

Decomposing Changes in Income Risk using Consumption Data*

Richard Blundell[†] Hamish Low[‡] Ian Preston[†]

July 2008

Abstract

This paper concerns the decomposition of income risk into permanent and transitory components using repeated cross-section data on income and consumption. Our focus is on the detection of changes in the magnitudes of variances of permanent and transitory risks. A new approximation to the optimal consumption growth rule is developed. Evidence from a dynamic stochastic simulation is used to show that this approximation can provide a robust method for decomposing income risk in a nonstationary environment. We examine robustness to unobserved heterogeneity in consumption growth and to unobserved heterogeneity in income growth. We use this approach to investigate the growth in income inequality in the UK in the 1980s.

JEL: C30, D52, D91.

Keywords: income risk, inequality, approximation methods, consumption

*We thank Orazio Attanasio, Simon Lee and Luigi Pistaferri for helpful comments. The paper is part of the program of research of the ESRC Centre for the Microeconomic Analysis of Public Policy at IFS. Financial support from the ESRC is gratefully acknowledged. All errors are ours.

[†]University College London and Institute for Fiscal Studies.

[‡]University of Cambridge and Institute for Fiscal Studies.

1 Introduction

The episodes of growth in income inequality over the last twenty five years in many modern economies have been well documented. Attention has focussed recently on the extent to which understanding this growth in inequality requires understanding how the distribution of permanent and transitory shocks to individual income processes have changed¹. The distinction between permanent and transitory risk is crucial to understanding the welfare consequences of income risk². Cross-sectional income surveys alone are of no help in getting at this distinction. However, the combination of consumption and income data can reveal much more³. The aim of this paper is to highlight the value of using repeated cross-section data on income and consumption in decomposing income variability into permanent and transitory shocks and to use this data to investigate the inequality boom of the 1980s in the UK. Our emphasis is on the detection of changes in the magnitudes of variances of permanent and transitory risks using consumption and income data.

Typically panel data surveys on consumption and income are unavailable but repeated cross-section household expenditure surveys that contain measurements on consumption and income are commonly available in many economies and over long periods of time⁴. 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 (CEX) has been available since 1980 (see Johnson, Smeeding and Torrey, 2005) and there are many other examples from other countries.

This paper makes three main contributions. It first examines assumptions

¹See Burkhauser and Poupore, 1997; Buchinsky and Hunt, 1999; Moffitt and Gottschalk, 2002; Meghir and Pistaferri, 2004; Storesletten, Telmer and Yaron, 2004.

²See, for example, the discussions in Blundell and Preston, 1998; Heathcote, Storesletten and Violante, forthcoming; and Low, Meghir and Pistaferri, 2006.

³see Blundell and Preston, 1998; Deaton and Paxson, 1994; Krueger and Perri, 2002.

⁴If panel data is available then alternative approaches become feasible, as explored for example in Blundell, Pistaferri and Preston, 2004

on intertemporal consumption choices under which repeated cross-section data can be used to identify the distribution of uninsured transitory and permanent shocks to income. Second, it assesses the accuracy with which components of income risk can be identified using survey data on income and consumption. In doing so it develops an approximation to the optimal consumption growth rule under CRRA preferences and, using a dynamic simulation, shows that this approximation can provide an accurate method for decomposing income risk. In particular, the approximation is found to separate accurately the variances of the permanent and transitory components of idiosyncratic uninsured shocks to income. Although we allow for common shocks our focus is on the identification of idiosyncratic risk. Finally, we use this methodology to unravel the persistence of the underlying shocks to income during the UK inequality boom of the 1980s.

The simplest approximation we examine is one in which individuals are unable to self-insure against permanent shocks but are able to insure fully against transitory shocks. This approximation, developed in Blundell and Preston (1998), implies that the cross-section variance of consumption will reflect only accumulated permanent shocks to income and further, the amount by which the cross-section variance of income exceeds the variance of consumption can be attributed to growth in the transitory variance. At a theoretical level, we show the order of the error of this approximation. Through simulation of individuals choosing consumption in an economy with permanent and transitory income shocks, we show that the approximation decomposes the income risk fairly accurately and correctly identifies changes in risk over time.

This initial approximation ignores precautionary saving and other forms of insurance. In particular, an error in the approximation arises through underpredicting self-insurance against permanent shocks and overpredicting self-insurance against transitory shocks. This error is shown to imply an underestimate of the actual risk to permanent income and an overestimate of the change in the variance of transitory shocks. We show that this er-

ror can be reduced through a more sophisticated approximation that uses information on the extent of self-insurance. Such information on the extent of insurance can come from asset data directly, if it is available. We develop an alternative strategy which estimates the degree of insurance using the covariance between consumption and income and the variance of consumption. This latter estimate of insurance requires that expenditure and income data are available in the same survey. The advantage of measuring insurance in this way is that, in principle, all forms of insurance against income shocks are captured, whereas using asset data only identifies insurance through precautionary saving. We evaluate both these new approximations and show important improvements in the accuracy of recovering the income risk decompositions.

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

The layout of the rest of the paper is as follows. In section 2 we derive the approximations which relate consumption inequality to income risk. The usefulness of having consumption and income data in the same survey and having measurements on financial wealth is explored in detail. In section 3 we develop an approach for idiosyncratic trends in consumption and income and also discuss the robustness of our approximation to liquidity constraints and heterogeneity in discount rates. Section 4 describes the environment we simulate and reports the results of our Monte Carlo experiments. Section 5 presents new estimates of the decomposition of income risk for Britain from

the inequality boom of the 1980s. 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

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

where Y_{it} represents the permanent component of income and u_{it} the transitory shock in period t . The final $T - R + 1$ periods of life are spent in mandatory retirement with no labour income.

The permanent component is assumed to follow a process

$$\Delta \ln Y_{it} = \eta_t + \omega_t + v_{it}$$

where η_t is a deterministic trend and ω_t a stochastic term, both common to the members of the cohort, while v_{it} is a permanent idiosyncratic shock⁵. The process for income can therefore be written

$$\Delta \ln y_{it} = \eta_t + \omega_t + \Delta u_{it} + v_{it}. \quad (1)$$

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 .

⁵In section 3, we consider extending analysis to the case where the income trend is individual specific and we show how the approximation can be used in the presence of this heterogeneity in income growth.

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

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

This is a popular specification compatible with an MA(1) process for idiosyncratic changes in log income⁶. We make no assumptions about the time series properties⁷ of the common shocks ω_t .

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 permanent shocks. These are the idiosyncratic components of permanent and transitory risk facing individuals.

Assuming the cross-sectional covariances of the shocks with previous periods' incomes to be zero, then

$$\Delta\text{Var}(\ln y_t) = \text{Var}(v_t) + \Delta\text{Var}(u_t) \tag{2}$$

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.

⁶See 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.

⁷Although we attach the common shock ω_t to the equation for permanent income, the lack of specificity about its time series properties means we should refrain from thinking of it as specifically permanent or transitory in nature.

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 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

$$E_t \left[\sum_{\tau=0}^{T-t} \frac{U(c_{it+\tau})}{(1+\delta)^\tau} \right]$$

where δ is a subjective discount factor, 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

$$\begin{aligned} U'(c_{it+\tau}) &= \lambda_{it+\tau} \\ E_t \lambda_{it+\tau} &= \left(\frac{1+\delta}{1+r} \right)^\tau \lambda_{it}, \quad \tau = 0, 1, \dots, T-t \end{aligned} \quad (4)$$

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

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

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,

$$\Delta \text{Var}(\ln c_t) = \text{Var}(\varepsilon_t) + \mathcal{O}(E_{t-1} |\varepsilon_{it}|^3) \quad (5)$$

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

$$\varepsilon_{it} = \pi_{it}(v_{it} + \alpha_t u_{it}) + \pi_{it} \Omega_t + \mathcal{O}_p \left(E_{t-1} \|\boldsymbol{\nu}_{it}^{R-1}\|^2 \right) \quad (6)$$

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 two additional parameters are introduced

- α_t : an annuitisation factor, common within a cohort, capturing the importance of transitory shocks to lifetime wealth relative to permanent shocks.

- π_{it} : a self-insurance factor capturing the significance of asset holdings as a component of current human and financial wealth.

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.

Let $\bar{\pi}_t$ and $\text{Var}_t(\pi_t)$ be the cross section mean and variance of π_{it} . Then 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 $\Delta \ln y_{it} = \eta_t + \omega_t + \Delta u_{it} + v_{it}$, then*

$$\begin{aligned}
\Delta \text{Var}(\ln y_t) &= \text{Var}(v_t) + \Delta \text{Var}(u_t) \\
\Delta \text{Var}(\ln c_t) &= (\bar{\pi}_t^2 + \text{Var}(\pi_t))\text{Var}(v_t) + (\bar{\pi}_t^2 + \text{Var}(\pi_t))\alpha_t^2 \text{Var}(u_t) \\
&\quad + \text{Var}(\pi_t)\Omega_t^2 + 2\text{Cov}(\pi_t, c_{t-1})\Omega_t + \mathcal{O}(E_{t-1}\|\nu_{it}\|^3) \\
\Delta \text{Cov}(\ln c_t, \ln y_t) &= \bar{\pi}_t \text{Var}(v_t) + \Delta[\bar{\pi}_t \alpha_t \text{Var}(u_t)] \\
&\quad + \mathcal{O}(E_{t-1}\|\nu_{it}\|^3).
\end{aligned} \tag{7}$$

Proof: See Appendix A1.

Taking income inequality together with consumption inequality and sufficient information on α_t and the distribution of π_{it} we are able to use the life-cycle model to separate the permanent income risk from the growth in transitory uncertainty.

From these expressions 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 π_{it} and α_t .

1. Particularly simple forms follow by allowing $\bar{\pi}_t \simeq 1$, $\text{Var}(\pi_t) \simeq 0$ and

$\alpha_t \simeq 0$, implying no self-insurance and a long horizon. Specifically

$$\begin{aligned}\Delta\text{Var}(\ln y_t) &= \text{Var}(v_t) + \Delta\text{Var}(u_t) \\ \Delta\text{Var}(\ln c_t) &\simeq \text{Var}(v_t) \\ \Delta\text{Cov}(\ln c_t, \ln y_t) &\simeq \text{Var}(v_t)\end{aligned}\tag{8}$$

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 (7). The evolution of the covariance should follow that of the consumption variance and this provides one testable overidentifying restriction per period of the data.

2. Relaxing the assumption that $\bar{\pi}_t \simeq 1$ but keeping $\text{Var}(\pi_t) \simeq 0$ and $\alpha_t \simeq 0$ implies

$$\begin{aligned}\Delta\text{Var}(\ln y_t) &= \text{Var}(v_t) + \Delta\text{Var}(u_t) \\ \Delta\text{Var}(\ln c_t) &\simeq \bar{\pi}_t^2 \text{Var}(v_t) \\ \Delta\text{Cov}(\ln c_t, \ln y_t) &\simeq \bar{\pi}_t \text{Var}(v_t).\end{aligned}\tag{9}$$

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

- With extraneous information on assets and incomes and assumptions about income growth, estimates of $\bar{\pi}_t$ 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, $\bar{\pi}_t$ could be estimated simultaneously with the variances of the shocks by, say,

minimum distance methods. In principle, sufficient degrees of freedom exist to estimate $\bar{\pi}_t$ separately for each period; in practice, it would make sense to impose some degree of smoothness on the path of $\bar{\pi}_t$ over time, for example by estimating a suitable parametric time path, thereby retaining some degrees of freedom for testing.

3. With sufficient information on the distribution of assets we could calculate $\bar{\pi}_t$, $\text{Var}(\pi_t)$ and $\text{Cov}(\pi_t, c_{t-1})$. If individuals differ in the insurance parameter π_{it} then common income shocks create heterogeneous consumption shocks so that it would be possible in principle, knowing $\text{Var}(\pi_t)$ and $\text{Cov}(\pi_t, c_{t-1})$, to recover estimates of common shocks Ω_t . In practice, information allowing us to enter a sensible value for the necessary moments of π_t is unlikely to be available.
4. With precise measurement of interest rates we could also allow $\alpha_t \neq 0$ and make full use of all terms in (7). In practice, the evolution of α_t is so gradual that identification using its changes over time would be tenuous and we do not discuss this further.

Cross section variances and covariances of log income and consumption 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 (7) after choosing or estimating values for $\bar{\pi}_t$, $\text{Var}(\pi_t)$ and α_t , 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.

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. Allowing for such heterogeneity by permitting for heterogeneous consumption trends Γ_{it}

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

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

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

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 $\bar{\pi}_t$ or we can only use our simplest approximation assuming $\bar{\pi}_t = 1$.

Income Trends: Individuals also differ in their expectations about income growth, particularly across occupations and across education groups. For example, Baker (1997) and Haider (2001) argue for the importance of

heterogeneity in income trends. 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

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

the equations for the evolution of the variances become

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

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 $\bar{\pi}_t$. 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, which is unlikely to hold in practice. 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 $\bar{\pi}_t$) and the permanent variance can be distinguished from expected variability.

Guvenen (2007) argues 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.

4 Monte-Carlo

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 Monte Carlo exercise is to examine the accuracy with which changes to the underlying structural variances can be recovered. To do this, we simulate the consumption behaviour of individuals in a life-cycle model allowing for heterogeneity in π_t and under a range of assumptions about discounting, risk aversion, liquidity constraints and the income process.

The specific Monte Carlo 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, especially in the US (see, for example, Moffitt and Gottschalk, 2002; Meghir and Pistaferri, 2004; Blundell, Pistaferri and Preston, 2005). 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 Numerical Model

Utility is given by the additive, constant elasticity of substitution form,

$$E_t \left[\sum_{\tau=0}^{T-t} \frac{1}{(1 + \delta_i)^\tau} \left(\frac{\gamma_i}{1 + \gamma_i} c_{it}^{1+1/\gamma_i} \right) \right]. \quad (10)$$

When we allow for preference heterogeneity, this enters through δ_i and through γ_i .

Transitory and permanent shocks to income are assumed to be log-normally

distributed,⁸

$$\begin{aligned}
\ln y_{it} &= \ln Y_{it} + u_{it}, & u_{it} &\sim \mathcal{I}\mathcal{N}(0, \sigma_{u_t}^2) \\
\ln Y_{it} &= \eta_{it} + \omega_t + \ln Y_{it-1} + v_{it}, & v_{it} &\sim \mathcal{I}\mathcal{N}(0, \sigma_{v_t}^2) \\
\omega_t &= -\frac{\sigma_{v_t}^2 + \Delta\sigma_{u_t}^2}{2}.
\end{aligned} \tag{11}$$

When we allow for heterogeneity in 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 .

$$\begin{array}{c}
\sigma_{v,L}^2 \quad \sigma_{v,H}^2 \\
\begin{array}{|c|}
\hline
\begin{array}{cc}
\sigma_{v,L}^2 & \begin{array}{cc} 1 - \beta & \beta \end{array} \\
\sigma_{v,H}^2 & \begin{array}{cc} \beta & 1 - \beta \end{array}
\end{array} \\
\hline
\end{array}
\end{array} \tag{12}$$

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.⁹

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. 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 .

⁸In 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.

⁹In solving the model for a particular individual, it is irrelevant whether a particular shock is idiosyncratic or common because the model is partial equilibrium.

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:

$$A_{it} \geq 0 \tag{13}$$

We set $T = 70$, with the last 10 years of life spent in mandatory retirement. Individuals can also use asset holdings to increase consumption in retirement. Parameters used in the baseline are summarised in table 1.

Table 1: **Baseline Parameter Values**

Description		Value
Discount Rate	δ	0.02
EIS	γ	-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
Switching Period	S	40
Retirement Age	R	60
Terminal Period	T	70

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.

Baseline parameter values are recorded in Table 1. The way in which these parameters are varied across experiments is described in Table 2. A first block of experiments considers the effect of higher and lower values for the discount rate, EIS and income growth, maintaining in each case similar values for all individuals. A second block of experiments then allows for cross-sectional heterogeneity in the values of these parameters to allow for idiosyncratic trends as discussed earlier. Finally three further experiments consider further modifications, setting the growth in transitory variance to zero, reducing the number of retirement years to discourage asset accumulation and increasing the probability of 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 (8), 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 partially insure 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 $\pi=1$.

We can improve on this simplest approximation by allowing for these insurance mechanisms. We do this in two ways. If we have information on asset holdings, then the approximation can be corrected to take account of the amount of self-insurance through saving and we would not expect differences in the utility cost of saving to affect the accuracy of the corrected estimates. The quality of the correction depends on the quality of information about assets. In the correction considered here, based on (9), we use an estimate of $\bar{\pi}_t$ using sample median values of assets and incomes assuming no anticipated growth in log incomes and known r . We label such estimates *Asset-based* π .

The final approximation estimates $\bar{\pi}_t$, and hence the amount of insurance,

Table 2: **Experiment Parameter Values**

Description	δ	γ	η	$\Delta\sigma_{u_t}^2$	T
Baseline	0.02	-0.67	0.0	0.1	70
High discount rate	0.04	-0.67	0.0	0.1	70
Low discount rate	0.01	-0.67	0.0	0.1	70
High EIS	0.02	-2.00	0.0	0.1	70
Low EIS	0.02	-0.20	0.0	0.1	70
High income growth	0.02	-0.67	0.02	0.1	70
Low income growth	0.02	-0.67	-0.01	0.1	70
Hetero discount rate	$\left\{ \begin{array}{c} 0.01 \\ 0.02 \\ 0.04 \end{array} \right\}$	-0.67	0.0	0.1	70
Hetero EIS	0.02	$\left\{ \begin{array}{c} -0.20 \\ -0.67 \\ -2.00 \end{array} \right\}$	0.0	0.1	70
Hetero income growth	0.02	-0.67	$\left\{ \begin{array}{c} -0.01 \\ 0.0 \\ 0.02 \end{array} \right\}$	0.1	70
Early hetero income growth	0.02	-0.67	$\left\{ \begin{array}{c} -0.01 \\ 0.0 \\ 0.02 \end{array} \right\}$	0.1	70
No transitory variance growth	0.02	-0.67	0.0	0.0	70
Liquidity constrained	0.02	-0.67	0.0	0.1	62
Social security	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.

jointly with the variances of the shocks by minimum distance assuming a linear path over time.¹⁰ We label such estimates *MDE* π .

In each case the moments (7) 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.

4.2 Results

4.2.1 Self insurance

Crucial to the approximations (7) are the means and variances of the self insurance parameters π_{it} . In Figure 1 we show the values of $\bar{\pi}_t$ 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 π_{it} and not the approximations used in estimation.

The baseline case gives a $\bar{\pi}_t$ 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. High income growth reduces the need for saving since individuals do not want to accumulate savings and move resources into the future when income is high. All of these cases therefore involve diminished self insurance and raise $\bar{\pi}_t$. Lower values of discount rates, EIS or income growth on the other hand all reduce $\bar{\pi}_t$.

¹⁰As compared to the estimate relying on the use of asset data, this estimate of $\bar{\pi}_t$ would in practice capture any type of insurance although there is, of course, only self-insurance in the actual simulations.

The experiments with heterogeneity in these parameters give similar mean values of $\bar{\pi}_t$ to the baseline case. Eliminating transitory variance growth raises $\bar{\pi}_t$ but not by very much.

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.

The final experiment considered introduces a social security pension equal to half of final income. Incentives to accumulate private assets for consumption in retirement are reduced. Moreover, in this case self insurance against permanent shocks is less effective for any future income path and 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 (6) unless π_{it} is modified to account for this fact¹¹. The values for $\bar{\pi}_t$ used in this case incorporate such a modification and are substantially higher, particularly at older ages, than in the base case.

The approximations (7) also include terms involving $\text{Var}(\pi_t)$ and $\text{Cov}(\pi_t, c_{t-1})$ which are neglected in the estimation methods applied below because of the likely absence in practice of any reliable method of estimating heterogeneity in π_{it} . It is nonetheless possible in the simulations to calculate the true variance of π_{it} to check on the likely magnitude of the omitted terms. Figure 2 shows the squared coefficient of variation in π_{it} for each of the simulations. The heterogeneity is very small at age 30 but grows as shocks to income accumulate and the variance of asset holdings grows. Even at age 50 however the variance is below ten per cent of $\bar{\pi}_t^2$ in all scenarios. As might be expected, the heterogeneity in π_{it} is greater in the cases where asset accumulation is greater but also in the cases where preference parameters or the

¹¹The 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.

income process are heterogeneous.

4.2.2 Estimating the permanent variance

Baseline simulations: Figure 3 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 approximations.

The estimates using $\bar{\pi}_t = 1$ consistently underestimate the permanent variance. This is because asset holdings enable partial self-insurance against the permanent shocks.¹² The cross-section variance of consumption reflects the uninsured part of the permanent shock and this is an underestimate of 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 whether we use sample medians of assets and income, or whether we estimate $\bar{\pi}_t$ alongside the variances. The advantage of the latter correction is that we can correct for self-insurance without relying on asset data and the evidence of this simulation is that the estimates using MDE-based $\bar{\pi}_t$ actually perform slightly better anyway than those based on calculations from median incomes and assets.

Sensitivity to discount rates, EIS and income growth: Our first sensitivity analysis concerns the sensitivity of the accuracy of our approximations to the rate of consumption growth and the rate of income growth, maintaining the assumption that growth rates are homogenous across individuals.

¹²This partial insurance against permanent shocks would not be feasible in an infinite horizon setting.

Figures 4, 5 and 6 show estimates of the permanent variance as the discount rate, EIS and income growth rates are varied. The estimates which make no correction for self insurance are least accurate in those scenarios which encourage asset accumulation since these are the cases in which $\bar{\pi}_t$ is furthest from 1. On the other hand, in all scenarios, correcting for self-insurance by estimating π through minimum distance leads to very accurate estimates of the permanent variance.

Using median assets to calculate $\bar{\pi}_t$ also secures a considerable improvement over assuming $\bar{\pi}_t = 1$ though as in the baseline case there is a slight tendency to overpredict the permanent variance, particularly later in life.

Heterogeneity in consumption and income growth: One potentially important limitation of our results so far is the assumption that all individuals are ex ante identical. By contrast, the focus of the next set of experiments is the implications of individuals being heterogeneous. In Figure 7 we explore the implications of heterogeneity in consumption growth induced by heterogeneity in discount rates and heterogeneity in the elasticity of intertemporal substitution whereas in 7 we explore the implications of heterogeneity in deterministic income growth rates.

As discussed in section 3, 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. Estimates of the variance of permanent variance can still be obtained by dropping the equation that exploits this relationship and using the information in the income-consumption covariance. Because of the reduction in number of moments we no longer have the degrees of freedom required to estimate $\bar{\pi}_t$ within the minimum distance calculation so a value either needs to be imposed or calculated, say, from asset data.

In Figure 7 we report two estimates that ignore the problem of heterogeneity - those using asset-based and MDE $\bar{\pi}_t$ as in other exercises - but add an estimate, labelled as “robust”, which drops the contaminated moment

condition and uses a value of $\bar{\pi}_t$ based on asset data.

The accuracy of the estimates using median assets to correct for self-insurance seems to be reduced somewhat by the presence of heterogeneity and dropping the moments using the variance of consumption does move the estimated permanent variance closer to the truth. On the whole, the most accurate estimate seems, despite the heterogeneity, to be those based on full minimum distance but without correcting for heterogeneity.

When we have heterogeneity in income growth, 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 $\bar{\pi}_t$. As with heterogeneity in consumption growth, we report, in Figure 8, estimates using asset-based and MDE $\bar{\pi}_t$ using all moments, including those no longer valid, and a third estimate correcting for self insurance with the asset based $\bar{\pi}_t$ and using only the valid moments.

We consider two separate types of heterogeneity in income growth. First, we consider heterogeneity which persists across the whole life-cycle, as in Guvenen (2007). Using median asset holdings to calculate $\bar{\pi}_t$ leads to an overestimate of the permanent variance. Similarly estimating $\bar{\pi}_t$ by minimum distance also 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.

Second, 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). When the heterogeneity is present only early in the life-cycle, 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.

Liquidity Constraints: 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 9 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 overestimate 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.

Social security: 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 and we modify the calculation of asset based $\bar{\pi}_t$ appropriately. With this modification the permanent variance is picked up fairly accurately by either of the methods allowing for self insurance as shown in Figure 9.

4.2.3 Estimating changes in the transitory variance

Estimates of the change in the transitory variance for all of the experiments discussed are shown in Figure 10, in all cases using MDE to estimate the self insurance parameter. In all cases discussed so far the growth in the

transitory variance is picked up with a high degree of accuracy. One scenario is added to this picture - a simulation in which the growth in the transitory variance is turned off. As is evident, the difference is clearly detected.

4.2.4 Overidentifying Restrictions

Table 3 reports mean values of the χ^2 tests of the overidentifying restrictions calculated with each set of estimates and the frequency of rejection at the 5% level. Across all experiments, tests of the restrictions with $\bar{\pi}_t = 1$ 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. When the correction is by using minimum distance to estimate $\bar{\pi}_t$, the distribution of the overidentification tests appears very close to the appropriate χ^2_{17} distribution with a mean close to degrees of freedom and size typically close to 5% (slightly overrejecting).

5 Decomposing Income Risk in the Inequality Boom

We now turn to apply the ideas and techniques outlined above to data from the Family Expenditure Survey. This is an annual continuous cross-sectional budget survey with detailed data on incomes and expenditures of UK households. The period chosen for the application is that from 1979 to 1992. This is the ‘inequality boom’ period in the UK in which there was rapid growth in income inequality, see Atkinson (1997), for example. 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 (2005).

Table 3: Tests of Overidentifying Restrictions

Description	χ^2_{19}	Rejection rate	χ^2_{19}	Rejection rate	χ^2_{17}	Rejection rate
Baseline	124.94	1.00	23.17	0.16	17.53	0.06
High discount rate	52.57	0.97	20.64	0.09	17.72	0.07
Low discount rate	161.76	1.00	24.7	0.23	17.52	0.06
High EIS	67.27	1.00	20.00	0.07	18.44	0.09
Low EIS	177.89	1.00	51.39	0.96	17.65	0.07
High income growth	38.97	0.77	20.52	0.09	18.86	0.09
Low income growth	181.09	1.00	24.74	0.24	18.34	0.09
Hetero discount rate	139.70	1.00	26.45	0.29	17.44	0.06
Hetero EIS	160.79	1.00	38.04	0.75	17.17	0.05
Hetero income growth	318.02	1.00	58.00	0.99	19.04	0.10
Early hetero income growth	122.67	1.00	22.96	0.15	17.75	0.06
No transitory variance growth	68.96	1.00	33.93	0.61	18.34	0.08
Liquidity constrained	39.38	0.81	27.41	0.33	17.74	0.06
Social security	79.48	1.00	24.72	0.24	18.42	0.09

The income measure used is equivalised household income after housing costs.¹³ Expenditure is equivalised household expenditure on nondurables and semi-durables, excluding expenditures on housing. 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.

Figures 11 and 12 show the variances and covariances of income and consumption over the period, pooled and separated by birth cohort. In all pictures we see a continually rising variance of income. This rising path is followed in the earlier years by the variance of consumption and the covariance but these paths flatten off in later years.

In the absence of asset data we present estimates for the variance of the permanent shocks and the changes in the variance of the transitory shocks using the approximation based on minimum distance estimation¹⁴. These are calculated with asymptotically optimal weighting. The estimated variances, smoothed using a fifth order moving average and shown with pointwise 95 per cent confidence bands, are presented in Figure 13.

In Figures 14 and 15 we separate the two cohorts but estimating jointly with a common insurance parameter. The dramatic growth in overall income inequality experienced in the 1980s is evident in the patterns for both birth cohorts. The variance of transitory shocks to income grows throughout the period whereas, for both cohorts, the permanent variance is high in the mid 1980s but then appears to fall back in the later years. This 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

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

¹⁴In all cases we estimate a flat profile for \bar{p}_t very close to unity.

early to mid period of the 1980s.¹⁵

6 Conclusions

Increases in cross-section measures of income inequality may reflect the variance of permanent shocks or increases in the variability of transitory shocks. However, the differing sources of risk have very different implications for welfare. In this paper, we have examined what can be learned about income risk from 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 simple approximations to consumption rules can be used to decompose income variability into its components using such data. In assessing the accuracy of this decomposition we show that it is able to map accurately the evolution of transitory and permanent variances of income shocks across a range of alternative parameterisations.

The usefulness of the approximation was shown clearly in a study of income risk in Britain during the inequality boom of the 1980s. We found that across the two main birth cohorts examined, there was a systematic rise in the variance of transitory shocks to income together with a sharp ‘spike’ in the variance of permanent shocks in the early 1980s.

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 uncer-

¹⁵See Gosling and Machin (1995) and Gosling, Machin and Meghir (2000) and references therein.

tainty from that variability which is due only to this heterogeneity in income paths.

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 long 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. In this direction, Blundell, Pistaferri and Preston (2005) create such a panel by combining the CEX and PSID in the US and establish the identification of additional transmission or ‘insurance’ parameters as well as the separate evolution of permanent and transitory income variances.

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.B. (1997), ‘Bringing Income Distribution in From the Cold’, *The Economic Journal*, March.
- [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 (2005), "Consumption Inequality and Partial Insurance", IFS Working Papers, W04/28, http://www.ucl.ac.uk/~uctp39a/aer_final.pdf, forthcoming *American Economic Review*.
- [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] Buchinsky, M. and J. Hunt (1999), "Wage mobility in the United States" *The Review of Economics and Statistics* 81(3): 351 – 368
- [13] 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
- [14] Campbell, J. Y., (1993) "Intertemporal asset pricing without consumption data", *American Economic Review*, **83**, 487-512.
- [15] Carroll, C. (1997), "Buffer Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics*, 107(1):1-56.
- [16] 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.
- [17] Deaton, A., and C. Paxson (1994), "Intertemporal choice and inequality", *Journal of Political Economy*, **102**, 384-94.
- [18] Goodman, L. A., (1960) "On the exact variance of products", *Journal of the American Statistical Association*, **55**, 708-713.
- [19] 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.

- [20] 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.
- [21] 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.
- [22] Gourinchas, P.-O. and Parker, J.A. (2002) “Consumption Over the Life Cycle” *Econometrica*, 70(1): 47-89
- [23] Guvenen, F. (2007), “Learning your Earning: Are Labor Income Shocks Really Very Persistent?”, *American Economic Review*, Volume 97, Number 3, June, 687-712.
- [24] Hall, R., and F. Mishkin (1982), “The sensitivity of consumption to transitory income: Estimates from panel data of households”, *Econometrica*, **50**, 261-81.
- [25] Haider, S.J. (2001), “Earnings Instability and Earnings Inequality in the United States, 1967-1991” *Journal of Labor Economics*, 19(4):799-836.
- [26] 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.
- [27] Heathcote, J., K. Storesletten and G. Violante (2008), “Insurance and Opportunities: A Welfare Analysis of Labor Market Risk”, *Journal of Monetary Economics*, forthcoming
- [28] 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.
- [29] Krueger, D., and F. Perri (2002), “Does income inequality lead to consumption inequality? Some empirical findings and a theoretical explanation”, NBER Working Paper 9202.
- [30] 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.

- [31] Mace, B. (1991), “Full insurance in the presence of aggregate uncertainty”, *Journal of Political Economy*, **99**, 928-56.
- [32] 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.
- [33] Mann, H. B. and Wald, A., (1943) “On stochastic limit and order relationships”, *The Annals of Mathematical Statistics*, **14**, 217-226.
- [34] Meghir, C. and L. Pistaferri (2004), “Income variance dynamics and heterogeneity”, *Econometrica*, **72**, 1-32.
- [35] Moffitt, R. and P. Gottschalk (1994), “Trends in the autocovariance structure of earnings in the US: 1969-1987”, Brown University, *mimeo*.
- [36] Moffitt, R. and P. Gottschalk (2002), “Trends in the Transitory Variance of Earnings in the United States”, *Economic Journal*, vol. 112(478), C68-C73
- [37] 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.

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 (4)

$$E_t U'(c_{it+1}) = U'(c_{it}) \left(\frac{1 + \delta}{1 + r} \right) = U'(c_{it} e^{\Gamma_{it+1}}) \quad (14)$$

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

$$\begin{aligned} U'(c_{it+1}) &= U'(c_{it} e^{\Gamma_{it+1}}) \left[1 + \frac{1}{\gamma(c_{it} e^{\Gamma_{it+1}})} [\Delta \ln c_{it+1} - \Gamma_{it+1}] \right. \\ &\quad \left. + \frac{1}{2} \beta(\tilde{c}, c_{it} e^{\Gamma_{it+1}}) [\Delta \ln c_{it+1} - \Gamma_{it+1}]^2 \right] \end{aligned} \quad (15)$$

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

$$\begin{aligned} E_t U'(c_{it+1}) &= U'(c_{it} e^{\Gamma_{it+1}}) \left[1 + \frac{1}{\gamma(c_{it} e^{\Gamma_{it+1}})} E_t [\Delta \ln c_{it+1} - \Gamma_{it+1}] \right. \\ &\quad \left. + \frac{1}{2} E_t \{ \beta(\tilde{c}, c_{it} e^{\Gamma_{it+1}}) [\Delta \ln c_{it+1} - \Gamma_{it+1}]^2 \} \right] \end{aligned} \quad (16)$$

Substituting for $E_t U'(c_{it+1})$ from (14),

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

and thus

$$\Delta \ln c_{it+1} = \Gamma_{it+1} - \frac{\gamma(c_{it}e^{\Gamma_{it+1}})}{2} E_t \left\{ \beta(\tilde{c}, c_{it}e^{\Gamma_{it+1}}) [\Delta \ln c_{it+1} e^{\Gamma_{it+1}}]^2 \right\} + \varepsilon_{it+1} \quad (17)$$

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

$$\Delta \ln c_{it+1} = \varepsilon_{it+1} + \Gamma_{it+1} + \mathcal{O}(E_t |\varepsilon_{it+1}|^2). \quad (18)$$

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

$$\Delta \ln c_{it+1} = \varepsilon_{it+1} + \Gamma_{t+1} + \mathcal{O}(E_t |\varepsilon_{it+1}|^2). \quad (19)$$

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).

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

$$\begin{aligned} F(\boldsymbol{\xi}) &= \ln \sum_{j=0}^N \exp \xi_j^0 + \sum_{j=0}^N \frac{\exp \xi_j^0}{\sum_{k=0}^N \exp \xi_k^0} (\xi_j - \xi_j^0) \\ &\quad + \frac{1}{2} \sum_{j=0}^N \sum_{k=0}^N \frac{\partial^2 F(\tilde{\boldsymbol{\xi}})}{\partial \xi_j \partial \xi_k} (\xi_j - \xi_j^0) (\xi_k - \xi_k^0) \end{aligned} \quad (20)$$

where $\tilde{\boldsymbol{\xi}}$ lies between $\boldsymbol{\xi}$ and $\boldsymbol{\xi}^0$ and is used to make the expansion exact. The coefficients in the remainder term are given by

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

where δ_{jk} denotes the Kronecker delta. These coefficients are bounded because $0 < \exp \tilde{\xi}_j / \sum_k \exp \tilde{\xi}_k < 1$.

Hence, taking expectations of (20) subject to information set \mathcal{I}

$$\begin{aligned} E_{\mathcal{I}}[F(\boldsymbol{\xi})] &= \ln \sum_{j=0}^N \exp \xi_j^0 + \sum_{j=0}^N \frac{\exp \xi_j^0}{\sum_{k=0}^N \exp \xi_k^0} (E_{\mathcal{I}} \xi_j - \xi_j^0) \\ &\quad + \frac{1}{2} \sum_{j=0}^N \sum_{k=0}^N E_{\mathcal{I}} \left(\frac{\partial^2 F(\tilde{\boldsymbol{\xi}})}{\partial \xi_j \partial \xi_k} (\xi_j - \xi_j^0) (\xi_k - \xi_k^0) \right) \end{aligned} \quad (21)$$

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

$$\begin{aligned} \xi_j &= \ln c_{it+j} - j \ln(1+r) \\ \xi_j^0 &= E_{t-1} \ln c_{it+j} - j \ln(1+r), \quad i = 0, \dots, T-t. \end{aligned} \quad (22)$$

Then, substituting equation (22) into equation (21) and noting only the order of magnitude for the remainder term,

$$\begin{aligned} E_{\mathcal{I}} \left[\ln \sum_{j=0}^{T-t} \frac{c_{it+j}}{(1+r)^j} \right] &= \ln \sum_{j=0}^{T-t} \exp[E_{t-1} \ln c_{it+j} - j \ln(1+r)] \\ &\quad + \sum_{j=0}^{T-t} \theta_{it+j} [E_{\mathcal{I}} \ln c_{it+j} - E_{t-1} \ln c_{it+j}] \\ &\quad + \mathcal{O}(E_{\mathcal{I}} \|\boldsymbol{\varepsilon}_{it}^T\|^2) \end{aligned} \quad (23)$$

where

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

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

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

We now apply the expansion (21) 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 $N = R - t$ and let

$$\begin{aligned} \xi_j &= \ln y_{it+j} - j \ln(1+r) \\ \xi_j^0 &= E_{t-1} \ln y_{it+j} - j \ln(1+r) \quad j = 0, \dots, R-t-1 \\ \xi_N &= \ln [A_{it} - A_{iT+1} (1+r)^{-(T-t)}] \\ \xi_N^0 &= E_{t-1} \ln [A_{it} - A_{iT+1} (1+r)^{-(T-t)}] \end{aligned} \quad (24)$$

Then, substituting equation (24) into equation (21), and again noting only the order of magnitude for the remainder term,

$$\begin{aligned}
E_{\mathcal{I}} \ln \left(\sum_{j=0}^{R-t-1} \frac{y_{it+j}}{(1+r)^j} + A_{it} - \frac{A_{iT+1}}{(1+r)^{T-t}} \right) \\
&= \ln \left[\sum_{j=0}^{R-t-1} \exp [E_{t-1} \ln y_{it+j} - j \ln(1+r)] + \exp E_{t-1} \ln \left[A_{it} - \frac{A_{iT+1}}{(1+r)^{T-t}} \right] \right] \\
&\quad + \pi_{it} \sum_{j=0}^{R-t-1} \alpha_{t+j} [E_{\mathcal{I}} \ln y_{it+j} - E_{t-1} \ln y_{it+j}] \\
&\quad + (1 - \pi_{it}) \left[E_{\mathcal{I}} \ln \left[A_{it} - \frac{A_{iT+1}}{(1+r)^{T-t}} \right] - E_{t-1} \ln \left[A_{it} - \frac{A_{iT+1}}{(1+r)^{T-t}} \right] \right] \\
&\quad + \mathcal{O} \left(E_{t-1} \|(\boldsymbol{\nu}_{it}^{R-1})\|^2 \right)
\end{aligned} \tag{25}$$

where

$$\begin{aligned}
\alpha_{t+j} &= \frac{\exp [E_{t-1} \ln y_{it+j} - j \ln(1+r)]}{\sum_{k=0}^{R-t-1} \exp [E_{t-1} \ln y_{it+k} - k \ln(1+r)]} \\
&= \frac{\exp \left[\sum_{k=0}^j (\eta_{t+k} + E_{t-1} \bar{v}_{t+k}) + E_{t-1} \bar{u}_{t+j} - j \ln(1+r) \right]}{\sum_{k=0}^{R-t-1} \exp \left[\sum_{l=0}^k (\eta_{t+l} + E_{t-1} \bar{v}_{t+l}) + E_{t-1} \bar{u}_{t+k} - k \ln(1+r) \right]}
\end{aligned}$$

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

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

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

We are able to equate the subjects of equations (23) and (25) because the realised budget must balance and $\sum_{j=0}^{R-t} \frac{c_{it+j}}{(1+r)^j}$ and $\sum_{j=0}^{R-t-1} \frac{y_{it+j}}{(1+r)^j} + A_{it} -$

$\frac{A_i T+1}{(1+r)^{T-t}}$ therefore have the same distribution. We use (23) and (25), taking differences between expectations at the start of the period, before the shocks are realised, and at the end of the period, after the shocks are realised. This gives

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

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

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

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

Squaring the two sides, taking expectations and inspecting terms reveals that the terms which are $\mathcal{O}(E_t \|\boldsymbol{\varepsilon}_{it}^T\|^2 + E_{t-1} \|\boldsymbol{\varepsilon}_{it}^T\|^2)$ are $\mathcal{O}\left(E_t \|\boldsymbol{\nu}_{it}^{R-1}\|^2 + E_{t-1} \|\boldsymbol{\nu}_{it}^{R-1}\|^2\right)$.

Furthermore, since, for all $j \geq 0$, $\|\nu_{it+j}\|^2 = \mathcal{O}_p\left(E_{t-1} \|\nu_{it+j}\|^2\right)$ by Chebyshev's inequality, $E_t \|\boldsymbol{\nu}_{it}^{R-1}\|^2 = \mathcal{O}_p\left(E_{t-1} \|\boldsymbol{\nu}_{it}^{R-1}\|^2\right)$.

Thus

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

and therefore

$$\Delta \ln c_{it} = \Gamma_t + \pi_{it}(v_{it} + \alpha_t u_{it}) + \pi_{it} \Omega_t + \mathcal{O}_p\left(E_{t-1} \|\boldsymbol{\nu}_{it}^{R-1}\|^2\right). \quad (26)$$

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 (26) and the equation driving the income process (1) 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 form¹⁶

$$\begin{aligned}
\Delta \text{Var}(\ln y_t) &= \text{Var}(v_t) + \Delta \text{Var}(u_t) \\
\Delta \text{Var}(\ln c_t) &= (\bar{\pi}_t^2 + \text{Var}(\pi_t)) \text{Var}(v_t) + (\bar{\pi}_t^2 + \text{Var}(\pi_t)) \alpha_t^2 \text{Var}(u_t) \\
&\quad + \text{Var}(\pi_t) \Omega_t^2 + 2 \text{Cov}(\pi_t, c_{t-1}) \Omega_t + \mathcal{O}(E_{t-1} \|\nu_{it}\|^3) \\
\Delta \text{Cov}(\ln c_t, \ln y_t) &= \bar{\pi}_t \text{Var}(v_t) + \Delta[\bar{\pi}_t \alpha_t \text{Var}(u_t)] \\
&\quad + \mathcal{O}(E_{t-1} \|\nu_{it}\|^3).
\end{aligned}$$

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

¹⁶Note that $\text{Cov}(\ln y_{t-1}, u_{t-1}) = \text{Var}(u_{t-1})$ and $\text{Cov}(\ln c_{t-1}, u_{t-1}) = \bar{\pi}_t \alpha_t \text{Var}(u_{t-1})$.

Figure 1: Mean self insurance: $\bar{\pi}_t$

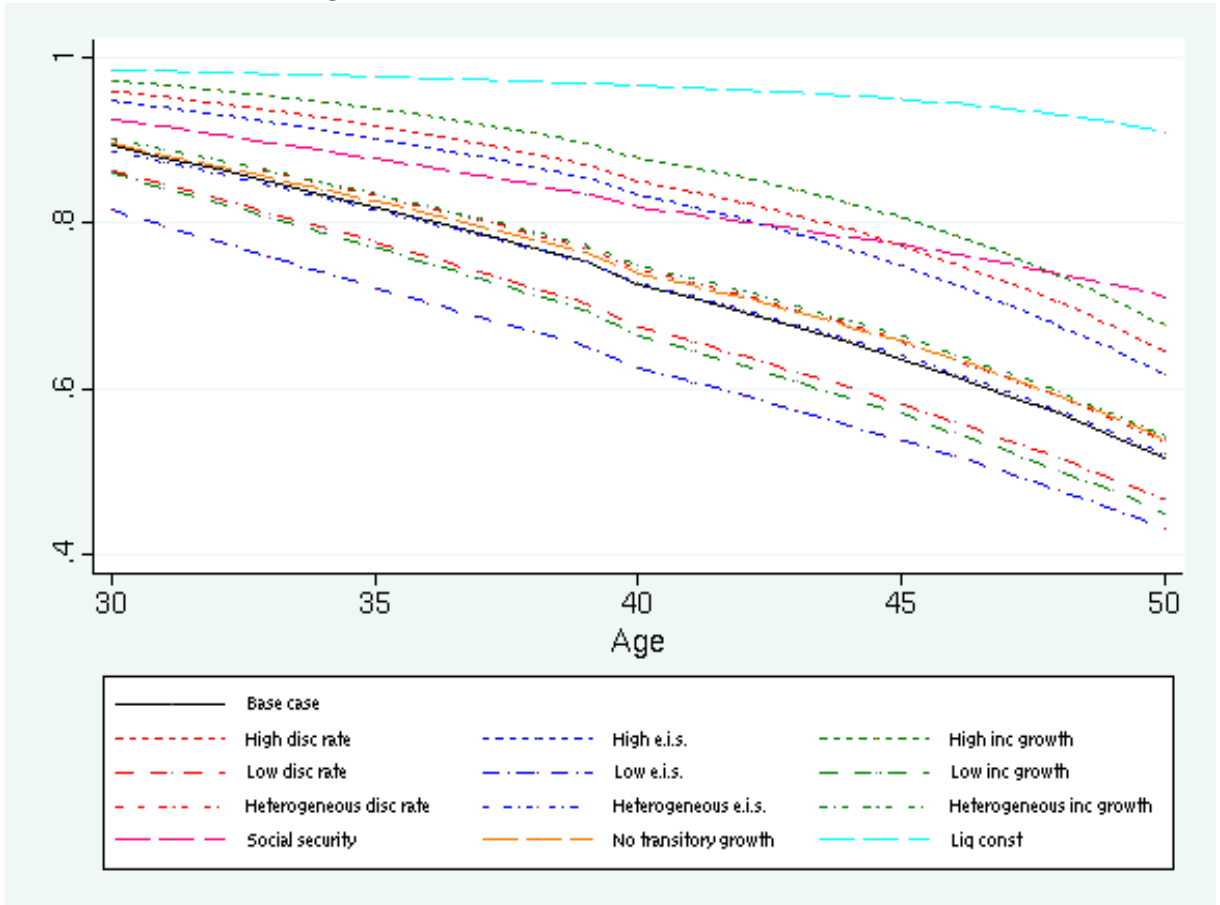


Figure 2: Variation in self insurance: $\text{Var}(\pi_t)/\bar{\pi}_t^2$

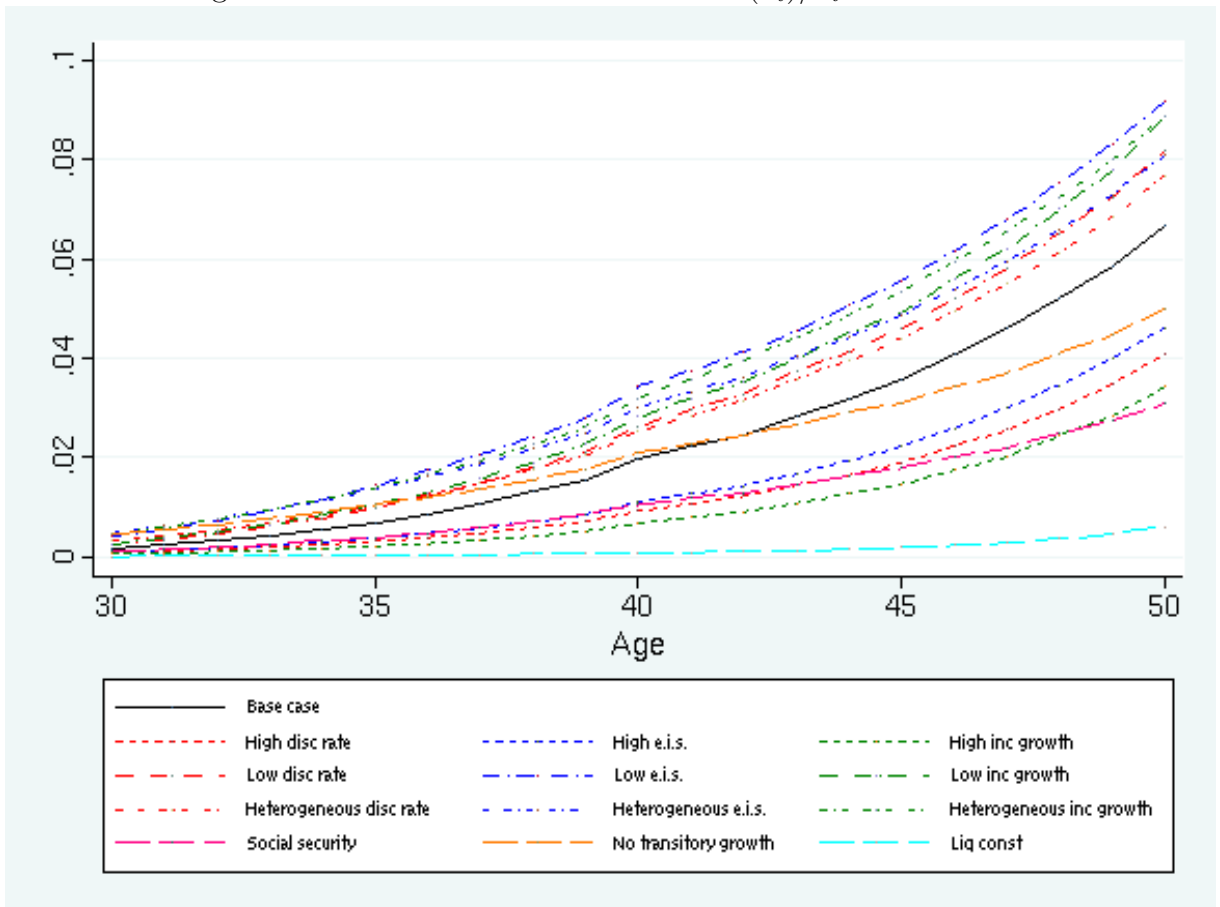


Figure 3: Permanent Variance: Base case

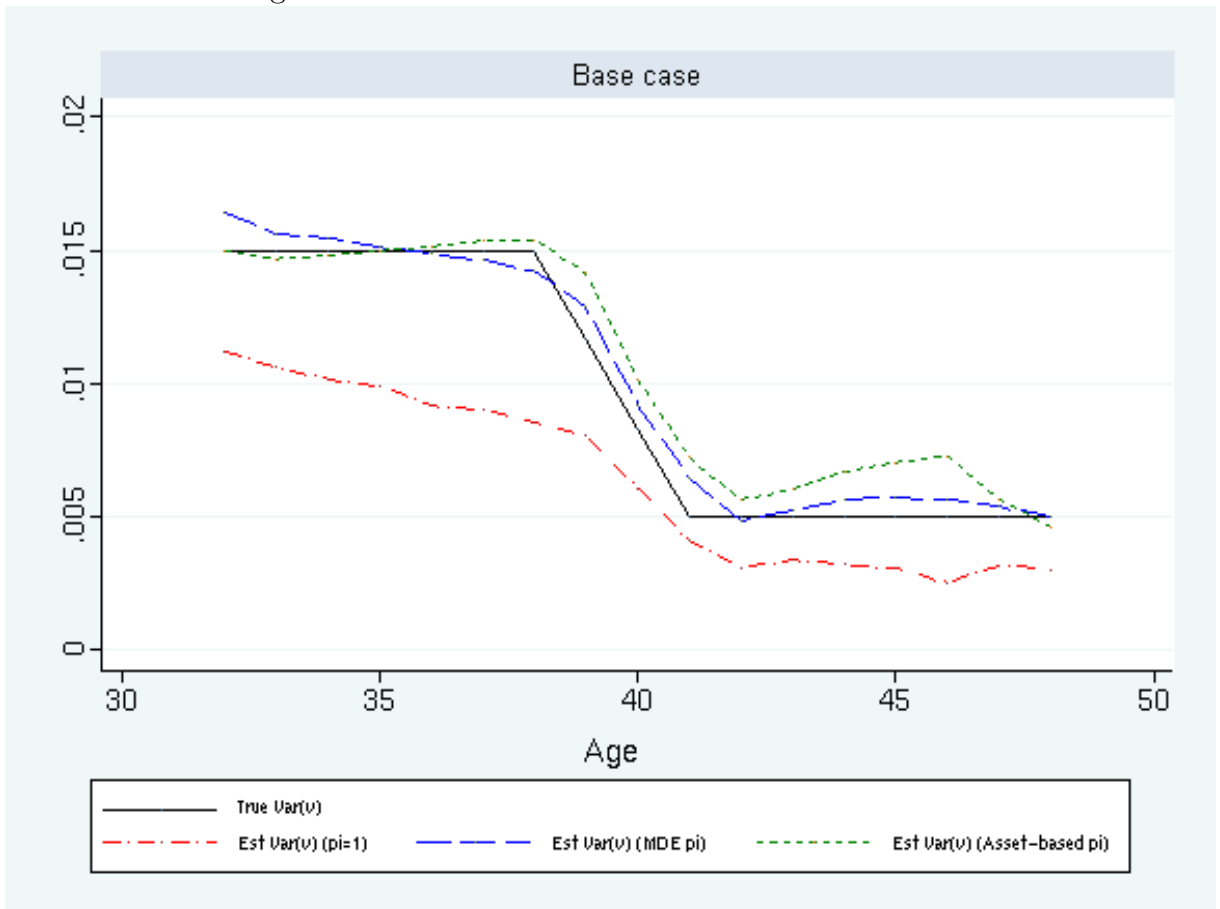


Figure 4: Permanent Variance: Effect of discount rate

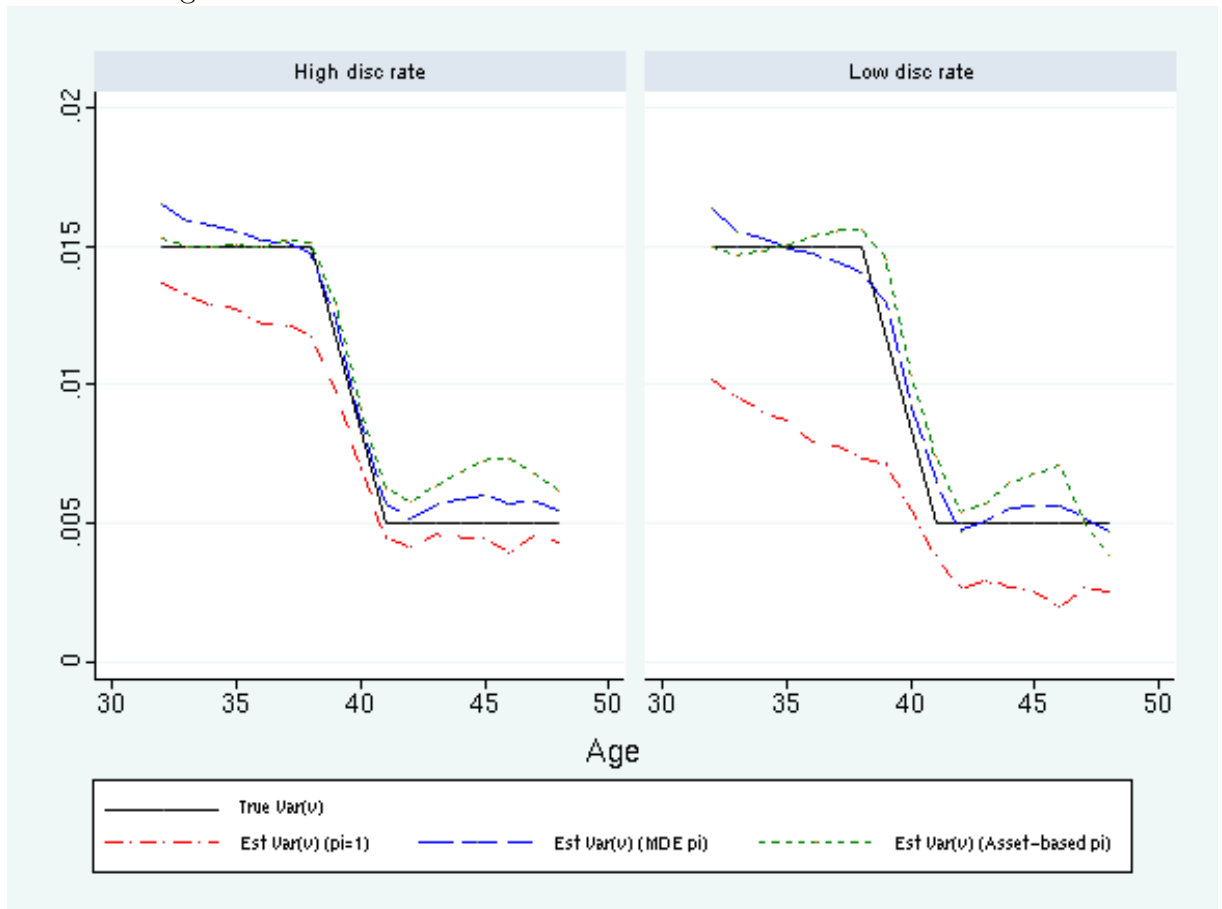


Figure 5: Permanent Variance: Effect of EIS

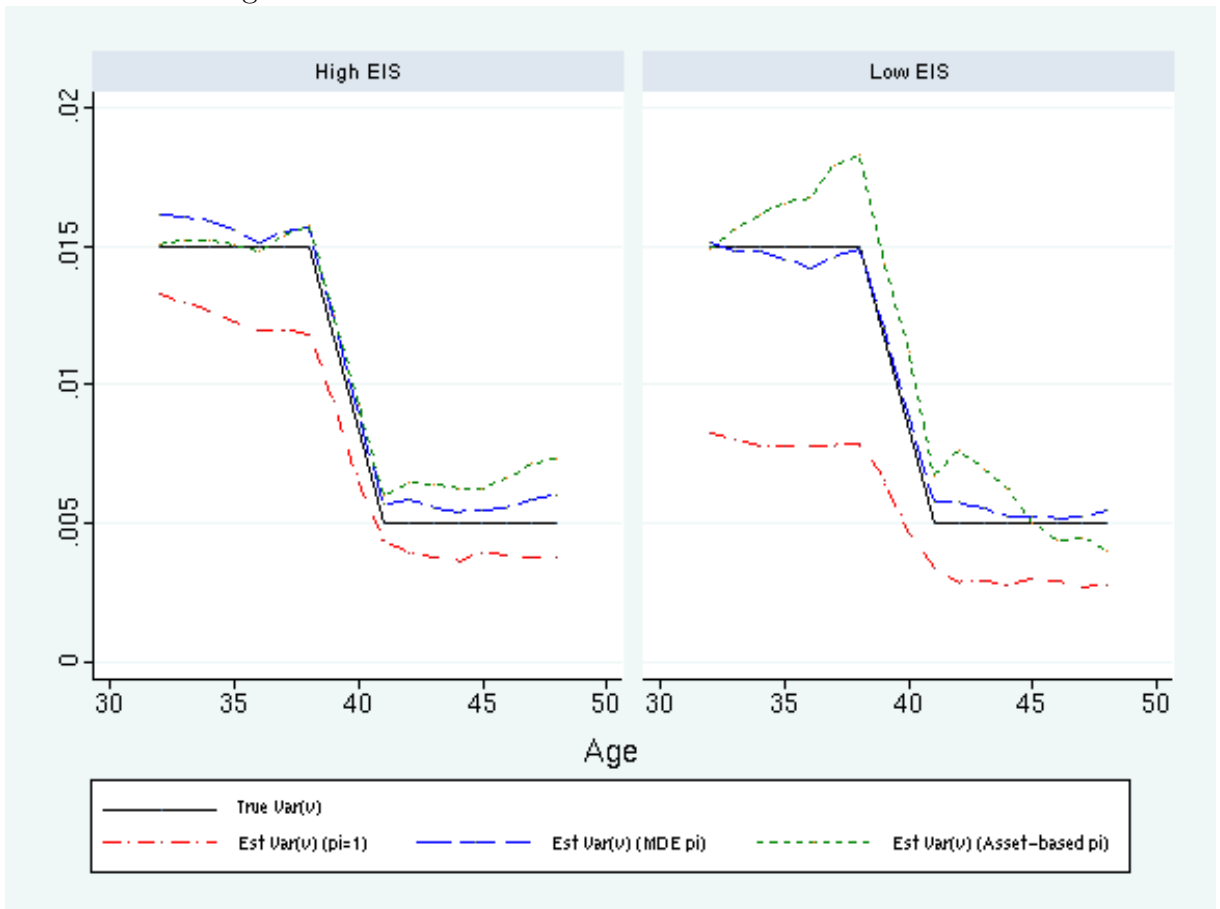


Figure 6: Permanent Variance: Effect of income growth

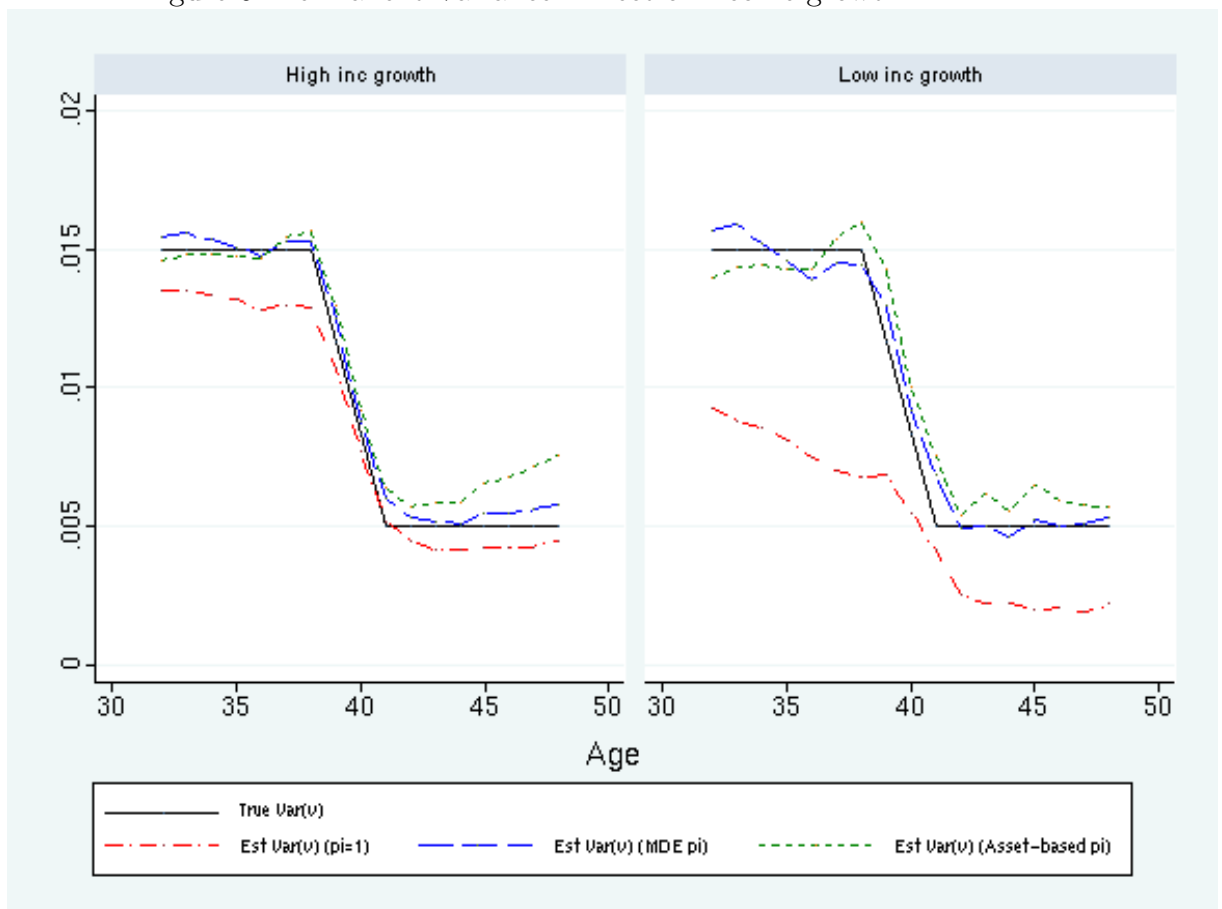


Figure 7: Permanent Variance: Effect of consumption growth heterogeneity

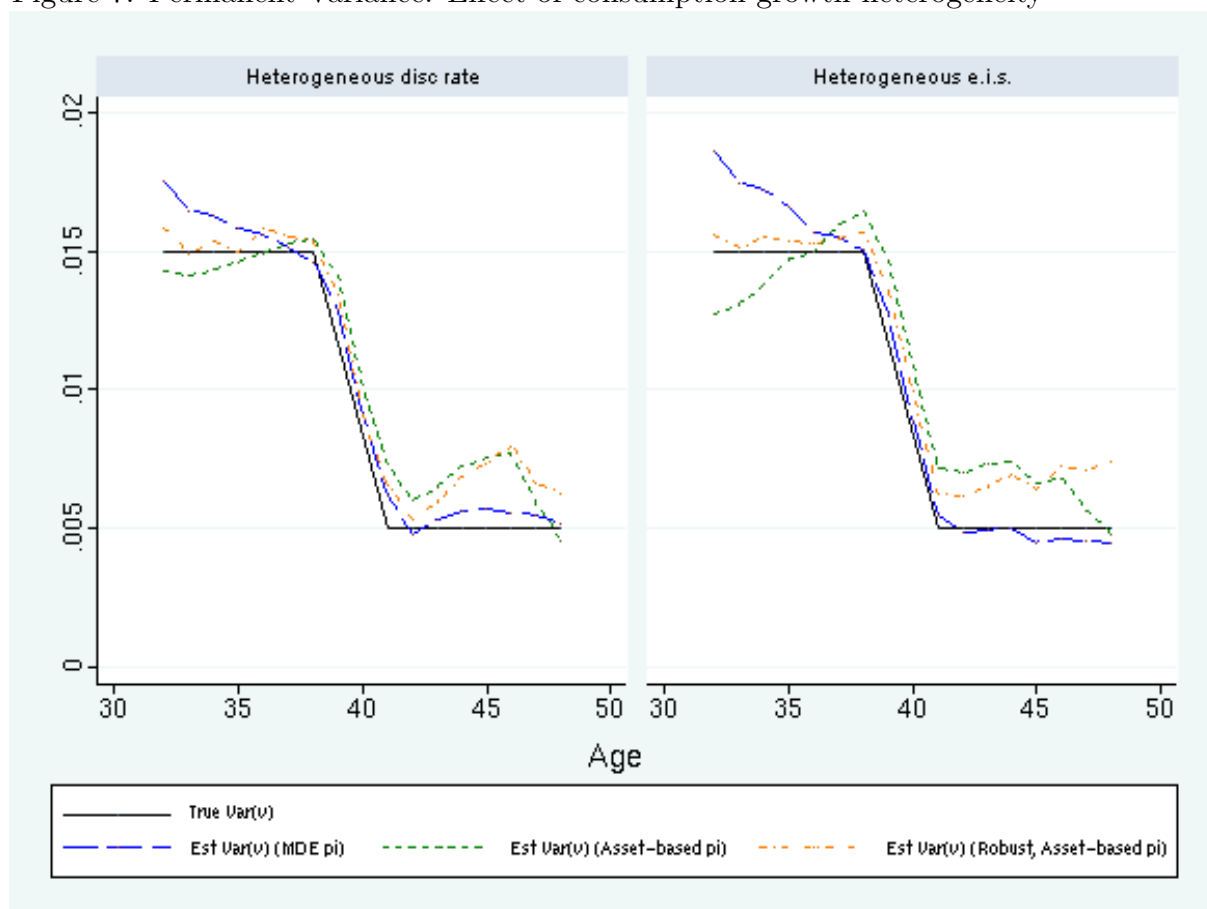


Figure 8: Permanent Variance: Effect of income growth heterogeneity

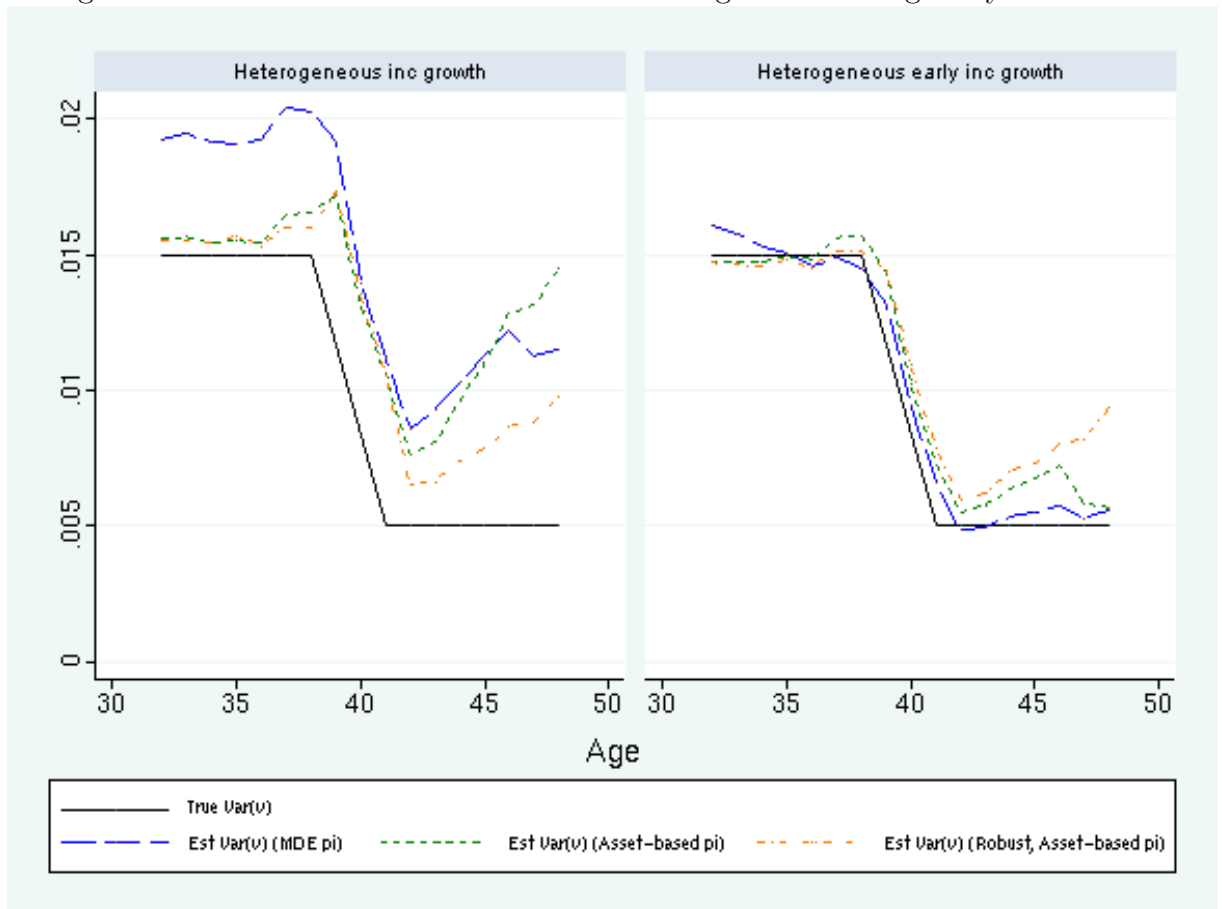


Figure 9: Permanent Variance: Effect of liquidity constraints and social security

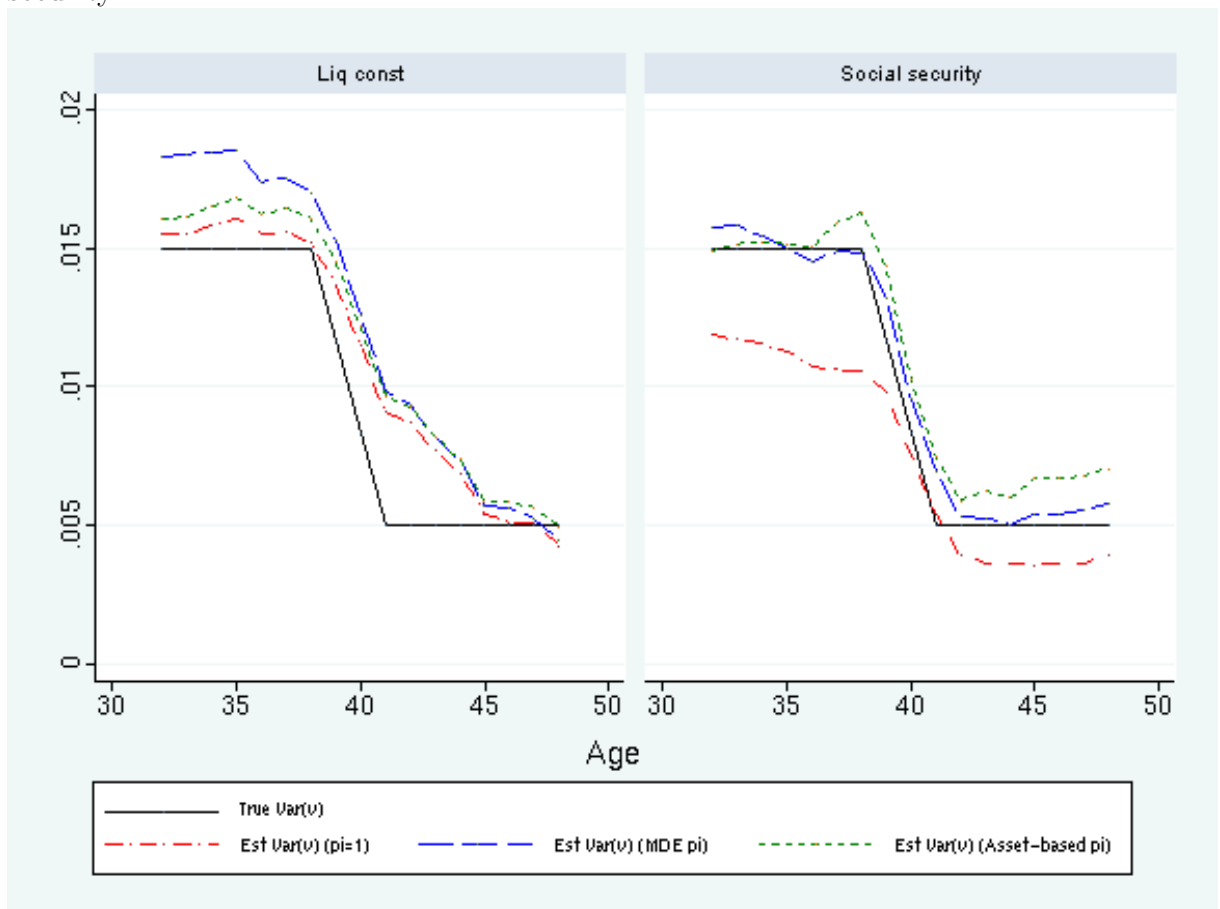


Figure 10: Change in Transitory Variance: MDE π

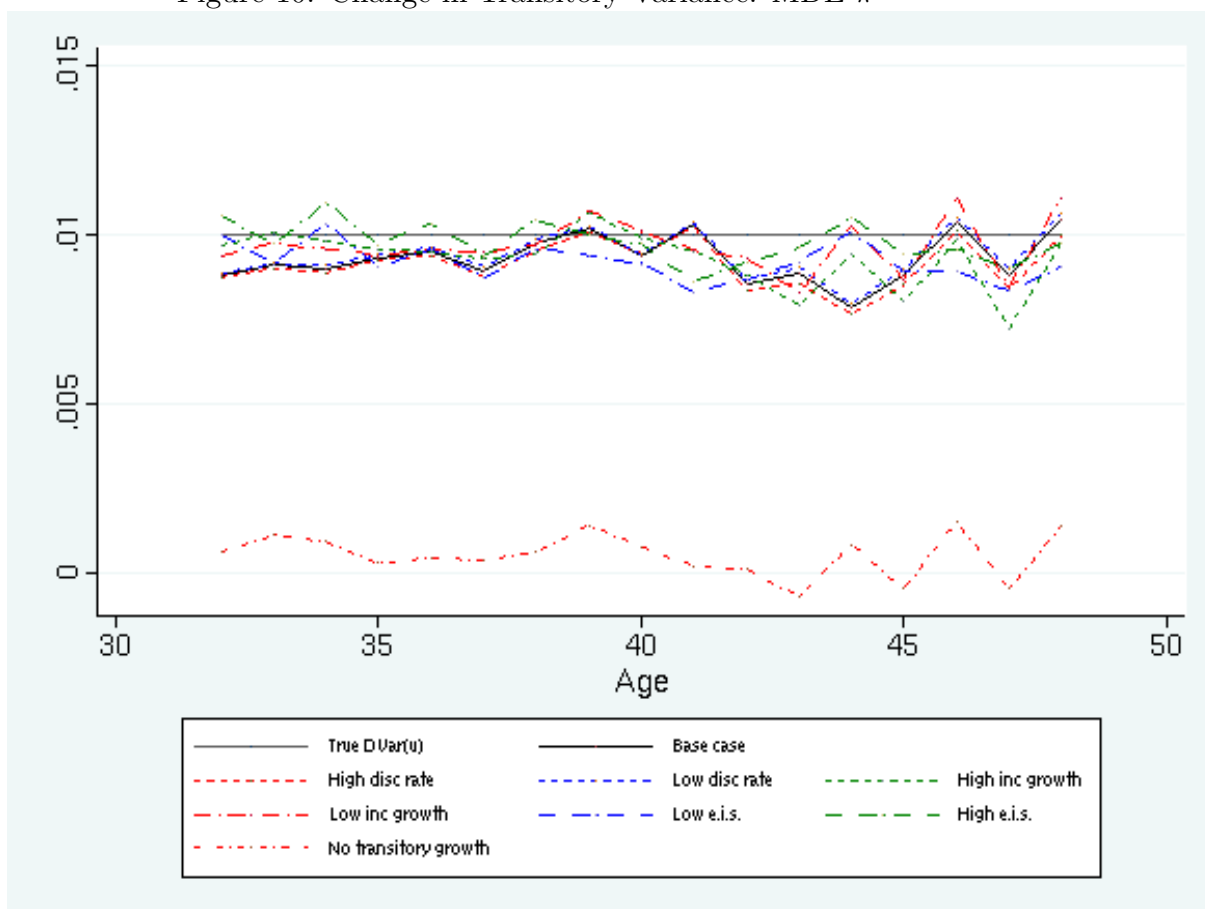


Figure 11: Variances: UK 1979-1992

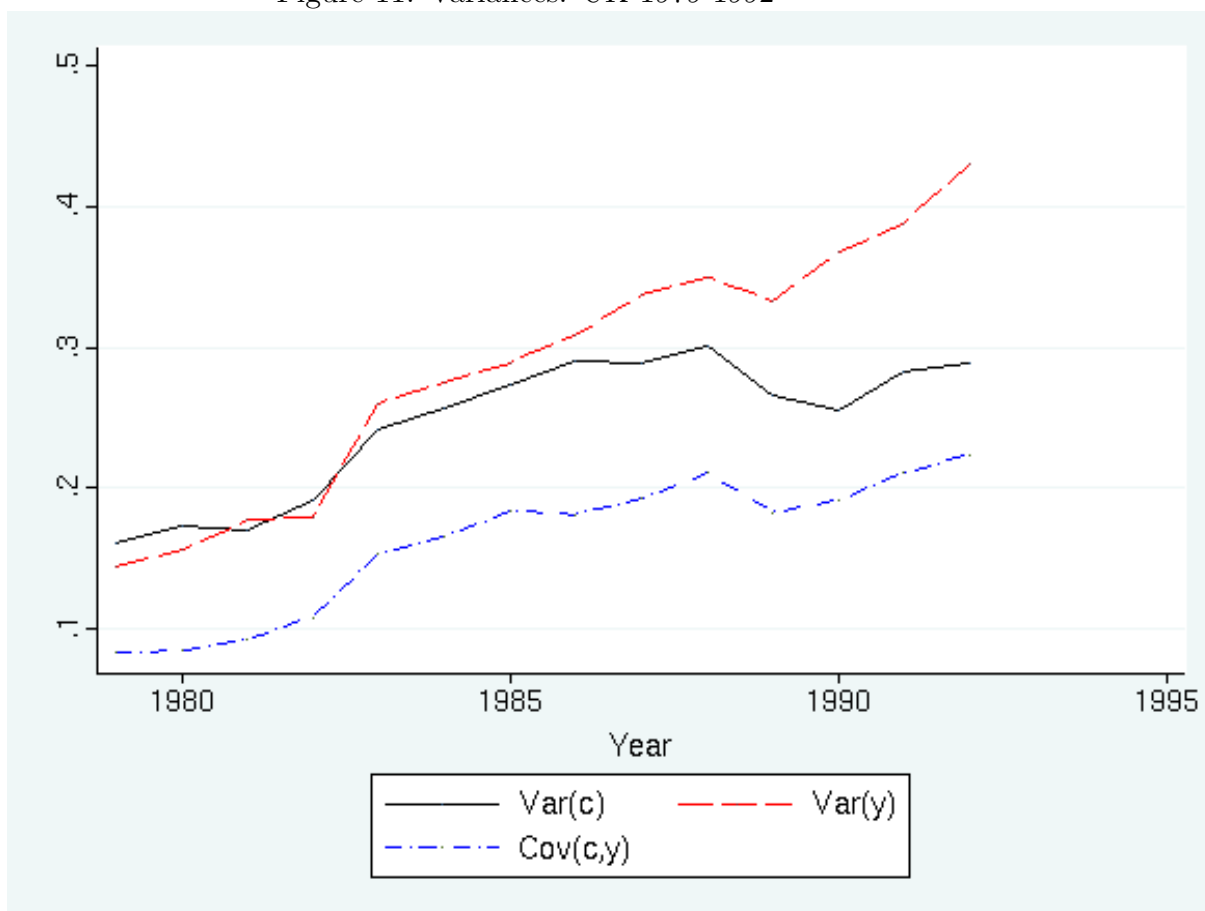


Figure 12: Variances by Cohort: UK 1979-1992

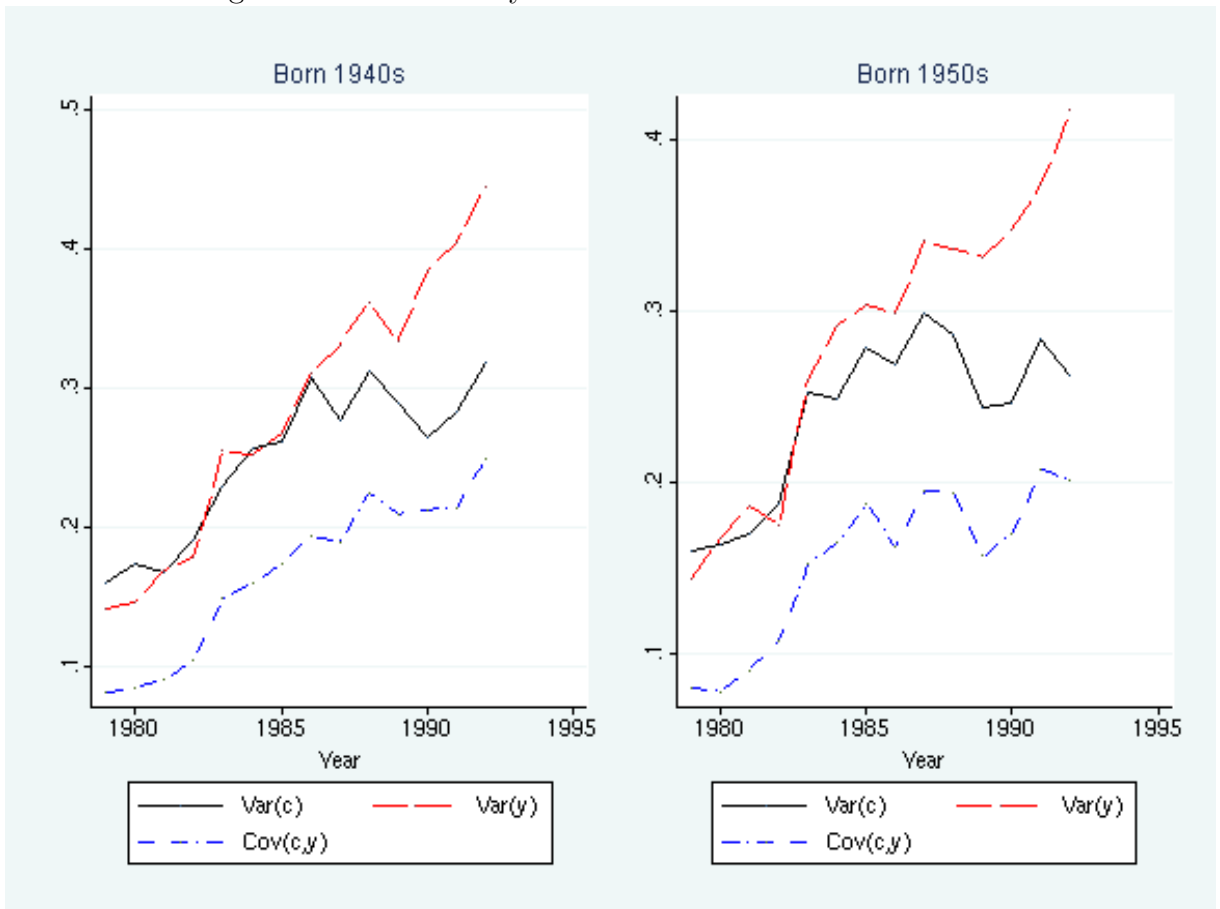


Figure 13: Change in Transitory Variance and Level of Permanent Variance:
UK 1979-1992

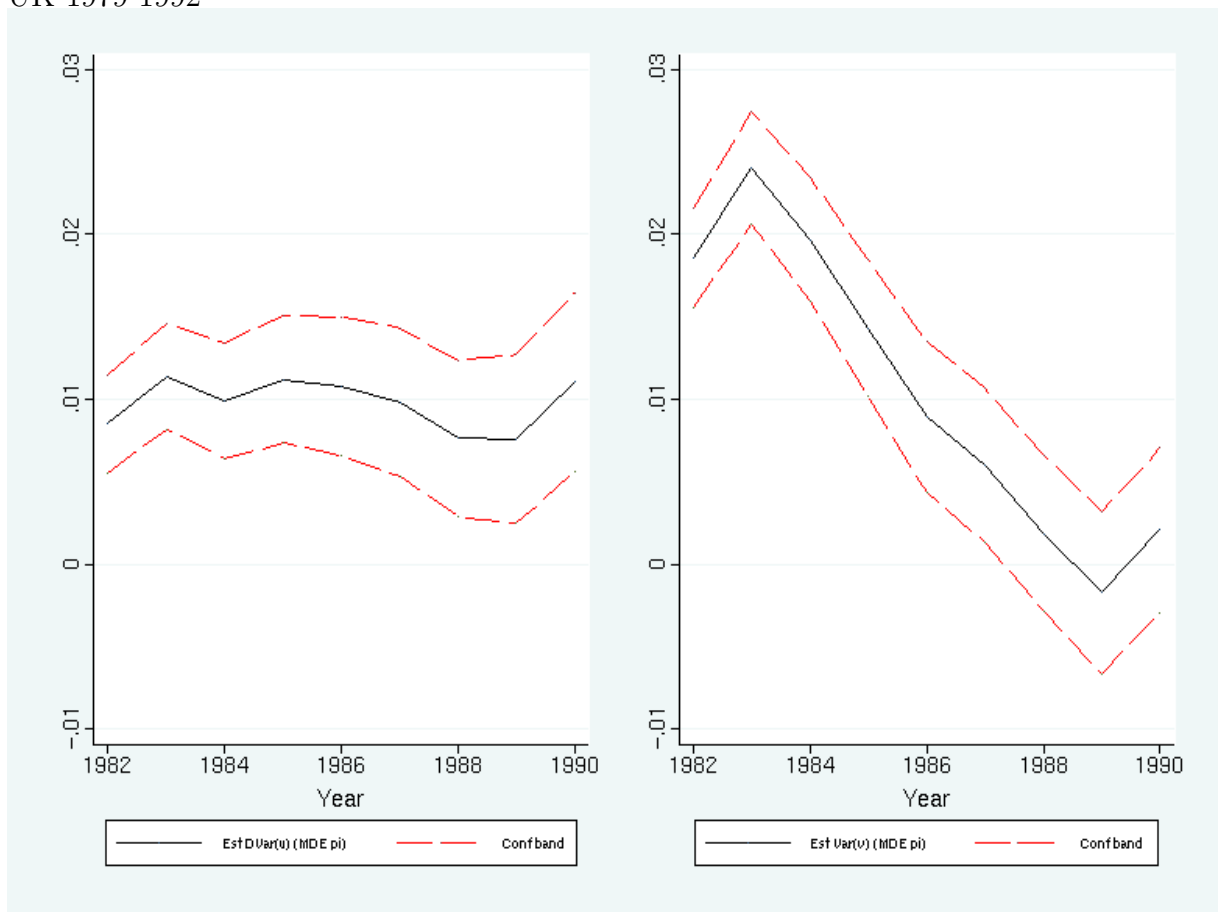


Figure 14: Change in Transitory Variance by Cohort: UK 1979-1992

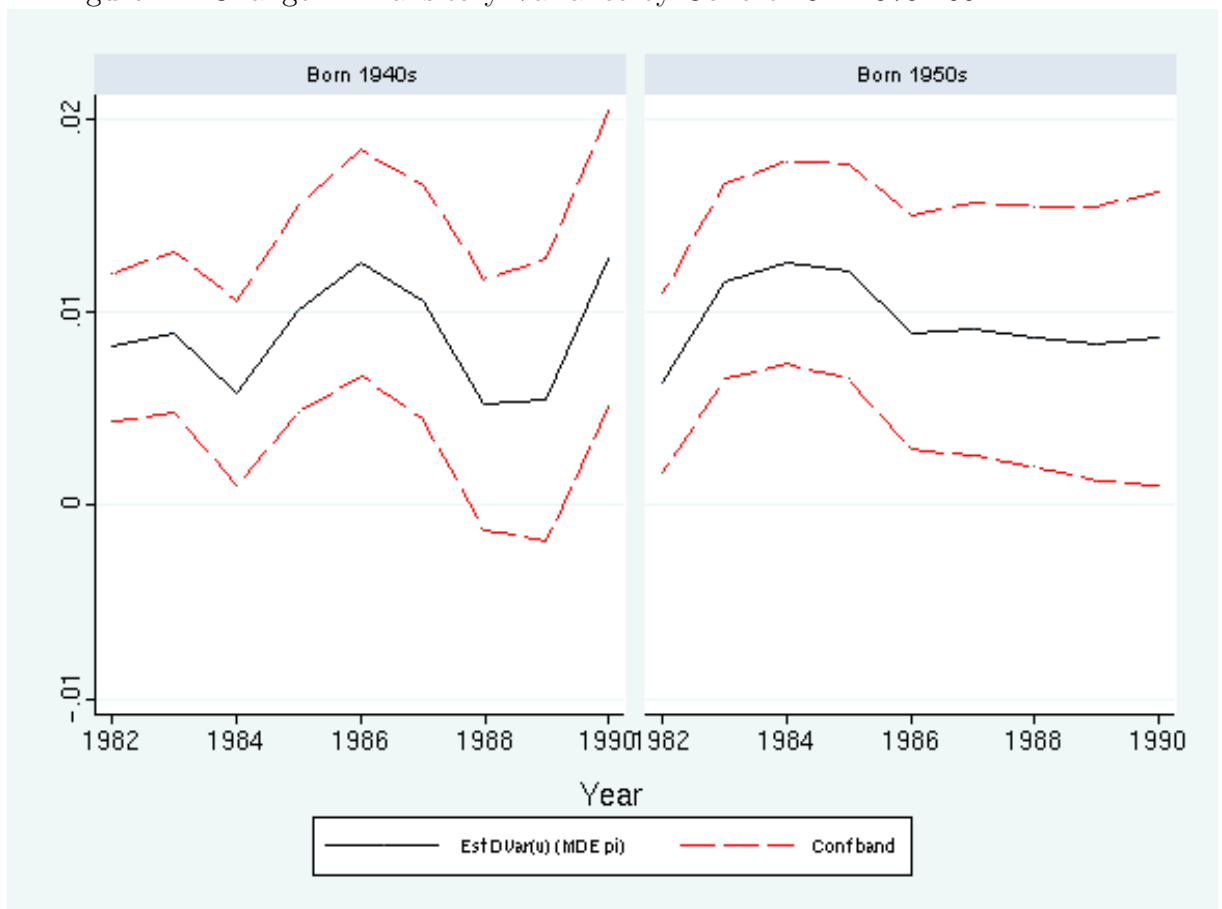


Figure 15: Permanent Variance by Cohort: UK 1979-1992

