An Empirical Investigation of Labor Income Processes*

Fatih Guvenen†

August 25, 2005

Abstract

The current literature offers two views on the nature of the income process. According to the first view, which we call the "Restricted Income Profiles" (RIP) model (MaCurdy, 1982), individuals are subject to large and very persistent shocks, while facing similar life-cycle income profiles (conditional on a few characteristics). According to the alternative view, which we call the "Heterogeneous Income Profiles" (HIP) model (Lillard and Weiss, 1979), individuals are subject to income shocks with modest persistence, while facing individual-specific income profiles. In this paper, we first show that ignoring profile heterogeneity, when in fact it is present, introduces an upward bias into the estimates of persistence. Second, we estimate a parsimonious parameterization of the HIP model that is suitable for calibrating economic models. The estimated persistence is about 0.8 in the HIP model compared to about 0.99 in the RIP model. Moreover, the heterogeneity in income profiles is estimated to be substantial, explaining between 65 to 80 percent of income inequality at the time of retirement. We also analyze the differences in the income process by education and find that profile heterogeneity is significantly larger among higher educated individuals. Finally we 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 model, casting doubt on the previous interpretation of this evidence.

JEL classification: C33, D31, J31.

*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, Per Krusell, Lee Ohanian, Kjetil Storesletten, Gianluca Violante, 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.

†Email: guvn@mail.rochester.edu; Website: http://www.econ.rochester.edu/guvenen/guvenen.htm

1 Introduction

When markets are incomplete, labor income risk plays a central role in many decisions that individuals make. Understanding the nature of income risk is thus an essential prerequisite for understanding a wide range of economic questions, such as the determination of wealth inequality (Hubbard, Skinner and Zeldes, 1995; Huggett, 1996), life-cycle consumption behavior (Carroll and Samwick, 1997; Gourinchas and Parker, 2002; Guvenen, 2005), the welfare costs of business cycles (Storesletten, Telmer and Yaron, 2001; Lucas, 2003), and the determination of asset prices (Constantinides and Duffie, 1996), among others. The conclusions one reaches in these analyses clearly depend on the properties of the labor income process used to calibrate these models.

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

$$\begin{aligned} y_h^i &= \beta^i h + z_h^i \\ z_h^i &= \rho z_{h-1}^i + \eta_h^i, \end{aligned} \tag{1}$$

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 we 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 $\rho \ll 1$ and $\sigma_\beta^2 \gg 0$ (Lillard and Weiss, 1979; Hause, 1980; and more recently Baker, 1997; and Haider, 2001). Thus according to this first view, which we call the "Heterogeneous Income Profiles" (HIP) model, individuals are subject to shocks with modest persistence, while facing lifecycle 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 (Becker, 1965; Ben-Porath, 1967).

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

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

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).[2] Thus, according to this alternative view, which we call the "Restricted Income Profiles" (RIP) model, individuals are subject to extremely persistent—nearly random walk—shocks, while facing similar life-cycle income profiles.

In this paper, we examine labor income data from several angles to help distinguish between these two income processes. We begin our analysis by showing that ignoring the heterogeneity in income growth rates (as is done in the RIP model), 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.

We study an example to show that this bias can be substantial: when labor income is generated from the HIP process given above with *i.i.d* shocks (equation (1)), the persistence parameter is estimated to be about 0.90 if RIP is assumed, instead of the true value of zero. This example thus suggests that allowing for heterogeneity in income growth rates is critical for the consistent estimation of the persistence parameter. It also explains why the previous studies estimating RIP models always obtained persistence parameters much higher than those implied by their HIP counterparts. In section 2.4 we also discuss more generally which features of labor income data allow us to distinguish between the RIP and HIP models.

We 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) process requires two state variables (the values of the AR and MA innovations) and the ARMA(1,2) requires three state variables to form optimal forecasts of the income process. Consequently, the majority of the existing life-cycle (or overlapping generations) models are instead calibrated using an income process fea-

[2] It is then curious that the conclusion reached by MaCurdy's test seems to contradict the direct estimation evidence supporting the HIP model. We discuss this point below.

turing an AR(1) component plus a transitory shock (among many others, Hubbard et al., 1995; Huggett, 1996; Campbell et al., 2001; Storesletten et al., 2004; Heathcote et al., 2004). 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 either an ARMA(1,1) or (1,2) process.[3] Thus, 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, we 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 model compared to about 0.99 when RIP is imposed. Together, these estimates imply that between 65 to 80 percent of income inequality at the age of retirement is due to heterogeneous profiles.

Third, we examine the differences in the income processes across education groups. While several studies have investigated this question in the context of RIP models (Hubbard et al., 1994; Carroll and Samwick, 1997), there exists no corresponding analysis in the context of HIP models. We find that in the HIP model the persistence parameter and innovation variance is very similar across education groups, but there is a major difference in a key dimension: 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 model, which implies similar income processes for both groups, with some mild evidence of larger innovation variances for lower educated individuals.

Fourth, we try to reconcile the test used by MaCurdy and others which does not reject the RIP model, with the direct estimation results which support the HIP model. 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 we emphasize a different point that applies even in large sample, where inflated size or low power are not relevant. We argue that the tests used by MaCurdy (1982) and Abowd and Card (1989) are not valid if the alternative hypothesis is a stochastic process (such as HIP) that contains an AR(1)

[3]Baker experiments with an AR(1) process to provide a comparison to Lillard and Weiss (1979). But as is well-known classical measurement error biases estimates of persistence downward when transitory shocks are not allowed.

component. To see this point, first it is easily shown from equation (1) that

$$cov(\Delta y^i_h, \Delta y^i_{h+n}) = \sigma^2_\beta - [\rho^{n-1}(\frac{1-\rho}{1+\rho})\sigma^2_\eta] \quad \text{for} \quad n \geq 2$$

Notice that the term in brackets vanishes as n gets large, so *higher order* autocovariances of income changes must be positive if indeed $\sigma^2_\beta > 0$. 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 model the first 11 covariances will be negative (see figure 6), 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 exactly what is implied by the HIP model. In Section 4 we show that the autocovariance and autocorrelation structures generated by the estimated HIP model are also quantitatively similar to their empirical counterparts. These results cast doubt on the previous interpretation of this evidence in the literature as supporting the RIP model.

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 reconciles the direct estimation evidence with earlier tests implemented in the literature, and Section 5 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. Our main sample consists of male head of households between the ages of 22 and 62. We 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). We 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).

These criteria leave us with our 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, we 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, we 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 our data covers 1967-92. The (potential) labor market experience of an individual is defined as $h = (age - max($years of schooling, $12) - 6)$. Further details of the selection criteria, variable definitions and some summary statistics for the primary sample are contained in Appendix A.

2.2 A Statistical Model

The process for *log earnings*, $\widetilde{y}_{it}^{h}$, of individual i with h years of labor market experience in year t is given by

$$\widetilde{y}_{it}^{h} = g\left(\theta_{t}^{0}, \mathbf{X}_{it}^{h}\right) + f\left(\theta^{i}, \mathbf{X}_{it}^{h}\right) + z_{it}^{h} + \phi_{t}\varepsilon_{it}^{h} \tag{2}$$

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

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 (hence the coefficient vector θ_{t}^{0} is not individual-specific) 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 our 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, we do not pursue this approach in the baseline specification (Instead we capture all the cross-sectional variation in income growth rate in f). Later in the paper, we will allow for 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.

The second function, f, is the centerpiece of our 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. We assume this function to be linear in experience: $f\left(\boldsymbol{\theta}^i, \mathbf{X}_{it}^h\right) = \alpha^i + \beta^i h$, where the random vector $\boldsymbol{\theta}^i \equiv \left(\alpha^i, \beta^i\right)$ is distributed across individuals with zero mean, variances of σ_α^2 and σ_β^2, and covariance of $\sigma_{\alpha\beta}$.[4]

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, 2005). The current specification provides a reasonable trade-off for this purpose.[5]

Modeling the Dynamics of Income.—The stochastic 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).[6] 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. Second, there have been some significant changes in the sizes of both persistent and transitory income shocks over the sample period under study (c.f., Moffitt

[4]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.

[5]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.

[6]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 our purposes.

and Gottschalk, 1995; Meghir and Pistaferri, 2004). To capture this non-stationarity, we write z_{it}^h as an AR(1) process with heteroskedastic shocks:

$$z_{it}^h = \rho z_{it-1}^{h-1} + \pi_t \eta_{it}^h, \quad z_{it}^0 = 0,$$

where π_t captures possible time-variation in the innovation variance. Similarly, the transitory shock in equation (2), ε_{it}^h, is scaled by ϕ_t to account for possible non-stationarity in that component. The innovations η_{it}^h and ε_{it}^h 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_η^2 and σ_ε^2 respectively. Furthermore, measurement error is a pervasive problem in micro data sets, and income data in PSID is no exception. This measurement error will be captured in the transitory component if it is serially independent, or will be included in z_{it}^h 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, y_{it}^h, is obtained by regressing $\widetilde{y}_{it}^h$ 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 moments of this residual for a cohort of a given age are:

$$\begin{aligned} var\left(y_{it}^h\right) &= \left[\sigma_\alpha^2 + 2\sigma_{\alpha\beta} h + \sigma_\beta^2 h^2\right] + var\left(z_{it}^h\right) + \phi_t^2 \sigma_\varepsilon^2 \\ cov\left(y_{it}^h, y_{it+n}^{h+n}\right) &= \left[\sigma_\alpha^2 + \sigma_{\alpha\beta}\left(2h+n\right) + \sigma_\beta^2 h\left(h+n\right)\right] + \rho^n var\left(z_{it+n}^{h+n}\right), \end{aligned} \tag{3}$$

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

$$\begin{aligned} var\left(z_{it}^1\right) &= \pi_t^2 \sigma_\eta^2, \\ var\left(z_{i1}^h\right) &= \pi_1^2 \sigma_\eta^2 \sum_{j=0}^{h-1} \rho^{2j}, && t = 1, h > 1 \\ var\left(z_{it}^h\right) &= \rho^2 var\left(z_{it-1}^{h-1}\right) + \pi_t^2 \sigma_\eta^2, && t > 1, h > 1 \end{aligned} \tag{4}$$

Note that in the first line we implicitly assume that the initial value of the persistent shock is zero for all individuals. In the second line we 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

Figure 1: Ignoring Profile Heterogeneity Results in an Upward Bias in Estimated Persistence

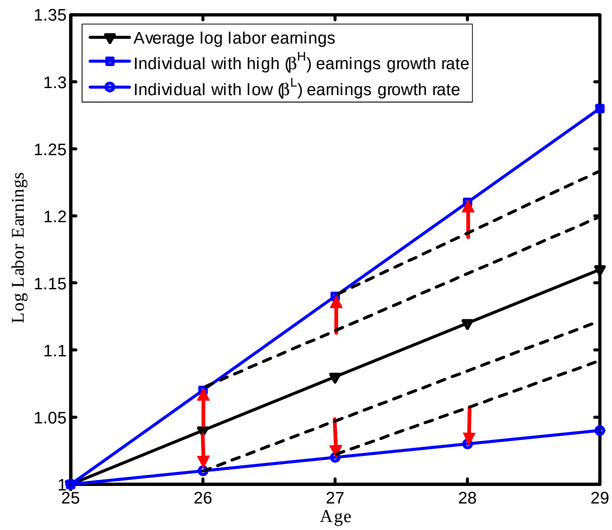

accumulated effect over the last h years.

Our estimation strategy 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 (Chamberlain (1984)). A typical element of $\mathbf{C}$ (at location $(\tau, \tau + n)$) is obtained by averaging $\left(y^h_{i,\tau} y^{h+n}_{i,\tau+n}\right)$ 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 (3) 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 we relegate its details (including the choice of weighting matrix, the exact formulas used, and related issues) to Appendix B.

2.3 Profile Heterogeneity and the Estimates of Persistence

Before proceeding further, we show that restricting income profiles across the population (as in the RIP model), 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 figure 1. Clearly, the income paths of both of these individuals will deviate from the average profile (denoted with "-^") 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.)

The resulting bias can be substantial as can be seen in the following example. Consider a simplified version of the income process given in (2): $\widetilde{y}_{it}^h = \alpha^i + \beta^i h + \varepsilon_{it}^h$, where β^i has mean $\overline{\beta}$, and ε_{it}^h 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 rates (assuming a life-cycle profile of $\alpha^i + \overline{\beta}h$ for all individuals). In this case, the income residuals are:

$$y_{it}^h \equiv \widetilde{y}_{it}^h - \left(\alpha^i + \overline{\beta}h\right) = \left(\beta^i - \overline{\beta}\right) h + \varepsilon_{it}^h$$

It easy to see that y_{it}^h 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 (we relax this assumption below). Then, under the (incorrect) assumption of RIP, a consistent estimator of the persistence of income shocks is the minimizer of $(1/I)\sum_{i=1}^{I}\left(y_{it+1}^{h+1} - \widehat{\rho} y_{it}^h\right)^2$, which is given by

$$\widehat{\rho} = \frac{h\,(h-1)\,\sigma_\beta^2}{(h-1)^2\,\sigma_\beta^2 + \sigma_\varepsilon^2}.$$

Notice that $\widehat{\rho}$ 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, we substitute some plausible values (that is, values consistent with our estimates in the next section) into this formula: $\sigma_\beta^2 = 0.0004$, and $\sigma_\varepsilon^2 = 0.03$. If the observed cohort is 44 years old $(h = 20)$ the estimated persistence is $\widehat{\rho} = 0.87$. Similarly, if $h = 30$, one obtains $\widehat{\rho} = 0.95$. 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 $\widehat{\rho} = 0.91$, even though the true persistence is, again, zero.

Furthermore, since this bias arises from heterogeneity in *growth rates*, the fact that we accounted for fixed effects in *levels*—as is commonly done in the literature—had no mitigating effects. In other words, if we also restrict α^i across individuals in the calculations above, the corresponding values of $\widehat{\rho}$ 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.

Figure 2: The Fits of the HIP and RIP Models to the Empirical Age-Inequality Profile of Income

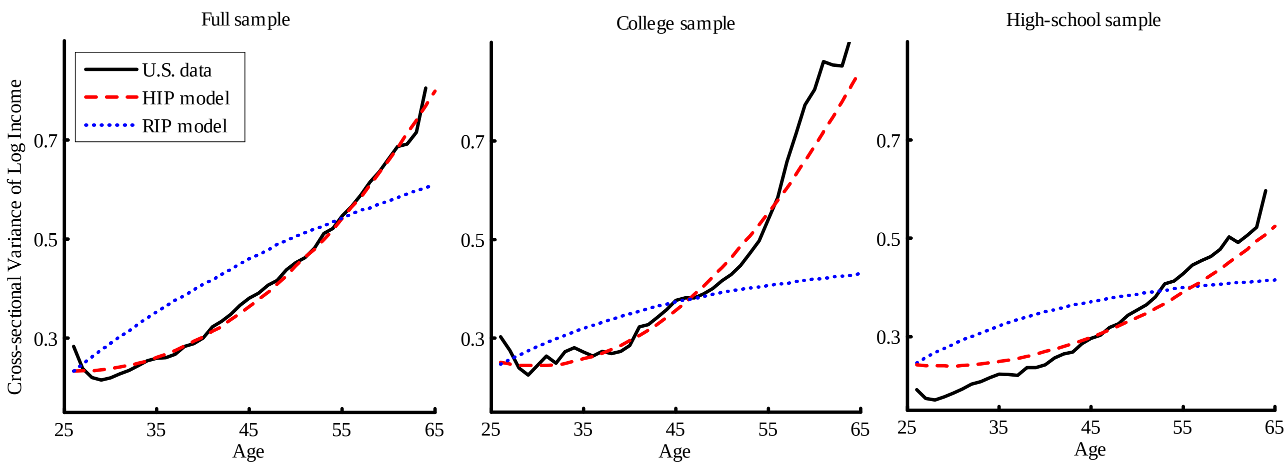

2.4 RIP versus HIP: Where Does Identification Come From?

The problem of distinguishing between the RIP and HIP models is reminiscent of the familiar debate in macroeconomics about whether GDP growth is better represented by a stochastic trend (RIP model), or by stationary shocks around a deterministic trend (HIP model). Given the well-known difficulties associated with distinguishing between those two hypotheses (c.f., Christiano and Eichenbaum (1990)), it seems reasonable to suspect a similar difficulty in the current context. Thus an important question to answer is the following: Where does identification between the RIP and HIP models 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, we can characterize the evolution of the cross-sectional distribution of income as a cohort gets older. As we explain below, it then becomes possible to distinguish between the RIP and HIP models by exploiting the different implications of each process for the evolution of this cross-sectional distribution.

To see this clearly, consider the case where the panel data set contains income observations on a single cohort over time. Furthermore, suppose that income shocks have stationary variances ($\phi^2 = \pi^2 \equiv 1$).[7] In this case, the second moments of the cross-sectional distribution for this cohort are given by:

$$\begin{aligned} var\left(y_i^h\right) &= \left[\sigma_\alpha^2 + 2\sigma_{\alpha\beta}h + \sigma_\beta^2 h^2\right] + var\left(z_i^h\right) + \sigma_\varepsilon^2 \\ cov\left(y_i^h, y_i^{h+n}\right) &= \left[\sigma_\alpha^2 + \sigma_{\alpha\beta}\left(2h+n\right) + \sigma_\beta^2 h\left(h+n\right)\right] + \rho^n var\left(z_i^h\right), \end{aligned} \tag{5}$$

where we eliminated the subscript t, since time and age are perfectly correlated within each cohort.

There are two sources of identification, which can be seen by inspecting these formulas. The first piece of information is provided by the change in the cross-sectional variance of income as the cohort ages, which is shown on the first line of (5). The terms in the square bracket capture the effect of profile heterogeneity, which is a *convex* increasing function of age. The second term captures the effect of the AR(1) shock, which is a *concave* increasing function of age as long as $\rho < 1$. Thus, if the variance of income in the data increases in a convex fashion as the cohort gets older, this would be captured by the HIP terms (notice that the coefficient on h^2 is σ_β^2), whereas a non-convex shape would be captured by the presence of AR(1) shocks. Figure 2 (solid lines) plots the age-variance profile of income residuals for the whole population as well as for college and high-school educated individuals separately. As can be seen here, the variance increases in a convex fashion in all three cases suggesting a need for profile heterogeneity.[8] (We discuss other aspects of this figure in the next section).

The second source of identification is provided by the autocovariances displayed in the second line. The covariance between ages h and $h+n$ is again composed of two parts. As before, the terms in the square bracket capture the effect of heterogeneous profiles and is a convex function of age. Moreover, the coefficients of the linear and quadratic terms depend both on h and n, which allows covariances to be decreasing, increasing or non-monotonic in the lag at each age. The second term captures the effect of the AR(1) shock, and notice that for a given h, it depends on the covariance lag n only through the

[7]The expressions in (3) and (4) make clear how the time-effects π_t and ϕ_t are identified: 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. This implies however that the two time-effects are not separately identified at the last date. To obtain identification at T we make the assumption that $\pi_{T-1}^2 = \pi_T^2$ following Heathcote et al. (2004).

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

Figure 3: The Covariance Structure of Income Residuals for College-Educated Individuals

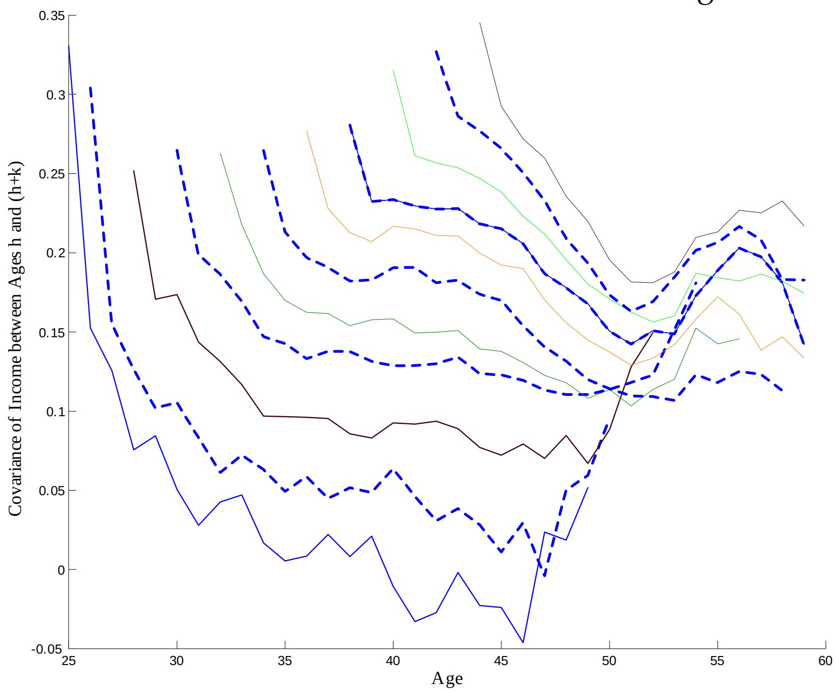

geometric discounting term ρ^n. The strong prediction of this form is that, starting at age h, covariances should decay geometrically at the rate ρ, regardless of the initial age. Thus, in the RIP model (which only has the AR(1) component) covariances are restricted to decay at the same rate at every age, and cannot be non-monotonic in n.

Figure 3 displays some representative elements of the empirical covariance matrix to make it easier to visualize. For example, the left most (solid) line plots $cov\left(y_i^{25}, y_i^{25+n}\right)$ for $n = 0, 1, .., 25$, and other lines plot the same for $h = 27, 29, ..., 45$ (subject to $h + n < 60$).[9] The key point to observe in this figure is that autocovariances are non-monotonic: they are typically convex (U-shaped), first decreasing and then increasing with the lag order.[10]

[9]To obtain this covariance matrix, we first regressed the raw covariances on cohort dummies to eliminate cohort effects. The covariances still contain the time-effects in variances, an issue that we address in the next footnote.

[10]As mentioned in the previous footnote, we do not eliminate time-effects when constructiving these covariances. However, note that the conclusions drawn about the shape of the covariance matrix in the previous paragraph are independent of the presence of time effects. This is because the covariances given in the second line of (5) would be exactly the same even if we do not restrict time effects $\phi^2 \neq 1$, and $\pi^2 \neq 1$ (compare the second lines of (3) and (5)). So the non-monotone shape of covariance structure cannot be explained by appealing to non-stationary shock variances in the present RIP framework.

There is however, an alternative way to introduce time effects, proposed by Moffitt and Gottschalk (1995), where π_t interacts with z_{it} rather than with its innovations, η_{it}. This specification in principle allows for

Table 1: ESTIMATING THE PARAMETERS OF THE LABOR INCOME PROCESS

	Group	Model	ρ	σ^2_α	σ^2_β	$corr_{\alpha\beta}$	σ^2_η	σ^2_ε
(1)	A	RIP	.988	.058	—	—	.015	.061
			(.024)	(.011)			(.007)	(.010)
(2)	A	HIP	.821	.022	.00038	$-.23$	.029	.047
			(.030)	(.074)	(.00008)	(.43)	(.008)	(.007)
(3)	C	RIP	.979	.031	—	—	.0099	.047
			(.055)	(.021)			(.013)	(.020)
(4)	C	HIP	.805	.023	.00049	$-.70$	.025	.032
			(.061)	(.112)	(.00014)	(1.22)	(.015)	(.017)
(5)	H	RIP	.972	.053	—	—	.011	.052
			(.023)	(.015)			(.007)	(.008)
(6)	H	HIP	.829	.038	.00020	$-.25$	.022	.034
			(.029)	(.081)	(.00009)	(.59)	(.008)	(.007)

Notes: Standard errors are in parentheses. In the second column, A = all individuals, C = college-educated group, and H = high school educated group. Time effects in the variances of persistent and transitory shocks are included in the estimation in all rows, but are not reported to save space. The reported variances are averages over the sample period.

Moreover, covariances do not appear to decay at the same rate regardless of age. Both of these observations are consistent with the richer structure allowed by the HIP model, but not with the more restrictive RIP model.

3 Empirical Findings

We first estimate the parameters of the process (2) by ignoring individual-specific variation in income growth rates ($\sigma^2_\beta \equiv 0$) but allowing for an individual fixed-effect, α^i (RIP model). The first row in Table 1 displays the results. 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 that in the long-run the persistent component dominates the cross-sectional distribution of income.

Starting in the second row, we allow for heterogeneity in income growth rates (HIP model). The first main finding is that the estimated persistence falls from 0.988 to 0.82.

non-monotonicity in the covariance structure. We have re-estimated the HIP model with this alternative parameterization yielding $\rho = 0.78$, and $\sigma^2_\beta = 0.00042$, which is very similar results to our estimates from the baseline model reported in the next section.

Figure 4: The Remaining Effect of an AR(1) Shock for Different Values of ρ

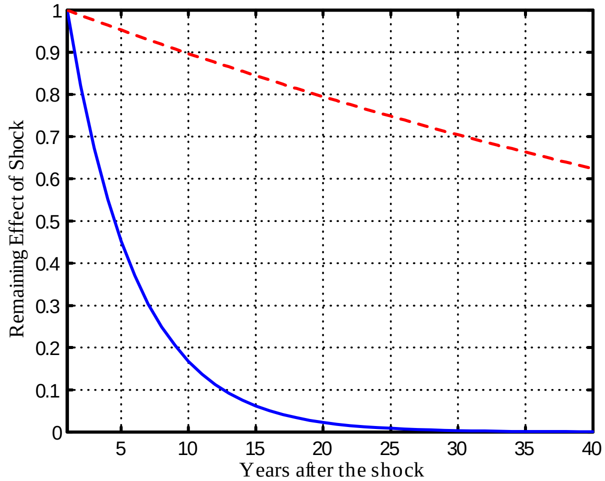

As is well-known, the difference between these two estimates is substantial (figure 4): when $\rho = 0.82$, the effect of an income shock is reduced to fourteen percent of its initial value in ten years, whereas for $\rho = 0.988$, it retains almost ninety percent of its initial value at the same horizon. 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 substantially different economic choices.

3.1 The labor income process by education group

We 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 models (Hubbard et al., 1994, and Carroll and Samwick, 1997). Thus, to provide a benchmark, we begin by estimating the RIP model for college- and high school-educated individuals. Rows 3 and 5 of table 1 report the parameter estimates for the two groups. The estimated persistence parameters are 0.979 and 0.972 for the college- and high school educated-groups respectively. Similarly, the innovation variances of the AR(1) shocks are 0.0099 and 0.0011 respectively. Overall, the estimated parameters reveal remarkably 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 our findings is that in the RIP model 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.

Now we re-estimate the income process of each group allowing for HIP (rows 4 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 a key dimension: the dispersion of income profiles is significantly larger for college-educated individuals ($\sigma_\beta^2 = .00049$) compared to high school-educated individuals ($\sigma_\beta^2 = .00020$). In fact, this difference could be partly anticipated from figure 2, which shows a larger increase in within-cohort income inequality among the former group than the latter group.

Figure 2 also shows how the estimated RIP and HIP models fit the age-inequality profile of income for the sample of all individuals (left panel) as well as for college- and high school-educated individuals (middle and right panels).[11] It is clear that allowing for heterogeneity in profiles helps the model better account for the slightly convex rise in dispersion over the life-cycle. Although it is certainly true that the fit of the RIP model would improve with a higher estimated persistence parameter, notice that both models are estimated to fit *all* the elements of the covariance matrix, and not just the variances plotted in the figure. A higher ρ results in a poorer fit for the off-diagonal elements in the RIP model, because these covariances show a steep decline as the lag order increases (see figure 3).

[11]The age-inequality profile is obtained by regressing the raw variances of each age-year cell, on age and cohort dummies following Deaton and Paxson (1994), and the graphs plot the coefficients on scaled age dummies

Table 2: ESTIMATES OF THE RIP MODEL BY EDUCATION LEVEL IN THE LITERATURE

Paper	Group	ρ	σ^2_η	σ^2_ε
Hubbard, Skinner,	<12 yrs of education	.955	.033	.040
and Zeldes (1994)		(.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	0-8 grades	1.0†	.0190	.0894
Samwick (1997)		–	(.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 ours is that these studies estimate income processes for household income whereas we estimate for individuals. †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 (Lillard and Weiss, 1979; Hause, 1980; Baker, 1997). A natural interpretation for 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 2 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, we rearrange the expression for cross-sectional income inequality (given in (3)) to obtain:

$$var_i(y_i^h) = \left(\sigma_\alpha^2 + \sigma_\varepsilon^2\right) + \left(\frac{1-\rho^{2h+1}}{1-\rho^2}\sigma_\eta^2\right) + \left[2\sigma_{\alpha\beta}h + \sigma_\beta^2 h^2\right],$$

where we substituted $var\left(z_i^h\right)$ from (4), and set π_t and ϕ_t equal to 1.

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

The last parenthesis contains 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,

Figure 5: Quantifying the Contribution of HIP to Income Inequality

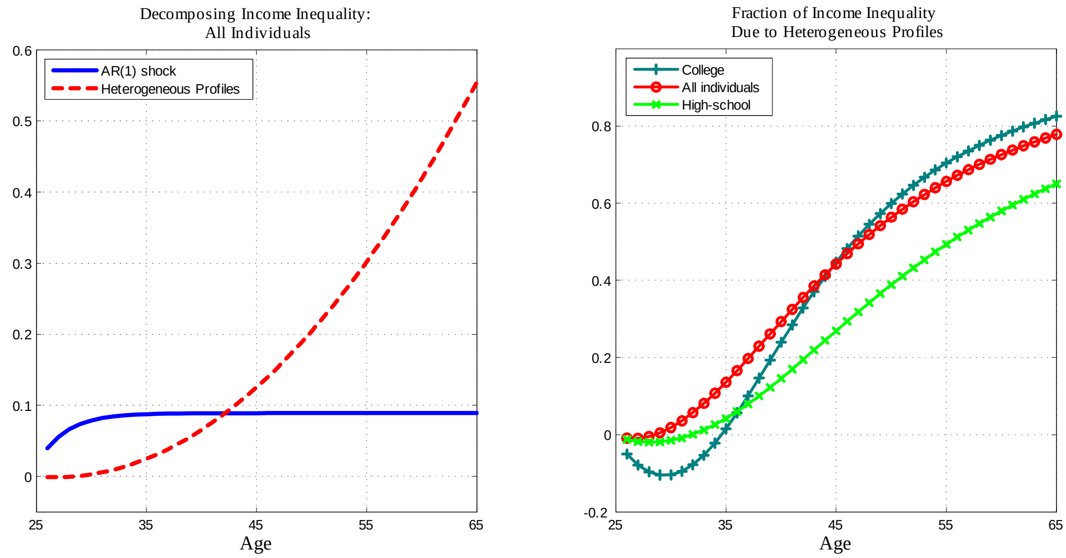

the effect of profile heterogeneity increases rapidly with age, and results in substantial inequality later in life.

The right panel of figure 5 plots the *fraction* of total inequality attributable to HIP. In the sample of all individuals (denoted "-o"), HIP accounts for 79 percent of inequality at retirement age. More importantly, HIP accounts for 81 percent of the inequality among college-educated individuals ("-+") and 64 percent of the inequality among high school-educated individuals ("-x").[12] 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.

[12]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 ($-.70$) 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.

4 A Comparison to the Existing Literature

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

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. From equation (2), the autocovariance of income growth at lag n is:

$$cov(\Delta y_i^h, \Delta y_i^{h+n}) = \sigma_\beta^2 - \rho^{n-1}\left(\frac{1-\rho}{1+\rho}\sigma_\eta^2\right), \tag{6}$$

for $n \geq 2$. Thus, covariances involve a positive constant term $\left(\sigma_\beta^2\right)$ 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 model 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 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 models is to test if higher order autocovariances are greater than zero.

The first row of Table 3 reports the results of this test conducted in Topel (1990) using PSID data covering $1968 - 83$. For completeness, the second row reports the same statistics using our longer panel. The same pattern can be seen in both samples. Starting from the second lag, there is no evidence of a positive covariance: they are all negative and statistically not different from zero, which appears to cast doubt on the HIP model.[13]

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 $n \geq 2$, their magnitude is very small (0.00038) making it difficult to distinguish it from a value of zero implied by the RIP model. Baker (1997) conducts a careful Monte Carlo analysis to examine the power

[13]The variance in Topel (1990) is about three times smaller than ours, probably because he looks at within-job wage changes. MaCurdy (1982) and Baker (1997) report variances closer to ours.

Table 3: THE COVARIANCE STRUCTURE OF INCOME GROWTH: A COMPARISON OF U.S. DATA AND THE HIP MODEL

Sample	Lag 0	1	2	3	4	5
	Autocovariance					
(1) Data (Topel)	.0476	−.0176	.00058	−.00166	−.00014	−.00067
	(.0019)	(.0014)	(.0008)	(.0007)	(.0008)	(.0007)
(2) Data (This paper)	.1215	−.0385	−.0031	−.0023	−.0025	−.00004
	(.0023)	(.0011)	(.0010)	(.0008)	(.0007)	(.0008)
(3) Model	.0840	−.0329	−.0014	−.0011	−.0007	−.0007
	(.0013)	(.0010)	(.0008)	(.0009)	(.0008)	(.0007)
	Autocorrelation					
(4) Data (Topel)	1.00	−.394	.013	−.039	−.003	−.016
(5) Data (This paper)	1.00	−.317	−.026	−.019	−.021	−.001
(6) Model	1.00	−.391	−.016	−.012	−.009	−.009
	(.000)	(.008)	(.009)	(.010)	(.008)	(.009)

Notes: Standard errors are in parenthesis. The statistics from Topel (1990) are from Table B1 in Appendix B, which are calculated from PSID 1968-83 with 8683 observations. The counterparts from simulated data are calculated using 12,000 observations.

of this test, and concludes that it does not provide strong evidence against HIP. Thus, we do not address this issue further in this paper.[14]

There is, however, a second concern about the use of this test, which is as critical as the first, and is independent of sample size. 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 all the studies mentioned above have focused on the first 5 to 10 lags. Figure 6 plots the autocovariances of income growth for the first 20 lags using our parameter estimates of the HIP model from Table 1. For the sample of all individuals (denoted "-o") autocovariances are negative *up to the 12th lag*, simply because σ^2_β 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 positive only at the 10^{th} and 15^{th} lags respectively. These calculations show that the findings of negative autocovariances reported in MaCurdy (1982), Topel (1990)

[14]Notice that even though McCurdy (1982)'s test cannot distinguish the autocovariances from zero, we are 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 we exploit the information in the entire variance covariance matrix which yields more precise information.

Figure 6: Evaluating MaCurdy's Test: The Covariance Structure of Income Growth

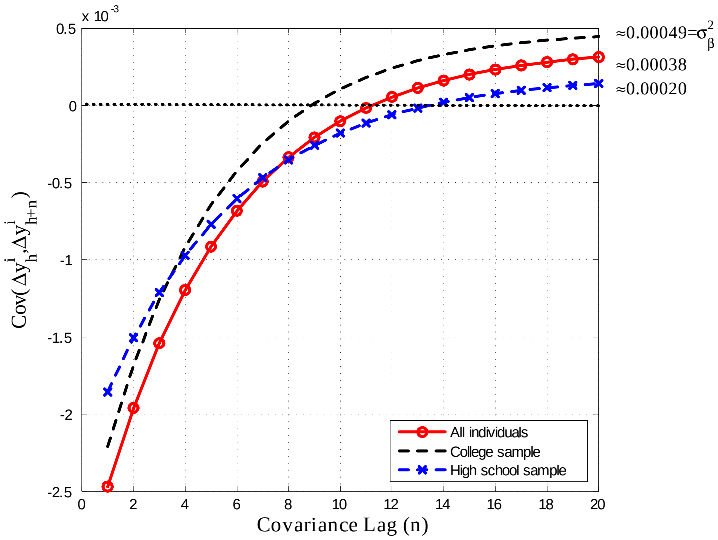

and Topel and Ward (1992) is exactly what is implied by the HIP model. Notice that this issue is separate from the power of the test, and suggests that even if an econometrician had access to a very large data set, the signs of these lower order covariances are not informative about HIP.

We conduct a Monte Carlo exercise to further explore this issue. We simulate income paths using the HIP model and the parameter values from the second row of Table 1.[15] The third row in Table 3 displays the averages of autocovariances over 500 replications along with the standard errors of the sampling distribution. The higher order covariances are negative, close to zero, and statistically insignificant after the second lag, similar to their empirical counterparts. Similarly, Panel B displays 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. Overall, these results suggest that tests based on the sign of covariances do not provide conclusive evidence on profile heterogeneity. The HIP model generates the same negative autocovariance structure that was previously used to reject it.

Finally, a similar concern applies to the variant of this test implemented by Abowd

[15]We first simulated income paths for 500,000 individuals. Then we drew 12,000 pairs of observations $(\Delta y_i^h, \Delta y_i^{h+k})$ without replacement for randomly selected initial age, h, and $k = 1, ..11$. The first ten autocovariances of income changes are then calculated using this sample and the exercise is repeated 500 times.

and Card (1989). 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 figure 6 the deviations of autocovariances from zero (between lags 2 and 10) are 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, 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.[16] In other words, the rejection of the null hypothesis would not necessarily support the HIP model either. Thus, 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.

5 Conclusion

In this paper we examined labor income data to distinguish between the RIP and HIP models. The existing evidence from labor income data has commonly been interpreted by macroeconomists as strongly in favor of the RIP model. Consequently, almost all life-cycle (or overlapping generations) models in the literature are calibrated using the RIP model as the income process.

We first argued that imposing a priori restrictions on income profiles as is done in the RIP model introduces an upward bias into the estimated persistence if the true data generating process features heterogeneous profiles. When we allow for HIP, the estimates we obtained indicate substantial heterogeneity in income profiles, and income shocks with modest persistence. Second, we also show that the HIP process we estimate generates small and negative autocovariances, quantitatively similar to their empirical counterpart, casing doubt on the previous interpretation of this finding as supporting the RIP model.

The HIP model 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 model which indi-

[16]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)

cates 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 model as the income process, which often implies puzzling differences in behavior by education level (see for example, Hubbard, et al., 1994; and Davis, Kubler and Willen, 2002).

It is important *not* to interpret these results—especially the low persistence of shocks in the HIP model—as suggesting that income uncertainty is not as large as that implied by the RIP model. The statistical analysis we conducted 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 (2005) we conduct such an analysis and argue that in fact a substantial part of this systematic variation is unknown to individuals early on, and is revealed very slowly, implying that the HIP model also features substantial income uncertainty. However, the nature of this risk and its distribution over the life-cycle is different than in the RIP model.

A Data Appendix

The data are drawn from the first 26 waves of the PSID. We include an individual into our sample if he satisfies the following criteria for a total of twenty years between 1968 and 1993. The individual

(i) is a male head of household

(ii) is between 22 and 62 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 we 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. Among these individuals 53 of them report a change in their education status during the sample period. We exclude these individuals from the two subsamples defined by education used in Section 3.1.

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. We 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 our 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 (Haider (2001); Storesletten et al. (2004)). Haider's estimates from the HIP specification are similar to ours (in particular, $\rho = 0.64$, and $\sigma^2_\beta = 0.00041$).

Table A1 reports some summary statistics for the primary sample.

B The Estimation Method

This appendix describes the minimum distance estimation (MDE) of the parameters of the income process given in equation (2). Let c_n be a typical element of the covariance matrix $\mathbf{C}$ of the income residuals, where $n = 1, .., N\ (= T(T+1)/2)$ enumerates unique elements of this matrix, and let $d_n(X_i, b)$ denote the corresponding model covariances given by equation (3), 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: $\mathbf{F}(b, X_i, \Upsilon_i) \equiv [F_1(b, X_i, \Upsilon_{i1}), ..., F_N(b, X_i, \Upsilon_{iN})]'$, where Υ_i is the indicator functions stacked into a vector conformably to $\mathbf{F}$. The moment conditions that we are estimating are of the form:

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

The MD estimator is the solution to

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

where $\widetilde{A}_N$ 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 we follow their recommendation. Notice however that because the panel is not balanced, each moment in the vector **F** is calculated using a different number of observations (determined by the non-zero elements of Υ_{in}). To adjust for this difference, we set $\widetilde{A}_N \equiv A_N A_N$, where A_N is a diagonal matrix with element (I/I_n) at the n^{th} 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 $\widehat{b}_N$ is consistent, asymptotically Normal with an asymptotic covariance matrix $\Sigma \equiv (D'D)^{-1} D'\Omega D (D'D)^{-1}$, 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 $E[\mathbf{F}(b, X_i, \Upsilon_i)\,\mathbf{F}(b, X_i, \Upsilon_i)\prime]$. Both expectations are replaced by sample averages when implemented.

References

[1] Abowd, John, and David Card (1989): "On the Covariance Structure of Earnings and Hours Changes," *Econometrica,* Vol. 57, pp.411-445.

[2] Altonji, Joseph G, and Lewis M Segal (1996): "Small Sample Bias in GMM Estimation of Covariance Structure," *Journal of Business and Economics Statistics,* Vol. 109, pp. 659-684.

[3] Baker, Michael (1997): "Growth-rate Heterogeneity and the Covariance Structure of Life-Cycle Earnings," *Journal of Labor Economics,* Vol. 15, pp. 338-375.

[4] Becker, Gary S. (1965): "Human Capital : A Theoretical and Empirical Analysis, with Special Reference to Education," reprinted 1994, University of Chicago Press.

[5] Ben-Porath, Yoram (1967): "The Production of Human Capital and the Life Cycle of Earnings," *Journal of Political Economy*,.Vol. 75.

[6] Bound, J. and A. B. Krueger, (1991), "The Extent of Measurement Error in Longitudinal Earnings Data: Do two wrongs make a right?", *Journal of Labor Economics* Vol. 9, pp. 1–24.

[7] Campbell John Y., Joao F. Cocco, Francisco J. Gomes, and Pascal J. Maenhout (2001): "Investing Retirement Wealth: A Life-Cycle Model," in *Risk Aspects of Investment-Based Social Security Reform,* John Campbell and Martin Feldstein eds., University of Chicago Press: Chicago, IL, pp. 439-473.

[8] Carroll, Christopher, and Andrew Samwick (1997): "The Nature of Precautionary Wealth," *Journal of Monetary Economics*, 40, No. 1.

[9] Chamberlain, Gary (1984): "Panel Data," in *Handbook of Econometrics*, eds. Z. Griliches and M. Intriligator. Amsterdam: North Holland, pp. 1274-1318.

[10] Christiano Lawrence, and Martin Eichenbaum (1990): "Unit Roots in Real GNP: Do We Know, and Do We Care?" Carnegie-Rochester Conference Series on Public Policy.

[11] Constantinides, George M., and Darrell Duffie (1996): "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy,* Vol. 104, pp. 219-240.

[12] Davis, Steve J., Felix Kubler and Paul Willen (2003): "Borrowing Costs and the Demand for Equity Over the Life-cycle," Working Paper, University of Chicago.

[13] Deaton, Angus, and Christina Paxson (1994): "Intertemporal Choice and Inequality," *Journal of Political Economy*, Vol. 102, pp. 437-467.

[14] Gourinchas Pierre-Olivier, and Jonathan A. Parker (2002): "Consumption Over the Life Cycle," *Econometrica,* Vol. 70, pp. 47-89.

[15] Guvenen, Fatih (2005): "Learning Your Earning: Are Labor Income Shocks Really Very Persistent?" Working Paper, University of Rochester.

[16] Haider Steve J. (2001): "Earnings Instability and Earnings Inequality of Males in the United Stated: 1967-1991, *Journal of Labor Economics,* Vol. 19, pp. 799-836.

[17] Hause, John C. (1980): "The Fine Structure of Earnings and the On-the-Job Training Hypothesis," *Econometrica,* Vol. 48, pp. 1013-29.

[18] Heathcote, Jonathan, Kjetil Storesletten, Gianluca Violante (2004): "The Macroeconomic Implications of Rising Wage Inequality in the United States," Working Paper, NYU.

[19] Hubbard, R Glenn, Jonathan Skinner and Steve P. Zeldes (1994): "The Importance of Precautionary Saving in Explaining Individual and Aggregate Savings," *Carnegie-Rochester Conference on Public Policy,* Vol. 40, pp. 59-105.

[20] Hubbard, R Glenn, Jonathan Skinner and Steve P. Zeldes (1995): "Precautionary Saving and Social Insurance," *Journal of Political Economy,* Vol. 103, pp. 360-399.

[21] Huggett, Mark (1996): "Wealth Distribution in Lifecycle Economies," Journal of Monetary Economics, Vol 38, pp. 469-94.

[22] Katz, Lawrence F., and David H. Autor (1999): "Changes in the Wage Structure and Earnings Inequality," in the *Handbook of Labor Economics,* Vol. 3, Amsterdam: North Holland.

[23] Katz, Lawrence F., and Kevin M. Murphy (1992): "Changes in Relative Wages, 1963-87: Supply and Demand Factors," *Quarterly Journal of Economics,* Vol 107, pp. 35-78.

[24] Lillard, Lee A., and Robert T Reville (1999): "Life-Cycle Human Capital Investment: New Evidence on an Old Issue," Working Paper, University of Michigan.

[25] Lillard, Lee A., and Yoram A Weiss (1979): "Components of Variation in Panel Earnings Data: American Scientists, 1960-70," *Econometrica*, Vol. 47, pp. 437-454.

[26] Lucas, Robert E., Jr. (2003): "Macroeconomic Priorities," *American Economic Review,*

[27] MaCurdy, Thomas (1982): "The Use of Time-Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis," *Journal of Econometrics,* Vol. 18, pp. 83-114.

[28] Meghir, Costas, and Luigi Pistaferri (2004): "Income Variance Dynamics and Heterogeneity," Econometrica, Vol. 72, pp. 1-32.

[29] Mincer, Jacob (1974): "Schooling, Experience and Earnings," New York: Columbia University Press.

[30] Moffitt, Robert A., and Peter Gottschalk (1995): "Trends in the Autocovariance Structure of Earnings in the US, 1969-1987, Working Paper, Brown University.

[31] Storesletten, Kjetil, Chris Telmer and Amir Yaron (2001): "The Welfare Costs of Business Cycles Revisited: Finite Lives and Cyclical Variation in Idiosyncratic Risk," *European Economic Review,* Vol. 45, pp. 1311-39.

[32] Storesletten, Kjetil, Chris Telmer and Amir Yaron (2004): "Cyclical Dynamics in Idiosyncratic Labor-Market Risk," *Journal of Political Economy*, Vol. 112, pp. 695-717.

[33] Topel, Robert H. (1990): "Specific Capital, Mobility and Wages: Wages rise with Job Seniority," NBER Working Paper No. 3294.

[34] Topel, Robert H., and Michael P. Ward (1992): "Job Mobility and the Careers of Young Men," *Quarterly Journal of Economics,* Vol. 107, pp. 439-79.