INCOME VARIANCE DYNAMICS AND HETEROGENEITY

Costas Meghir
Luigi Pistaferri

THE INSTITUTE FOR FISCAL STUDIES
WP01/07

INCOME VARIANCE DYNAMICS AND HETEROGENEITY*

By Costas Meghir† and Luigi Pistaferri‡

Abstract

Recent theoretical work has shown the importance of measuring microeconomic uncertainty for models of both general and partial equilibrium under imperfect insurance. In this paper the assumption of i.i.d. income innovations used in previous empirical studies is removed and the focus of the analysis placed on models for the conditional variance of income shocks, which is related to the measure of risk emphasized by the theory. We first discriminate amongst various models of earnings determination that separate income shocks into idiosyncratic transitory and permanent components. We allow for education- and time-specific differences in the stochastic process for earnings and for measurement error. The conditional variance of the income shocks is modelled as a parsimonious ARCH process with both observable and unobserved heterogeneity. The empirical analysis is conducted on data drawn from the 1967-1992 Panel Study of Income Dynamics. We find strong evidence of sizeable ARCH effects as well as evidence of unobserved heterogeneity in the variances.

Keywords: Microeconomic uncertainty, Earnings, ARCH.
JEL Classification: D80; J30.

*We are grateful to two anonymous referees, and the editor Glenn Ellison for constructive and helpful comments. We also thank Joe Altonji, Orazio Attanasio, James Banks, Richard Blundell, David Card, Richard Dickens, Tullio Jappelli, Hamish Low, Tom MaCurdy, Thierry Magnac, Ian Preston, Danny Quah, Frank Windmejier and seminar participants at several universities in Europe and in the US for useful comments. Financial support was provided by the ESRC Centre for Fiscal Policy at the Institute for Fiscal Studies (IFS) and the "Taube Faculty Research Fund" at the Stanford Institute for Economic and Policy Research. The usual disclaimers apply.

†*Institute for Fiscal Studies, and Dept. of Economics, University College London, Gower Street, London WC1E 6BT, England.*

‡*Dept. of Economics, Stanford University, Stanford CA 94305, U.S.A., and C.E.P.R.*

1 Introduction

While a number of papers have focused on modelling the time series properties of the mean of earnings the modelling of the variance has been neglected. Seminal papers in this area have been Lillard and Willis (1978), MaCurdy (1982), and Abowd and Card (1989) who fit ARMA type processes, using panel data. However, the properties of the variance of income are also important for our understanding of behavior. For example results reported in Deaton (1992) relating to consumption and savings, rely on the assumption that income shocks are independently and identically distributed, while Caballero (1990) discusses the potential importance of relaxing the i.i.d. assumption. Some recent studies allow for heterogeneity when measuring the variance of income (Carroll and Samwick, 1997; Hubbard, Skinner and Zeldes, 1994), while Alvarez, Browning and Ejrnaes (2001) emphasize heterogeneity in many aspects of the income process including the mean and the variance.[1]

The assumption of i.i.d. income innovations has probably been used for its convenience since it simplifies the analysis and search for the numerical solutions in simulations of say the consumption plan. Nevertheless, it can lead to wrong conclusions about individual behavior. In particular, the presence of ARCH effects or stochastic volatility has implications for the study of life-cycle consumption and savings, for the welfare effects of uncertainty, as well as for income mobility and poverty, among other issues. The importance of considering the evolution of the variance of earnings over the business cycle, as well as acknowledging differences across individuals (under imperfect insurance) has been emphasized recently by Browning, Hansen and Heckman (2000).

In this paper we address some of these issues and model earnings as the sum of a martingale component and a (possibly persistent) transitory disturbance. Since the presence of permanent shocks is central for a number of economic questions, we investigate the validity of this representation by testing whether the variance of the permanent shock is zero and consider the implications of some alternative specifications. Since part of the transitory fluctuation is in reality measurement error, we tackle explicitly the issue of reporting errors in

[1] Other authors rely on subjective expectation data (Guiso, Jappelli and Pistaferri, 2002). See also Zeldes (1992) for a discussion of these issues.

earnings data. We allow for the possibility that individuals with different education levels face different income processes, both at the mean and at the variance level, thus allowing for the changing returns to observable skills (see Juhn, Murphy and Pierce, 1993 for the US and Gosling, Machin and Meghir, 2000 for the UK).

In a second step, we model the conditional variance of the earnings shocks as an ARCH process with observed and unobserved heterogeneity.[2] Our task is complicated by the fact that in the context of a model with permanent and transitory income shocks, these are not separately observable. We show how moment conditions can be derived to estimate the dynamic properties of the conditional variance of income and to test for permanent unobserved heterogeneity in the variance of both the transitory and the permanent innovation.[3] Our data draw from the Panel Study of Income Dynamics (PSID) for the years 1967-92.

We find that permanent and transitory shocks are important components of income shocks. We also find strong evidence of state dependence in the variance of both permanent and transitory components. Finally we find that earnings variances are heterogeneous across individuals. Thus the i.i.d. nature of income innovations is rejected by our data which points to both heteroskedasticity and stochastic earnings risk as important facts.

The paper has five more sections. In Section 2 we introduce the data used in the empirical application. In sections 3 and 4 we discuss and motivate our approach for modelling the conditional mean and the conditional variance of earnings, respectively, and present and discuss our empirical findings. Section 5 analyzes some implications of the empirical results. Section 6 concludes.

[2] See among others Engle (1982).

[3] Banks, Blundell and Brugiavini (2001) is the only attempt we know of to estimate the conditional variance of income shocks according to an ARCH process. However, the authors do not allow for the distinction between transitory and permanent disturbances and do not present ARCH estimates. Moreover, their analysis is based on cohort data and misses truly idiosyncratic uncertainty. Meghir and Windemejier (2000) recover orthogonality conditions for the estimation of an ARCH process but do not allow for a distinction between transitory shocks and permanent shocks. Chamberlain and Hirano (1999) estimate earnings dynamics allowing for volatility heterogeneity.

2 The data

The data used in this study are drawn from the 1968-1993 family and individual-merged files of the PSID (waves I through XXVI).[4]

The PSID started in 1968 collecting information on a sample of roughly 5,000 households. Of these, about 3,000 were representative of the US population as a whole (the core sample), and about 2,000 were low-income families (the Census Bureau's SEO sample). Thereafter, both the original families and their split-offs (children of the original family forming a family of their own) have been followed. In the empirical analysis we use both the core sample and the SEO sample. Some authors (Lillard and Willis, 1978) suggest dropping the SEO low-income sample because of endogenous selection. In other words, an initial condition problem arises. However, given linearity, the initial condition problem is taken care of by the presence of the permanent component. To put it differently, we deal with the problem by estimating models for the growth rate rather than specifications in levels. This is also true for the model of the variance where we include fixed effects.

Questions referring to labor income are retrospective; thus, those asked in 1968, say, refer to the 1967 calendar year. The earnings variable is the labor portion of money income from all sources; the variable name in the PSID tapes is "head's money income from labor" and includes the labor part of farm income and business income, wages, bonuses, overtime, commissions, professional practice, labor part of income from roomers and boarders or business income.[5] We deflate the nominal measure of earnings by the GNP personal consumption expenditure deflator (using 1992 as the base year). We use information on the highest grade completed to allocate individuals in our sample to three education groups: High School dropouts (those with less than 12 grades of schooling), High School graduates (those with at least a High School diploma, but no College degree), and College graduates (those with a College degree or more).

Step-by-step details on sample selection are reported in the Appendix. Briefly, we select

[4] See Hill (1992) for more details about the PSID.

[5] As noted by Gottshalk and Moffitt (1993), the measure of labour income available in the PSID has sources that may reflect capital income, such as the labour part of farm income and roomers and boarders. We do not account for this problem.

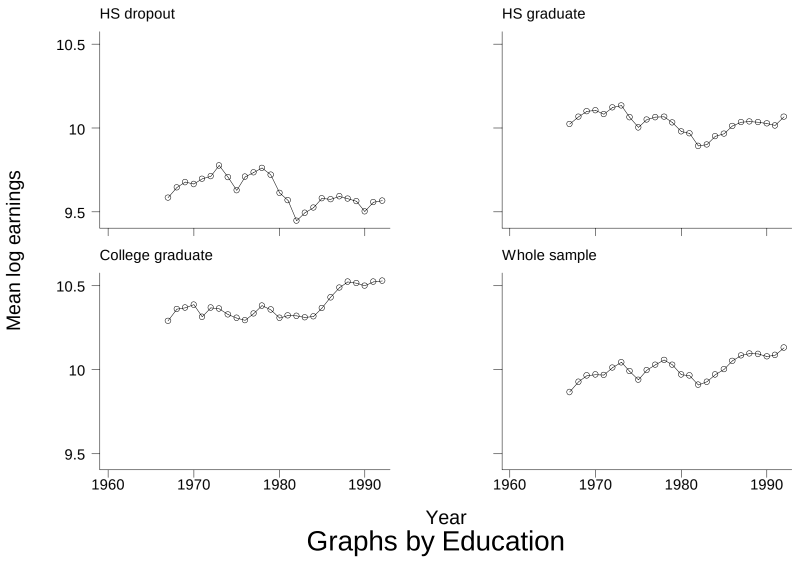

Figure 1: The mean of log real earnings.

male heads aged 25 to 55 with at least nine years of usable earnings data. The selection process leads to a sample of 2,069 individuals and 31,631 individual-year observations. Relevant sample statistics are presented in Tables A1 to A3 (sample composition by year and by education, and demographic characteristics). Figures 1 and 2 plot the mean and the variance of log real earnings against time for each education group and for the whole sample. These figures reproduce well know facts about the distribution of male earnings in the US (see Levy and Murnane, 1992).

3 The conditional mean of earnings

As in previous empirical work, we posit the following model for the conditional mean of log earnings:

$$y_{it} = m_t^e + \beta_t^{e\prime} Z_{it} + u_{it} \tag{1}$$

where y_{it} is the logarithm of real annual measured earnings, the superscript "e" stands for education, m_t^e is a calendar year effect, Z_{it} a vector of observable characteristics, and u_{it} the stochastic component of earnings. Aggregate shocks specific to an education group are captured by m_t^e (aggregate risk). A distinguishing feature of this model is the assumption that

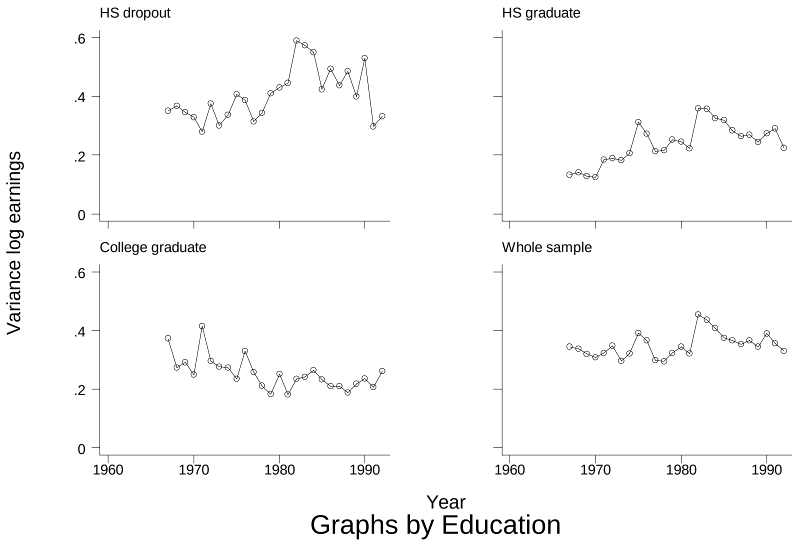

Figure 2: The variance of log real earnings.

the return to education varies with calendar time while the returns to observable attributes vary with time as well as by education in a very general way. This characterization is in line with e.g. Card and Lemieux (2001), who model the last three decades' shifts in the income distribution with changes in the prices of skills.

We assume that the unexplained component of income can be decomposed into a measurement error, a transitory innovation with low persistence and a martingale permanent component. Therefore:

$$u_{it} = r_{it} + e_{it} + p_{it} \tag{2}$$

with r_{it} being a classical i.i.d. measurement error, e_{it} the transitory shock and p_{it} the permanent component of income which follows the process:

$$p_{it} = p_{it-1} + \zeta_{it} \tag{3}$$

We assume that the transitory shock follows an MA(q) process (i.e. $e_{it} = \Theta^e(L,q)\varepsilon_{it} = \varepsilon_{it} - \sum_{j=1}^{q} \theta_j^e \varepsilon_{it-j}$), with the order q of the process to be determined empirically, and assume that the permanent and the transitory shocks are uncorrelated at all leads and lags.[6]

We define the unexplained component of the rate of growth of earnings as:

$$g_{it} = \Delta u_{it} = \Delta r_{it} + \Theta^e(L, q)\Delta\varepsilon_{it} + \zeta_{it} \tag{4}$$

This is a composite MA$(q+1)$ process. Equation 4 implies that g_{it} will be orthogonal to variables dated $t-q-2$ or earlier.

We recover the order of the MA process for the transitory shock from the estimated autocovariances of g_{it}. For instance, if transitory shocks are not serially correlated (i.e., if $q = 0$) one should find $E(g_{it}g_{it-j}) = 0$ for $j > 1$.[7]

Since we use estimated residuals we need to assume that the underlying processes, including the measurement error, have distributions such that the cross sectional moments of g_{it} exist up to the fourth order. However, since we are dealing with finite lived individuals we do not necessarily require that these moments exist as the time dimension goes to infinity.

In what follows we compute all the standard errors using the block bootstrap procedure (see Hall and Horowitz, 1996 and Horowitz, 2002). In this way we account for serial correlation of arbitrary form, heteroskedasticity, as well as for the fact that we use pre-estimated residuals.[8] We should point out that this procedure is conservative, since it allows for more serial correlation than that implied by the moment conditions we use. Hence the bootstrap standard deviations will be using only the N dimension of the sample and the precision of our parameters is likely to be underestimated.

[6] Some examples of permanent innovations are associated to job mobility, long-term unemployment, health shocks, promotions and even demotions. Transitory shocks to individual earnings include overtime labor supply, piece-rate compensation, bonuses and premia, etc.; in general, such shocks are mean reverting and their effect does not last long.

[7] The restriction can be tested computing the χ^2 zero restrictions test described by Abowd and Card (1989).

[8] Later on we will be using GMM to estimate processes for the mean and the variance of income. Windmeijer (2000) provides a small sample correction for the standard errors of two step GMM procedures, whose asymptotic approximation is known to be heavily biased. However, in our case the bootstrap procedure of Hall and Horowitz (1996) not only corrects for this but also accounts for pre-estimated residuals and hence we have preferred it.

3.1 The conditional mean and the structure of the error term

We first estimate the conditional mean process. We start by regressing, for each year and each education group, log real earnings on a constant term, a quadratic in age, dummies for race (white), region of residence and residence in a SMSA. This first stage allows the effects of all characteristics to be education specific and to vary over time. Moreover the returns to education also vary over time. For the remaining analysis of the stochastic properties of income we use the residuals from these regressions, i.e. $(\widehat{g}_{it} = \Delta y_{it} - \Delta\widehat{m}_t^e - \Delta\widehat{\beta}_t^{e\prime} Z_{it})$ in the place of g_{it}, and replace theoretical moments with sample analogs.

The next step is to evaluate whether the data conform to our hypotheses concerning the stochastic structure of the error term. We estimate the autocovariances of g_{it} using standard methods (Abowd and Card, 1989).

Table 1
The autocovariances of the unexplained growth of earnings

Order	Pooled sample	High school dropout	High school graduate	College graduate
0	0.109 (0.0051)	0.164 (0.014)	0.103 (0.0068)	0.065 (0.0058)
1	−0.0303 (0.0021)	−0.0535 (0.0055)	−0.028 (0.0027)	−0.0105 (0.0018)
2	−0.0079 (0.0014)	−0.0134 (0.0046)	−0.0077 (0.0015)	−0.0025 (0.0013)
3	−0.0024 (0.0011)	−0.0018 (0.0034)	−0.0026 (0.0014)	−0.0024 (0.0011)
4	0.0007 (0.0011)	0.0074 (0.0035)	−0.0015 (0.0013)	−0.0017 (0.0012)

Note: Asymptotic standard errors are reported under the coefficient estimate. Values are pooled over all years and individuals.

The estimated unconditional autocovariances up to order four are presented in Table 1 for the pooled sample and separately for the three education groups we focus on. For simplicity, we report values pooled over time.[9] In Table 2 we report the test of zero restrictions for the

[9] The estimated matrices of autocovariances and the associated standard errors are available on request for the whole sample and separately for each education group. Each matrix contains 325 unique elements.

null hypothesis that $E(g_{it}g_{it-j}) = 0$ (with $1 \leq j \leq 4$) allowing the estimated autocovariances to differ over time.

In the pooled sample, unexplained earnings growth rates appear correlated up to the second order. Autocovariances at the third order and beyond are small and statistically insignificant or of borderline significance. This is confirmed also by the analysis conducted on different education groups (see Table 1). In the pooled sample the test that $E(g_{it}g_{it-3}) = 0$ has a p-value of 19 percent (see Table 2). The null is similarly not rejected when we stratify the sample by schooling. The statistical implication is that $q = 1$.[10] We impose this restriction thereafter. The economic implication is that transitory shocks are somewhat persistent: it takes at least one period for the full impact of the transitory shock to be felt.

Table 2
Tests of zero restrictions

	Order of autocovariance			
	1	2	3	4
Pooled Sample	**320** df 24 p-0	**61.2** df 23 p-0	**27.7** df 22 p-0.19	**20.2** df 21 p-0.52
High School Dropout	**159** df 24 p-0	**38.0** df 23 p-0.03	**31.7** df 22 p-0.09	**23.9** df 21 p-0.30
High School Graduate	**193** df 24 p-0	**57.3** df 23 p-0.00	**23.8** df 22 p-0.36	**28.1** df 21 p-0.14
College	**73.7** df 24 p-0	**28.3** df 23 p-0.22	**22.7** df 22 p-0.43	**27.3** df 21 p-0.17

Note: In this table we present tests for zero autocovariance of order 1-4. We provide the test statistic for the hypothesis that the respective autocovariance is zero in all time periods; the degrees of freedom of the test (df), which is determined by the number of time periods for which we can estimate the autocovariance and the p-value for the hypothesis that the autocovariances are zero. We have also tested the hypothesis that all autocovariances of order 3 or higher are jointly zero as in Abowd and Card (1989). The test statistic is 279.83 with 253 degrees of freedom and a p-value of 12%.

3.2 Testing for the absence of permanent shocks

Whether permanent shocks are present and the magnitude of their variance is an issue of great importance for an number of economic questions, such as consumption. This is

[10] For the college graduate $q = 0$ cannot be rejected; we estimate the ARCH parameters below for both $q = 0$ and $q = 1$.

established in number of papers in the literature. Thus we next estimate the variance of the permanent shock and test the null hypothesis that it is zero.[11] To do this we use the sample autocovariances and impose the restriction that they are generated by (4). We use Equally Weighted Minimum Distance (EWMD) for reasons explained in Altonji and Segal (1996).

The key moment condition that identifies the variance of the permanent shock is

$$E\left[g_{it}\left(\sum_{j=-(1+q)}^{(1+q)} g_{it+j}\right)\right] = E\left(\zeta_{it}^2\right) \tag{5}$$

$E\left(\zeta_{it}^2\right)$ being the unconditional variance of the permanent shock. In the Appendix we present EWMD estimates that allow for non-stationarity (see Table A4). We also present an estimate of $E\left(\zeta_{it}^2\right)$ in Table 3 (along with its standard error) under the assumption that this is constant over time, which gives us a one degree of freedom test for the null of no permanent shock.

The test statistic, which is equal to the pooled estimate of the permanent shock divided by it standard error, is asymptotically (for large N) distributed standard normal. The standard error is computed using the block bootstrap, allowing for pre-estimated residuals as well as for serial correlation and heteroskedasticity.

There are a number of advantages of the test we use. First, it does not hinge on the assumption of covariance stationarity. Stationarity is in fact often rejected with PSID data on earnings. Second, it can be generalized to any form of serial correlation (of the MA type) in the transitory component. Finally, it is robust to the presence of measurement error (either classical or with MA-type serial correlation).

3.2.1 Results

The pooled variances of the permanent shock are estimated to be 0.0313 (with a bootstrap standard error of 0.0026) in the whole sample, 0.0331 (0.0067) for the High School dropouts, 0.0277 (0.0039) for the High School graduates, and 0.0437 (0.0068) for the College graduates (the results are reported at the bottom of Table 3). The hypothesis of no permanent shock

[11] The null hypothesis encompasses two different assumptions concerning the structure of the error term: either that a permanent component is absent altogether, or that it is time invariant (e.g. a random growth model). The alternative hypothesis is that the permanent component follows a martingale process. We ignore the problem associated with the fact that q is pre-estimated rather than known.

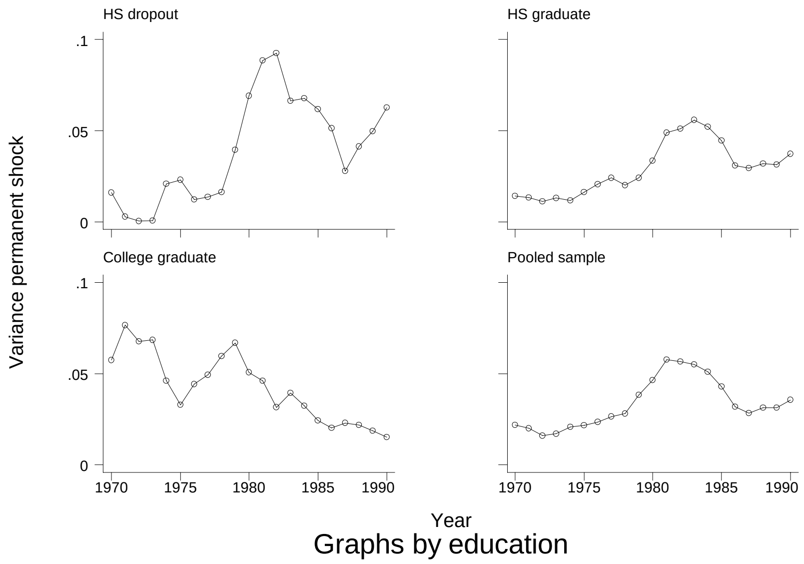

Figure 3: The variance of the permanent shock to earnings.

is strongly rejected for all education groups with p-values well below 1 percent in all cases. Henceforth we assume that log earnings for each education group follows the mean process described in equations 1, 2 and 3 with the transitory shock following an MA(1) process.

Assuming this variance represents uninsurable risk, the high variance for the College graduates is consistent with the idea that the higher returns emanating from increased education come at the cost of higher earnings risk. Our estimates are close to those found elsewhere in the empirical literature (see Carroll and Samwick, 1997), although our focus is on earnings rather than family income. Given that family income may include some form of implicit or explicit insurance, our estimates are likely to be larger.

Figure 3 plots the estimates of the unconditional variance of the permanent shock against time, for the whole sample and by education (see Table A4 in the Appendix for the standard errors of these estimates). Overall the variance increases throughout the 1970s and in the early 1980s, it declines after 1984 with a slight tendency to increase at the end of the survey period. This evidence is very similar to that reported by Moffitt and Gottshalk (1994), who note that "the permanent variance grows, on average, through about 1982 or 1983, [while] it levels off or falls subsequently".[12] A similar pattern holds for the least educated and for the High School graduates. Finally, for the College graduates there is not much evidence of a monotonic increase in permanent income variance in the 1970s; on the other hand, the decline in variance that occurs in the 1980s is much more pronounced than for the other two groups.

3.3 The variance of the transitory shock, the variance of measurement error and the MA coefficients

Although we can identify the variance of the permanent shock it is not possible to disentangle the unconditional variance of the transitory shock, the variance of the measurement error and the MA coefficients. Earlier studies have ignored this point. However, this may be important to the extent that the transitory shock reflects uncertainty and induces economic responses, while measurement error is noise due to imperfect data.

We follow two approaches to deal with this issue. The first strategy is to obtain bounds for the unidentified measures. The second is to use an external estimate of the measurement error in earnings.[13] Both approaches rely on the assumption that measurement error is classical $-$ an assumption that is not universally accepted.[14]

For illustrative purposes take the case with an MA(1) transitory shock and assume invertibility of the MA process. Thus $e_{it} = \varepsilon_{it} - \theta \varepsilon_{it-1}$ with $|\theta| < 1$. The autocovariances of earnings growth of order one and two can be used to derive the following two equations in three unknowns:

$$\begin{aligned} \sigma_\varepsilon^2 &= \tfrac{E(g_{it}g_{it-2})}{\theta} && I \\ \sigma_r^2 &= -E\left(g_{it}g_{it-1}\right) - \tfrac{(1+\theta)^2}{\theta} E\left(g_{it}g_{it-2}\right) && II \end{aligned} \tag{6}$$

The sign of $E\left(g_{it}g_{it-2}\right)$ defines the sign of θ. In our case we can conclude that $\theta < 0$

[12] See Moffitt and Gottshalk (1994), page 12.

[13] A natural extension of our framework is a multiple indicator model that can be used to identify the sources of permanent and transitory earnings variation (e.g., wage changes, unemployment, etc.), and to distinguish between transitory shocks and measurement error. This strategy is pursued in Altonji, Martins and Siow (2001).

[14] As we discuss below, other aspects of our approach, such as the identification of ARCH effects do not depend on whether the measurement error is classical.

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

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

$$\begin{aligned} E\left[g_{it}^2 - 2\overline{\sigma_r^2}\right] &= \sigma_\zeta^2 + 2\left(1+\theta+\theta^2\right)\sigma_\varepsilon^2 \\ E\left(g_{it}g_{it-1} + \overline{\sigma_r^2}\right) &= -\left(1+\theta\right)^2\sigma_\varepsilon^2 \\ E\left(g_{it}g_{it-2}\right) &= \theta\sigma_\varepsilon^2 \end{aligned}$$

where $\overline{\sigma_r^2}$ is available externally. The three moments above depend only on θ, σ_ζ^2 and σ_ε^2. We can then estimate these parameters using EWMD.

3.3.1 Results

Table 3 reports the results of the two approaches we follow to bound or point estimate the MA coefficient and the variances.

In the pooled sample we find that the lowest θ (in absolute value) that satisfies $\sigma_r^2 \geq 0$ is -0.18. The variance of the transitory shock is in the range $\sigma_\varepsilon^2 \in [0.0079, 0.0439]$, while $\sigma_r^2 \in [0, 0.0303]$. For the High School dropouts we find $\sigma_\varepsilon^2 \in [0.0130, 0.0767]$. The corresponding MA parameter is $\theta \in [-1, -0.17]$. For the High School graduates, we find $\sigma_\varepsilon^2 \in [0.0077, 0.0407]$ corresponding to $\theta \in [-1, -0.19]$. Finally, for the College graduates, $\sigma_\varepsilon^2 \in [0.0027, 0.0160]$ with $\theta \in [-1, -0.17]$. The upper bound of the variance of measurement error is 0.0530 for the High School dropouts, 0.0278 for the High School graduate, and 0.0114 for the College graduate.

Bound and Krueger (1994) conduct a validation study of the CPS data on earnings and conclude that measurement error explains 28 percent of the overall variance of the rate of growth of earnings in the CPS. Bound, Brown, Duncan and Rodgers (1994) find a value of 22 percent using the PSID-Validation Study.[15] We assume an intermediate value of 25 percent. Since in our data the earnings growth variances are 0.1651, 0.1033, and 0.0650, respectively for the High School dropouts, the High School graduates, and the College graduates, we calculate (separately for the three education groups): $\overline{\sigma_r^2} = 0.0206, 0.0129$, and 0.0081.[16] We use EWMD to estimate θ, σ_ε^2 and σ_ζ^2 (conditioning on $\overline{\sigma_r^2}$) and find that θ ranges between -0.25 (High School dropout) and -0.51 (College graduate). The variance of the transitory shock is 0.0548 for the High School dropout, 0.0267 for the High School graduate, and 0.0049 for the College graduate.

Table 3
The unconditional variance of income shocks

	Pooled sample	High-school dropout	High-school graduate	College graduate
	Conditioning on the feasible range of θ			
Transitory shock	0.0439	0.0767	0.0407	0.0160
(upper bound)	(0.0070)	(0.0221)	(0.0079)	(0.0070)
Transitory shock	0.0079	0.0130	0.0077	0.0027
(lower bound)	(0.0012)	(0.0038)	(0.0015)	(0.0012)
Measurement error	0.0303	0.0530	0.0278	0.0114
(upper bound)	(0.0021)	(0.0055)	(0.0030)	(0.0018)
θ	−0.18	−0.17	−0.19	−0.17
(upper bound)				
	Using an external estimate of σ_r^2			
Measurement error	0.0138	0.0206	0.0129	0.0081
Transitory shock	0.0300	0.0548	0.0267	0.0049
	(0.0031)	(0.0100)	(0.0039)	(0.0030)
θ	−0.2566	−0.2535	−0.2674	−0.5101
	(0.0339)	(0.0610)	(0.0400)	(0.2901)
Permanent shock	0.0313	0.0331	0.0277	0.0437
	(0.0026)	(0.0067)	(0.0039)	(0.0068)

Note: Equally weighted minimum distance estimates. Standard errors reported in parenthesis.

The above can be viewed as unconditional averages of the underlying (changing) variances and θ. It is however possible to allow for non-stationarity and still be able to identify the

[15] See Bound, Brown and Mathiowetz (2001) for a recent survey of the growing literature on measurement error in micro data.

[16] In the absence of better information, we assume that the fraction of earnings growth variance due to measurement error is the same across education groups (25 percent). Bound and Krueger (1994) provide evidence that measurement error is not correlated with education.

parameters of interest. These are reported, for the pooled sample only, in Table A5 in the Appendix.

4 The conditional variance of earnings

We now specify the conditional variance of the transitory and the permanent shock, thus allowing for non-i.i.d. income innovations. In both cases we specify an ARCH(1) structure of the form

$$\begin{aligned} E_{t-1}\left(\varepsilon_{it}^{2}\right) &= \kappa_{t}^{e} + \gamma^{e}\varepsilon_{it-1}^{2} + \lambda_{i} \quad \textit{Transitory} \\ E_{t-1}\left(\zeta_{it}^{2}\right) &= \phi_{t}^{e} + \varphi^{e}\zeta_{it-1}^{2} + \eta_{i} \quad \textit{Permanent} \end{aligned} \tag{7}$$

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

The terms λ_i and η_i are fixed effects that capture all those elements that are invariant over time and reflect long term occupational choices, etc. The latter reflects permanent variability of income due to factors unobserved by the econometrician. Such variability may in part have to do with the particular occupation or job that the individual has chosen. This variability will be known by the individuals when they make their occupational choices and hence it also reflects preferences. Whether this variability reflects permanent risk or not is of course another issue which cannot be answered here.[17]

[17] An interesting possibility allowed in ARCH models for time-series data is that of asymmetry of response to shocks. In other words, the conditional variance function is allowed to respond asymmetrically to positive and negative past shocks. This could be interesting here as well, for a considerable amount of asymmetry in the distribution of earnings is related to unemployment. Caballero (1990) shows that asymmetric distributions enhance the need for precautionary savings. In our case, however, models embedding the notion of asymmetry are not identifiable. The reason is that we do not observe separately the transitory and permanent shock. Finally, our results are consistent with the presence, in addition to ARCH effects, of a stochastic

4.1 The conditional variance of the transitory shock

In the general MA(q) case the conditional autocovariance of the unexplained component of earnings growth of order $q+1$ identifies the conditional variance of the transitory shock up to the constant MA(q) parameter θ_q^e:

$$E_{t-1}\left(g_{it+q+1}g_{it}\right)=\theta_q^e E_{t-1}\left(\varepsilon_{it}^2\right) \tag{8}$$

Taking the first lag of the transitory variance function in (7), pre-multiplying by the MA coefficient θ_q^e, using (8), and then applying the law of iterated expectations, yields:

$$E_{t-2}\left(g_{it+q+1}g_{it}-\theta_q^e\kappa_t^e-\gamma^e g_{it+q}g_{it-1}-\theta_q^e\lambda_i\right)=0 \tag{9}$$

Two important points can be made about this result. First, identification of the ARCH coefficient requires only the knowledge of the *order* of the MA process *not* the value of the parameters θ_q^e. Second, the ARCH coefficients are identified in the presence of classical or even serially correlated measurement error (so long as that has no ARCH component).

One way to eliminate unobserved heterogeneity is to use first differences. However, this gives rise to weak instruments since it is hard to predict changes in the autocovariances using lagged ones; weak instruments can lead to substantial small sample biases (see Bound, Jaeger and Baker, 1995). We use within groups combined with instrumental variables, which for moderate to a large time series dimension is likely to behave much better. Neglecting time effects for simplicity, the ex-post equivalent of (9) is:

$$g_{it+q+1}g_{it}=\gamma^e g_{it+q}g_{it-1}+\theta_q^e\lambda_i+\omega_{it}$$

with $E_{t-2}\left(\omega_{it}\right)=0$. Even if $\mathrm{var}\left(\lambda_i\right)=0$, an instrumental variable procedure is still required for identification. We then remove individual specific means by

$$\left(g_{it+q+1}g_{it}-\overline{g_{i(+q+1)}g_i}\right)=\gamma^e\left(g_{it+q}g_{it-1}-\overline{g_{i(+q)}g_{i(-1)}}\right)+\left(\omega_{it}-\overline{\omega_i}\right) \tag{10}$$

volatility term. This would only be detectable if it is sufficiently serially correlated.

where $\overline{x_{i(+j)}} = (T-j)^{-1}\sum_{s=1+j}^{T} x_{is}$ and $\overline{x_{i(-j)}} = (T-j)^{-1}\sum_{s=1}^{T-j} x_{is}$. This procedure eliminates the fixed effect but not the endogeneity of $g_{it+q}g_{it-1}$, so we still need to use an IV procedure, and more precisely we need instruments lagged $t-2$, for instance $g_{it-2}g_{it-4}$. It is well known (Nickell, 1981) that the within-group estimator for a dynamic panel data model is biased for short T. However, the bias decreases when T gets large and disappears asymptotically when $T \to \infty$. In our case, T is large enough to make the within group bias probably negligible: T ranges between 9 and 26 for each individual. Moreover, as a sensitivity test both for attrition bias and for this within groups issue, we also report estimates based on individuals observed 16 years or more.

4.2 The conditional variance of the permanent shock

To identify the parameters of the variance function of the permanent shock in equation (7) we make use of the fact that for an MA(q)-transitory shock:

$$E_{t-q-2}\left[g_{it}\left(\sum_{j=-(1+q)}^{(1+q)} g_{it+j}\right)\right] = E_{t-q-2}\left(\zeta_{it}^2\right) \tag{11}$$

It is then easy to prove that the relevant orthogonality condition for this problem is:

$$E_{t-q-3}\left[g_{it}\left(\sum_{j=-(1+q)}^{(1+q)} g_{it+j}\right) - \phi_t^e - \varphi^e g_{it-1}\left(\sum_{j=-(1+q)}^{(1+q)} g_{it+j-1}\right) - \eta_i\right] = 0 \tag{12}$$

Again we only require to know the order of the MA process, not the values of its parameters. Moreover, as for the transitory shock case, the presence of measurement error (even if serially correlated) is allowed for in this moment condition. Thus our estimated ARCH coefficients are robust to the presence of measurement errors.

As before we apply the within groups transformation to eliminate unobserved heterogeneity. We use autocovariances lagged $t-q-3$ or more as instruments.

Estimation requires the availability of a panel data set with a large enough time period of observation for each individual. Data requirements become increasingly stringent as the order of the moving average process for the transitory shock increases. For our case with $q = 1$, the estimation of (9) and (12) requires at least nine and eleven years of data respectively. In this respect the PSID, which is our data source, is ideal for estimation purposes.

Obviously there is some attrition in the sample and by having to use individuals with at least nine observations each we loose those who stay in the data for less. Thus, although we do not need to assume that panel attrition is random, we do need to assume that attrition only depends on *fixed* unobserved and observed characteristics (e.g. on the initial conditions and on the heterogeneity in the variances) and not on the actual shocks. In the empirical section we provide a test for this hypothesis by comparing our main results to those obtained by limiting the sample to those who are observed for more than 16 periods.

4.3 Alternative specifications

The way that the mean process is specified affects the conditional variance estimates for the shocks. A number of papers before us have found that earnings data are consistent with a process comprising a martingale component and a transitory shock (MaCurdy, 1982 and Abowd and Card, 1989 are two of the most prominent examples). However, this approach is not uncontroversial.

A more general model that nests ours is one where the unobservables components of log earnings have a time varying effect as in $u_{it} = d_t^e(k_i + p_{it} + e_{it}) + r_{it}$. The analysis we carry out on the conditional variances could also be carried out on the basis of this mean specification, with a modification to the moment conditions. In other words, this richer specification does not lead to identification problems. Under this model, the long term autocovariances of the residual log earnings growth should not be zero. The results of Table 1 show not only that the autocovariances become insignificant after the second lag but also that they decline rapidly to zero. Both facts are consistent with our specification.

We can also test the null of our model against this one directly. By quasi-differencing we obtain (see Holtz-Eakin, Newey and Rosen, 1988):

$$y_{it} = m_t^e - a_t^e m_{t-1}^e + a_t^e y_{it-1} + \beta_t^{e\prime} Z_{it} - a_t^e \beta_t^{e\prime} Z_{it-1} + d_t^e(\Delta e_{it} + \zeta_{it}) + r_{it} - a_t^e r_{it-1} \qquad (13)$$

where $a_t^e = d_t^e / d_{t-1}^e$ and all parameters in (13) can be identified. Testing that the coefficient on the lagged log earnings is constant over time and equal to one takes us back to the original specification. The joint test that this is true jointly for all three education groups has a p-value of 13%, which confirms the adequacy of our specification.[18]

In the case discussed by Lillard and Reville (1999), $u_{it} = tk_i + e_{it} + r_{it}$ and the residuals of log-income in first differences take the form

$$g_{it} = \Delta r_{it} + \Theta^e(L,q)\Delta\varepsilon_{it} + \zeta_{it} + k_i \tag{14}$$

If the variance of k_i, σ_k^2, is small, it may go statistically undetected in high-order earnings growth autocovariances. However, even a small σ_k^2 may have important implications for the evolution of earnings (Baker, 1997). In this case the unconditional variance of the permanent shock is overstated (and the unconditional variance of the transitory shock understated). However, the orthogonality condition (9) still has the same form, with a modification for the heterogeneity term:

$$E_{t-2}\left(g_{it+q+1}g_{it} - \theta_q^e \kappa_t^e - \gamma^e g_{it+q}g_{it-1} - \Lambda_i\right) = 0$$

where $\Lambda_i = (1-\gamma^e)\, k_i^2 + \theta_q^e \lambda_i$. Thus the ARCH parameters are still identified with the conditions stated earlier, even if we ignore this linear trend.

Lillard and Reville (1999) also point out that the rate of return to experience as measured by age may be heterogeneous. The previous model allows for that through the linear term, since, in the presence of fixed effects a linear time trend and a linear age term are not distinguishable. However, if a quadratic term in age or experience were heterogeneous residual growth will be

$$g_{it} = \Delta r_{it} + \Theta^e(L,q)\Delta\varepsilon_{it} + \zeta_{it} + k_i\left(\alpha_0 + \alpha_1 X_{it}\right)$$

where X_{it} is a measure of experience. In this case (5) takes the form:

$$E\left[g_{it}\left(\sum_{j=-(1+q)}^{(1+q)} g_{it+j}\right)\right] = E\left(\zeta_{it}^2\right) + (3+2q)\left(\alpha_0 + \alpha_1 X_{it}\right)^2 \sigma_k^2 \tag{15}$$

[18]This will also be a test of the special case where $u_{it} = d_t^e k_i + e_{it} + r_{it}$.

resulting in overestimation of the variance of the permanent shock that declines with labor market experience if earnings profiles are concave. Again this model is not consistent with a zero autocovariances of order three or over. In addition, the test of overidentifying restrictions which we use when estimating the ARCH model should reject the ARCH specifications with additive unobserved heterogeneity (equation 9 or 12). Note that in principle one could generalize our approach to allow for this.

4.4 The estimation of the ARCH coefficients

4.4.1 Checking identification conditions

We now proceed to the estimation of the processes for the conditional variance of the permanent and the transitory shocks. We start our analysis by examining whether the instruments we use have sufficient explanatory power to estimate the ARCH model. This is particularly important in our case since for identification we have to depend on information lagged several periods. The recent literature on weak instruments has emphasized the importance of such an exercise (see, Staiger and Stock (1994), Bound, Jaeger and Baker (1995)). To examine this issue, we compute the bootstrap p-value of the χ^2-test for the significance of the excluded instruments in the reduced form. Results are reported in Table IV, based on 300 bootstrap replications.

In the first difference specification for the transitory shock, the reduced forms being examined refer to a linear model predicting $\Delta g_{it+1}g_{it-1}$ using as instruments $g_{it-3}g_{it-5}$ and $g_{it-4}g_{it-6}$, which are dated $t-3$ or earlier. For the within group specification we denote by a bar over a variable its mean over time for individual i. Then the reduced forms refer to a linear model predicting $\left(g_{it+1}g_{it-1} - \overline{g_{i(+1)}g_{i(-1)}}\right)$ in equation 10 using as instruments $g_{it-2}g_{it-4}$ and $g_{it-3}g_{it-5}$. For the permanent shock model in the first difference case we are predicting $\Delta\left[g_{it-1}\left(\sum_{j=-2}^{2} g_{it+j-1}\right)\right]$ using $g_{it-7}\left(\sum_{j=-2}^{2} g_{it+j-7}\right)$ and $g_{it-8}\sum_{j=-2}^{2} g_{it+j-8}$. In the within groups case we are predicting $\left[g_{it-1}\left(\sum_{j=-2}^{2} g_{it+j-1}\right) - \overline{g_{i(-1)}\left(\sum_{j=-2}^{2} g_{i(j-1)}\right)}\right]$ using $g_{it-6}\left(\sum_{j=-2}^{2} g_{it+j-6}\right)$ and $g_{it-7}\sum_{j=-2}^{2} g_{it+j-7}$. All reduced forms and instrument sets include age and time dummies.

For the first difference model the only case that the instruments have some predictive power is for the transitory model when we pool across all education groups. In all other cases the instruments have little predictive power. By contrast, in the reduced forms corresponding to within groups we find very strong predictive power: a p-value well below 1 percent in all cases.

It is very well known that in a dynamic equation, for sufficiently high autoregressive parameters (γ^e or φ^e in equation 7), lags have little or no predictive power for current changes (see Arellano and Bond (1991), Blundell and Bond (1999)). The lack of predictive power we find could thus imply a high autoregressive coefficient (and even perhaps a unit root) in the conditional variance. In principle one of the models could still be estimated in first differences. However, Staiger and Stock (1994) show that even with significant F statistics GMM could be biased considerably when the fit of the reduced form is bad. Our assessment of the results leads us to use the within groups procedure. The latter is likely to be less biased given the weakness of the instruments for the first differenced model and the length of our panel.

Table 4

	Pooled sample	High school dropout	High school graduate	College graduate
		Transitory shock		
First Differences				
χ^2 statistic, 2 degrees of freedom	18.84	17.40	26.56	0.96
Bootstrap p-value reduced form	0.0000	0.1700	0.3633	0.6233
Within Groups				
χ^2 statistic, 2 degrees of freedom	245.52	42.83	267.88	66.75
Bootstrap p-value reduced form	0.0000	0.0033	0.0000	0.0000
		Permanent shock		
First Differences				
χ^2 statistic, 2 degrees of freedom	2.57	1.10	0.10	4.39
Bootstrap p-value reduced form	0.2367	0.8833	0.9333	0.2033
Within Groups				
χ^2 statistic, 2 degrees of freedom	360.73	65.31	289.26	86.45
Bootstrap p-value reduced form	0.0000	0.0000	0.0000	0.0000

We now proceed to the estimation of the processes for the conditional variance of the permanent and the transitory shocks. We start our analysis by examining whether the instruments we use have sufficient explanatory power to estimate the ARCH model. This is particularly important in our case since for identification we have to depend on information lagged several periods. The recent literature on weak instruments has emphasized the importance of such an exercise (see, Staiger and Stock (1994), Bound, Jaeger and Baker (1995)). To examine this issue, we compute the bootstrap p-value of the χ^2-test for the significance of the excluded instruments in the reduced form. Results are reported in Table IV, based on 300 bootstrap replications.

In the first difference specification for the transitory shock, the reduced forms being examined refer to a linear model predicting $\Delta g_{it+1}g_{it-1}$ using as instruments $g_{it-3}g_{it-5}$ and $g_{it-4}g_{it-6}$, which are dated $t-3$ or earlier. For the within group specification we denote by a bar over a variable its mean over time for individual i. Then the reduced forms refer to a linear model predicting $\left(g_{it+1}g_{it-1} - \overline{g_{i(+1)}g_{i(-1)}}\right)$ in equation 10 using as instruments $g_{it-2}g_{it-4}$ and $g_{it-3}g_{it-5}$. For the permanent shock model in the first difference case we are predicting $\Delta\left[g_{it-1}\left(\sum_{j=-2}^{2} g_{it+j-1}\right)\right]$ using $g_{it-7}\left(\sum_{j=-2}^{2} g_{it+j-7}\right)$ and $g_{it-8}\sum_{j=-2}^{2} g_{it+j-8}$. In the within groups case we are predicting $\left[g_{it-1}\left(\sum_{j=-2}^{2} g_{it+j-1}\right) - \overline{g_{i(-1)}\left(\sum_{j=-2}^{2} g_{i(j-1)}\right)}\right]$ using $g_{it-6}\left(\sum_{j=-2}^{2} g_{it+j-6}\right)$ and $g_{it-7}\sum_{j=-2}^{2} g_{it+j-7}$. All reduced forms and instrument sets include age and time dummies.

For the first difference model the only case that the instruments have some predictive power is for the transitory model when we pool across all education groups. In all other cases the instruments have little predictive power. By contrast, in the reduced forms corresponding to within groups we find very strong predictive power: a p-value well below 1 percent in all cases.

It is very well known that in a dynamic equation, for sufficiently high autoregressive parameters (γ^e or φ^e in equation 7), lags have little or no predictive power for current changes (see Arellano and Bond (1991), Blundell and Bond (1999)). The lack of predictive power we find could thus imply a high autoregressive coefficient (and even perhaps a unit root) in the conditional variance. In principle one of the models could still be estimated in first differences. However, Staiger and Stock (1994) show that even with significant F statistics GMM could be biased considerably when the fit of the reduced form is bad. Our assessment of the results leads us to use the within groups procedure. The latter is likely to be less biased given the weakness of the instruments for the first differenced model and the length of our panel.

4.4.2 The conditional variance of the earnings shock

The estimates of the conditional variance function for the transitory shock are based on the orthogonality condition (9) adapted to the within group case. The estimates in the permanent shock case are based on the within group version of the orthogonality condition (12). The within group-instrumental variable estimator we use is the GMM estimator of Hansen (1982) adapted to our problem.

We take $q = 1$ from our earlier results and estimate separately for each education group the following specification:

$$\widetilde{\psi_{it}} = \xi_t + \beta_1 \widetilde{age_{it}} + \beta_2 \widetilde{age^2_{it}} + \gamma \widetilde{\psi_{it-1}} + \widetilde{\omega_{it}} \tag{16}$$

where $\widetilde{x_{it}} = x_{it} - \overline{x_i}$, $\overline{x_i} = T_i^{-1} \sum_{s=1}^{T_i} x_{is}$. For the transitory variance $\psi_{it} = \widehat{g}_{it+2}\widehat{g}_{it}$ and for the permanent variance $\psi_{it} = [\hat{g}_{it}(\hat{g}_{it-2} + \hat{g}_{it-1} + \hat{g}_{it} + \hat{g}_{it+1} + \hat{g}_{it+2})]$; $\widehat{g}_{it}$ is the residual from the estimated mean log earnings process in first differences for individual i in period t..

The instruments are lags of ψ_{it} and are set up as in Arellano and Bond (1991). To avoid overfitting in the reduced forms we do not exploit all the available linear orthogonality conditions. In fact, we truncate the set of available instruments at the first two available lags. Moreover, this should improve the power of the overidentifying restrictions test.[19]

4.4.3 The results

Table 6 reports the WG-GMM results for the coefficients in the conditional variance, as well as a number of diagnostic tests. We report bootstrap standard errors based on 1000 replications as well as the level of significance of the ARCH coefficients based on the bootstrap critical values. The results are presented by education group. We then consider pooling over the whole sample and test whether the ARCH effects vary by education. All specifications

[19] This bias arises from inflating the degrees of freedom of the test with the inclusion of irrelevant instruments.

include time dummies and a quadratic in age. However neither time or age effects are significant in this conditional specification.

Turning now to the ARCH coefficients themselves the highest persistence parameter for both the permanent and the transitory shock is obtained for the High School graduates. For this group, which is by far the largest, the coefficients are unambiguously significant and high, with a 0.67 value for the transitory shock case and nearly 0.9 for the permanent shock. For the High school dropouts the ARCH coefficient for the transitory shock is not significant (p-value of 13%) but for the permanent shock the effect is quite large and has a p-value of 7%, for a sample that is about half that of the high school graduates.

For the College graduates the coefficient on the permanent shock is essentially zero, while for the transitory shock the coefficient is not significant and imprecisely estimated, making it difficult to draw a firm conclusion. We also impose the (acceptable) hypothesis that for this group the transitory shocks follow an MA(0) process ($q = 0$). In this case the estimate of γ^e is 0.2004 (bootstrap s.e. 0.2435) and that of φ^e 0.0221 (0.2045). Thus for the coefficient of the permanent shock the estimation precision goes up substantially but both coefficients remain insignificant.[20]

In Table 6 we also present tests of the overidentifying restrictions (OID). These are all acceptable, except for the model of the variance of the permanent shock for the College Graduates, where the p-value is 2%. This can occur if the variance includes a persistent stochastic volatility component, which could be correlated with the instruments. To address this issue we lagged the instruments one more period. The OID test now has a p-value of 14%. The coefficient φ^e became 0.17 with a standard error of 0.29. Hence, the results are not significantly different to what we obtained before.

[20]The estimated age and age square effects in the transitory shock case are 0.0011 (0.0029) and -0.0008 (0.0034), while in the permanent shock case they are -0.0086 (0.0080) and 0.0107 (0.0093).

Table 6
The conditional variance of earnings shocks

	High school dropout	High school graduate	College graduate
		Transitory shock	
γ^e	0.1918 (0.14)[13%]	0.6672 (0.152)[0.4%]	0.3878 (0.192)[10.3%]
Age	−0.0030 (0.0065)	−0.0005 (0.0029)	−0.0006 (0.0027)
$Age^2/100$	0.0028 (0.0082)	0.0010 (0.0036)	0.0007 (0.0036)
P-value OID test	19%	14%	15%
P-value for time effects	95%	99%	99%
P-value for time and age effects	98%	96%	98%
P-value for unobs. heterogeneity	34.4%	2%	4.3%
γ^e **Pooled (all education groups)**		0.404 (0.096)[0%]	
		Permanent shock	
φ^e	0.3299 (0.21)[7.4%]	0.8912 (0.18)[0.0%]	0.0283 (0.33)[90.4%]
Age	−0.0076 (0.0056)	0.0017 (0.0034)	−0.0065 (0.0055)
$Age^2/100$	0.0090 (0.0063)	−0.0008 (0.0040)	0.0083 (0.0066)
P-value OID test	14%	40%	2%
P-value for time effects	93%	89%	95%
P-value for time and age effects	93%	93%	91%
P-value test for unobs. heterogeneity	0%	2.7%	37.2%
φ^e **Pooled (all education groups)**		0.56 (0.126)[0.3%]	

Notes: Bootstrap standard errors based on 1000 replications reported in round brackets and p-values in square brackets. The OID test is a Sargan test for the null of instruments validity. It is distributed χ^2 with degrees of freedom equal to the number of exclusion restrictions. The test for the null of no unobserved heterogeneity is based on 500 replications.

We now consider a pooled model where the ARCH coefficients are restricted to be equal across education groups. In fact the differences across education groups are not significant. Based on the estimates of Table 6 the χ^2 test that the coefficients are the same for the transitory shock has a p-value of 15%; for the permanent shock the p-value is 6.6%.

When we restrict the three coefficients to be the same, while leaving all the other regressors in the variance to have education specific coefficients, we obtain 0.404 persistence for the transitory shock (standard error 0.096) and 0.56 for the permanent one (standard error 0.128), with important improvements in the precision. Both the restricted coefficients are highly significant and we take these as the most reliable estimates of the ARCH coefficients. These estimates for persistence (*vis-á-vis* the 0.89 coefficient for the High School graduates) imply lower fourth moments for the marginal distribution of the innovations to log earnings for our finitely lived individuals.[21]

A key issue is whether unobserved heterogeneity in the variance is an important factor, i.e., whether there are permanent differences in the volatility of income among individuals, over and above what is accounted for by observables, such as education, time and age. To carry out this test we estimated the model both in levels and within groups; we then used the bootstrap to derive critical values for the $\chi^2_{(1)}$ statistic for the equality of the two coefficients.

For all groups there is evidence of unobserved heterogeneity, at least for one of the two variances. For the High school graduates unobserved heterogeneity is evident in both the variance of the permanent shock and the variance of the transitory shock. For the dropouts, we find unobserved heterogeneity for the variance of the permanent shock only. Finally for the College graduates heterogeneity seems to be important only for the variance of the transitory shock.

A final issue is the extent to which attrition from the PSID has biased our results. We have assumed that attrition is all accounted for by permanent characteristics that are eliminated by our estimation procedure (first differencing for the mean and within groups for the variances). To provide some evidence for this we compare our estimates to those obtained by using only those individuals who are 16 or more years in the sample. This kind of selection mimics attrition bias since it cuts out individuals observed for a shorter time period. The estimates based on this restricted sample are reported in Table A6 of the Appendix. There appears to be some differences, with the coefficients on this selected sample being slightly lower. This is an indication that, if anything, attrition may have

[21] Engle (1982) has shown that if a random variable is conditionally normal with an ARCH process for the variance, the fourth moment of the unconditional distribution is finite only if the ARCH coefficient is lower than $3^{-0.5}$. In our case, with finitely lived individuals the implication of higher levels of variance persistence is a growing fourth order moment of the cross sectional distribution of income innovations over the life-cycle.

biased the ARCH coefficients up. However, when we compute a test statistic comparing the ARCH coefficients in the two data sets we find that none of the differences are significant at conventional levels as reported in the Table.[22]

5 Implications

Given the results above we now provide two examples which illustrate the effects that ARCH can have in explaining a number of important phenomena.

5.1 Income mobility

Allowing for ARCH effects will have implications for earnings mobility since state dependence becomes richer relative to allowing for dynamics in the mean only. As an illustration we focus on the probability of being below the poverty line q_L conditioning on being below q_L in the previous period, i.e. $\Pr(y_{it} < q_L | y_{it-1} < q_L)$.

We compute this measure of poverty persistence using a distribution of income that allows for ARCH effects and a distribution of income that would have been fitted on the same data, but ignoring ARCH effects. In the ARCH case log earnings are generated by the sum of a random walk process where the innovations ζ_{it} are normal and have conditional variance $E_{t-1}(\zeta_{it}^2) = (1-\varphi)\sigma_\zeta^2 + \varphi\zeta_{it-1}^2$ and a transitory component $e_{it} = \varepsilon_{it} - \theta\varepsilon_{it-1}$ where the innovations ε_{it} are also normal and have conditional variance $E_{t-1}(\varepsilon_{it}^2) = (1-\gamma)\sigma_\varepsilon^2 + \gamma\varepsilon_{it-1}^2$. Alternatively we use i.i.d. innovations whose variances are equal to the unconditional variances of the previous process, namely σ_ζ^2 and σ_ε^2.

We focus on a single cohort of individuals, assume a life-cycle of 40 periods and set the parameters σ_ζ^2, σ_ε^2, φ and γ to the values estimated in Table 3 and Table 6. To avoid an initial conditions problem in the variance we draw initial shocks for both cases from the unconditional distributions.

Figure 1 plots the difference between $\Pr(y_{it} < q_L | y_{it-1} < q_L)$ calculated under ARCH-

[22] Fitzgerald, Gottschalk and Moffitt (1998) examine the effect of attrition in the PSID and conclude that attrition propensities are correlated with individual-specific levels of earnings instability, but that their effects on attrition are not large, suggesting that "they are unlikely to induce significant bias in studies which have [...] dynamic measures as outcome variables" (p. 296).

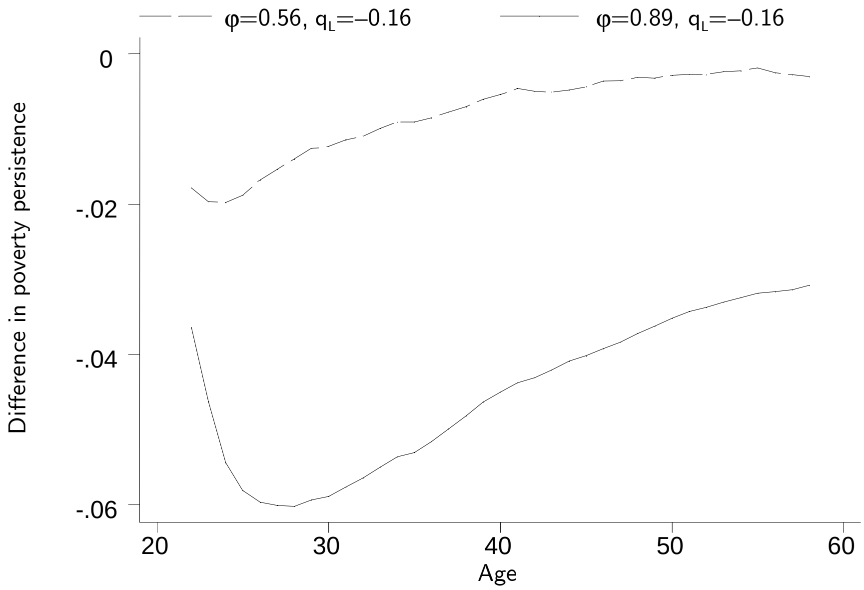

Figure 4: Difference in the persistence of poverty between the ARCH and the i.i.d. case.

type shocks and the same probability calculated under i.i.d. income shocks over the life cycle. We set the poverty line at $q_L = -0.16$ corresponding to the bottom 25^{th} percentile of the income distribution,[23] and $\varphi = \{0.89, 0.56\}$ corresponding to the ARCH estimates for the permanent shock in the High school graduate and Pooled sample case (see Table 6).

The difference in poverty persistence is negative throughout the life cycle. The intuition is that in the lower and upper tails of the distribution individuals with large shocks in period $t-1$ are over-represented. According to the ARCH specification, they have a higher conditional variance (relative to the i.i.d. case) and thus a higher probability of a large shock which will remove them from the tail and push them back to the centre. The reduced persistence on the two extremes of the income distribution is compensated by greater persistence in the central part of the distribution.

Two more facts are worth noting. First, the difference in poverty persistence is higher for the young and tends to decline as time goes by. Second, it is accentuated by higher values of the ARCH parameter φ. Similar results are obtained when setting the poverty line at the 10^{th} percentile.

[23]Since there is no aggregate growth in this model, we set q_L to a constant value.

5.2 Precautionary savings

As Caballero (1990) has stated, the presence of stochastic higher order moments in income, to the extent that they relate to uninsurable uncertainty, is capable of explaining a number of important phenomena that have been observed in consumption data, including excess smoothness, excess growth and excess sensitivity – one mechanism for doing this is the impact of the stochastic process of income on the volatility of consumption growth and on precautionary savings. His analytical results, although suggestive, are not well suited for our income process, which is in logs.

To see the effect of introducing ARCH shocks we consider an approximation to consumption growth. Preferences for consumption c_{it} are CRRA ($u(c_{it}) = \frac{1}{\rho+1} c_{it}^{\rho+1}$, $\rho < 0$), and income is $\log y_{it} = \log y_{it-1} + \zeta_{it}$, where ζ_{it} is conditionally normal with mean zero[24] and $E_{t-1}\left(\zeta_{it}^2\right) = (1-\varphi)\sigma_\zeta^2 + \varphi\zeta_{it-1}^2$. The interest rate is taken equal to the intertemporal discount rate. Adapting the results of Banks, Blundell and Brugiavini (2001) we can write consumption growth as[25]

$$\Delta \log c_{it} = \frac{(1-\rho)}{2}\pi_{it}^2 \left\{ E_{t-1}\left(\zeta_{it}^2\right) + \frac{1}{4}E_{t-1}\left[\left(\zeta_{it}^2 - E_{t-1}\left(\zeta_{it}^2\right)\right)^2\right]\right\} + \Delta \log w_{it} \qquad (17)$$

where π_{it} is the ratio of lagged income to consumption y_{it-1}/c_{it-1} which for the purposes of illustration we take to be 1 and $\Delta \log w_{it}$ is the change in life-cycle wealth from one period to the next.

Individual consumption growth due to precautionary savings is given by the first term on the RHS of (17). For small realizations of ζ_{it-1} (in the $[-\sigma_\zeta, \sigma_\zeta]$ range) consumption growth in the ARCH case is lower than in the i.i.d. case. This is because people forecast a lower variance than the unconditional one – on which precautionary behavior in the i.i.d. case depends upon. This can help explain a high degree of heterogeneity in savings behavior.

In the aggregate there is no such ambiguity however and ARCH always predicts more saving. To see this note that when $\varphi < 3^{-0.5}$ the fourth moment of the ARCH process converges to $3\sigma_\zeta^4\left(\frac{1-\varphi^2}{1-3\varphi^2}\right)$ which is larger than the fourth moment of the long run cross-sectional

[24] We neglect the effect of transitory shocks given that they are smoothed away through savings.

[25] Detailed calculations are available from the authors on request.

income distribution in the absence of ARCH effects (i.e. $3\sigma_{\zeta}^4$). Thus under ARCH, aggregate consumption growth due to precautionary savings increases by $\frac{(1-\rho)}{2}\frac{\varphi^2}{1-3\varphi^2}\sigma_{\zeta}^4$ relative to the i.i.d. case. Using our estimates of φ and σ_{ζ}^2, this is roughly 0.8% when $\rho = -2$ (Attanasio and Weber, 1995) and 1.6% when $\rho = -5$. With aggregate consumption growth averaging 3% in the last few decades, these are sizeable changes.

When the persistence coefficient $\varphi \geq 3^{-0.5}$ the fourth order moment does not converge. With finite lived consumers, the moments never become infinite of course, but this means that aggregate consumption growth within a cohort can be very high. In other words ARCH effects can rationalize steep consumption growth over the life-cycle, even with moderate or small levels of variance in the income shocks.

Both examples are just illustrative of the potential importance of understanding the form of the income process and should not be taken as a full blown analysis, which would need to take into account of a number of other important issues in the respective fields.

6 Conclusions

In this paper we have estimated an income process for individual annual earnings in the US, allowing for differences across education groups and taking into account changes over time. In line with earlier studies we have confirmed that earnings are best described as being driven by a permanent income shock and a serially correlated transitory shock. We have also emphasized the importance of measurement error. We test our specification against a number of alternatives and find ours to be a good description of the data.

We then estimate conditional variance processes for both the transitory and the permanent component, having shown that these processes are separately identified, even in the presence of non-classical measurement error. We find that there is strong evidence of sizeable ARCH effects for both the variances of the transitory and permanent shocks. We also find strong evidence for fixed differences across individuals in the variance of shocks. We then illustrate some potential implications that ARCH effects may have in the field of income mobility and savings. A more complete analysis of these issues, with our richer income process is of course an important area for further work.

Finally there are two issues (among others) that require further research: The extent to which the income shocks we have been modelling represent uninsurable income uncertainty (see Blundell and Preston, 1998); and distinguishing between employment and wage risk separately.

A Appendix

A.1 Step-by-step details on sample selection

The 1968-1993 PSID individual file contains information on 53,013 individuals (all those ever present in the sample). We drop members of the Latino sample added in 1990 (10,022 individuals), and those who are never heads of their household (26,962 individuals). This reduces the sample to 16,029 individuals. We keep only those who are *continuously* heads of their household, who are in the sample for nine years or more, and are aged 25 to 55 over this period. This leaves us with a sample of 4,539 individuals.

We then drop female heads and remain with a sample of 3,663 male heads. We eliminate those with a spell of self-employment over the sample period, missing earnings, and unusable (zero or top-coded) earnings data. This leaves 2,340 individuals. We also drop those with missing education and race records, and those with inconsistent education records. We are left with 2,153 individuals. Finally, we eliminate individuals with outlying earnings records, defined as a change in log earnings greater than 5 or less than -1. The final sample includes 2,069 individuals.

The composition of the sample by year and by education is reported in Tables A1 and A2, respectively. Selected demographic and socio-economic characteristics are reported in Table A3 for selected sample years (1967, 1979, and 1991).

Table A1
Distribution of observations, by year

Year	Number of observations	Year	Number of observations
1967	746	1980	1393
1968	787	1981	1425
1969	831	1982	1455
1970	878	1983	1501
1971	923	1984	1539
1972	1008	1985	1491
1973	1078	1986	1436
1974	1163	1987	1383
1975	1234	1988	1329
1976	1258	1989	1280
1977	1300	1990	1230
1978	1319	1991	1177
1979	1365	1992	1102

Table A2
Distribution of observations, by education

Number of years	Number of individuals			
	Pooled sample	High-school dropout	High-school graduate	College graduate
9	231	64	131	36
10	185	50	96	39
11	182	51	92	39
12	154	28	93	33
13	144	42	81	21
14	142	35	80	27
15	139	34	83	22
16	121	27	67	27
17	123	40	54	29
18	86	20	43	23
19	109	21	60	28
20	93	22	49	22
21	92	16	48	28
22	54	12	25	17
23	52	16	26	10
24	45	12	20	13
25	37	10	17	10
26	80	19	39	22

Table A3
Descriptive statistics: Demographic characteristics

	1967	1979	1992
Age	36.91 (6.34)	37.24 (9.35)	41.21 (5.64)
High school dropout	0.46 (0.50)	0.25 (0.43)	0.12 (0.32)
High school graduate	0.38 (0.49)	0.53 (0.50)	0.61 (0.49)
Hours	2,250 (558)	2,134 (534)	2,151 (540)
Married	0.96 (0.21)	0.88 (0.33)	0.87 (0.33)
White	0.66 (0.48)	0.67 (0.47)	0.70 (0.46)
Children	2.81 (2.11)	1.40 (1.31)	1.33 (1.23)
Family size	4.95 (2.06)	3.58 (1.58)	3.51 (1.44)
Family income	27,146 (13,254)	36,216 (17,542)	43,181 (24,062)
North-East	0.20 (0.40)	0.17 (0.37)	0.16 (0.37)
North-Central	0.25 (0.43)	0.26 (0.44)	0.24 (0.43)
South	0.41 (0.49)	0.42 (0.49)	0.43 (0.50)
SMSA	0.70 (0.46)	0.67 (0.47)	0.53 (0.50)

Table A4
Estimated variances of the permanent shock

Year	Pooled sample	High school dropout	High school graduate	College graduate
1969	0.0280 (0.0125)	0.0221 (0.0155)	0.0244 (0.0241)	0.0576 (0.0321)
1970	0.0242 (0.0124)	0.0119 (0.0105)	0.0117 (0.0042)	0.0847 (0.0657)
1971	0.0136 (0.0077)	0.0144 (0.0163)	0.0067 (0.0043)	0.0296 (0.0189)
1972	0.0223 (0.0135)	−0.0175 (0.0251)	0.0222 (0.0107)	0.1155 (0.0413)
1973	0.0124 (0.0073)	0.0045 (0.0145)	0.0049 (0.0092)	0.0579 (0.0182)
1974	0.0165 (0.0061)	0.0154 (0.0147)	0.0127 (0.0082)	0.0321 (0.0123)
1975	0.0332 (0.0112)	0.0434 (0.0234)	0.0175 (0.0174)	0.0482 (0.0150)
1976	0.0155 (0.0089)	0.0105 (0.0158)	0.0189 (0.0149)	0.0186 (0.0091)
1977	0.0222 (0.0092)	−0.0165 (0.0153)	0.0259 (0.0112)	0.0662 (0.0282)
1978	0.0421 (0.0124)	0.0470 (0.0333)	0.0282 (0.0118)	0.0634 (0.0280)
1979	0.0199 (0.0079)	0.0191 (0.0198)	0.0061 (0.0069)	0.0489 (0.0224)
1980	0.0533 (0.0178)	0.0524 (0.0356)	0.0382 (0.0119)	0.0881 (0.0624)
1981	0.0660 (0.0204)	0.1355 (0.0763)	0.0567 (0.0191)	0.0150 (0.0097)
1982	0.0540 (0.0122)	0.0778 (0.0371)	0.0521 (0.0161)	0.0354 (0.0133)
1983	0.0495 (0.0101)	0.0645 (0.0301)	0.0443 (0.0132)	0.0445 (0.0129)
1984	0.0616 (0.0158)	0.0567 (0.0359)	0.0710 (0.0237)	0.0386 (0.0126)
1985	0.0413 (0.0101)	0.0824 (0.0467)	0.0412 (0.0106)	0.0142 (0.0081)
1986	0.0253 (0.0093)	0.0462 (0.0495)	0.0213 (0.0080)	0.0205 (0.0096)
1987	0.0289 (0.0077)	0.0251 (0.0182)	0.0304 (0.0100)	0.0260 (0.0159)
1988	0.0309 (0.0095)	0.0128 (0.0200)	0.0371 (0.0146)	0.0227 (0.0089)
1989	0.0340 (0.0087)	0.0865 (0.0428)	0.0283 (0.0101)	0.0173 (0.0054)
1990	0.0293 (0.0070)	0.0498 (0.0254)	0.0290 (0.0095)	0.0159 (0.0048)
1991	0.0438 (0.0209)	0.0516 (0.0248)	0.0550 (0.0335)	0.0123 (0.0083)
χ^2 $(d.o.f)$	322.81 (275)	1648.65 (275)	354.89 (275)	1784.31 (275)

Note: In Table A4 and A5 we impose equality of the permanent shock variances in the first two and in the last two years of the sample period. This is to avoid instability when few moments are used for identification. The value reported at the bottom of each table is a goodness of fit statistic for the estimated model. It is defined as $(\mathbf{c}-\mathbf{f}(\boldsymbol{\beta}))'\mathbf{R}^{-}(\mathbf{c}-\mathbf{f}(\boldsymbol{\beta}))$, where $\mathbf{c}$ is the vector of estimated moments, $f(\boldsymbol{\beta})$ is the theoretical counterpart we attempt to fit, and $\mathbf{R}^{-}$ is a generalized inverse of $\mathbf{R}=\mathbf{WVW}$ with $\mathbf{W}=\mathbf{I}-\mathbf{G}(\mathbf{G}'\mathbf{G})^{-1}\mathbf{G}'$ and $\mathbf{V}$ the matrix of empirical fourth moments. The test statistic is asymptotically distributed χ^2 with degrees of freedom equal to the difference between the number of elements in $\mathbf{c}$ and the rank of the gradient matrix $\mathbf{G}$. See Newey (1985) for more details.

Table A5
Conditioning on an external estimate of σ_r^2

Year	$E\left(\zeta_{it}^2\right)$	$E\left(\varepsilon_{it}^2\right)$
1967	-.-	0.0119 (0.0212)
1968	-.-	0.0074 (0.0064)
1969	0.0298 (0.0134)	0.0132 (0.0069)
1970	0.0274 (0.0124)	0.0153 (0.0068)
1971	0.0173 (0.0081)	0.0352 (0.0126)
1972	0.0248 (0.0133)	0.0605 (0.0210)
1973	0.0136 (0.0071)	0.0242 (0.0067)
1974	0.0158 (0.0062)	0.0313 (0.0125)
1975	0.0334 (0.0112)	0.0819 (0.0219)
1976	0.0166 (0.0091)	0.0618 (0.0286)
1977	0.0218 (0.0098)	0.0136 (0.0064)
1978	0.0417 (0.0125)	0.0113 (0.0067)
1979	0.0213 (0.0079)	0.0267 (0.0094)
1980	0.0552 (0.0178)	0.0156 (0.0062)
1981	0.0652 (0.0200)	0.0168 (0.0106)
1982	0.0525 (0.0123)	0.0509 (0.0151)
1983	0.0500 (0.0100)	0.0430 (0.0153)
1984	0.0623 (0.0158)	0.0345 (0.0153)
1985	0.0429 (0.0099)	0.0230 (0.0074)
1986	0.0269 (0.0094)	0.0266 (0.0084)
1987	0.0312 (0.0077)	0.0090 (0.0057)
1988	0.0316 (0.0095)	0.0224 (0.0122)
1989	0.0368 (0.0087)	0.0005 (0.0063)
1990	0.0309 (0.0069)	0.0212 (0.0111)
1991	0.0483 (0.0210)	0.0058 (0.0057)
1992	-.-	0.0227 (0.0237)
θ	-0.2030 (0.0308)	
χ^2 $(d.o.f.)$	325.06 (275)	

Table A6
ARCH estimates using individuals observed 16 years or more

	HS dropout	HS graduate	College graduate
		Transitory shock	
γ^e	0.215 (0.119)	0.190 (0.102)	0.232 (0.156)
P-value attrition test	84.2%	5.4%	51%
		Permanent shock	
φ^e	0.289 (0.216)	0.584 (0.214)	−0.041 (0.417)
P-value attrition test	30.8%	19.2%	81.2%

References

[1] Abowd, John and David Card (1989), "On the covariance structure of earnings and hours changes", *Econometrica*, **57**, 411-45.

[2] Altonji, Joseph G., Ana Paula Martins, and Aloysius Siow (2001), "Dynamic factor models of consumption, hours and income", *Research in Economics* (forthcoming).

[3] Altonji, Joseph G., and Lewis M. Segal (1996), "Small-sample bias in GMM estimation of covariance structures", *Journal of Business and Economic Statistics,* **14**, 353-66.

[4] Alvarez, Javier, Martin Browning, and Mette Ejrnaes (2002), "Modelling income processes with lots of heterogeneity", *mimeo,* Universidad de Alicante and University of Copenhagen.

[5] Arellano, Manuel, and Stephen Bond (1991), "Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations", *Review of Economic Studies*, **58**, 277-97.

[6] Attanasio, Orazio, and Guglielmo Weber (1995), "Is consumption growth consistent with intertemporal optimization? Evidence from the Consumer Expenditure Survey", *Journal of Political Economy*, **103**, 1121-57

[7] Baker, Michael (1997), "Growth-rate heterogeneity and the covariance structure of life-cycle earnings", *Journal of Labour Economics*, **15**, 338-375.

[8] Banks, James, Richard Blundell and Agar Brugiavini (2001), "Risk pooling and consumption growth", *Review of Economic Studies.*

[9] Blundell, Richard W., and Stephen Bond (1999), "Initial conditions and moment restrictions in dynamic panel data models", *Journal of Econometrics,* **87**, 115-43.

[10] Blundell, Richard W., and Ian Preston (1998), "Consumption inequality and income uncertainty", *Quarterly Journal of Economics,* **113**, 603-640.

[11] Bound, John, Charles Brown, Greg J. Duncan, and Willard L. Rodgers (1994), "Evidence on the validity of cross-sectional and longitudinal labor market data", *Journal of Labor Economics,* **12**, 345-68.

[12] Bound, John, Charles Brown and Nancy Mathiowetz (2001), "Measurement error in survey data", in James Heckman and Edward Leamer, eds., *Handbook of Econometrics,* vol. 5, Elsevier Science. Amsterdam: North Holland.

[13] Bound, John, David Jaeger and Regina Baker (1995) "Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak", *Journal of the American Statistical Association,* **90**, 443-50.

[14] Bound, John, and Alan Krueger (1994), "The extent of measurement error in longitudinal earnings data: Do two wrongs make a right?", *Journal of Labor Economics,* **9**, 1-24.

[15] Browning, Martin, Lars P. Hansen, and James J. Heckman (2000), "Micro data and general equilibrium models", in John B. Taylor and Michael Woodford, eds., *Handbook of Macroeconomics*, Amsterdam: North Holland.

[16] Caballero, Ricardo J. (1990), "Consumption puzzles and precautionary savings", *Journal of Monetary Economics*, **25**, 113-36.

[17] Card, David, and Thomas Lemieux (2001), "Can falling supply explain the rising return to college for younger men? A cohort-based analysis", *Quarterly Journal of Economics*, **116**, 705-46.

[18] Carroll, Cristopher and Andrew Samwick (1997), "The nature of precautionary saving", *Journal of Monetary Economics*, **40**, 41-71.

[19] Chamberlain, Gary, and Keisuke Hirano (1999), "Predictive distributions based on longitudinal earnings data", *Annales d'Economie et de Statistique*, **55-56**, 211-42.

[20] Deaton, Angus (1992), *Understanding consumption.* Oxford: Oxford University Press.

[21] Engle, Richard F. (1982), "Autoregressive conditional heteroscedasticity with estimates of the variance of the United Kingdom inflation", *Econometrica*, **50**, 987-1007.

[22] Fitzgerald, John, Peter Gottschalk, and Robert Moffitt (1998), "An analysis of sample attrition in panel data: The Michigan Panel Study of Income Dynamics", *Journal of Human Resources*, **33**(2), 251-99.

[23] Gosling, Amanda, Stephen Machin, and Costas Meghir (2000), "The changing distribution of male wages in the UK", *The Review of Economic Studies*, **67**, 635-66.

[24] Gottschalk, Peter and Robert Moffitt (1993), "The growth of earnings instability in the US labour market", *Brooking Papers on Economic Activity*, 217-272.

[25] Guiso, Luigi, Tullio Jappelli, and Luigi Pistaferri (2002), "An empirical analysis of earnings and employment risk", *Journal of Business and Economic Statistics,* **20**, 241-53.

[26] Hall, Peter, and Joel L. Horowitz (1996), "Bootstrap critical values for tests based on Generalized-Method-of-Moments estimators", *Econometrica*, **64**, 891-916.

[27] Hansen, Lars P. (1982), "Large sample properties of generalized method of moments estimators", *Econometrica,* **50**, 1029-54.

[28] Hill, Martha S. (1992), *The Panel Study of Income Dynamics: A user's guide*, Newbury Park, California: Sage Publications.

[29] Holtz-Eakin, Daniel, Withney Newey and Harvey Rosen (1988), "Estimating Vector Autoregressions with Panel Data" *Econometrica*, **56**, 1371-95.

[30] Horowitz, Joel L. (2002), "The bootstrap", in James Heckman and Edward Leamer, eds., *Handbook of Econometrics,* vol. 5, Elsevier Science. Amsterdam: North Holland.

[31] Hubbard, Glenn, Jonathan Skinner, and Stephen Zeldes (1994), "The importance of precautionary motives in explaining individual and aggregate saving", *Carnegie Rochester Series on Public Policy*, **40**, 59-125.

[32] Juhn, Chinhui, Kevin Murphy, and Brooks Pierce (1993), "Wage inequality and the rise in returns from skills", *Journal of Political Economy,* **101**, 410-42.

[33] Levy, Frank, and Richard J. Murnane (1992), "U.S. earnings levels and earnings inequality: A review of recent trends and proposed explanations", *Journal of Economic Literature*, **30**, 1333-81.

[34] Lillard, E. Lee, and Robert T. Reville (1999), "Life cycle human capital investment: New evidence on an old issue", University of Michigan and RAND, mimeo.

[35] Lillard, Lee, and Robert Willis (1978), "Dynamic aspects of earnings mobility", *Econometrica*, **46**, 985-1012.

[36] MaCurdy, Thomas E. (1982), "The use of time series processes to model the error structure of earnings in a longitudinal data analysis", *Journal of Econometrics*, **18**, 82-114.

[37] Meghir, Costas, and Frank Windemejier (1999), "Moment conditions for dynamic panel data models with multiplicative individual effects in the conditional variance", *Annales d'Economie et de Statistique*, **55-56**, 317-30.

[38] Moffitt, Robert and Peter Gottschalk (1994), "Trends in the autocovariance structure of earnings in the US: 1969-1987", Brown University, Department of Economics, mimeo.

[39] Newey, Whitney K. (1985), "Generalized method of moments specification tests", *Journal of Econometrics,* **29**, 229-56.

[40] Nickell, Stephen (1981), "Biases in dynamic models with fixed effects", *Econometrica,* **49**, 1417-26.

[41] Staiger, Douglas, and James H. Stock (1994), "Instrumental variables regression with weak instruments", *Econometrica*, **65**, 557-86.

[42] Windmeijer, Frank (2000), "A finite sample correction for the variance of linear two-step GMM estimators", IFS Working Paper 00/19.

[43] Zeldes, Stephen (1992), "Comment" to C. Carroll, "Buffer stock saving and the Life Cycle/Permanent Income Hypothesis", *Brookings Papers on Economic Activity*, **2**, 61-156.