Social security	1.0	0.02	0.67	0.0	0.1	70

For experiments with heterogeneity, one half of each sample have the middle value of the heterogeneous parameter and one quarter of the sample have each of the extreme values. For the experiment with early heterogeneity in income growth, the heterogeneity is present only up to age 30, after which income grows at a common rate of 0. For the experiment with social security, individuals enjoy an additional retirement income equal to one half of income in the final period of working life.

4.2 Simulation Results

4.2.1 Self insurance

Crucial to the accuracy of the approximations (11) are the extent of self insurance, captured by ϕ_{it} . In Figure 1 we show the values of $\bar{\phi}_t$ for each of the simulations across the twenty years over which we follow individuals. Note that these are the means of the distribution of the true ϕ_{it} and not the approximations used in estimation. A value of $\phi = 1$ indicates no assets are being held and permanent shocks pass through one-for-one into consumption. A value less than 1 indicates partial insurance, whereas a value greater than 1 would mean an individual is borrowing, and is over-exposed to income shocks.

The baseline case gives a $\bar{\phi}_t$ declining, as future labour income diminishes and assets are built up, from a little below 0.9 at age 30 to a little more than 0.5 at age 50.

A high discount rate discourages saving since it is more costly in terms of utility for individuals to self-insure. A high elasticity of intertemporal substitution also discourages saving. The CRRA specification implies that a high γ means individuals have low risk aversion and low prudence and this means savings are less valuable and there is less precautionary saving and self-insurance. These cases therefore involve diminished self insurance and raise $\bar{\phi}_t$. Lower values of discount rates or the EIS on the other hand reduce $\bar{\phi}_t$. The experiments with heterogeneity in these parameters give similar mean values of ϕ_t to the baseline case. Eliminating the growth in the transitory variance reduces saving and raises $\bar{\phi}_t$ but not by very much.

Figure 1 reports the range of $\bar{\phi}_t$ for some of the experiments considered. The values of $\bar{\phi}_t$ for the other experiments depend on how the cost of saving varies. For example, reducing life expectancy after retirement reduces the motive to accumulate assets during working life and this is combined with an explicit borrowing constraint in the liquidity constrained experiment. In this case, unsurprisingly, asset accumulation is heavily reduced and self insurance

is the lowest of any of the scenarios considered.

Finally, introducing a social security pension equal to half of final income. reduces the incentives to accumulate private assets for consumption in retirement. Moreover, in this case self insurance against permanent shocks is less effective for any future income path, given current asset holdings, because the influence of shocks to income carry on into retirement. The relation between shocks to income and consumption is no longer captured accurately by (8) unless the formula for ϕ_{it} (equation 9) is modified to account for this.¹⁰ The values for $\bar{\phi}_t$ used in this case incorporate such a modification and are substantially higher, particularly at older ages, than in the base case.

This discussion of how the transmission of permanent shocks into consumption via ϕ_{it} varies in different scenarios highlights that motives for holding assets are not additive: assets held for retirement can be used to smooth shocks if necessary, and similarly precautionary balances which are not used can then be consumed in retirement. Our approach does not have to model this fungibility of assets because the approximations estimate ϕ_{it} directly and thus provides an estimate of the pass-through without modelling the source.

4.2.2 Estimating the permanent variance

Baseline simulations Figure 2 shows estimates of the permanent variance by age of the cohort for our baseline case. We report the true path of the variance and the alternative approximation.

The estimates assuming $\bar{\phi}_t = 1$ consistently underestimate the permanent variance. This is because asset holdings enable partial self-insurance against the permanent shocks.¹¹ The cross-section variance of consumption reflects the uninsured part of the permanent shock and this is an underestimate of

¹⁰The correct coefficient treats the anticipated social security receipts as part of labour income, weighted according to the proportion of final salary to which individuals are entitled.

¹¹This partial insurance against permanent shocks would not be feasible in an infinite horizon setting.

the actual permanent shock. Nonetheless the change in the value of the variance $\text{Var}(v_t)$ is clearly picked up.

Further, correcting for self-insurance possibilities secures a considerable improvement in estimates with the means across Monte Carlo replications very close to the true values in the simulations and no evident deterioration in quality with age. This improvement is observed when we estimate $\bar{\phi}_t$ alongside the variances. This correction for partial-insurance does not rely on asset data or specifying the mechanism through which insurance occurs. This finding is the key conclusion of our simulations: we can recover accurately the variance of shocks to income if we estimate the transmission parameter alongside the variances. As we show below, the accuracy of this process does depend on some assumptions, such as the absence of binding liquidity constraints.

In addition to this comparison with the true value of the variance, we can test the overidentifying restrictions by calculating the frequency of rejection at the 5% level across the simulations. In the baseline, and across all experiments, tests of the restrictions with $\bar{\phi}_t = 1$ always reject strongly. Given that these estimates of the permanent variance are systematically downward biased, this rejection is not surprising. By contrast, when we correct for self-insurance, rejections are much less frequent, and the distribution of the overidentification tests appears very close to the appropriate χ^2_{17} distribution with a mean close to the degrees of freedom and with size typically close to 5% (slightly over-rejecting).

Sensitivity to the Preference Parameters In Figure 3, we consider how accurately our approximation estimates the permanent variance when we vary the discount rate (the left-hand column) and the elasticity of intertemporal substitution (the right-hand column). In each case, the first two rows consider high and low values of the relevant preference parameter, maintaining the assumption that preferences and hence consumption growth rates are homogenous across individuals. The third row allows for heterogeneity

in preferences and so heterogeneity in consumption growth.

When preferences are homogenous, the estimates which assume there is no self insurance against permanent shocks are most accurate in those scenarios where savings is costly. In these cases, for example when discounting is high or the EIS is high, there is less asset accumulation and $\bar{\phi}_t$ is closest to 1. On the other hand, in all scenarios, correcting for self-insurance by estimating $\bar{\phi}_t$ through minimum distance leads to very accurate estimates of the permanent variance.

Heterogeneity in preferences and the resulting heterogeneity in consumption paths means that the change in the cross-section variance of consumption should no longer be used to identify the variance of permanent income shocks. The third row of figure 3 shows that simple estimates of the permanent variance that use this information, despite the presence of heterogeneity, lead to downward bias, although of similar magnitude to the cases with homogenous preferences. However, as discussed in section 3, we can drop this contaminated moment and use only the information in the income-consumption covariance to estimate the permanent variance. This estimate is labelled “Robust, $\bar{\phi}_t = 1$ ” in the third row of figure 3. Because of the reduction in number of moments we no longer have the degrees of freedom required to estimate $\bar{\phi}_t$ within the minimum distance calculation so a value either needs to be imposed or calculated, say, from asset data.

In figure 3, we impose $\bar{\phi}_t = 1$, and this moves the estimated permanent variance closer to the truth. On the other hand, the most accurate estimate seems, despite the heterogeneity, to be those based on full minimum distance, estimating $\bar{\phi}_t$ but without correcting for heterogeneity.

Sensitivity to the Specification of the Income Process We show how two types of variation of the income process affect our ability to estimate the permanent variance. First, we allow for heterogeneous income profiles, as in section 3. Second, we allow income to follow an AR(1) process as in our general theorem 10, rather than a unit root.

The top row of Figure 4 shows estimates of the permanent variance when income growth is heterogenous. On the left-hand side, we allow for heterogeneity in the expected income paths that individuals face across their whole lives, as in Guvenen (2007). On the right hand side, we allow for heterogeneity in paths only for the first 10 years, as in Haider and Solon (2006). The bottom row of Figure 4 shows estimates of the permanent variance when income follows an AR(1) process, with persistence equal to 0.95 on the left-hand side and also with persistence equal to 0.90 on the right hand side.

When we have heterogeneity in income growth over the whole lifetime, information on the variance of log income and the covariance between log income and consumption will be contaminated by variability due to this heterogeneity. Nonetheless the permanent variance remains estimable from the consumption variance given a suitable estimate for $\bar{\phi}_t$. We report estimates using our simple approximation ignoring the possibility of heterogeneity and assuming $\bar{\phi} = 1$. We also report estimates where $\bar{\phi}_t$ is estimated by minimum distance alongside the permanent variance but again maintaining the false assumption that there is no heterogeneity.

Finally, we allow for heterogeneity and following approximation 18, we report estimates of the permanent variance using only the evolution of the variance of consumption and imposing $\bar{\phi} = 1$. Our simple approximation underestimates the permanent variance, and this arises, as before, because the variability of consumption is dampened by self-insurance and this is interpreted as indicating a lower permanent variance. On the other hand, estimating $\bar{\phi}_t$ by minimum distance over-predicts the permanent variance. This arises because variability in income due to heterogeneity is being attributed to the permanent shock. Correcting for the heterogeneity by dropping the moments using the variability in income reduces the estimates of the permanent variance, although there is still some over-prediction.

When we consider heterogeneity in income growth rates which lasts only until age 30, which is more in keeping with the results of Haider and Solon (2006), and if we use data after that heterogeneity is resolved, then our results

look very similar to the baseline and the use of the moments involving the variance of income do not introduce evident bias.

The graphs for different values of ρ report two methods of estimating the permanent variance: first, maintaining the false assumption that $\rho = 1$ and the assumption that $\bar{\phi}_t = 1$; second, estimating $\bar{\phi}_t$ and ρ alongside the permanent variance. Lower values of ρ have a direct effect reducing the value of $\bar{\phi}$ below 1, and so our estimates of the permanent variance using the first method become worse as ρ decreases, as shown in the bottom two graphs in figure 4. When we estimate $\bar{\phi}_t$ and ρ alongside the permanent variance, the underestimate of the permanent variance is corrected, and indeed over-corrected when ρ becomes as low as 0.9. Table 3 reports the estimates of ϕ and ρ from using this second method. The method estimates accurately the degree of persistence of shocks across a range of values of ρ .

Other Sensitivity Tests In our baseline estimates and the sensitivity analysis so far, individuals do not face explicit borrowing constraints. Further, the need to save for retirement means that individuals do not have a strong desire to borrow except when very impatient. In Figure 5, we show the estimates of the permanent variance when individuals have a strong incentive to borrow, but face an explicit borrowing constraint. We generate this scenario by drastically cutting the length of the retirement period so that individuals behave as buffer stock consumers (as in Carroll, 1997).

When individuals are liquidity constrained, they are no longer able to insure transitory shocks fully and transitory shocks will generate extra variability in the cross-section variance of consumption. Since our simplest approximation assumes that transitory shocks are fully insured, this extra variability in the consumption data is interpreted as variability in permanent income leading to an over-estimate of the permanent variance. Our corrections for self-insurance make little difference to this bias because the bias in this case is not due to underestimating the extent of self-insurance against permanent shocks. On the other hand, our approximation continues to capture much of

the true decline in the permanent variance.¹²

A final experiment modifies the basic set-up by giving individuals a social security income in retirement equal to one half of income in period $R - 1$. As discussed earlier this changes the relation between income shocks and consumption. Despite this, the permanent variance is picked up fairly accurately by either of the methods allowing for self insurance as shown in Figure 5.

4.2.3 Estimating changes in the transitory variance

Estimates of the change in the transitory variance for the main experiments discussed are shown in Figure 6, in all cases using MDE to estimate the self insurance parameter. In the cases discussed so far the growth in the transitory variance is picked up with a high degree of accuracy. The exception to this is the case with liquidity constraints, where the approximation mis-identifies transitory shocks as permanent shocks because transitory shocks do get transmitted into consumption.

5 Income Inequality and Income Risk: Results for Britain, 1979-1997.

We now turn to apply the ideas and techniques to the study of inequality in Britain over the period 1979 - 1997, covering two business cycles. We use data from the Family Expenditure Survey. This is an annual continuous cross-sectional budget survey with detailed data on incomes and consumption expenditures of British households.

The period chosen covers the deep recession in the early 1980s following by the ‘inequality boom’ period in Britain in which there was rapid growth

¹²It is important to stress that this bias arises because liquidity constraints do actually bind. The presence of future potentially binding constraints does not cause the same issue because the Euler equation will still hold.

in income inequality; see Atkinson (1999), for example. The British economy then fell in to a further recession in the early 1990s followed by a period of sustained growth. Over this period there was also growth in consumption inequality, especially in the early to mid 1980s; see Blundell and Preston (1998). These patterns in consumption and income inequality match many of the features observed in the US over this period, see Johnson, Smeeding and Torrey (2005) and Blundell, Pistaferri and Preston (2008).

The income measure used is equivalised household income after housing costs.¹³ Expenditure is equivalised household expenditure on nondurables.¹⁴ In each year we trim from the sample households with either income or expenditure in the highest or lowest 0.5 per cent of the survey. Households are classified into cohorts according to ten year bands for date of birth of head of household. We focus our attention here on households headed by individuals in two central birth cohorts for which there is a reasonable sample across the whole of the period - those born in the 1940s and 1950s - keeping only those households with heads aged between 25 and 60.

Estimated variances and covariances of income and consumption over the period are calculated for each year by pooling data on all households headed by an individual born in the appropriate cohort and sampled within a three year band centred on the year in question. Figure 7 shows these estimated variances and covariances of income and consumption over the period for the two birth cohorts combined. The variance of income is rising over the sample period though the rise flattens off towards the end. This rising path is followed in the earlier years by the variance of consumption but the path flattens off somewhat in the early 1990s. The covariance between income and consumption tends to increase throughout the period. The right panel of the

¹³This is a standard UK definition for disposable household income, see Brewer, Goodman, Muriel and Sibiet (2007). The equivalence scale used is the OECD scale.

¹⁴Expenditure on durable goods does not coincide with consumption of durable services. The possibility of delaying replacement of durable or semidurable goods as a way to manage income shocks (see Browning and Crossley, 2009) makes a definition including durable expenditures inappropriate.

same figure shows the year-on-year changes in the three series (smoothed by a three year moving average).

Using decomposition (15) we cannot reject the hypothesis of a unit root for the persistent component in the income process.¹⁵ Initially we estimate $\bar{\pi}_t$ jointly with $\text{Var}(v_t)$ and $\Delta\text{Var}(u_t)$ by minimum distance estimation on the pooled data with asymptotically optimal weighting. Assuming a constant $\bar{\pi}_t$ we find an estimated value of 0.816 (with a standard error of 0.031) seems plausible: this is a value not only well within the range of values simulated in Figure 1 but also not unlike estimates for US data¹⁶ in Blundell, Pistaferri and Preston (2008). We relax the constancy over time of $\bar{\pi}_t$ below so that $\bar{\pi}_t$ follows a linear time trend.

The estimated variances for the persistent and transitory components of income, smoothed using a third order moving average, are shown with pointwise 95 per cent confidence bands Figure 8. The first panel presents the estimated permanent variance. The evidence points towards two peaks, the first as the economy emerged from recession in the 1980s and a second during the recession of the early 1990s. The earlier period of high permanent variance corresponds to the period of key labour market reforms and the strong growth in returns to education which also occurred in the early to mid period of the 1980s.¹⁷

The estimated growth rate in the transitory variance peaks alongside the permanent variance as the economy emerges from recession in the 1980s, but the growth rate declines through the 1990s recession. As discussed in section 4.2.2, the presence of liquidity constraints can lead our method to over-estimate the extent of the permanent variance and under-estimate the growth in the transitory variance. The growth in permanent variance could therefore

¹⁵The estimate of ρ is 1.004 with a 95% bootstrap confidence interval that contains unity.

¹⁶That paper, for instance, estimates a comparable value of 0.793 for the cohort born in the 1940s over the period 1979-92.

¹⁷See Gosling and Machin (1995) and Gosling, Machin and Meghir (2000) and references therein.

be picking up a growth in transitory variance among liquidity constrained households.

The remaining figures present estimates for the two birth cohorts separately though with a common linearly age-dependent path for $\bar{\pi}_t$ for the two cohorts. The estimated path¹⁸ for $\bar{\pi}_t$ declines from 0.849 at age 30 to 0.783 at age 50. Again, this fits extremely well with the earlier simulations. The estimated variances, smoothed using a third order moving average and shown with pointwise 95 per cent confidence bands, are presented in Figures 11 and 12.

For both cohorts the period of highest permanent variance is the mid 1980s. For the older cohort the permanent variance falls away from then on with only slight evidence of a second peak. For the younger cohort the evidence of the second peak is quite pronounced with permanent variance in the early 1990s rising more or less to the level of the mid 1980s. The transitory variance, on the other hand, appears to be growing through most years for both cohorts, except towards the very end of the period covered.

The source of these differences across the two cohorts can be seen in the different paths of the variance and covariance shown in Figure 9, and shown more explicitly in Figure 10 which displays the changes over time. In the 1980s the change in the variances for both cohorts look similar: there is a spike in the variance of income which is matched by a spike in the consumption variance and covariance.

These results suggest a spike in the permanent variance for both cohorts as they move through the 1980s. By contrast, in the 1990s, the cohorts look different to each other. The older cohort has a second spike in the variance of income, but very little change in the consumption variance indicating that growth in income uncertainty for this cohort during the early 1990s largely reflected a growth in the variance of transitory shocks. For the younger cohort there is a clear spike in the variance of consumption, but much less

¹⁸The slope of -0.0034 per year, though plausible, has a standard error of 0.0066 and can not be considered statistically significantly different from zero.

of a spike in income. This suggests that the younger cohort experienced a spike in permanent variance coincident with a decline in transitory variance as indicated in Figures 11 and 12.

6 Conclusions

Increases in cross-section measures of income inequality may reflect the variance of persistent shocks or increases in the variability of transitory shocks. However, the differing sources of risk have very different implications for welfare,¹⁹ and the importance of these different sources of risk is likely to change over individuals lifetimes and over the business cycle. In this paper, we provide a way of identifying how these risks evolve by using repeated cross-section data on income and consumption. This is the type of data typically available in consumer expenditure surveys.

Using a dynamic stochastic simulation framework we have shown that approximations to consumption rules that allow for the nonstationary environment faced by households can be used to decompose income variability into its components. In assessing the accuracy of this decomposition we show that it is able to map accurately the evolution of transitory and persistent variances of income shocks across a range of alternative parameterisations. Our results allow for an autoregressive process for persistent income shocks and we have shown how to recover the degree of persistence using cross-section moments.

We use this method to map out the evolution of income risk in Britain through the recessions of the early 1980s and early 1990s. We show the different patterns of income persistence in the aftermath of recessions across different birth cohorts. The early 1980s coincided with a spike in the variance of permanent shocks across all birth cohorts, while for the early 1990s recession, the spike in the variance of permanent shocks can only be detected

¹⁹See, for example, the discussions in Blundell and Preston, 1998; Heathcote, Storesletten and Violante, forthcoming; and Low, Meghir and Pistaferri, 2009.

among the younger birth cohort.

In the standard decomposition any unobserved heterogeneity in income paths will be labelled as unexplained variability in the growth in income and be defined as risk. Panel data on income can be used to explore the degree of heterogeneity, as discussed in Baker (1997) and Guvenen (2007), although typically long panels are required to clearly identify heterogeneity in income paths. We have shown that the approximation developed here can accommodate such heterogeneity in income paths. Further, with additional assumptions, we can use the variance of consumption to separate out uncertainty from that variability which is due only to this heterogeneity in income paths.

The approach developed here relies on the assumptions of optimising behaviour and of individuals having preferences with a constant relative risk aversion form. However, unlike direct solutions using dynamic programming (as in Gourinchas and Parker, 2002), we do not have to specify (or estimate) the shape of the consumption function or the values for the discount rate, risk aversion or elasticity of intertemporal substitution. Further, we show how to allow for idiosyncratic trends in consumption and income.

As a final point it is worth emphasizing that repeated cross-sections alone, even with accurate measures on income and consumption, have their limitations. A longer term goal would be to establish accurate measures of consumption in panel surveys of income dynamics. This would allow the identification of richer models and a more accurate distinction between alternative specifications. Such a panel could identify additional transmission or ‘insurance’ parameters as well as the separate evolution of permanent and transitory income variances.

A.1 Appendix: Proof of Theorem 1

The approximation in section 2 uses the Euler equation to relate consumption growth to innovations. These innovations are related to income shocks

through an approximation to the budget constraint. The validity of the approximation depends on the order of the error in approximations to the Euler equation and to the budget constraint. The aim of this appendix is firstly to show how the approximation relating consumption variance to income variance is derived and secondly to show the order of the error of this approximation.

A.1.1 Approximating the Euler Equation

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

By (5)

$$E_t U'(c_{it+1}) = U'(c_{it}) \left(\frac{\delta_{t+1}}{\delta_t(1+r)} \right) = U'(c_{it} e^{\Gamma_{it+1}}) \quad (22)$$

for some Γ_{it+1} .

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

$$\begin{aligned} U'(c_{it+1}) &= U'(c_{it} e^{\Gamma_{it+1}}) \left[1 + \frac{1}{\gamma(c_{it} e^{\Gamma_{it+1}})} [\Delta \ln c_{it+1} - \Gamma_{it+1}] \right. \\ &\quad \left. + \frac{1}{2} \beta(\tilde{c}, c_{it} e^{\Gamma_{it+1}}) [\Delta \ln c_{it+1} - \Gamma_{it+1}]^2 \right] \end{aligned} \quad (23)$$

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

Taking expectations

$$\begin{aligned} E_t U'(c_{it+1}) &= U'(c_{it} e^{\Gamma_{it+1}}) \left[1 + \frac{1}{\gamma(c_{it} e^{\Gamma_{it+1}})} E_t [\Delta \ln c_{it+1} - \Gamma_{it+1}] \right. \\ &\quad \left. + \frac{1}{2} E_t \{ \beta(\tilde{c}, c_{it} e^{\Gamma_{it+1}}) [\Delta \ln c_{it+1} - \Gamma_{it+1}]^2 \} \right] \end{aligned} \quad (24)$$

Substituting for $E_t U'(c_{it+1})$ from (22),

$$\frac{1}{\gamma(c_{it} e^{\Gamma_{it+1}})} E_t [\Delta \ln c_{it+1} - \Gamma_{it+1}] + \frac{1}{2} E_t \{ \beta(\tilde{c}, c_{it} e^{\Gamma_{it+1}}) [\Delta \ln c_{it+1} - \Gamma_{it+1}]^2 \} = 0$$

and thus

$$\Delta \ln c_{it+1} = \Gamma_{it+1} - \frac{\gamma(c_{it}e^{\Gamma_{it+1}})}{2} E_t \left\{ \beta(\tilde{c}, c_{it}e^{\Gamma_{it+1}}) [\Delta \ln c_{it+1} e^{\Gamma_{it+1}}]^2 \right\} + \varepsilon_{it+1} \quad (25)$$

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

$$\Delta \ln c_{it+1} = \varepsilon_{it+1} + \Gamma_{it+1} + \mathcal{O}(E_t |\varepsilon_{it+1}|^2). \quad (26)$$

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

$$\Delta \ln c_{it+1} = \varepsilon_{it+1} + \Gamma_{t+1} + \mathcal{O}(E_t |\varepsilon_{it+1}|^2). \quad (27)$$

A.1.2 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).

Let $\Xi^{N+2} = \left\{ \boldsymbol{\xi} \in \mathbb{R}^{N+2} \mid \sum_{j=0}^N \exp \xi_j + \xi_{N+1} > 0 \right\}$ and define a function $F : \Xi^{N+2} \rightarrow \mathbb{R}$ by $F(\boldsymbol{\xi}) = \ln \left[\sum_{j=0}^N \exp \xi_j + \xi_{N+1} \right]$. By exact Taylor expansion around an arbitrary point $\boldsymbol{\xi}^0 \in \Xi^{N+2}$ there exists a $\tilde{\boldsymbol{\xi}}$ between $\boldsymbol{\xi}$ and $\boldsymbol{\xi}^0$ such that

$$\begin{aligned} F(\boldsymbol{\xi}) &= \ln \left[\sum_{j=0}^N \exp \xi_j^0 + \xi_{N+1}^0 \right] + \sum_{j=0}^N \frac{\exp \xi_j^0}{\sum_{k=0}^N \exp \xi_k^0 + \xi_{N+1}^0} (\xi_j - \xi_j^0) \\ &\quad + \frac{\xi_{N+1} - \xi_{N+1}^0}{\sum_{k=0}^N \exp \xi_k^0 + \xi_{N+1}^0} \\ &\quad + \frac{1}{2} \sum_{j=0}^{N+1} \sum_{k=0}^{N+1} \frac{\partial^2 F(\tilde{\boldsymbol{\xi}})}{\partial \xi_j \partial \xi_k} (\xi_j - \xi_j^0) (\xi_k - \xi_k^0) \end{aligned} \quad (28)$$

The coefficients in the remainder term are given by

$$\begin{aligned}
\frac{\partial^2 F(\tilde{\xi})}{\partial \xi_j \partial \xi_k} &= \frac{\exp \tilde{\xi}_j}{\sum_k^N \exp \tilde{\xi}_k + \tilde{\xi}_{N+1}} \left(\delta_{jk} - \frac{\exp \tilde{\xi}_j}{\sum_k^N \exp \tilde{\xi}_k + \tilde{\xi}_{N+1}} \right) & j, k < N+1 \\
&= - \frac{\exp \tilde{\xi}_j}{\left[\sum_k^N \exp \tilde{\xi}_k + \tilde{\xi}_{N+1} \right]^2} & j < N+1, k = N+1 \\
&= - \frac{1}{\left[\sum_k^N \exp \tilde{\xi}_k + \tilde{\xi}_{N+1} \right]^2} & j = k = N+1
\end{aligned}$$

where δ_{jk} denotes the Kronecker delta.

Taking expectations of (28) subject to arbitrary information set $\mathcal{I}$

$$\begin{aligned}
E_{\mathcal{I}}[F(\xi)] &= \ln \left[\sum_{j=0}^N \exp \xi_j^0 + \xi_{N+1}^0 \right] + \sum_{j=0}^N \frac{\exp \xi_j^0}{\sum_{k=0}^N \exp \xi_k^0 + \xi_{N+1}^0} (E_{\mathcal{I}} \xi_j - \xi_j^0) \\
&\quad + \frac{\xi_{N+1} - \xi_{N+1}^0}{\sum_{k=0}^N \exp \xi_k^0 + \xi_{N+1}^0} \\
&\quad + \frac{1}{2} \sum_{j=0}^{N+1} \sum_{k=0}^{N+1} E_{\mathcal{I}} \left[\frac{\partial^2 F(\tilde{\xi})}{\partial \xi_j \partial \xi_k} (\xi_j - \xi_j^0)(\xi_k - \xi_k^0) \right] \quad (29)
\end{aligned}$$

We now apply this expansion to the two sides of the budget constraint, expanding on each side around the paths that would be taken by the variables in the event that all shocks are realised as zero, $u_{it} = v_{it} = \omega_t = 0$.

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

$$\begin{aligned}
\xi_j &= \ln c_{it+j} - j \ln(1+r) \\
\xi_j^0 &= \ln C_{it+j} - j \ln(1+r), \quad j = 0, \dots, T-t \quad (30)
\end{aligned}$$

and

$$\xi_{T-t+1} = \xi_{T-t+1}^0 = 0$$

where $\ln C_{it+j} = \ln c_{i0} + \sum_{\tau=1}^{t+j} \Gamma_{\tau}$ is the path followed by consumption in the absence of income risk. Then, substituting equation (30) into equation (29)

and noting only the order of magnitude for the remainder term,

$$\begin{aligned}
E_{\mathcal{I}} \left[\ln \sum_{j=0}^{T-t} \frac{c_{it+j}}{(1+r)^j} \right] &= \ln \sum_{j=0}^{T-t} \frac{C_{it+j}}{(1+r)^j} \\
&+ \sum_{j=0}^{T-t} \theta_{it+j} [E_{\mathcal{I}} \ln c_{it+j} - \ln C_{it+j}] \\
&+ \mathcal{O}(E_{\mathcal{I}} \|\boldsymbol{\varepsilon}_{i0}^T\|^2)
\end{aligned} \tag{31}$$

where

$$\theta_{it+j} = \frac{\exp \xi_j^0}{\sum_{k=0}^N \exp \xi_k^0} = \frac{C_{it+j}/(1+r)^j}{\sum_{k=0}^{T-t} C_{it+k}/(1+r)^k},$$

and $\boldsymbol{\varepsilon}_{i0}^T$ denotes the vector of lifetime consumption innovations $(\varepsilon_{i0}, \varepsilon_{i1}, \dots, \varepsilon_{iT})'$.

The term θ_{it+j} can be seen as an annuitisation factor for consumption.

We now apply the expansion (29) similarly to the expected present value of resources, $\sum_{j=0}^{R-t-1} (1+r)^{-j} y_{it+j} + A_{it} - A_{iT+1} (1+r)^{-(T-t)}$. Let $N = R-t-1$ and let

$$\begin{aligned}
\xi_j &= \ln y_{it+j} - j \ln(1+r) \\
\xi_j^0 &= \ln Y_{it+j} - j \ln(1+r) \quad j = 0, \dots, R-t-1 \\
\xi_{R-t} &= [A_{it} - A_{iT+1} (1+r)^{-(T-t)}] \\
\xi_{R-t}^0 &= [A_{it}^0 - A_{iT+1}^0 (1+r)^{-(T-t)}]
\end{aligned} \tag{32}$$

where $A_{it}^0 = (1+r)^t A_{i0} + \sum_{j=1}^t (1+r)^{t-j} [Y_{ij} - C_{ij}]$ is the path followed by assets in the event that all realised shocks are zero. Then, substituting equation (32) into equation (29), and again noting only the order of magnitude

for the remainder term,

$$\begin{aligned}
E_{\mathcal{I}} \ln \left(\sum_{j=0}^{R-t-1} \frac{y_{it+j}}{(1+r)^j} + A_{it} - \frac{A_{iT+1}}{(1+r)^{T-t}} \right) \\
= \ln \left[\sum_{j=0}^{R-t-1} \frac{Y_{it+j}}{(1+r)^j} + A_{it}^0 - \frac{A_{iT+1}^0}{(1+r)^{T-t}} \right] \\
+ \pi_{it} \sum_{j=0}^{R-t-1} \alpha_{t+j} [E_{\mathcal{I}} \ln y_{it+j} - \ln Y_{it+j}] \\
+ \frac{E_{\mathcal{I}} \left[A_{it} - \frac{A_{iT+1}}{(1+r)^{T-t}} \right] - \left[A_{it}^0 - \frac{A_{iT+1}^0}{(1+r)^{T-t}} \right]}{\sum_{j=0}^{R-t-1} Y_{it+j}/(1+r)^j + A_{it}^0 - A_{iT+1}^0/(1+r)^{T-t}} \\
+ \mathcal{O} \left(E_{\mathcal{I}} \left\| (\boldsymbol{\nu}_{i0}^{R-1}) \right\|^2 \right) \tag{33}
\end{aligned}$$

where

$$\alpha_{t+j} = \frac{Y_{it+j}/(1+r)^j}{\sum_{k=0}^{R-t-1} Y_{it+k}/(1+r)^k}$$

can be seen as an annuitisation factor for income and

$$\begin{aligned}
\pi_{it} &= 1 - \frac{\exp \xi_N^0}{\sum_{k=0}^N \exp \xi_k^0} \\
&= \frac{\sum_{j=0}^{R-t-1} Y_{it+j}/(1+r)^j}{\sum_{j=0}^{R-t-1} Y_{it+j}/(1+r)^j + A_{it}^0 - A_{iT+1}^0/(1+r)^{T-t}}
\end{aligned}$$

is (roughly) the share of expected future labor income in current human and financial wealth (net of terminal assets) and $\boldsymbol{\nu}_{i0}^{R-1}$ denotes the vector of lifetime income shocks $(\nu'_{i0}, \nu'_{i1}, \dots, \nu'_{iR-1})'$.

We are able to equate the subjects of equations (31) and (33) because the realised budget must balance and $\sum_{j=0}^{R-t} \frac{c_{it+j}}{(1+r)^j}$ and $\sum_{j=0}^{R-t-1} \frac{y_{it+j}}{(1+r)^j} + A_{it} - \frac{A_{iT+1}}{(1+r)^{T-t}}$ therefore have the same distribution. We use (31) and (33), 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.

Noting the income process

$$\Delta \ln y_{it} = \rho \Delta \ln y_{it-1} + [\eta_t - \rho \eta_{t-1}] + \Delta[\omega_t - \rho \omega_{t-1}] + \Delta v_{it} + \Delta[u_{it} - \rho u_{it-1}], \quad (34)$$

this gives²⁰

$$\varepsilon_{it} + \mathcal{O}_p(E_{t-1} \|\boldsymbol{\varepsilon}_{i0}^T\|^2) = \phi_{it} v_{it} + \psi_{it} u_{it} + \pi_{it} \Omega_t + \mathcal{O}_p\left(E_{t-1} \|\boldsymbol{\nu}_{i0}^{R-1}\|^2\right)$$

where

$$\begin{aligned} \phi_{it} &= \pi_{it} \sum_{j=0}^{R-t-1} \alpha_{t+j} \rho^j, \\ \psi_{it} &= \pi_{it} \alpha_t, \end{aligned}$$

the main term on the left hand side is the innovation to the expected present value of consumption and the main terms on the right hand side comprise the innovation to the expected present value of income and

$$\Omega_{it} = \sum_{j=0}^{R-t-1} \alpha_{t+j} (E_t - E_{t-1}) \omega_{t+j},$$

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 $\mathcal{O}_p\left(E_{t-1} \|\boldsymbol{\varepsilon}_{i0}^T\|^2\right)$ are $\mathcal{O}_p\left(E_{t-1} \|\boldsymbol{\nu}_{i0}^{R-1}\|^2\right)$ and thus

$$\varepsilon_{it} = \phi_{it} v_{it} + \psi_{it} u_{it} + \pi_{it} \Omega_t + \mathcal{O}_p\left(E_{t-1} \|\boldsymbol{\nu}_{i0}^{R-1}\|^2\right)$$

and therefore

$$\Delta \ln c_{it} = \Gamma_t + \phi_{it} v_{it} + \psi_{it} u_{it} + \pi_{it} \Omega_t + \mathcal{O}_p\left(E_{t-1} \|\boldsymbol{\nu}_{i0}^{R-1}\|^2\right). \quad (35)$$

²⁰We use the fact that, by Chebyshev's inequality, terms which are $\mathcal{O}\left(\|\nu_{it}\|^2\right)$ and $\mathcal{O}\left((E_t - E_{t-1}) \|\boldsymbol{\nu}_{it-1}^{R-1}\|^2\right)$ are $\mathcal{O}_p\left(E_{t-1} \|\boldsymbol{\nu}_{it-1}^{R-1}\|^2\right)$ and terms which are $\mathcal{O}\left(\|\varepsilon_{it}\|^2\right)$ and $\mathcal{O}\left((E_t - E_{t-1}) \|\boldsymbol{\varepsilon}_{it-1}^T\|^2\right)$ are $\mathcal{O}_p\left(E_{t-1} \|\boldsymbol{\varepsilon}_{it-1}^T\|^2\right)$.

A.1.3 Cross Section Variances

We assume that the variances of the shocks v_{it} and u_{it} are the same in any period for all individuals in any cohort, that shocks are uncorrelated across individuals and that the cross-sectional covariances of the shocks with previous periods' incomes are zero.

Using equation (35) and the equation driving the income process (34) and noting terms that are common within a cohort, the growth in the cross-section variance and covariances of income and consumption can now be seen to take the form

$$\begin{aligned}
\Delta \text{Var}(\ln y_t) &= \rho^2 \Delta \text{Var}(\ln y_{t-1}) + \Delta \text{Var}(u_t) - \rho^2 \Delta \text{Var}(u_{t-1}) + \Delta \text{Var}(v_t) \\
\Delta \text{Var}(\ln c_t) &= (\bar{\phi}_t^2 + \text{Var}(\phi_t)) \text{Var}(v_t) + (\bar{\psi}_t^2 + \text{Var}(\psi_t)) \text{Var}(u_t) \\
&\quad + \text{Var}(\pi_t) \Omega_t^2 + 2 \text{Cov}(\pi_t, \ln c_0) \Omega_t + \mathcal{O}(E_{t-1} \|\boldsymbol{\nu}_{i0}^{R-1}\|^3) \\
\Delta \text{Cov}(\ln c_t, \ln y_t) &= (\rho - 1) \text{Cov}(\ln c_{t-1}, \ln y_{t-1}) + \bar{\phi}_t \text{Var}(v_t) + \bar{\psi}_t \text{Var}(u_t) - \rho \bar{\psi}_{t-1} \text{Var}(u_{t-1}) \\
&\quad + \text{Cov}(\pi_t, \ln y_0) \Omega_t - \rho \text{Cov}(\pi_{t-1}, \ln y_0) \Omega_{t-1} + \mathcal{O}(E_{t-1} \|\boldsymbol{\nu}_{i0}^{R-1}\|^3)
\end{aligned}$$

using the formula of Goodman (1960) for variance of a product of uncorrelated variables.

References

- [1] Abowd, J. M. and D. Card (1989), “On the covariance structure of earnings and hours changes”, *Econometrica*, **57**, 411-45.
- [2] Aiyagari, R. (1994), “Uninsured risk and aggregate saving”, *Quarterly Journal of Economics*, **109**, 659-84
- [3] Altug, S., and R. Miller (1990), “Household choices in equilibrium”, *Econometrica*, **58**, 543-70.
- [4] Atkinson, A. (1999), ‘The distribution of income in the UK and OECD countries in the twentieth century’, *Oxford Review of Economic Policy*, vol. 15, no. 4, pp. 56–75.
- [5] Attanasio, O., and S. Davis (1996), “Relative wage movements and the distribution of consumption”, *Journal of Political Economy*, **104**, 1227-62.
- [6] Attanasio, O., and G. Weber (1995), “Is consumption growth consistent with intertemporal optimization: Evidence from the Consumer Expenditure Survey”, *Journal of Political Economy*, **103**, 1121-57.
- [7] Baker, M. (1997), “Growth-rate Heterogeneity and the Covariance Structure of Life-Cycle Earnings”, *Journal of Labor Economics*, 15, 338-375.
- [8] Blundell, R., L. Pistaferri and I. Preston (2008), “Consumption Inequality and Partial Insurance”, *American Economic Review*, **98**, 1887-1921.
- [9] Blundell, R., and I. Preston (1995), ‘Income, Expenditure and the Living Standards of UK Households’ *Fiscal Studies*, Vol. 16, No.3, 40-54, 1995.
- [10] Blundell, R., and I. Preston (1998), “Consumption inequality and income uncertainty”, *Quarterly Journal of Economics* **113**, 603-640.
- [11] Brewer, M., A. Goodman, A. Muriel and L. Sibieta (2007) “Poverty and inequality in the UK”, IFS Briefing Notes, BN7, March http://www.ifs.org.uk/publications.php?publication_id=3932

- [12] Browning, M. and T. F. Crossley (2009) “Shocks, Stocks and Socks: Smoothing Consumption over a Temporary Income Loss”, forthcoming *Journal of the European Economic Association*.
- [13] Buchinsky, M. and J. Hunt (1999), “Wage mobility in the United States” *The Review of Economics and Statistics* 81(3): 351 – 368
- [14] Burkhauser, R.V. and J. G. Poupore (1997) “A cross-national comparison of permanent inequality in the United States and Germany” *The Review of Economics and Statistics*, 79(1): 10-17
- [15] Campbell, J. Y., (1993) “Intertemporal asset pricing without consumption data”, *American Economic Review*, **83**, 487-512.
- [16] Carroll, C. (1997), “Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis.” *Quarterly Journal of Economics*, 107(1):1-56.
- [17] Cutler, D., and L. Katz (1992), “Rising inequality? Changes in the distribution of income and consumption in the 1980s”, *American Economic Review*, **82**, 546-51.
- [18] Deaton, A., and C. Paxson (1994), “Intertemporal choice and inequality”, *Journal of Political Economy*, **102**, 384-94.
- [19] Goodman, L. A., (1960) “On the exact variance of products”, *Journal of the American Statistical Association*, **55**, 708-713.
- [20] Gosling, A., and S. Machin (1995) “Trade Unions and the Dispersion of Earnings in British Establishments, 1980-90”, *Oxford Bulletin of Economics and Statistics*, vol. 57(2), pages 167-84, May.
- [21] Gosling, A., S. Machin, and C. Meghir (2000), “The Changing Distribution of Male Wages in the U.K”, *Review of Economic Studies*, vol. 67(4), pages 635-66, October.
- [22] Gottschalk, P and R. Moffitt (1994), “The Growth of Earnings Instability in the U.S. Labor Market” *Brookings Papers on Economic Activity*, Vol. 1994, No. 2, pp. 217-272.

- [23] Gourinchas, P.-O. and Parker, J.A. (2002) “Consumption Over the Life Cycle” *Econometrica*, 70(1): 47-89
- [24] Guvenen, F. (2007), “Learning your Earning: Are Labor Income Shocks Really Very Persistent?”, *American Economic Review*, Volume 97, Number 3, June, 687-712.
- [25] Hall, R., and F. Mishkin (1982), “The sensitivity of consumption to transitory income: Estimates from panel data of households”, *Econometrica*, **50**, 261-81.
- [26] Haider, S.J. (2001), “Earnings Instability and Earnings Inequality in the United States, 1967-1991” *Journal of Labor Economics*, 19(4):799-836.
- [27] Haider, S. J. and G. Solon (2006) “Life-Cycle Variation in the Association between Current and Lifetime Earnings”, *American Economic Review*, 96(4):1308-1320.
- [28] Heathcote, J., K. Storesletten and G. Violante (2008), “Insurance and Opportunities: A Welfare Analysis of Labor Market Risk”, *Journal of Monetary Economics*, forthcoming
- [29] Johnson, D.S., T.M. Smeeding and B.R.Torrey (2005), United States Inequality through the Prisms of Income and Consumption’, *Monthly Labor Review*, 11-24, April.
- [30] Krueger, D., and F. Perri (2002), “Does income inequality lead to consumption inequality? Some empirical findings and a theoretical explanation”, NBER Working Paper 9202.
- [31] Low, H., C. Meghir and L. Pistaferri (2006) “Wage Risk and Employment Risk over the Life Cycle”, Institute for Fiscal Studies Working Paper W06/27.
- [32] Mace, B. (1991), “Full insurance in the presence of aggregate uncertainty”, *Journal of Political Economy*, **99**, 928-56.
- [33] MaCurdy, T. E. (1982), “The use of time series processes to model the error structure of earnings in a longitudinal data analysis”, *Journal of Econometrics*, **18**, 83-114.

- [34] Mann, H. B. and Wald, A., (1943) “On stochastic limit and order relationships”, *The Annals of Mathematical Statistics*, **14**, 217-226.
- [35] Meghir, C. and L. Pistaferri (2004), “Income variance dynamics and heterogeneity”, *Econometrica*, **72**, 1-32.
- [36] Moffitt, R. and P. Gottschalk (1994), “Trends in the autocovariance structure of earnings in the US: 1969-1987”, Brown University, *mimeo*.
- [37] Moffitt, R. and P. Gottschalk (2002), “Trends in the Transitory Variance of Earnings in the United States”, *Economic Journal*, vol. 112(478), C68-C73
- [38] Saez, E. & Piketty, T. (2003). Income inequality in the United States: 1913-1998. *Quarterly Journal of Economics*, 118(1), 1-39.
- [39] Storesletten, K., C. Telmer and A. Yaron (2004), “Consumption and Risk Sharing over the Life Cycle”. *Journal of Monetary Economics*, vol. 51, pp. 609-633.

Table 3: **Minimum Distance Estimates of ϕ and ρ**

True ρ	Estimate ϕ	Estimate ρ
1.0	0.850 (0.049)	1.001 (0.002)
0.95	0.550 (0.038)	0.955 (0.005)
0.90	0.347 (0.041)	0.898 (0.015)

Figure 1: Transmission parameter: ϕ_t




Figure 2: Permanent Variance: Base case




Figure 3: Permanent Variance: Effect of discount rate and EIS




Figure 4: Permanent Variance: Effect of consumption and income growth heterogeneity and $\rho < 1$




Figure 5: Permanent Variance: Effect of liquidity constraints and social security




Figure 6: Change in Transitory Variance: MDE π




Figure 7: Variances and Change in Variances: UK 1979-1997




Figure 8: Permanent Variance and Change in Transitory Variance: UK 1979-1997




Figure 9: Variances by Cohort: UK 1979-1997




Figure 10: Change in Variances by Cohort: UK 1979-1997




Figure 11: Permanent Variance by Cohort: UK 1979-1997




Figure 12: Change in Transitory Variance by Cohort: UK 1979-1997

