Review of Economic Dynamics 12 (2009) 58–79

Contents lists available at ScienceDirect

Review of Economic Dynamics

www.elsevier.com/locate/red

An empirical investigation of labor income processes

Fatih Guvenen [a,b]

[a] University of Minnesota, Department of Economics, 4-101 Hanson Hall, Minneapolis, USA
[b] NBER, USA

ARTICLE INFO

Article history:
Received 20 January 2006
Revised 18 June 2008
Available online 1 July 2008

JEL classification:
C33
D31
J31

Keywords:
Idiosyncratic income risk
Incomplete markets models
Heterogeneous income profiles
Random walk model

ABSTRACT

In this paper, I reassess the evidence on labor income risk. There are two leading views on the nature of the income process in the current literature. The first view, which I call the "Restricted Income Profiles" (RIP) process, holds that individuals are subject to large and very persistent shocks, while facing similar life-cycle income profiles. The alternative view, which I call the "Heterogeneous Income Profiles" (HIP) process, holds that individuals are subject to income shocks with modest persistence, while facing individual-specific income profiles. I first show that ignoring profile heterogeneity, when in fact it is present, introduces an upward bias into the estimates of persistence. Second, I estimate a parsimonious parameterization of the HIP process that is suitable for calibrating economic models. The estimated persistence is about 0.8 in the HIP process compared to about 0.99 in the RIP process. Moreover, the heterogeneity in income profiles is estimated to be substantial, explaining between 56 to 75 percent of income inequality at age 55. I also find that profile heterogeneity is substantially larger among higher educated individuals. Third, I discuss the source of identification—in other words, the aspects of labor income data that allow one to distinguish between the HIP and RIP processes. Finally, I show that the main evidence against profile heterogeneity in the existing literature—that the autocorrelations of income changes are small and negative—is also replicated by the HIP process, suggesting that this evidence may have been misinterpreted.

1. Introduction

The nature of labor income risk plays a central role in many economic decisions that individuals make. Among many economic questions that hinge on income risk are the life-cycle consumption and portfolio choice behavior (Carroll and Samwick, 1997; Campbell et al., 2001; Gourinchas and Parker, 2002; Guvenen, 2007), the determination of wealth inequality (Huggett, 1996; Castaneda et al., 2003), the welfare costs of business cycles (Storesletten et al., 2001; Lucas, 2003), and the determination of asset prices (Constantinides and Duffie, 1996). The conclusions that a researcher reaches in these analyses, clearly, depend on the properties of the labor income process used to calibrate these models.

There are two leading views about the nature of the income process in the current literature. To provide context for the following discussion, suppose that the log labor income of individual i with h years of labor market experience is given by[1]:

E-mail address: fatihguvenen@gmail.com.
URL: http://www.econ.umn.edu/~guvenen.

[1] This income process is a simplified version of the models estimated in the literature, but still captures the components necessary for the present discussion. I study more general processes in Section 2.

1094-2025/$ – see front matter
doi:10.1016/j.red.2008.06.004

$$y_h^i = \beta^i h + z_h^i,$$
$$z_h^i = \rho z_{h-1}^i + \eta_h^i, \quad (1)$$

where β^i is the individual-specific income growth rate with cross-sectional variance σ_β^2; and η_h^i is the innovation to the AR(1) process with variance σ_η^2. In this preliminary discussion I abstract from heterogeneity in the intercept of income.

The early papers on income dynamics estimated versions of the process given in (1) from labor income data and found: $0.5 < \rho < 0.7$, and $\sigma_\beta^2 \gg 0$ (cf., Lillard and Weiss, 1979; Hause, 1980; and more recently Baker, 1997; and Haider, 2001). Thus according to this first view, which I call the "Heterogeneous Income Profiles" (HIP) model, individuals are subject to shocks with modest persistence, while facing life-cycle profiles that are individual-specific (and hence vary significantly across the population). One theoretical motivation for this specification is the human capital model, which implies differences in income profiles, for example, if individuals differ in their ability level.[2]

In an influential paper, MaCurdy (1982) cast doubt on these findings. He tested—and did not reject—the restriction $\sigma_\beta^2 = 0$ against the more general alternative of HIP. He then estimated versions of the income process given in (1) by *imposing* $\sigma_\beta^2 \equiv 0$, and found $\rho \approx 1$ (see also Abowd and Card, 1989; Topel, 1990; and Topel and Ward, 1992). Therefore, according to this alternative view, which I call the "Restricted Income Profiles" (RIP) model, individuals are subject to extremely persistent—nearly random walk—shocks, while facing similar life-cycle income profiles (conditional on some observable characteristics).

In this paper, I examine labor income data from several angles to help distinguish between these two income processes. I begin the analysis by showing that assuming away the heterogeneity in income growth rates (as is done in the RIP process), when in fact it is present, biases the estimated persistence parameter upward. It is easy to see why this happens: an individual with high (alternatively, low) income growth rate will systematically deviate from the average profile. Ignoring this fact will then lead the econometrician to interpret this *systematic* fanning out as the result of persistent positive (or negative) income shocks every period. I study an example which shows that this bias can be substantial: when labor income is generated by the HIP process given above (Eq. (1)) with i.i.d shocks, the persistence parameter is estimated to be about 0.90 if RIP is assumed, instead of the true value of zero. This example, therefore, suggests that allowing for heterogeneity in income growth rates can be critical for the consistent estimation of the persistence parameter.

I next estimate the HIP and RIP versions of a general labor income process. The stochastic component of the income process has typically been modeled in one of two ways in the literature. Following MaCurdy (1982), several studies have modeled the dynamics with an ARMA(1,1) or (1,2) process (among others, Abowd and Card, 1989; Meghir and Pistaferri, 2004). While this specification is quite flexible and provides a good description of income dynamics, it has one obvious drawback when used as input into an economic model: the ARMA(1,1) and (1,2) processes require two and three state variables, respectively, to form optimal forecasts of the income process. Consequently, the majority of the existing life-cycle (or overlapping generations) models are instead calibrated using a somewhat simpler specification—one that features an AR(1) component plus a transitory shock.[3] This specification introduces only one state variable into a dynamic programming problem, and provides a good compromise between fit and parsimony. However, the existing estimates of the HIP process in the literature also feature ARMA processes.[4] Therefore, a second contribution of this paper is to estimate a HIP process, where the stochastic component is modeled parsimoniously as an AR(1) process plus a transitory shock, making it suitable as a basis for calibration.

Using data from the Panel Study of Income Dynamics (PSID) covering 1968 to 1993, I find statistically and quantitatively significant heterogeneity in income profiles. Furthermore, the persistence of income shocks is estimated to be about 0.8 in the HIP process, implying that between 65 to 80 percent of income inequality at the age of retirement is due to heterogeneous profiles. Instead, when the specification is restricted to be a RIP process, the estimated persistence is 0.99 (and a random walk specification could not be rejected).

Third, I examine the differences in income processes across education groups. While several studies have investigated this question in the context of RIP processes (Hubbard et al., 1994; Carroll and Samwick, 1997), there exists no corresponding analysis in the context of HIP processes. I find that in the HIP process there is a major difference between the two groups: the dispersion of income growth rates, σ_β^2, is more than twice as large for college graduates than it is for high school graduates. This is in contrast to the estimates from the RIP process, which implies similar income processes for both groups (with some mild evidence of larger innovation variances for lower educated individuals).

I next turn to identification. In particular, I examine what features of labor income data help us distinguish between the HIP and RIP processes. The panel structure is essential in this respect, because it allows us to characterize the evolution of the *cross-sectional* distribution of income as a cohort gets older. As I explain in Section 4, the autocovariance structure implied by the two income processes differ in important ways, which make it possible to distinguish between them. While

[2] Becker (1964) and Ben-Porath (1967) contain classic treatments of these models. More recently, Guvenen and Kuruscu (2007) apply a human capital model with ability heterogeneity to understand the trends in wage inequality in the US data since the 1970s.

[3] Among many others, see Hubbard et al. (1995), Huggett (1996), Campbell et al. (2001), Storesletten et al. (2001), Heathcote et al. (2004).

[4] Baker experiments with an AR(1) process to provide a comparison to Lillard and Weiss (1979). But in this specification he does not (nor do Lillard and Weiss) allow for a separate serially independent shock. As is well known, classical measurement error biases estimates of persistence downward when transitory shocks are not allowed.

this analysis clarifies how theoretical identification is obtained, it also highlights some empirical difficulties with identification: basically, higher order autocovariances provide valuable information for identification, but because of sample attrition, fewer and fewer individuals contribute to these moments. This not only increases the noise in these autocovariances, but perhaps more importantly, raises questions about potential selectivity bias. As a result, it is not clear that income data alone can provide a definitive verdict on the nature of income risk. An alternative approach would be to exploit the information embedded in individuals' economic choices (which contain information about how they perceive future income risks) and use them in conjunction with income data.

Finally, I try to reconcile the test used by MaCurdy and others which does not reject the RIP process, with the direct estimation results which lend support to the HIP process. In related work, Baker (1997) has conducted a careful Monte Carlo study and argued that the test lacks power in small sample against the alternative of HIP. Here I emphasize a different point that applies even in large sample, where inflated size or low power are not relevant. I argue that the tests used by MaCurdy (1982) and Abowd and Card (1989) are not appropriate for distinguishing between the RIP and HIP processes if the true data generating process (such as HIP) contains an AR(1) component with $\rho < 1$. To see this point, first it is easily shown from Eq. (1) that

$$\operatorname{cov}\left(\Delta y_h^i, \Delta y_{h+n}^i\right) = \sigma_\beta^2 - \left[\rho^{n-1}\left(\frac{1-\rho}{1+\rho}\right)\sigma_\eta^2\right] \quad \text{for } n \geqslant 2.$$

Notice that the term in brackets vanishes as n gets large, so *higher order* autocovariances of income changes must be positive if indeed $\sigma_\beta^2 > 0$. This observation forms the basis of MaCurdy's test. A key question however is: What is the lowest lag at which the covariances should become positive? This is important because the aforementioned studies have focused on the first 5 to 10 lags. By substituting the parameter values estimated in Section 3 into the expression above, one can easily show that in the HIP process the first 11 covariances will be negative (see Fig. 8), despite the fact that those estimates imply substantial heterogeneity in income profiles. This point suggests that the negative covariances of income changes reported in the literature is also implied by the HIP process. In Section 5, I show that the autocovariance and autocorrelation structures generated by the estimated HIP process are also quantitatively similar to their empirical counterparts. Moreover, even though autocovariances should eventually become positive according to the HIP process, in sample sizes close to those used in the literature (less than 30,000 observations) even the 20th autocovariance will not be significantly positive. These results cast doubt on the previous interpretation of this evidence in the literature as supporting the RIP process.

The rest of the paper is organized as follows. The next section describes the data and the estimation method. Section 3 presents the empirical results and quantifies the heterogeneity in income growth rates. Section 4 discusses identification. Section 5 reconciles the direct estimation evidence with earlier tests implemented in the literature, and Section 6 concludes.

2. Empirical analysis

2.1. The PSID data

This section briefly describes the data and the variables used in the empirical analysis. The labor earnings data are drawn from the first 26 waves of PSID covering the period from 1968 to 1993. The main sample consists of male head of households between the ages of 20 and 64. I include an individual into the sample if he satisfies the following conditions for twenty (not necessarily consecutive) years: the individual has (1) reported positive labor earnings and hours; (2) worked between 520 and 5110 hours in a given year; (3) had an average hourly earnings between a preset minimum and a maximum wage rate (to filter out extreme observations). I also exclude individuals who belong to the poverty (SEO) subsample in 1968. These criteria are similar to the ones used in previous studies (Abowd and Card, 1989; Baker, 1997; and Heathcote et al., 2004, among others). Further details of the selection criteria are contained in Appendix A.

These criteria leave us with a main sample of 1270 individuals with at least twenty years of data on each. To study the labor income processes of different education groups separately, I further draw two subsamples: the first contains 335 individuals with at least a four-year college degree (sixteen years of education or more), and the second contains 882 individuals with at most a high school degree (fifteen years of education or less). To make the text more readable, I will refer to the former group as "college-educated" and the latter as "high school educated," at the expense of a slight abuse of language. The measure of labor income includes wage income, bonuses, commissions, plus the labor portions of several types of income such as farm income, business income, etc. Labor income in PSID refers to the previous year, so the data covers 1967–1992.

Finally, following the bulk of the existing literature, the measure of labor market experience of an individual I use is "potential" experience defined as $h = (\text{age} - \max(\text{years of schooling}, 12) - 6)$. Although, ideally, it would be preferable to use a measure of actual experience, constructing a reliable measure using PSID data is not straightforward because it requires either observing individuals for all years since they enter the labor market which would substantially reduce the sample size or relying on retrospective questions that ask workers to recall how many years they have worked since entering labor market, which may not be reliable. Partly reflecting these concerns, all existing studies on income dynamics using the PSID data uses potential experience as I do here (including MaCurdy, 1982; Abowd and Card, 1989; Baker, 1997; Meghir and Pistaferri, 2004; Heathcote et al., 2005; among others). A potential caveat is that if high skill workers have more actual experience at a given age compared to low-skill workers (perhaps because they work longer hours or have a lower unemployment rate), using potential experience may underestimate their actual experience compared to low-skill workers. However, the fact that I am focusing on workers who remain in the sample for a relatively long period of time (and therefore have more stable employment histories) should partly mitigate this problem. Further details on variable definitions and some summary statistics for the primary sample are contained in Appendix A.

2.2. A statistical model

In this section, I generalize the income process in Eq. (1) to make it suitable for empirical analysis. Specifically, the process for *log labor earnings*, $y^i_{h,t}$, of individual i with h years of labor market experience in year t is given by

$$y^i_{h,t} = g(\boldsymbol{\theta}^0_t, \mathbf{X}^i_{h,t}) + f(\alpha^i, \beta^i, \mathbf{X}^i_{h,t}) + z^i_{h,t} + \phi_t \varepsilon^i_{h,t} \tag{2}$$

where $i = 1, \ldots, I$; $h = 1, \ldots, H$, and $t = 1, \ldots, T$.

The functions g and f denote the "life cycle" components of earnings. The first one, g, captures the part of variation that is common to all individuals (which is why the coefficient vector $\boldsymbol{\theta}^0_t$ is not indexed by i) and is assumed to be a cubic polynomial in experience, h. Notice that the coefficients of this polynomial are allowed to be time-varying. In addition to the standard time effects (aggregate shocks) in labor income movements captured by year-to-year variations in the intercept of g, this flexible specification also allows us to model some important changes that took place in the labor market during the sample period. For example, changes in the return to experience that took place during this period (Katz and Autor, 1999) are accounted for by the time-varying higher order terms in experience. Although, it is also possible to capture the rise in the skill premium during this period (Katz and Murphy, 1992) by adding an education dummy into g, I do not pursue this approach in the baseline specification. (Instead I capture all the cross-sectional variation in income growth rates in the f function). Later in the paper, I will estimate a separate income process for each education group to fully control for the effect of education on the life-cycle profiles as well as its effect on the persistence and variance of income shocks.

2.2.1. Heterogeneity in income profiles

The second function, f, is the centerpiece of the present analysis, and captures the component of life-cycle earnings that is individual- or group-specific. For example, if the growth rate of earnings varies with the ability of a worker, or is different across occupations, this variation will be reflected in an individual- or occupation-specific slope coefficient in f. I assume this function to be linear in experience: $f(\alpha^i, \beta^i, \mathbf{X}^i_{h,t}) \equiv \alpha^i + \beta^i h$, where the random vector (α^i, β^i) is distributed across individuals with zero mean, variances of σ^2_α and σ^2_β, and covariance of $\sigma_{\alpha\beta}$.[5]

Although it is straightforward to generalize f to allow for heterogeneity in higher order terms, Baker (1997, p. 373) finds that this extension does not noticeably affect parameter estimates or improve the fit of the model. In addition, recall that one goal of this study is to estimate an income process that is parsimonious enough to be used for calibrating macroeconomic models. However, each additional term introduced into f will appear as an additional state variable in a dynamic programming problem (see, for example, Guvenen, 2007). The current specification provides a reasonable trade-off for this purpose.[6]

2.2.2. Modeling the dynamics of income

The dynamic component of income is modeled as an AR(1) process plus a purely transitory shock. This specification is fairly common in the literature and, despite its parsimonious structure, it appears to provide a good description of income dynamics in the data (Topel, 1990; Hubbard et al., 1994; Moffitt and Gottschalk, 1995; Storesletten et al., 2004).[7] The AR(1) process can capture mean-reverting shocks, such as human capital innovations that depreciate over time, or long-term nominal wage contracts whose value decreases over time in real terms, as well as fully permanent shocks as a special case. Furthermore, there have been some significant changes in the sizes of both persistent and transitory income shocks over the sample period under study (cf., Moffitt and Gottschalk, 1995; Meghir and Pistaferri, 2004). To capture this non-stationarity, I write $z^i_{h,t}$ as an AR(1) process with heteroskedastic shocks:

$$z^i_{h,t} = \rho z^i_{h-1,t-1} + \pi_t \eta^i_{h,t}, \quad z^i_{0,t} = 0, \tag{3}$$

where π_t captures possible time-variation in the innovation variance. Similarly, the transitory shock in Eq. (2), $\varepsilon^i_{h,t}$, is scaled by ϕ_t to account for possible non-stationarity in that component. The innovations $\eta^i_{h,t}$ and $\varepsilon^i_{h,t}$ are assumed to be independent of each other and over time (and independent of α^i and β^i), with zero mean, and variances of σ^2_η and σ^2_ε respectively. Furthermore, measurement error is a pervasive problem in micro data sets, and income data in PSID is no

[5] The zero-mean assumption is merely a normalization since g already includes an intercept and a linear term. Thus, in any given year, the population averages of the intercept and slope are given by the first two coefficients of g.

[6] Lillard and Reville (1999) on the other hand, provide some evidence suggesting that the quadratic term may be important so this seems to be an extension worth considering in future work.

[7] As noted earlier, although it is also possible to model dynamics using an unrestricted ARMA (1, 1) or (1, 2) process, the resulting specification introduces additional state variables into dynamic programming problems, making it unsuitable for the present purpose.

exception. This measurement error will be captured in the transitory component if it is serially independent, or will be included in $z^i_{h,t}$ if it has an autoregressive component (Bound and Krueger, 1991). It is important to keep this point in mind when interpreting the empirical findings in the next section.

The income *residual*, $\hat{y}^i_{h,t}$, is obtained by regressing $y^i_{h,t}$ on the polynomial g. Since the individual-specific parameters, α^i and β^i, are not observable, f is treated as part of the random component of the income process and is included in the residual. For a given year, the cross-sectional second-order moments of this residual for a cohort of a given age are:

$$\begin{aligned}
&\operatorname{var}(\hat{y}^i_{h,t}) = \left[\sigma^2_\alpha + 2\sigma_{\alpha\beta}h + \sigma^2_\beta h^2\right] + \operatorname{var}(z^i_{h,t}) + \phi^2_t\sigma^2_\varepsilon, \\
&\operatorname{cov}(\hat{y}^i_{h,t}, \hat{y}^i_{h+n,t+n}) = \left[\sigma^2_\alpha + \sigma_{\alpha\beta}(2h+n) + \sigma^2_\beta h(h+n)\right] + \rho^n \operatorname{var}(z^i_{h,t}),
\end{aligned} \tag{4}$$

where $n = 1, \ldots, \min(H-h, T-t)$, and the variance of the AR(1) component is obtained recursively:

$$\begin{aligned}
&\operatorname{var}(z^i_{1,t}) = \pi^2_t\sigma^2_\eta, \\
&\operatorname{var}(z^i_{h,1}) = \pi^2_1\sigma^2_\eta\sum_{j=0}^{h-1}\rho^{2j}, \quad t=1,\ h>1, \\
&\operatorname{var}(z^i_{h,t}) = \rho^2\operatorname{var}(z^i_{h-1,t-1}) + \pi^2_t\sigma^2_\eta, \quad t>1, h>1.
\end{aligned} \tag{5}$$

Note that in the first line I implicitly assume that the initial value of the persistent shock is zero for all individuals. In the second line I assume that the innovation variance was constant over time before the sample started in 1968, so that the cross-sectional variance for a cohort aged h in the first year of the sample can be determined by the accumulated effect over the last h years.[8]

My estimation strategy (first proposed by Chamberlain, 1984) is based on minimizing the "distance" between the elements of the $(T \times T)$ empirical covariance matrix of income residuals (denote it by $\mathbf{C}$) and its counterpart implied by the statistical model described above. A typical element of $\mathbf{C}$ (at location $(\tau, \tau+n)$) is obtained by averaging $(\hat{y}^i_{h,\tau}\hat{y}^i_{h+n,\tau+n})$ across individuals of all ages who were present in these two years. The theoretical counterpart is calculated by aggregating over h the formulas for the covariances given in (4) for each (h,t) cell. This estimation method has been used extensively in the literature (including most of the studies referenced in this paper), so it is familiar enough that I relegate its details (including the choice of weighting matrix, the exact formulas used, and related issues) to Appendix C.

2.3. Profile heterogeneity and the estimates of persistence

Before proceeding further, I show that restricting income profiles across the population (as in the RIP process), when in fact such heterogeneity is present, leads to inconsistent estimates of the persistence parameter. To see this point, consider two individuals with different income growth rates, $\beta^H > \beta^L$, whose income profiles are plotted in Fig. 1. Clearly, the income paths of both of these individuals will deviate from the average profile (denoted with "-ˆ") in a systematic way over time. Ignoring this fact (by assuming $\beta^H = \beta^L \equiv \overline{\beta}$) will then lead the econometrician to interpret this *systematic* fanning out as the result of a sequence of persistent positive (or negative) income shocks to these individuals, indicated by the up and down arrows in the figure.

To give an idea about the potential magnitude of this bias, a quantitative example will be helpful. Consider a simplified version of the income process given in (2): $y^i_{h,t} = \alpha^i + \beta^i h + \varepsilon^i_{h,t}$, where β^i has population mean $\overline{\beta}$, and $\varepsilon^i_{h,t}$ is serially independent with zero mean. In addition, suppose that the econometrician allows for a fixed effect in the intercept, but not in the growth rate (assuming a life-cycle profile of $\alpha^i + \overline{\beta}h$ for all individuals). In this case, the income residuals are:

$$\hat{y}^i_{h,t} \equiv y^i_{h,t} - (\alpha^i + \overline{\beta}h) = (\beta^i - \overline{\beta})h + \varepsilon^i_{h,t}.$$

It is easy to see that $\hat{y}^i_{h,t}$ does not have zero mean for a given individual over time; instead it will either trend up or down. Finally, suppose that the econometrician observes a single cohort, and only when they are h and $h+1$ years old (I relax this assumption below). Then, under the (incorrect) assumption of RIP, a consistent estimator of the persistence of income shocks is the minimizer of $(1/I)\sum_{i=1}^{I}(\hat{y}^i_{h+1,t+1} - \tilde{\rho}\hat{y}^i_{h,t})^2$, which has a probability limit given by

$$\tilde{\rho} = \frac{h(h+1)\sigma^2_\beta}{h^2\sigma^2_\beta + \sigma^2_\varepsilon}.$$

Notice that $\tilde{\rho}$ is increasing in h, and approaches 1 in the limit, when in fact the true persistence is zero. To get a quantitative sense of the potential bias, I substitute some plausible values (that is, values consistent with my estimates in the next section) into this formula: $\sigma^2_\beta = 0.0004$, and $\sigma^2_\varepsilon = 0.03$. If the observed cohort is 44 years old ($h = 20$) the

[8] The expressions in (4) and (5) make clear how the time-effects π_t and ϕ_t will be identified in the estimation: π_t has a lasting effect on subsequent covariances (that is, it shifts the entire covariance structure after date t) whereas ϕ_t *only* affects the variance at time t. One implication of this, however, is that π_t and ϕ_t are not separately identified at the last date. To obtain identification at T I make the assumption that $\pi^2_{T-1} = \pi^2_T$.

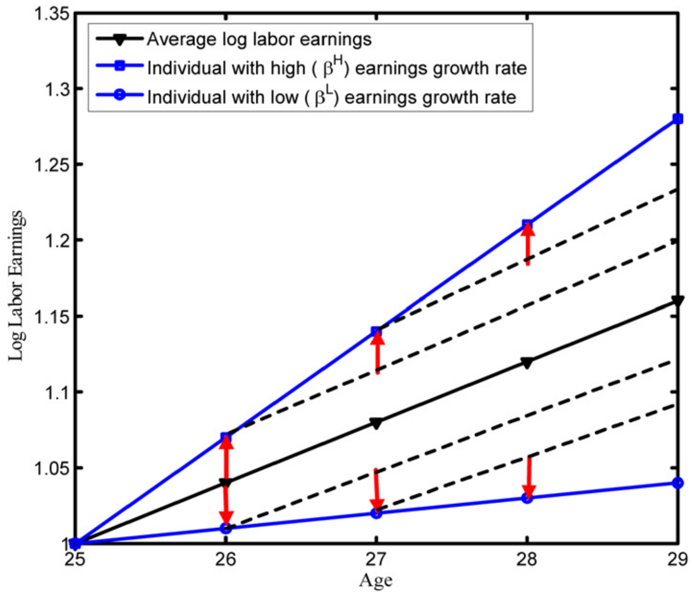

Fig. 1. Ignoring profile heterogeneity results in an upward bias in estimated persistence.

estimated persistence is $\tilde{\rho} = 0.87$. Similarly, if $h = 30$, one obtains $\tilde{\rho} = 0.95$. This calculation can be easily extended to show that when there is a population of individuals uniformly distributed from 25 to 64 years of age ($h = 1$ to 40), the estimated persistence would be $\tilde{\rho} = 0.91$, even though the true persistence is, again, zero. It is also easy to show that, under the same assumptions, the innovation variance of this *perceived* AR(1) process will be estimated to be $\sigma_\beta^2[h(1-\tilde{\rho})+1]^2 + (1+\tilde{\rho}^2)\sigma_\varepsilon^2$, which equals 0.058 at age 44 and 0.060 at age 54. This is about twice the variance of the actual innovation variance, $\sigma_\varepsilon^2 = 0.03$, used to generate the data.

Finally, since this bias arises from heterogeneity in *growth rates*, the fact that I accounted for fixed effects in *levels*—as is commonly done in the literature—had no mitigating effects. In other words, if I also restrict α^i across individuals in the calculations above, the corresponding values of $\tilde{\rho}$ remain almost unchanged. This simple example illustrates the close link between profile heterogeneity and the estimated persistence, and suggests that modeling the former could be critical for a consistent estimation of the latter.

3. Empirical findings

I first estimate the parameters of the process (2) by ignoring individual-specific variation in income growth rates—that is, by restricting $\sigma_\beta^2 \equiv 0$—but allowing for an individual fixed-effect, α^i (RIP process). In all rows, time effects in the variances of the persistent and transitory shocks are included in the estimation. The first row in Table 1 displays the estimates of the main parameters of interest (estimates of time effects are reported in Appendix B).

The estimate of ρ is 0.988, and one cannot statistically reject that income shocks are permanent at conventional significance levels. The innovation standard deviation of z is also large—about 12 percent per year—so in the long-run the persistent component dominates the cross-sectional distribution of income.

In panel B (row 4), I allow for heterogeneity in income growth rates. The first main finding is that the estimated persistence falls from 0.988 to 0.82. As is well known, the difference between these two estimates is substantial (Fig. 2): when $\rho = 0.82$, the effect of an income shock is reduced to fourteen percent of its initial value in ten years, whereas for $\rho = 0.988$, it retains almost ninety percent of its initial value at the same horizon. After twenty years, the effect of the former shock almost vanishes whereas the latter shock still keeps eighty percent of its initial impact. As can be anticipated from these comparisons, individuals facing each of these processes are likely to make very different economic choices. Similarly, Fig. 2 also plots the impulse response for $\rho = 0.95$ and 1.0, which are two values commonly used in the literature to calibrate economic models with RIP processes. Again, the remaining impact of an autoregressive shock is very different in each case.

As noted above, if measurement error follows an autoregressive process we have to be careful about interpreting ρ purely as the persistence of income shocks. For example, if measurement error has a lower persistence than income, the estimated ρ will understate the true persistence of income shocks. Nevertheless, notice that ρ is estimated to be almost a unit root under the RIP specification, which suggests that such a downward bias is not likely to be quantitatively large. Since the HIP process is estimated from the same income data, and therefore, contains the same measurement error, it seems unlikely that the low estimate of ρ in the HIP specification is due to measurement error. Of course, if measurement error is classical,

Table 1
Estimating the parameters of the labor income process

	Sample	ρ	σ_α^2	σ_β^2	$corr_{\alpha\beta}$	σ_η^2	σ_ε^2
			Panel A: σ_β^2 *restricted* to be zero (RIP process)				
(1)	All	.988	.058	–	–	.015	.061
		(.024)	(.011)			(.007)	(.010)
(2)	College	.979	.031	–	–	.0099	.047
		(.055)	(.021)			(.013)	(.020)
(3)	High-school	.972	.053	–	–	.011	.052
		(.023)	(.015)			(.007)	(.008)
			Panel B: σ_β^2 *unrestricted* (HIP process)				
(4)	All	.821	.022	.00038	−.23	.029	.047
		(.030)	(.074)	(.00008)	(.43)	(.008)	(.007)
(5)	College	.805	.023	.00049	−.70	.025	.032
		(.061)	(.112)	(.00014)	(1.22)	(.015)	(.017)
(6)	High-school	.829	.038	.00020	−.25	.022	.034
		(.029)	(.081)	(.00009)	(.59)	(.008)	(.007)
(7)	All (large sample)	.842	.072	.00043	−.33	.032	.044
		(.024)	(.055)	(.00007)	(.40)	(.006)	(.008)
(8)	All (first 10 cov.)	.899	.055	.00055	−.73	.016	.047
		(.042)	(.060)	(.00013)	(.38)	(.010)	(.009)

Notes. Standard errors are in parentheses. Time effects in the variances of persistent and transitory shocks are included in the estimation in all rows. The estimated time effects for rows 1 and 4 are reported in Table 5, others are omitted to save space. The reported variances are averages over the sample period.

as is commonly assumed, then this problem does not arise (it would be captured in the transitory component, σ_ε^2, and would have no effect on ρ).

An important question is whether the estimates are sensitive to the sample selection criteria used to obtain the primary sample. In particular, recall that I require an individual to satisfy these criteria for twenty years to be included in the sample. Although, as I discuss further below, there are good reasons for this requirement, one could be concerned that we are eliminating individuals with unstable jobs who might be facing more persistent shocks than the rest of the population. Including these individuals could therefore result in a higher estimated persistence as well as maybe a different estimate for the dispersion of income growth rates. To explore this possibility, I draw a new subsample using the same criteria as before except that I now require individuals to stay in the sample for at least four years instead of twenty. The resulting sample has 4381 individuals with at least four observations each. Row 7 displays the results. The estimated persistence goes up slightly compared to the baseline, from 0.82 to 0.84, and the innovation variance also increases from 0.029 to 0.032. The

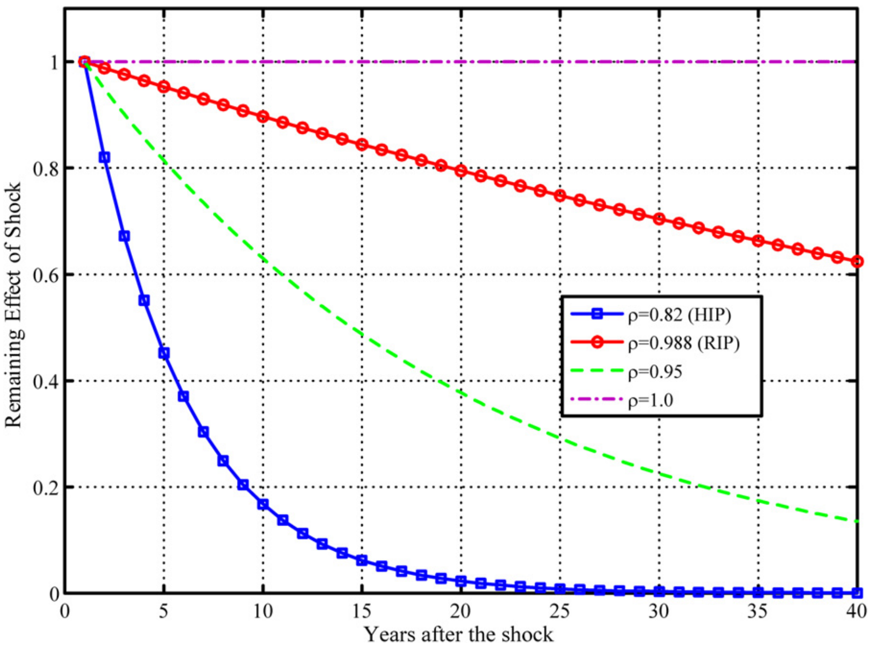

Fig. 2. The remaining effect of an AR(1) shock for different values of ρ.

dispersion of income growth rates is also higher at 0.00043. While these results are consistent with the fact that the new sample contains more heterogeneous households, the difference in estimates is not large.

Before closing this section, it is useful to compare these results to earlier work that estimates alternative versions of the HIP process using representative samples of US households. Among these, Haider (2001) employs a rotating panel design using PSID data and includes individuals satisfying sample selection criteria for three years or more. He also finds evidence of large heterogeneity in income growth rates ($\sigma_\beta^2 = 0.00041$) and estimates ρ to be 0.64. Notice that while his estimate of σ_β^2 is very close to what I find here, the persistence parameter is significantly lower. This is also the case in Baker (1997, Table 4) who uses a fully balanced panel from PSID (1968–1988) and obtains: $\sigma_\beta^2 = 0.00039$ and $\rho = 0.67$. One reason for the lower estimates of ρ found by these authors could be that they model the dynamics of income as an unrestricted ARMA (1, 1) or (1, 2) process, compared to the more parsimonious specification adopted here. To sum up, the estimates of ρ obtained in this paper are substantially lower than a unit root. At the same time, they still represent an upper bound of the values found in the literature using HIP processes. However, the estimates of σ_β^2 appear to be very similar across these three studies.

3.1. The labor income process by education group

I next examine if, and how, the labor income process differs by education group. This question has so far only been investigated in the context of RIP processes (Hubbard et al., 1994, and Carroll and Samwick, 1997). Thus, to provide a benchmark, I begin by estimating the RIP process for college- and high school-educated individuals. Rows 2 and 3 of Table 1 report the parameter estimates for the two groups: ρ is estimated to be 0.979 and 0.972 for the college- and high school educated-groups, respectively. Similarly, the innovation variances of the AR(1) shocks are very close to each other: 0.0099 and 0.011 respectively. Overall, the estimated parameters reveal very similar income processes for the two education groups.

Although this finding may seem surprising (given the many differences one could think of between the labor market risks faced by different education groups), it is in fact consistent with the results obtained in previous studies. Table 2 displays the estimated income processes from two studies that are most often used for calibrating macroeconomic models. In Hubbard et al. (1994), the estimated persistence ranges from 0.946 to 0.955 but shows no systematic pattern with education. The innovation variance seems to go down with higher education, but the difference is not statistically significant. Carroll and Samwick (1997) impose the further restriction that income shocks are permanent for all groups ($\rho \equiv 1$), and only estimate the variances. They find innovation variances to be increasing with education at lower levels, but then fall back at higher education levels. The differences between groups are again not statistically significant. They find some evidence that transitory shock variances get smaller with education. The conclusion that emerges from these studies and my findings is that in the RIP process income risk does not vary substantially by education level. If anything, there is some evidence that income risk is somewhat greater for lower educated individuals.

I next estimate the income process of each group allowing for HIP (rows 5 and 6 of Table 1). The estimated persistence is now significantly lower for both groups ($\rho^C = 0.81$ versus $\rho^H = 0.83$), but there is still little difference across education groups. However, there *is* now a major difference in an important dimension: the dispersion of income profiles is significantly larger for college-educated individuals ($\sigma_\beta^2 = 0.00049$) compared to high school-educated individuals ($\sigma_\beta^2 = 0.00020$). In fact, this difference could be partly anticipated from Fig. 4, which shows a larger increase in within-cohort income inequality among the former group than the latter.

Table 2
Estimates of the RIP model by education level in the literature

Paper	Group	ρ	σ_η^2	σ_ε^2
Hubbard et al. (1994)	<12 yrs of education	.955	.033	.040
		(.106)	(.076)	(.075)
	12–15 yrs of education	.946	.025	.021
		(.129)	(.063)	(.054)
	16+ yrs of education	.955	.016	.014
		(.121)	(.040)	(.033)
Carroll and Samwick (1997)	0–8 grades	1.0[a]	.0190	.0894
		–	(.0137)	(.0256)
	9–12 grades	1.0	.0214	.0658
		–	(.0090)	(.0168)
	High school diploma	1.0	.0277	.0431
		–	(.0069)	(.0129)
	Some college	1.0	.0238	.0342
		–	(.0047)	(.0088)
	College graduates	1.0	.0146	.0385
		–	(.0068)	(.0126)

Notes. One difference between these studies and the present one is that these studies estimate income processes for household income whereas I estimate for individuals.
[a] The persistence parameter is restricted to 1.0 (random walk shocks) in Carroll and Samwick and hence is not estimated.

Finally, the correlation between the slope and the intercept is negative in all rows of Table 1 (although not precisely estimated), consistent with earlier work. A natural interpretation of this negative correlation is suggested by the human capital model: individuals who invest more early in life—perhaps in response to higher learning ability—and suffer from lower income are compensated by higher income growth. Moreover, the correlation is more negative for the college-educated group (−0.70) compared to the rest (−0.25), suggesting that human capital accumulation could be more important for wage growth in high-skill occupations (Mincer, 1974; Hause, 1980).

3.2. Quantifying the heterogeneity in income profiles

The second main finding (in row 4 of Table 1) is that the heterogeneity in income growth rates measured by σ_β^2 is (statistically) significant. To show that this estimate is also economically significant, I rearrange the expression for cross-sectional income inequality (given in (4)) to obtain:

$$\mathrm{var}_i(\hat{y}_h^i) = (\sigma_\alpha^2 + \sigma_\varepsilon^2) + \left(\frac{1-\rho^{2h+1}}{1-\rho^2}\sigma_\eta^2\right) + [2\sigma_{\alpha\beta}h + \sigma_\beta^2 h^2],$$

where I substituted $\mathrm{var}(z_h^i)$ from (5), and set π_t and ϕ_t equal to 1.

This expression provides a useful decomposition of inequality into its components. The first set of parentheses contain terms that do not depend on age (and hence make up the intercept of the age-inequality profile). The second set of parentheses capture the rise in inequality due to the accumulated effect of AR(1) shocks. The solid line in the left panel of Fig. 1 plots the magnitude of this term over the life-cycle of a cohort. For the estimated value of $\hat{\rho} = 0.82$, this component increases slightly in the first seven years and then remains roughly constant.

The last set of parentheses contain terms that capture the effect of HIP on inequality. It consists of a decreasing linear term (since $\sigma_{\alpha\beta} < 0$), and an increasing quadratic term, in h. It is easy to see that even when σ_β^2 is very small, the effect of profile heterogeneity on income inequality will grow rapidly with h^2, as the cohort gets older. As the dashed line in the left panel shows, early in the life-cycle the contribution of profile heterogeneity to income inequality is very small. In fact, until about age 47 more than half of the income inequality is generated by the fixed effect, and transitory and persistent shocks. However, the effect of profile heterogeneity increases rapidly with age, and results in substantial inequality later in life.

The right panel of Fig. 3 plots the *fraction* of total inequality attributable to HIP. In the sample of all individuals (denoted "-o"), HIP accounts for 70 percent of inequality at age 55 (33 years of experience). More importantly, at the same age, HIP accounts for 75 percent of the inequality among college-educated individuals ("-+") and 56 percent of the inequality

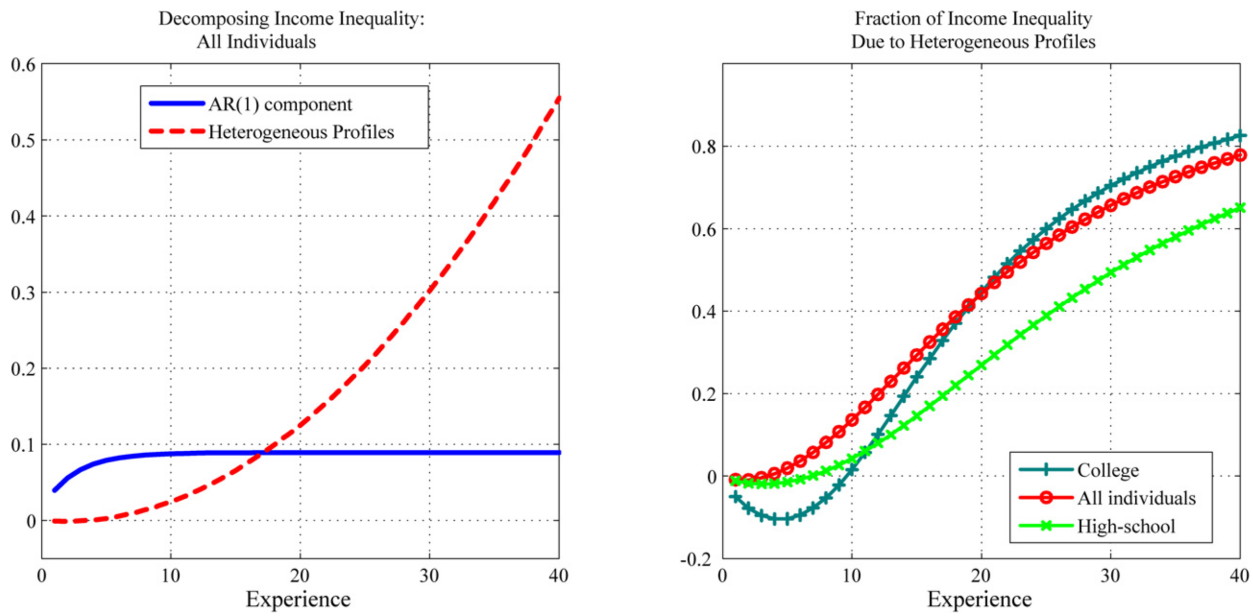

Fig. 3. Quantifying the contribution of HIP to income inequality.

among high school-educated individuals ("-x").[9] The fact that heterogeneity in income profiles is substantial even within these education groups has an important implication for calibrating macroeconomic models. It suggests that the common practice of allowing for a different income profile for each education group, *while omitting within-group variation*, captures only a small part of the profile heterogeneity in the population.

4. What is the source of identification?

The problem of distinguishing between the RIP and HIP processes is reminiscent of the familiar debate in macroeconomics about whether GDP growth is better represented by a stochastic trend (RIP process), or by stationary shocks around a deterministic trend (HIP process). Given the well-known difficulties associated with distinguishing between those two hypotheses (cf., Christiano and Eichenbaum, 1990), it seems reasonable to suspect a similar difficulty in the current context.[10] Thus an important question to answer is the following: Where does identification between the RIP and HIP processes come from?

The main difference between the present problem and the debate in macroeconomics is that in our case we have access to *panel data* on labor income, unlike macroeconomists who had to rely on a single time-series of GDP observations. With panel data, I can characterize the evolution of the cross-sectional distribution of income as a cohort gets older. As I explain below, it then becomes possible to distinguish between the RIP and HIP processes by exploiting the different implications of each process for the evolution of this cross-sectional distribution.

More specifically, consider the covariance matrix of income residuals for a given cohort. The diagonal elements of this matrix correspond to the cross-sectional variance of log income, whereas the off-diagonal elements correspond to autocovariances at various lags. For a cohort with a working life of 40 years, there are 40 variance terms and, many more, up to 780 autocovariance terms. To understand identification it is thus instructive to examine the covariance matrix by focusing on the diagonal and off-diagonal elements separately.[11] This is what I do next.

4.1. Age structure of variances

The first piece of information is provided by the change in the cross-sectional variance of income as the cohort ages. The formula for these variances (Eq. (4)) is reproduced here for convenience:

$$\operatorname{var}(\hat{y}^i_{h,t}) = \underbrace{\left[\sigma^2_\alpha + 2\sigma_{\alpha\beta}h + \sigma^2_\beta h^2\right]}_{\text{HIP component}} + \underbrace{\operatorname{var}(z^i_{h,t})}_{\text{AR(1) component}} + \phi^2_t\sigma^2_{\varepsilon,t}. \tag{6}$$

For the clarity of this discussion, I first consider the case where the panel data on income is from a single cohort *and* income shocks have stationary variances over time: $\phi^2_t = \pi^2_t \equiv 1$ for all t. I relax these assumptions in a moment.

The terms in square brackets capture the effect of profile heterogeneity, which is a *convex* function of age (notice that the coefficient on h^2 is σ^2_β). The second term captures the effect of the AR(1) shock, which is a *concave* increasing function of age as long as $\rho < 1$ and becomes linear in age when $\rho = 1$. Thus, if the within-cohort variance of income increases in a convex fashion in the data as the cohort gets older, this would be captured by the HIP terms, whereas a non-convex shape (including a linear one) would be captured by the presence of AR(1) shocks.

To construct the empirical counterpart of the age–variance profile, however, I do need to account for time-variation in shock variances (i.e., *time*-effects), and also account for the fact that the present panel data pertains to more than one cohort who could differ in their income variances (i.e., *cohort*-effects). Nevertheless, it is well-known that attempting to simultaneously identify time, cohort, and age effects is problematic, since any one of these effects can be written as a linear combination of the other two.[12] Therefore, I follow the bulk of the literature and control for cohort and time effects in turn.

4.1.1. Cohort effects in variances

I first generate the empirical graphs by controlling for cohort effects only (following Deaton and Paxson, 1994 and Storesletten et al., 2004), which is the more common approach of the two. To this end, I first construct five-year overlapping

[9] Notice that in the college sample the contribution of HIP to inequality is negative (i.e., HIP *reduces* inequality) in the first 10 years of the life-cycle. This is due to the large negative correlation (−0.70) between the slope and intercept of income profiles in this group. As a result, early in life individuals with low initial income but fast income growth catch up with those with slow income growth but high initial income, which reduces inequality early on.

[10] This resemblance is on the methodological level. Substantively, the two questions are quite different: in the macroeconomics literature most researchers agreed that GDP movements were extremely persistent, but the question was whether the autoregressive parameter was equal to 1 or slightly lower than that. Instead here, the RIP model implies an autoregressive parameter close to 1, whereas the HIP model implies *substantially* less persistence. Therefore, the distinction between the two income processes is not merely a technical curiosity, but has important substantive implications.

[11] Notice that the estimation uses the $T \times T$ matrix of covariances over *time*. The elements of that matrix are obtained by aggregating the underlying covariances of each cohort studied here appropriately as explained in Section 2.

[12] Hall (1971) is one of the early papers to notice this problem. For a detailed treatment of the identification of time, cohort, and age effects see Heckman and Robb (1985). As these authors show even if one allows for a higher order polynomial in these three effects, it is only possible to identify a subset of the polynomial coefficients. In principle, it is possible to distinguish between the three effects by imposing further restrictions (such as assuming that time effects are orthogonal to a linear time trend; see, for example, Deaton and Paxson, 1994). Of course, in many cases these additional restrictions may be hard to justify, and even in those cases identification typically remains weak empirically.

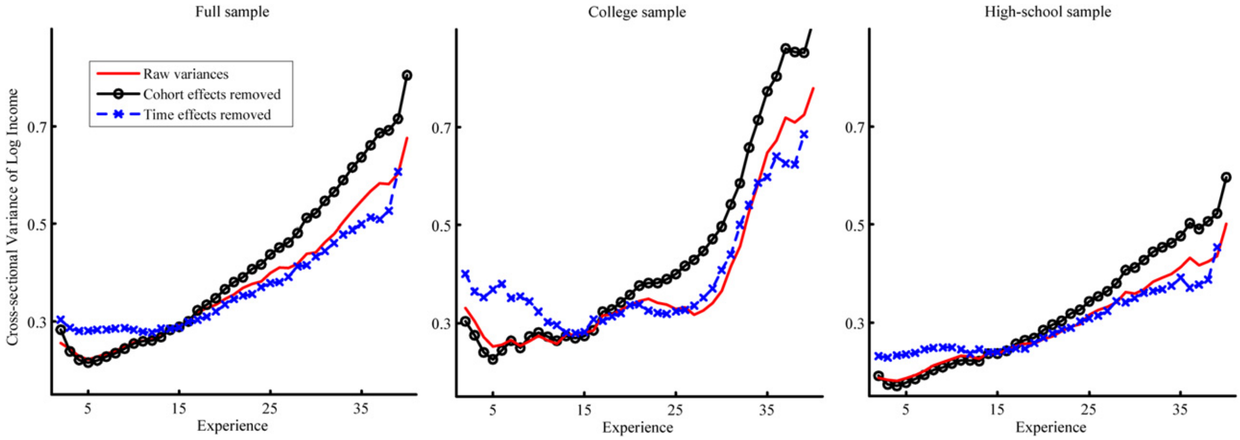

Fig. 4. Age–variance profile of log income in PSID.

experience bands by grouping all individuals who have experience level $h-2$ to $h+2$ and assign the mid-point (h) as the experience of that group. This prevents having a very small number of observations in a given experience cell thereby mitigating the noise that results from sampling variation. Because my sample consists of individuals from age 20 to 64, there are 41 experience bands ranging from (mid-point) age 22 to 62. Next, I group individuals by experience and (year-of-birth) cohort and compute the variances in each experience–cohort cell. I then regress these raw variances on a full set of experience and cohort dummies, and report the coefficients on the experience dummies in Fig. 4. To preserve the overall level of inequality in the data, the dummies are scaled to match the level of variance in the data at $h=15$.

The lines marked with circles in the three panels show the resulting experience–variance profile for the whole population (left) and for college (middle) and high-school (right) graduates. As can be seen in these graphs, variances increase in a slightly convex fashion in all three cases consistent with the heterogeneity in income profiles found in the estimation. Moreover, the experience–variance profile rises the most, and appears to be the most convex, for the college sample, consistent with the large estimate of σ_β^2 for this group. The opposite is true for the high school sample, again consistent with the small estimate of σ_β^2 for this group.

4.1.2. Time effects in variances

In a recent paper, Heathcote et al. (2005) argue that it might be important to control for time effects when constructing the experience–*variance* profile since income inequality has increased substantially during the 1980s (which is included in my sample period). To address this issue, I now control for time effects only. More specifically, after constructing the raw variances for all experience–cohort cells described above, I regress them on a full set of experience and *time* dummies. The dashed lines marked with x's plot the coefficients on experience dummies.

A couple of remarks are in order. First, comparing these "cleaned variances" to the raw data (solid lines), notice that controlling for time effects raises the experience–variance profile early in life (up to about 15 years of experience), whereas controlling for cohort effects raises them later in life. Second, while I confirm Heathcote et al.'s (2005) finding that controlling for time effects results in a smaller overall rise in the variance of log income over the life-cycle, the profile continues to be slightly convex for the full sample as well as for the high-school sample. For the college sample, the experience–variance profile is still convex but now goes down until about 15 years of experience after which point it starts to increase first slowly and then more rapidly. Overall, while controlling for cohort versus time effects could be important for some questions, the slight convexity of the experience–variance profile seems to be robust to whichever route is chosen. It should also be stressed that in both cases the convexity is typically moderate in the data (except for the college sample).

In Fig. 5, for comparison, I plot (solid line) the theoretical experience–variance profile implied by the HIP process (using Eq. (4) and the parameter estimates in rows 4, 5 and 6 of Table 1) which is slightly convex as in the data. Because cohort and time effects in turn have been removed from variances in constructing Fig. 4, the theoretical counterparts here are also constructed without time effects (i.e., $\pi_t = \phi_t = 1$) for a meaningful comparison to that figure. Similarly, the dashed lines plot the variances implied by the RIP process (using the parameters in rows 1, 2, and 3 of the same table) which shows a slightly concave shape instead. Notice that, in principle, a RIP process could also generate a convex pattern if the AR(1) process have age-dependent innovation variances. However, this would require innovation variances to be *increasing* with age, which is at odds with the empirical evidence presented in Baker and Solon (2003, Fig. 2) and Meghir and Pistaferri (2004, Table 5) who find a decreasing (or U-shaped) pattern over the life-cycle.[13] To sum up, the experience–variance profile is *slightly* convex in the US data, which is one reason the estimation in the previous section revealed evidence of HIP.

[13] Clearly, AR(1) shocks can also generate a convex profile if $\rho > 1$. But, as I discuss below, this would imply that covariances *increase* with the lag order, which is at odds with empirical evidence.

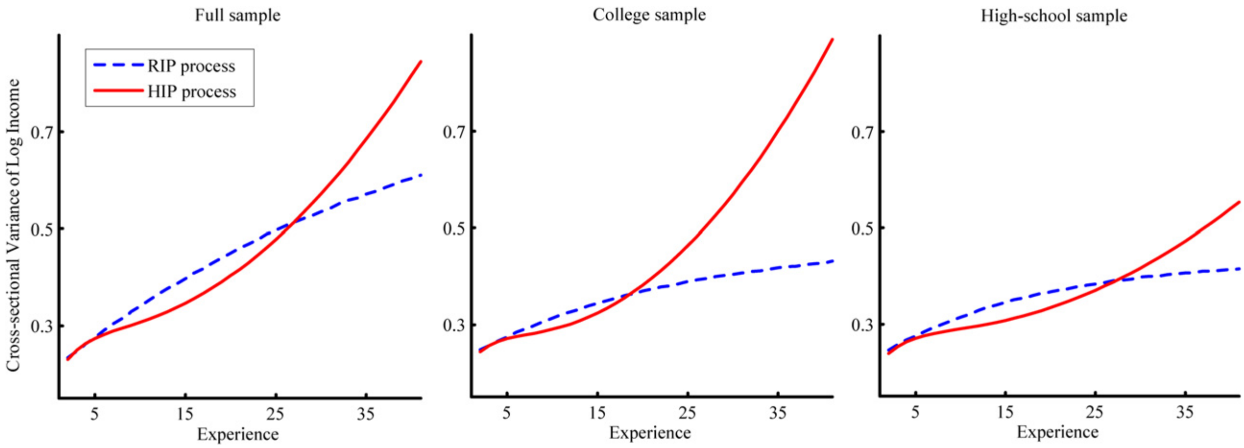

Fig. 5. Theoretical experience–variance profile of log income implied by HIP and RIP processes.

4.2. Age structure of autocovariances

The second source of identification is provided by the off-diagonal elements of the covariance matrix. As before, I begin by assuming that the econometrician observes a single cohort over time, in which case the formula is

$$\operatorname{cov}(\hat{y}^i_h, \hat{y}^i_{h+n}) = \underbrace{\left[\sigma^2_\alpha + \sigma_{\alpha\beta}(2h+n) + \sigma^2_\beta h(h+n)\right]}_{\text{HIP component}} + \underbrace{\rho^n \operatorname{var}(z^i_h)}_{\text{AR(1) component}}, \tag{7}$$

where I dropped the t subscript since it is perfectly correlated with h.

As shown in this equation, the covariance between experience h and $h+n$ is again composed of two parts. As before, the terms in square brackets capture the effect of heterogeneous profiles and is a convex function of experience. Moreover, and more importantly, the coefficients of the linear and quadratic terms depend on both h and n, which allows covariances to be decreasing, increasing, or non-monotonic in the lag order for each experience level. The second term captures the effect of AR(1) shocks, and notice that for a given h, it depends on the covariance lag n only through the geometric discounting term ρ^n. The strong prediction of this form is that, starting at h, covariances should decay geometrically at the rate ρ, regardless of h. Thus, in the RIP process (which only has the AR(1) component) covariances are restricted to decay at the same rate at every experience level, and cannot be non-monotonic in n.

To construct the empirical counterpart I proceed as before. I first compute the entire structure of raw autocovariances for each cohort in the sample (using the age-cohort cells constructed above). Then I regress these raw autocovariances on a full set of cohort dummies and average the resulting residuals across all cohorts. Fig. 6 plots these "cleaned" autocovariances separately for the college (top panel) and high-school (bottom) samples. For example, the leftmost (solid) line plots $\operatorname{cov}(\hat{y}^i_1, \hat{y}^i_{1+n})$ for $n = 1, 2, \ldots, 20$, and other lines plot the same for $h = 2, 3, \ldots, 25$, subject to $h + n \leqslant 34$, which corresponds to a real life age of about 55. (The autocovariances that start from even (odd) numbered ages are shown with solid (dashed) lines to make the graph easier to read.)

In the top panel of Fig. 6, the first observation is that the autocovariance structure displays rich patterns that change over the life-cycle. For example, autocovariances decay monotonically and rather rapidly with the lag order for small hs (less than about 6 or 7). However, as h grows, the autocovariance structure appears to become flatter first, and then becomes U-shaped at older ages: it decreases, and then increases, steeply with lag order. Notice that for $h + n \geqslant 25$ almost all covariances are increasing with lag order. Turning to the high school sample in the bottom panel, the covariance structure looks quite different. Although we do see a flattening of covariance profiles up to about $h = 15$ similar to that in the college sample, there is no tendency of covariances to become U-shaped or increasing at older ages, unlike in the college sample. In contrast, they appear to steepen again at higher experience levels.

To understand how this information could help distinguish between the two income processes, I now turn to the theoretical covariance structures implied by the HIP and RIP processes (shown in Fig. 7). These graphs are plotted using the parameter estimates for the college sample from rows 5 and 2 of Table 1 respectively. The top panel plots the autocovariance structure generated by the HIP process (lines marked circles), as well as the separate contributions of the AR(1) component and HIP component. As noted earlier the AR(1) component (dashed lines) generates a geometric decay in autocovariances with the lag order at the same rate for all experience levels. The HIP component (solid lines) instead generates autocovariances that are downward sloping in lag order (n) at young ages, whereas they become upward sloping at older ages. As a result, the HIP process generates autocovariances (which is the sum of the two components just mentioned) that is downward sloping at young ages, but becomes upward sloping (and even slightly U-shaped) at older ages. Comparing this figure to the empirical counterpart for college graduates, it is fair to say that the HIP process is broadly consistent with the pattern in the data, while missing on some important features, such as the deep U-shape observed at later ages. Although it seems

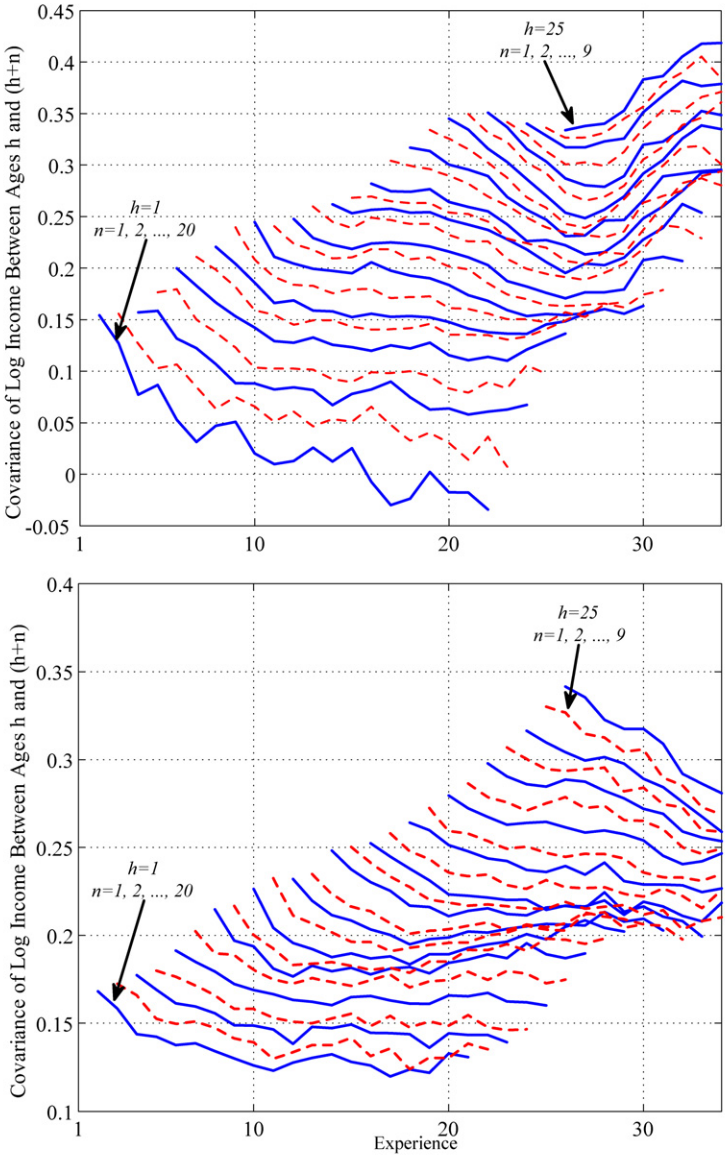

Fig. 6. Covariance structure of log income for college- (top) and high school- (bottom) educated individuals in PSID.

possible to accommodate this pronounced U-shape in the HIP framework by allowing an individual-specific quadratic term (in addition to the linear term), this would introduce an additional state variable into a dynamic programming problem. Therefore, I do not pursue this extension here.

I next turn to the RIP process shown in the lower panel. Since the autocovariance structure is generated entirely by the AR(1) component, they all decay towards zero at the same geometric rate ρ regardless of h. However, this does not imply that the *slope* of the autocovariance profile is the same for all h. This is because the slope is:

$$\operatorname{cov}(\hat{y}_h^i, \hat{y}_{h+n+1}^i) - \operatorname{cov}(\hat{y}_h^i, \hat{y}_{h+n}^i) = \rho^n(\rho - 1)\operatorname{var}(z_h^i).$$

In the RIP process, $\operatorname{var}(z_h^i)$ increases over the life cycle (due to the accumulation of highly persistent shocks), implying a more negative slope, and therefore an autocovariance profile that is steeper, as individuals get older.[14] This feature does not fit well with the generally flattening pattern of autocovariances observed in the college sample. In contrast, recall that in the high school sample, covariances do become somewhat steeper at older ages. This is broadly consistent with a RIP process and visually it is difficult to see the effect of a large HIP component on covariances for this group (except for some

[14] In other words, even though the *ratio* of subsequent autocovariances are constant for all hs, the *difference* is getting more negative as h gets larger.

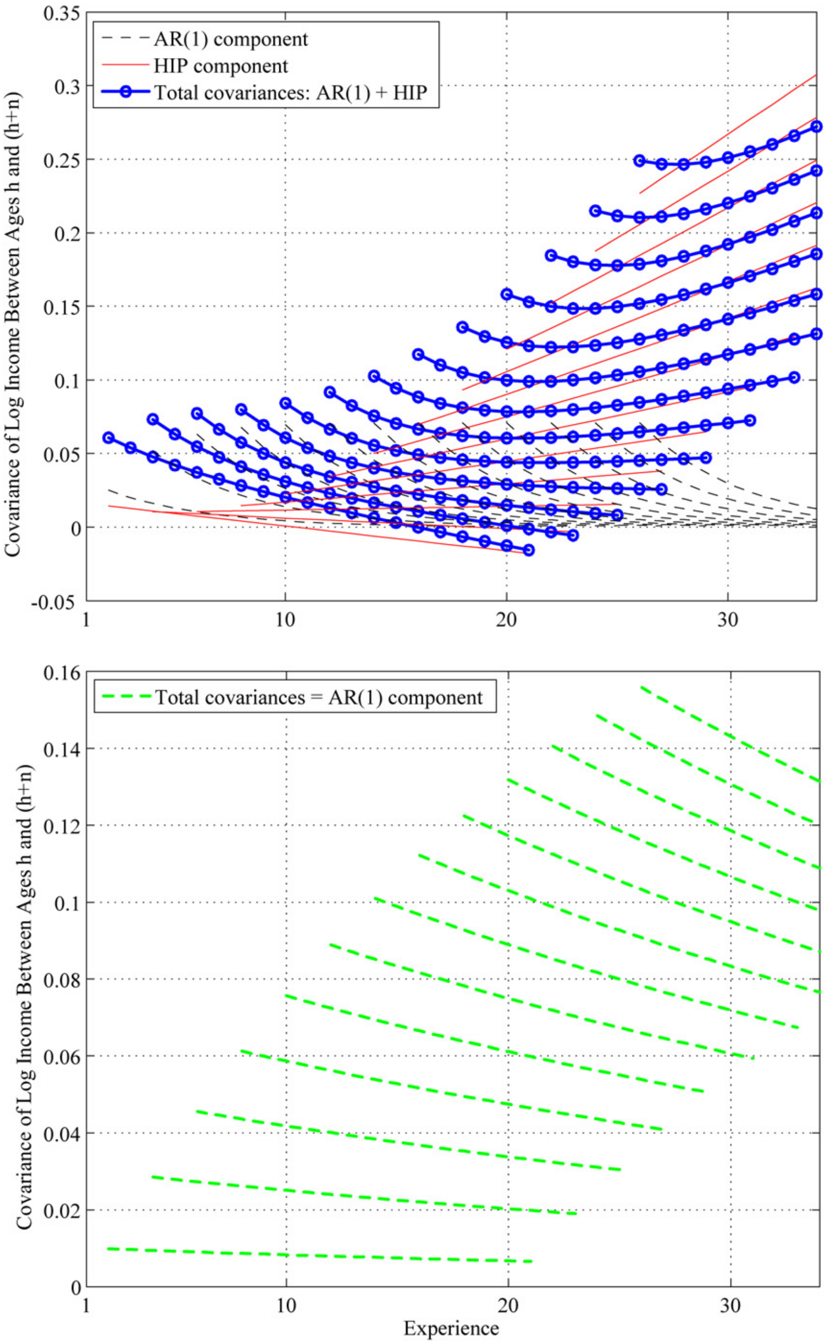

Fig. 7. Theoretical covariance structure of log income for college-educated individuals: HIP (top) and RIP (bottom) processes.

flattening and rising up to $h = 15$). Consistent with this observation, the formal estimation reveals a much smaller estimate of σ_β^2 for the high school graduates compared to college graduates.[15]

So far I have only controlled for cohort effects when constructing the autocovariance structure. But notice that the formula for autocovariances in (7) does not explicitly depend on time effects in shock variances (i.e., ϕ_t and π_t). Of course, time-variation would affect the *level* of $\text{var}(z_h^i)$ and, therefore, the level of the autocovariance structure for each h, but *not*

[15] The features of the shape of the autocovariance matrix discussed here are preserved in the larger sample described earlier (which includes all individuals that satisfy the selection criteria for four years or more). I omit these results for brevity but they are available upon request.

how covariances change with the lag order n, which is the focus of the preceding analysis. This implies that, unlike the age–variance structure, I do not need to control for time-effects in autocovariances.[16]

To sum up, the empirical autocovariance structure of income residuals display some rich patterns that change as an individual gets older. The autocovariance structure implied by AR(1) shocks alone, as in the RIP process, appears too restrictive. Instead, the HIP process, while not capturing some important aspects of the autocovariance structure in the data, allows for more flexibility and seems to a better job of fitting the autocovariance matrix of the college sample. As a result, the minimum distance estimation in the previous section finds evidence of large heterogeneity in income growth rates to fit this matrix. In contrast, the estimate of heterogeneity in the high school sample is smaller consistent with what one could conjecture by studying the covariance structure.

4.2.1. Some empirical difficulties in identification: a discussion

The preceding analysis of the covariance matrix also highlights some empirical difficulties in distinguishing between the HIP and RIP processes. The main difficulty is that while higher order autocovariances contain information that is valuable for identification, fewer and fewer individuals contribute to these covariances because of sample attrition, raising concerns about potential selectivity bias. In most of the analysis I required individuals to be present for twenty years. Although this requirement creates a subsample that may not exactly be a random sample of US households, it has the important advantage that autocovariances at different lags are computed for roughly the same groups of households. Therefore, it is possible to interpret the covariance matrix (and the resulting estimates) as valid for this subsample. If, instead, I include all individuals who are in the sample for, say, three years or more, then the number of individuals contributing to the 20th autocovariance will be about a quarter of the number of individuals contributing to the 3rd autocovariance. To the extent that these individuals are not a completely random subsample of the original sample, covariances at different lags will have variation due to sample selection that can confound the identification between HIP and RIP processes.

While the similarity between the estimates obtained from the primary sample and the larger sample is somewhat reassuring (rows 4 and 7 in Table 1), this potential difficulty should not be easily dismissed. To further illustrate this point, I re-estimated the HIP process using only the autocovariances up to the 10th order. This can be thought of as an extreme case where higher order autocovariances are so noisy that they are completely uninformative. As can be seen in row 8 of Table 1, the estimate of ρ rises to 0.90, although the heterogeneity in income growth rates also increases from 0.00038 to 0.00055. This result suggests that while the evidence on the existence of heterogeneity in income growth rates is less sensitive to the inclusion of higher covariances, this information is important for the estimate of ρ.

Overall, the importance of higher order autocovariances also underscores the limitation of relying on income data alone for distinguishing between these alternative income processes. Another, and arguably better, approach would be to base inference on individuals' economic choices, which contain valuable information about the environment faced by individuals, including the future income risks they perceive. For example, the response of forward looking choices, such as consumption and savings, to movements in income can be exploited to distinguish between different views about the income process. This is the approach taken in Guvenen (2007) who studies some well-known empirical facts about consumption behavior to learn about the nature of income risk. Guvenen and Smith (2007) take this one step further and conduct a full blown econometric analysis in an attempt to fully use the joint dynamics of consumption and income.

5. A comparison to the existing literature

In this section I try to reconcile the direct estimation results of the previous section supporting the HIP process with some previous tests used in the literature, which have been interpreted as supporting the RIP process (among others, MaCurdy, 1982; and Abowd and Card, 1989).

5.1. MaCurdy (1982)

The basic idea of these tests is based on the simple observation that with profile heterogeneity, individual income *growth* should be positively autocorrelated. This can be shown easily. Using Eqs. (3) and (5) (i.e., $z_0^i = 0$ for all i and no time effects in variances), the autocovariance of income growth at lag n is:

$$\operatorname{cov}\left(\Delta\hat{y}_h^i, \Delta\hat{y}_{h+n}^i\right) = \sigma_\beta^2 - \rho^{n-1}\left(\frac{1-\rho}{1+\rho}\sigma_\eta^2\right), \tag{8}$$

for $n \geqslant 2$. Thus, covariances involve a positive constant term (σ_β^2) that arises from the presence of HIP, and a negative term, which goes to zero as a geometric function of n. According to the HIP process then covariances should be positive—after a certain lag—if σ_β^2 is positive after all. Moreover, if $\rho = 1$ (income shocks are permanent) the negative term disappears and autocovariances should always be positive at any lag greater than 1. On the other hand, it is also easy to see that in the

16 In other words, time effects would shift the covariance structure starting for a given h up or down but would have no effect on how covariances change for a given h as we vary n. When I regress the autocovariances on time effects, this is exactly what I find. These results are not reported for brevity but are available from the author upon request.

Table 3
Covariance structure of income growth: US data versus the HIP model

Lag	Autocovariances				Autocorrelations	
	Data	Model	Model	Model with $\mathrm{var}(z_0^i) = 0.3$	Data	Model
$N \rightarrow$	27,681	27,681	500,000	27,681	27,681	27,681
0	.1215	.1136	.1153	.1149	1.00	1.00
	(.0023)	(.00088)	(.00016)	(.00113)	(0.000)	(.000)
1	−.0385	−.04459	−.04826	−.04352	−.3174	−.3914
	(.0011)	(.00077)	(.00017)	(.00087)	(0.010)	(.0082)
2	−.0031	−.00179	−.00195	−.00077	−.0261	−.0151
	(.0010)	(.00075)	(.00018)	(.00086)	(0.008)	(.0084)
3	−.0023	−.00146	−.00154	−.00081	−.0192	−.0128
	(.0008)	(.00079)	(.00020)	(.00074)	(0.009)	(.0087)
4	−.0025	−.00093	−.00120	−.00029	−.0213	−.0080
	(.0007)	(.00074)	(.00019)	(.00092)	(0.010)	(.0083)
5	−.0001	−.00080	−.00093	−.00019	−.0012	−.0071
	(.0008)	(.00081)	(.00020)	(.00082)	(0.007)	(.0090)
6	−.0000	−.00073	−.00067	−.00025	−.0020	−.0063
	(.0008)	(.00076)	(.00018)	(.00093)	(0.007)	(.0085)
7	.0001	−.00046	−.00049	−.00021	.0004	−.0043
	(.0010)	(.00077)	(.00019)	(.00084)	(0.009)	(.0086)
8	.0004	−.00030	−.00033	−.00017	.0036	−.0027
	(.0008)	(.00080)	(.00020)	(.00085)	(0.010)	(.0091)
9	−.0010	−.00027	−.00018	−.00007	−.0085	−.0024
	(.0007)	(.00074)	(.00018)	(.00082)	(0.011)	(.0084)
10	−.0017	−.00003	−.00010	.00015	−.0143	−.0003
	(.0006)	(.00072)	(.00019)	(.00085)	(0.009)	(.0081)
11	.0012	.00013	−.00001	.00022	.0103	.0011
	(.0008)	(.00075)	(.00017)	(.00085)	(0.008)	(.0085)
12	−.0009	.00011	.00005	.00021	−.0078	.0010
	(.0010)	(.00076)	(.00020)	(.00077)	(0.011)	(.0086)
13	.0006	.00009	.00012	.00034	.0051	.0083
	(.0009)	(.00079)	(.00021)	(.00079)	(0.009)	(.0088)
14	−.0023	.00024	.00017	.00028	−.0188	.0021
	(.0008)	(.00080)	(.00019)	(.00076)	(0.010)	(.0089)
15	.0053	.00017	.00021	.00031	.0438	.0015
	(.0007)	(.00076)	(.00020)	(.00081)	(0.008)	(.0086)
16	−.0045	.00020	.00024	.00035	−.0372	.0018
	(.0008)	(.00077)	(.00018)	(.00092)	(0.009)	(.0088)
17	.0014	.00029	.00027	.00039	.0116	.0025
	(.0009)	(.00083)	(.00017)	(.00092)	(0.010)	(.0092)
18	.0012	.00036	.00030	.00043	.0094	.0032
	(.0009)	(.00076)	(.00018)	(.00075)	(0.011)	(.0087)

Notes. N denotes the sample size (number of individual-years) used to compute the statistics. Standard errors are in parenthesis. The statistics in the "data" columns are calculated from the primary sample. The counterparts from simulated data are calculated using 27,681 observations, which is the total number of observations in the primary sample. Standard deviations are calculated using bootstrap and 1000 replications.

absence of HIP, autocovariances should be either negative or zero (depending on whether $\rho < 1$ or $\rho = 1$). This suggests that one way to distinguish between HIP and RIP processes is to test if higher order autocovariances are greater than zero. The first column of Table 3 reports the results of this test using the primary sample. As seen here, starting from the second lag, there is no evidence of a positive covariance: they are mostly negative and statistically not different from zero, which seemingly casts doubt on the HIP process.

There are two separate issues about the use of this test. The first one is that the non-rejection may be due to the low power of the test. To address this issue, consider the case where the covariances are most likely to be positive, that is, when $\rho = 1$. But note that while in this case covariances must be positive for all $n \geqslant 2$, their magnitude is very small (0.00038) making it difficult to distinguish it from a value of zero implied by the RIP process.[17]

There is, however, a second important concern about the use of this test. To see this point, recall that if in fact $\rho < 1$, the second (negative) term is present, so the covariances are not positive up to a certain lag. A key question then is the following: what is the lowest lag at which the covariances should be expected to become positive? This is critical because both studies mentioned above have focused on the first 5 to 10 lags. Fig. 8 plots the autocovariances of income growth for the first 20 lags using the parameter estimates of the HIP process from Table 1. For the sample of all individuals (denoted "-o") autocovariances are negative *up to the 12th lag*, simply because σ_β^2 is quantitatively so small compared to the term in brackets. Similar calculations for individuals with college- and high school-education show that the covariances become

[17] Notice that even though MaCurdy's (1982) test cannot distinguish the autocovariances from zero, I am able to get statistically significant estimates of profile heterogeneity. The reason is that taking the averages of first order autocorrelations results in the loss of useful information. Instead I exploit the information in the entire variance covariance matrix which yields more precise information.

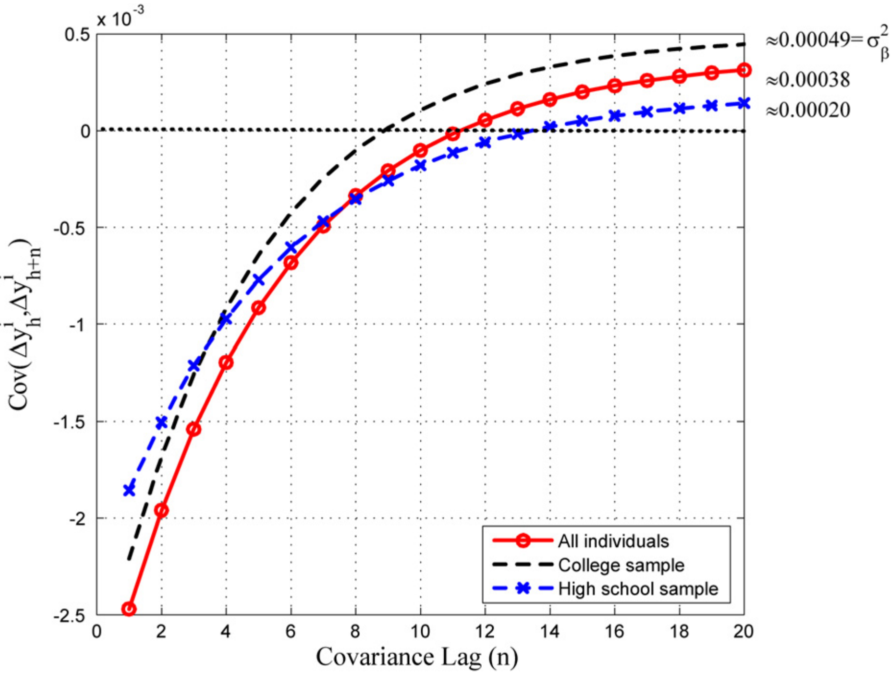

Fig. 8. Evaluating MaCurdy's test: the covariance structure of income growth.

positive only at the 10th and 15th lags respectively. These calculations show that the findings of negative autocovariances reported in MaCurdy (1982), Topel (1990) and Topel and Ward (1992) is also generated by the HIP process. Notice that this issue is separate from the power of the test, and suggests that even if an econometrician had access to a very large data set, the signs of these lower order covariances are dominated by the AR(1) component and are not very informative about HIP versus RIP.

I conduct a Monte Carlo exercise to further explore this issue. I simulate income paths using the HIP process and the parameter values from the second row of Table 1.[18] The second column in Table 3 displays the averages of autocovariances over 1000 replications along with the standard errors of the sampling distribution. The key point to notice from this table is that even though the average value of the autocovariances become positive after the 10th lag, none of the autocovariances up to the 18th lag are statistically significantly different from zero. In other words, for the sample size used in this paper, one should not expect to find empirical covariances to be statistically different from zero up to the 18th lag even if the true data generating process is HIP. In column 3, I repeat the Monte Carlo experiment assuming a sample size of 500,000, which is about 18 times larger than the primary sample in this paper, and close to 100 times larger than MaCurdy's sample (of 5130 observations). Even in this case, the 17th autocovariance is not significant while the 18th becomes barely significant at 5 percent level. Finally, the next two columns display the autocorrelation structure of income changes using simulated data along with its empirical counterpart. Again, the same pattern is apparent here: very weak negative autocorrelations, not significant after the first lag.

5.2. Initial conditions of the AR(1) process

A final issue to address is whether these results are sensitive to the assumption I made regarding the initial conditions of the stochastic process for income. In particular, Eq. (8) is derived (and the simulated data used in Table 3 is generated) assuming that the initial value of the persistent shock is the same for all individuals. However, if instead $\operatorname{var}(z_0^i) > 0$ then I would have:

$$\operatorname{cov}\left(\Delta\hat{y}_h^i, \Delta\hat{y}_{h+n}^i\right) = \sigma_\beta^2 - \rho^{n-1}\left(\frac{1-\rho}{1+\rho}\sigma_\eta^2\right) + \rho^{2(h-1)+n}\left[\operatorname{var}\left(z_0^i\right) - \frac{1-\rho}{1+\rho}\sigma_\eta^2\right], \tag{9}$$

where last set of brackets contain the new terms that are not present in Eq. (8). If the initial dispersion is sufficiently large, this last term would be positive, which in turn could weaken the argument made above that autocovariances up to a certain lag should be negative even when the true data generating process is HIP with an AR(1) component. To explore if this effect

[18] I first simulated income paths for 500,000 individuals. Then I drew 27,681 pairs of observations $(\Delta y_h^i, \Delta y_{h+n}^i)$ without replacement for randomly selected initial age, h, and $n = 1, \ldots, 18$. The first eighteen autocovariances of income changes are then calculated using this sample and the exercise is repeated 1000 times.

could be empirically important, I repeat the Monte Carlo exercise above, but now allow for a large initial dispersion in z. Specifically, I set $\text{var}(z_0^i) = 0.3$. To put this number in context, notice that the total variance of log labor income at age 25 is less than 0.3 (Fig. 4). Therefore, with this choice we are effectively attributing *all* inequality at that age to z_0^i and none to α^i or ε_0^i. As a different comparison, notice that this assumed variance of z_0^i is more than three times larger than the variance of the *stationary distribution* of z_h^i, which is $\sigma_\eta^2/(1-\rho^2) = 0.088$. Therefore, 0.3 is a generous upper bound, and this choice is intended to demonstrate the maximum impact initial conditions can plausibly have on the autocovariances of income growth.

Column 4 in Table 3 displays the resulting covariances, which are all higher (i.e., shifted up) from column 2, although the difference is quantitatively small. The fact that autocovariances are shifted up could be anticipated from Eq. (9), because the last term is positive for my choice of $\text{var}(z_0^i)$. But because the first 10 autocovariances in column 2 were negative, adding a small positive amount to each of them in fact makes them *closer* to zero, and therefore, statistically less significant than before. This result suggests that the existence of initial conditions provides one more reason for why MaCurdy's test may fail to reject the null of RIP even when the true data generating process is HIP. Turning to the higher order covariances (11th to 18th) that were already positive, these now become more positive than before, but are still very far from being significant at any conventional significance level. A closer inspection of Eq. (9) shows that the small effect of initial conditions on these autocovariances is not very surprising: the terms in brackets are multiplied by $\rho^{2(h-1)+n}$, which quickly vanishes as the cohort gets older and/or the lag order increases.[19] Overall, I conclude that the existence of dispersion in initial values of z does not affect the substantive conclusions about the expected sign and significance of the covariances of income growth.

5.3. Abowd and Card (1989)

A similar concern applies to the variant of this test implemented by Abowd and Card (1989, pp. 427–428). As an extension to MaCurdy's idea, these authors proposed to test if all higher order autocovariances are *jointly* equal to zero. The test essentially entails computing a weighted sum of *squared* autocovariances from lags 2 to 10, and comparing it to the corresponding critical value from the χ^2 distribution. However, as shown in Fig. 8 the deviations of autocovariances from zero up to the 10th lag are mainly due to the AR(1) component, and is in negative direction, rather than being due to HIP, and in the positive direction. But because covariances are squared, the test does not distinguish between negative and positive deviations. Therefore, with a large enough sample size, Abowd and Card's test would reject the null of zero even when the income process contains only an AR(1) component and there was no profile heterogeneity.[20] In other words, if one rejects this null hypothesis (of zero joint covariances), that would not necessarily support the HIP process either. Therefore, the interpretation of the results of this test is not straightforward when the data generating process has an AR(1) component. This is true regardless of sample size.

A second difficulty with interpreting Abowd and Card's test results as providing evidence against the HIP process is that they *jointly* test if both labor earnings *and* labor hours contain heterogeneous profiles. More specifically, they stack the autocovariances of income change (28 moments), hours change (28 moments), and the cross-covariances between the two variables at all lags and leads (56 moments) into a vector and test if all moments are jointly equal to zero (see Table 8, columns 1 and 2 in their paper). But it seems unlikely that labor hours will display significant heterogeneity in growth rates over the life-cycle (if they grow at all), in which case 84 of these moments should be expected to be zero. If this is the case, including these moments will bias their joint test towards non-rejection even when the first 28 moments are in fact non-zero. Notice that this problem is independent of the concerns with the AR(1) component raised above.

To summarize, these results suggest that the tests based on the *sign of autocovariances* used in the previous literature do not provide evidence on the absence or existence of profile heterogeneity. The HIP process generates the same negative and statistically insignificant autocovariance structure that was previously used to reject it.

6. Conclusion

The existing evidence from labor income data has commonly been interpreted by macroeconomists as strongly in favor of the RIP process. Consequently, almost all life-cycle (or overlapping generations) models in the literature are currently calibrated using the RIP process as the income process. In this paper I have reassessed the existing evidence and found that there are several pieces of evidence that lend support the HIP process. However, I have also discussed some issues that makes it difficult to definitively distinguish between the two hypothesis using income data alone.

I first argued that imposing a priori restrictions on income growth rate heterogeneity, as is done in the RIP process, introduces an upward bias into the estimated persistence if the true data generating process features heterogeneous profiles. When I allow for HIP, the estimates I obtained indicate substantial heterogeneity in income profiles, and income shocks with modest persistence. Second, I also show that the HIP process I estimate generates small and negative autocovariances, quan-

[19] For example, for $n = 10$, and $h = 10, 20, 30$ we get $\rho^{2(h-1)+n} = 0.0038$, 7.3×10^{-5}, and 1.37×10^{-6}.

[20] In fact, in this case the null would be rejected more easily because the covariance structure would shift downward making the lower order covariances more negative (and hence their squared value farther away from zero).

titatively similar to their empirical counterpart, casting doubt on the previous interpretation of this finding as supporting the RIP process.

The HIP process also implies that the income processes of high and low educated individuals differ in a key dimension: the dispersion of income growth rates is much larger for the former group than the latter. This is in contrast to the RIP process which indicates similar income processes for both groups (or more uncertainty for the latter group). This finding has potentially important implications for life-cycle studies which attempt to understand certain differences in economic behavior among education groups. Existing studies have used the RIP process as the income process, which often implies puzzling differences in behavior by education level (see, for example, Hubbard et al., 1994; and Davis et al., 2003).

I also show that identification between the two income processes relies on information throughout the covariance matrix, including the higher order covariances. However, because computing these higher order covariances require us to observe individuals at far apart points in their life-cycle, fewer and fewer individuals contribute to these moments, raising issues with selection and smaller sample sizes. A fruitful complementary approach could be to study forward looking economic choices which would contain valuable information about perceived future income risks.

It is important *not* to interpret the results of this paper (to the extent they are viewed as lending more support to the HIP process) as suggesting that income uncertainty is not as large as that implied by the RIP process. The statistical analysis conducted in this paper reveals an important systematic component by examining realized (ex post) wages, but is silent about whether (or how much) each individual knows about his own profile ex ante. The latter cannot be determined by examining labor income data alone, though it could in principle be inferred by studying the choices made by individuals. In Guvenen (2007) I conduct such an analysis and argue that in fact a substantial part of this systematic variation is likely to be unknown to individuals early on, and is revealed very slowly, implying that the HIP process also features substantial income uncertainty. However, the nature of this risk and its distribution over the life-cycle is different than in the RIP process.

Appendix A. Data appendix

The data are drawn from the first 26 waves of the PSID. I include an individual into the sample if he satisfies the following criteria for a total of twenty (not necessarily consecutive) years between 1968 and 1993. The individual (i) is a male head of household, (ii) is between 20 and 64 years old (inclusive), (iii) is not from the SEO sample (which oversamples poor households), (iv) has positive hours and labor income, (v) has hourly labor earnings more than $W_{\min}$ and less than $W_{\max}$, where I set $W_{\min}$ to $2 and $W_{\max}$ to $400 in 1993 and adjust them for previous years using the average growth rate of nominal wages obtained from BLS, (vi) worked for more than 520 hours (10 hours per week) and less than 5110 hours (14 hours a day, everyday)

There were a total of 1270 individuals satisfying these conditions for at least twenty years who comprise the primary sample. If, instead, I required these criteria to be satisfied for twenty *consecutive* years, there would be 210 fewer observations resulting in a sample size of 1060, so this flexibility in selection criteria increases the sample size by about 20 percent. When constructing the two subsamples defined by education used in Section 3.1, I exclude 53 individuals who have inconsistencies in their education variables over time.

A.1. Variable definitions

Age of the head is constructed by taking the first report of age by the individual and by adding the necessary number of years to obtain the age in other years (variable name V16631 in 1989). This is done to eliminate the occasional non-changes or two-year jumps in the age variable between consecutive interviews as a result of interviews not being conducted exactly one year apart.

Head's total labor income measure is comprehensive and includes salary income, bonuses, overtime, commissions, and the labor part of farm, business, market gardening, and roomers and boarders income, as well as income from professional practice or trade (variable name V17534 in 1989).

Annual labor hours of the head is the self-reported annual hours worked by the individual (variable name V16335 in 1989).

Head's average hourly earnings is calculated by the PSID as the ratio of total labor income to annual labor hours (variable name V17536 in 1989).

Education is based on the categorical education variable in the years it is available (variable name V17545 in 1989), and on years of schooling completed when this variable was not available (variable name V30620 in 1989). Potential labor market experience is constructed from this latter variable.

The traditional approach to panel construction (Lillard and Weiss, 1979; MaCurdy, 1982; Abowd and Card, 1989; and Baker, 1997, among others) requires an individual to satisfy the selection criteria for every year of the sample period to be included in the panel. Although this condition has the advantage of creating a balanced panel, it also has the drawback of reducing the sample size significantly as the time horizon expands, since individuals with even one year of missing data are excluded. I also require the individuals to be present in the sample for a long period of time while allowing for up to six missing observations for each individual. This is intended to make my panel construction comparable to earlier studies, while at the same time keeping a reasonably large number of observations. An alternative approach pursued by some recent studies is to include an individual into the panel if certain criteria are satisfied for a few—usually two or three—years

Table 4
Summary statistics for the primary sample, PSID 1968–1983

Year	Mean age	Mean years of education	Mean wage	Median wage	Variance of log(wage)	Mean earnings	Median earnings	Variance of log(earnings)	Number of observations
1968	31.460	11.88	14.60	13.09	0.2480	33,977	31,105	0.2916	604
1969	32.080	11.70	15.34	13.46	0.2621	35,472	32,319	0.3058	624
1970	32.536	12.62	16.07	14.21	0.2704	36,825	32,976	0.3091	663
1971	32.486	10.99	15.84	14.21	0.2746	35,785	31,896	0.3114	742
1972	32.689	12.64	15.78	14.16	0.2756	35,439	32,110	0.2959	809
1973	32.735	12.51	16.26	14.52	0.2632	36,968	33,312	0.2798	892
1974	32.885	12.65	16.57	14.88	0.2621	37,630	34,226	0.2884	974
1975	33.107	12.77	16.72	15.09	0.2652	37,181	32,772	0.3047	1040
1976	33.542	13.20	16.25	14.56	0.2673	35,733	31,726	0.3047	1096
1977	33.940	13.12	16.70	15.23	0.2830	37,413	34,043	0.3091	1149
1978	34.439	13.10	17.06	15.48	0.2621	38,412	35,116	0.2905	1209
1979	35.051	13.09	17.71	16.00	0.2820	39,748	35,952	0.2970	1248
1980	36.080	13.09	18.05	16.73	0.2746	40,536	36,364	0.2777	1249
1981	37.049	13.11	18.33	17.01	0.2777	40,719	36,816	0.2873	1256
1982	38.065	13.10	18.69	16.88	0.3272	40,223	36,877	0.3411	1256
1983	39.094	13.11	18.82	17.34	0.3181	41,078	36,393	0.3636	1236
1984	40.027	13.25	18.99	17.18	0.3434	41,855	37,140	0.3982	1244
1985	41.018	13.24	19.56	17.27	0.3684	43,961	38,280	0.3982	1245
1986	42.022	13.51	20.15	17.86	0.3795	45,416	37,997	0.4343	1233
1987	43.039	13.51	20.44	17.93	0.3856	45,863	38,835	0.4369	1241
1988	43.787	13.54	20.57	17.89	0.3844	47,629	38,670	0.4225	1210
1989	44.507	13.58	21.07	18.03	0.4096	48,628	38,649	0.4651	1185
1990	45.111	13.58	20.62	17.57	0.4020	48,067	38,429	0.4802	1149
1991	45.559	13.61	20.66	17.41	0.3919	47,535	39,302	0.4789	1097
1992	46.200	13.67	21.14	17.34	0.4251	47,927	37,810	0.4775	1054
1993	46.689	13.70	22.93	18.79	0.4476	51,082	39,965	0.5314	976

(Haider, 2001; Storesletten et al., 2004). Haider's estimates from the HIP specification are similar to mine (in particular, $\rho = 0.64$, and $\sigma_\beta^2 = 0.00041$). Table 4 reports some summary statistics for the primary sample.

Appendix B. Estimated time effects in variances (all individuals)

Table 5
Estimated variance of the persistent and transitory shocks by year

Year	RIP (all individuals)		HIP (all individuals)	
	$\text{var}(\eta_t)$ Estimate (S.E)	$\text{var}(\varepsilon_t)$ Estimate (S.E)	$\text{var}(\eta_t)$ Estimate (S.E)	$\text{var}(\varepsilon_t)$ Estimate (S.E)
1968	.0098 (.0036)	.0518 (.019)	.0249 (.0074)	.0432 (.015)
1969	.0081 (.0032)	.0449 (.016)	.0212 (.0059)	.0376 (.012)
1970	.0118 (.0039)	.0549 (.023)	.0245 (.0065)	.0453 (.017)
1971	.0097 (.0032)	.0388 (.013)	.0155 (.0039)	.0338 (.010)
1972	.0152 (.0041)	.0468 (.017)	.0207 (.0048)	.0397 (.013)
1973	.0088 (.0024)	.0577 (.019)	.0265 (.0059)	.0568 (.016)
1974	.0077 (.0021)	.0613 (.018)	.0206 (.0043)	.0364 (.009)
1975	.0106 (.0033)	.0484 (.015)	.0221 (.0048)	.0313 (.008)
1976	.0091 (.0023)	.0593 (.020)	.0254 (.0054)	.0391 (.011)
1977	.0126 (.0034)	.0512 (.016)	.0165 (.0038)	.0572 (.017)
1978	.0141 (.0044)	.0484 (.019)	.0238 (.0056)	.0436 (.012)
1979	.0178 (.0051)	.0568 (.021)	.0205 (.0042)	.0514 (.015)
1980	.0150 (.0048)	.0623 (.024)	.0280 (.0065)	.0560 (.018)
1981	.0197 (.0058)	.0659 (.018)	.0343 (.0079)	.0488 (.014)
1982	.0226 (.0061)	.0528 (.017)	.0386 (.0084)	.0454 (.016)
1983	.0251 (.0052)	.0592 (.022)	.0449 (.0080)	.0379 (.011)
1984	.0307 (.0077)	.0633 (.023)	.0360 (.0074)	.0417 (.014)
1985	.0259 (.0069)	.0638 (.020)	.0441 (.0086)	.0363 (.012)
1986	.0210 (.0048)	.0786 (.025)	.0350 (.0077)	.0568 (.017)
1987	.0221 (.0060)	.0693 (.024)	.0425 (.0084)	.0454 (.015)
1988	.0154 (.0049)	.0812 (.028)	.0385 (.0071)	.0563 (.018)
1989	.0207 (.0043)	.0823 (.025)	.0305 (.0065)	.0620 (.019)
1990	.0192 (.0065)	.0736 (.022)	.0244 (.0053)	.0531 (.016)
1991	.0173 (.0049)	.0692 (.026)	.0327 (.0069)	.0469 (.012)
1992	.0173	.0849 (.028)	.0327	.0657 (.019)

Appendix C. The estimation method

This appendix describes the minimum distance estimation (MDE) of the parameters of the income process given in Eq. (2). Let c_n be a typical element of the covariance matrix $\mathbf{C}$ of the income residuals, where $n = 1, \ldots, N$ $(= T(T+1)/2)$ enumerates unique elements of this matrix, and let $d_n(X_i, b)$ denote the corresponding model covariances given Eq. (4), where b denote the parameters of the income process. Define $F_n(b, X_i, \Upsilon_{in}) = \Upsilon_{in}[c_n - d_n(X_i, b)]$, where Υ_{in} is an indicator function that is equal to 1 if individual i contributes to moment condition n, and zero otherwise. Finally, stack all moment conditions into an $(N \times 1)$ vector: $\mathbf{F}(b, X_i, \Upsilon_i) \equiv [F_1(b, X_i, \Upsilon_{i1}), \ldots, F_N(b, X_i, \Upsilon_{iN})]'$, where Υ_i is the indicator functions stacked into a vector conformably to $\mathbf{F}$. The moment conditions that I am estimating are of the form:

$$E_i\big[\mathbf{F}(b, X_i, \Upsilon_i)\big] = 0.$$

The MD estimator is the solution to

$$\min_b \left[I^{-1}\sum_{i=1}^{I}\mathbf{F}(b, X_i, \Upsilon_i)\right]' \tilde{A}_N \left[I^{-1}\sum_{i=1}^{I}\mathbf{F}(b, X_i, \Upsilon_i)\right]$$

where $\tilde{A}_N$ is a positive definite matrix. Chamberlain (1984) discusses the choice of the asymptotically optimal weighting matrix. However, Altonji and Segal (1996) provide Monte Carlo evidence showing that the optimal weighting matrix often results in substantial small sample bias and recommend the use of an identity matrix instead, and I follow their recommendation. Notice however that because the panel is not balanced, each moment in the vector $\mathbf{F}$ is calculated using a different number of observations (determined by the non-zero elements of Υ_{in}). To adjust for this difference, I set $\tilde{A}_N \equiv A_N A_N$, where A_N is a diagonal matrix with element (I/I_n) at the nth diagonal, where I_n is obtained by summing Υ_{in} over i. This choice implies that each moment is calculated using all available observations and the resulting moments are weighted with an identity matrix.

Finally, the estimator $\hat{b}_N$ is consistent, asymptotically Normal with an asymptotic covariance matrix $\Sigma \equiv (D'D)^{-1}D'\Omega D \times (D'D)^{-1}$, where D is the Jacobian of the vector of moments, $E[\partial\mathbf{F}(b, X_i, \Upsilon_i)/\partial b']$, and Ω is the covariance matrix $E[\mathbf{F}(b, X_i, \Upsilon_i)\mathbf{F}(b, X_i, \Upsilon_i)']$. Both expectations are replaced by sample averages when implemented.

Acknowledgments

I would like to thank Daron Acemoglu, Joe Altonji, Orazio Attanasio, Michael Baker, Mark Bils, Richard Blundell, Jeremy Greenwood, James Heckman, Hugo Hopenhayn, Lawrence Katz, Narayana Kocherlakota (the editor), Per Krusell, Lee Ohanian, Gianluca Violante, an associate editor, two anonymous referees, the seminar participants at Harvard, MIT, Northwestern, NYU, New York FED, Rochester, UCLA, UPenn, USC, UT-Austin, Western Ontario, Yale, the 2004 SED conference in Florence, and the 2004 NBER Summer Institute for useful comments. I would also like to thank Damba Lkhagvasuren for excellent research assistance, and the National Science Foundation for financial support under grant SES-0351001. This paper originally started as a joint project with Shamim Mondal, and special thanks go to him. The usual disclaimer applies.

References

Abowd, J., Card, D., 1989. On the covariance structure of earnings and hours changes. Econometrica 57, 411–445.

Altonji, J.G., Segal, L.M., 1996. Small sample bias in GMM estimation of covariance structure. Journal of Business and Economics Statistics 109, 659–684.

Baker, M., 1997. Growth-rate heterogeneity and the covariance structure of life-cycle earnings. Journal of Labor Economics 15, 338–375.

Baker, M., Solon, G., 2003. Earnings dynamics and inequality among Canadian men, 1976–1992: Evidence from longitudinal income tax records. Journal of Labor Economics 21, 289–321.

Becker, G.S., 1964. Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education. reprinted 1993. Univ. of Chicago Press.

Ben-Porath, Y., 1967. The production of human capital and the life cycle of earnings. Journal of Political Economy 75.

Bound, J., Krueger, A.B., 1991. The extent of measurement error in longitudinal earnings data: Do two wrongs make a right? Journal of Labor Economics 9, 1–24.

Campbell, J.Y., Cocco, J.F., Gomes, F.J., Maenhout, P.J., 2001. Investing retirement wealth: A life-cycle model. In: Campbell, J., Feldstein, M. (Eds.), Risk Aspects of Investment-Based Social Security Reform. Univ. of Chicago Press, Chicago, IL, pp. 439–473.

Carroll, C., Samwick, A., 1997. The nature of precautionary wealth. Journal of Monetary Economics 40 (1), 41–72.

Castaneda, A., Diaz-Jimenez, J., Rios-Rull, J.V., 2003. Accounting for earnings and wealth inequality. Journal of Political Economy 111 (4), 818–857.

Chamberlain, G., 1984. Panel data. In: Griliches, Z., Intriligator, M. (Eds.), Handbook of Econometrics. Amsterdam, North-Holland, pp. 1274–1318.

Christiano, L., Eichenbaum, M., 1990. Unit roots in real GNP: Do we know, and do we care? Carnegie-Rochester Conference Series on Public Policy 32, 7–61.

Constantinides, G.M., Duffie, D., 1996. Asset pricing with heterogeneous consumers. Journal of Political Economy 104, 219–240.

Davis, S.J., Kubler, F., Willen, P., 2003. Borrowing costs and the demand for equity over the life-cycle. Working paper, University of Chicago, 2003.

Deaton, A., Paxson, C., 1994. Intertemporal choice and inequality. Journal of Political Economy 102, 437–467.

Gourinchas, P.-O., Parker, J.A., 2002. Consumption over the life cycle. Econometrica 70, 47–89.

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

Guvenen, F., Kuruscu, B., 2007. A quantitative analysis of the evolution of the US wage distribution: 1970–2000, Working paper No. 13095, National Bureau of Economic Research.

Guvenen, F., Smith, A., 2007. Inferring labor income risk from economic choices: An indirect inference approach. Working paper, University of Texas at Austin.

Haider, S.J., 2001. Earnings instability and earnings inequality of males in the United States: 1967–1991. Journal of Labor Economics 19, 799–836.

Hall, R.E., 1971. The measurement of quality change from vintage price data. In: Griliches, Z. (Ed.), Price Indexes and Quality Change. Harvard Univ. Press, Cambridge, MA.
Hause, J.C., 1980. The fine structure of earnings and the on-the-job training hypothesis. Econometrica 48, 1013–1029.
Heathcote, J., Storesletten, K., Violante, G., 2004. The macroeconomic implications of rising wage inequality in the United States. Working paper, NYU.
Heathcote, J., Storesletten, K., Violante, G., 2005. Two views of inequality over the life-cycle. Journal of the European Economic Association P&P 3, 765–775.
Heckman, J.J., Robb, R., 1985. Using longitudinal data to estimate age, period, and cohort effects in earnings equations. In: Mason, W., Fienberg, S. (Eds.), Cohort Analysis in Social Research: Beyond the Identification Problem. Springer-Verlag, New York.
Hubbard, R.G., Skinner, J., Zeldes, S.P., 1994. The importance of precautionary saving in explaining individual and aggregate savings. Carnegie-Rochester Conference Series on Public Policy 40, 59–105.
Hubbard, R.G., Skinner, J., Zeldes, S.P., 1995. Precautionary saving and social insurance. Journal of Political Economy 103, 360–399.
Huggett, M., 1996. Wealth distribution in life-cycle economies. Journal of Monetary Economics 38, 469–494.
Katz, L.F., Autor, D.H., 1999. Changes in the wage structure and earnings inequality. Handbook of Labor Economics, vol. 3. North-Holland, Amsterdam.
Katz, L.F., Murphy, K.M., 1992. Changes in relative wages, 1963–87: Supply and demand factors. Quarterly Journal of Economics 107, 35–78.
Lillard, L.A., Reville, R.T., 1999. Life-cycle human capital investment: New evidence on an old issue. Working paper, University of Michigan.
Lillard, L.A., Weiss, Y.A., 1979. Components of variation in panel earnings data: American scientists, 1960–70. Econometrica 47, 437–454.
Lucas Jr., R.E., 2003. Macroeconomic priorities. American Economic Review 93, 1–14.
MaCurdy, T., 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.
Meghir, C., Pistaferri, L., 2004. Income variance dynamics and heterogeneity. Econometrica 72, 1–32.
Mincer, J., 1974. Schooling, experience and earnings. Columbia Univ. Press, New York.
Moffitt, R.A., Gottschalk, P., 1995. Trends in the autocovariance structure of earnings in the US, 1969–1987. Working paper, Brown University.
Storesletten, K., Telmer, C., Yaron, A., 2001. The welfare costs of business cycles revisited: Finite lives and cyclical variation in idiosyncratic risk. European Economic Review 45, 1311–1339.
Storesletten, K., Telmer, C., Yaron, A., 2004. Cyclical dynamics in idiosyncratic labor-market risk. Journal of Political Economy 112, 695–717.
Topel, R.H., 1990. Specific capital, mobility and wages: Wages rise with job seniority. Working paper No. 3294, NBER.
Topel, R.H., Ward, M.P., 1992. Job mobility and the careers of young men. Quarterly Journal of Economics 107, 439–479.