

**Deutsche Bundesbank's 9th Spring Conference:
"Microdata Analysis and Macroeconomic Implications"**
Eltville, 27-28 April 2007

Anthony A Smith

University of Yale

Fatih Guvenen

University of Texas at Austin

**„Inferring Labour Income Risk From Economic Choices: An
Indirect Inference Approach “**

Inferring Labor Income Risk From Economic Choices: An Indirect Inference Approach

Fatih Guvenen Anthony A. Smith

VERY PRELIMINARY AND INCOMPLETE

April 25, 2007

Abstract

This paper sheds light on the nature of labor income risk by exploiting the information contained in the joint dynamics of households' labor earnings and consumption-choice decisions. In particular, this paper attempts to discriminate between two leading views on the nature of labor income risk: the “restricted income profiles” (RIP) model—in which individuals are subjected to large and persistent income shocks but face similar life-cycle income profiles—and the “heterogeneous income profiles” (HIP) model—in which individuals are subjected to income shocks with modest persistence but face individual-specific income profiles. Although these two different income processes have vastly different implications for economic behavior, earlier studies have found that labor income data alone is insufficient to distinguish between them. This paper, therefore, brings to bear the information embedded in consumption data. Specifically, we use the powerful new tools of indirect inference to estimate, using rich panel data on consumption and labor earnings, a structural consumption-savings model with borrowing constraints and risk-averse households. The method we develop is in fact more general than this particular application and allows the estimation of income processes from economic decisions in the presence of non-separability between consumption and leisure, partial insurance of income shocks, missing observations, among others. In this estimation, we use an auxiliary model—which forms the bridge between the data and the consumption-savings model—that provides a sharp distinction between the RIP and HIP models. Finally, we conduct formal statistical tests to assess the extent to which the RIP and HIP models find support in the data.

1 Introduction [Incomplete]

The goal of this paper is to elicit information about the nature of labor income risk from individuals' economic decisions (such as consumption-savings choice), which contain valuable information about the environment faced by individuals, including the future (income) risks they perceive.

To provide a framework for this discussion, consider the following process for log labor income of individual i with t years of labor market experience:

$$y_t^i = \underbrace{[a_0 + a_1t + a_2t^2 + a_3Educ + \dots]}_{\text{common life-cycle component}} + \underbrace{[\alpha^i + \beta^i t]}_{\text{profile heterogeneity}} + \underbrace{[z_t^i + \varepsilon_t^i]}_{\text{stochastic component}} \quad (1)$$

$$\text{where} \quad z_t^i = \rho z_{t-1}^i + \eta_t^i, \quad \text{and} \quad \eta_t^i, \varepsilon_t^i \sim iid$$

The terms in the first bracket capture the life-cycle variation in labor income that is *common* to all individuals with given observable characteristics. The second component captures potential *individual-specific* differences in income growth rates (as well as in levels, which is less important). Such differences would be implied for example by a human capital model with heterogeneity in learning ability.¹ Finally, the terms in the last bracket represent the stochastic variation in income, which is written here as the sum of an AR(1) component and a purely transitory shock. This is a specification commonly used in the literature.

A vast empirical literature has estimated various versions of (1) in an attempt to answer the following two questions:

1. Do individuals differ *systematically* in their income growth rates? If such differences exist, are they quantitatively important? i.e., is $\sigma_\beta^2 \gg 0$?
2. How large and how persistent are income shocks? i.e., what is σ_η^2 and ρ ?

Existing studies in the literature can be broadly categorized into two groups based on the conclusions they reach regarding these questions. The first group of papers *impose* $\sigma_\beta^2 \equiv 0$

¹See for example, the classic paper by Ben-Porath (1967). For a more recent example of such a human capital model, see Guvenen and Kuruscu (2006).

based on outside evidence,² and with this restriction estimate ρ to be close to 1. We refer to this version of the process in (1) as the “*Restricted* Income Profiles” (*RIP*) model. The second group of papers do not impose any restrictions on (1) and find that ρ is significantly less than 1 and σ_β^2 is large. We refer to this version of (1) as the “*Heterogeneous* Income Profiles” (*HIP*) model. In other words, according to the *RIP* view, most of the rise in within-cohort income inequality over the life-cycle is due to large and persistent shocks, whereas in the *HIP* view, it is due to systematic differences in income growth rates. While overall we interpret the results of these studies, and especially those of the more recent papers, as more supportive of the *HIP* model, it is fair to say that this literature has not produced an unequivocal verdict.³

A key point to observe is that these existing studies do not utilize the information revealed by individuals’ consumption-savings choice to distinguish between the *HIP* and *RIP* models.⁴ But endogenous choices, such as consumption and savings, contain valuable information about the environment faced by individuals, including the future risks they perceive. Therefore, the main purpose of this paper is to use the restrictions imposed by the *RIP* and *HIP* processes on consumption data—in the context of a life-cycle model—to bring more evidence to bear on this important question. We elaborate further below on the advantages of focusing on consumption-savings choice (instead of using labor income data in isolation or using other endogenous choices, such as labor supply) for drawing inference about the labor income process.

In a sense, the two questions discussed so far only scratch the surface of “the nature of income risk.” This is because those two questions are statistical in nature, i.e., they relate to how the income process is viewed by the *econometrician* who studies past observations

²The outside evidence refers to a test proposed by MaCurdy (1982) in which he failed to reject the null of *RIP* against the alternative of *HIP*. Two recent papers, Baker (1997) and Guvenen (2005), argue that MaCurdy’s test lacks power, and therefore, the lack of rejection of the *RIP* null does not provide evidence against the *HIP* model.

³A short list of these studies includes MaCurdy (1982), Abowd and Card (1989), and Topel (1990), which find support for the *RIP* model; Lillard and Weiss (1979), Hause (1980), and especially the more recent studies such as Baker (1997), Haider (2001), and Guvenen (2005a) which find support for the *HIP* model.

⁴Two recent papers do use consumption data but in a more limited fashion than this paper intends to do. In a recent paper, Huggett, Ventura and Yaron (2006) study a version of the Ben-Porath model and make some use of consumption data to measure the relative importance of persistent income shocks versus heterogeneity in learning ability. Although the income process generated by their model does not exactly fit into the specification in equation (1) their results are informative. Second, Guvenen (2007) uses consumption data to investigate if a *HIP* model estimated from income data is consistent with some stylized consumption facts. While both of these papers are informative about the *HIP* versus *RIP* debate, they make limited use of consumption data, especially of the *dynamics* of consumption behavior.

on individual income. But it is quite plausible that individuals may have more, or less, information about their income process than the econometrician at different points in their lifecycle, which raises two more questions:

3. If individuals indeed differ in their income growth rate as suggested by the HIP model, how much do individuals *know* about their β^i at *different points* in their life-cycle? In other words, what fraction of the heterogeneity in β^i constitutes “uncertainty” on the part of individuals as opposed to simply being some “known heterogeneity”?
4. What fraction of income movements measured by z_t^i and ε_t^i are really “unexpected shocks” as opposed to being “anticipated changes”?

These questions are inherently different than the first two in that they pertain to how *individuals* perceive their income process. As such, they cannot be answered using income data alone, but the answers can be teased out, again, from individuals’ economic decisions. To give one example (to question 4), consider a married couple who jointly decide that they will both work up to a certain age and then will have children at which time one of the spouses will quit his/her job to take care of the children. The ensuing large fall in household income will appear as a large permanent shock to the econometrician using labor income data alone, but consumption (and savings) data would reveal that this change has been anticipated.

Several papers have used consumption data and shed light on various properties of income processes (among others, Hall and Mishkin (1982), Deaton and Paxson (1994), Blundell and Preston (1998), and Blundell, Pistaferri and Preston (2006a)). This paper aims to contribute to this literature in the following ways. First and foremost, existing studies consider only versions of the RIP model (i.e., they set $\sigma_\beta^2 \equiv 0$ at the outset), whereas our goal is to distinguish between HIP and RIP models as well. Second, and furthermore, these studies also *impose* $\rho \equiv 1$, and only estimate the innovation variances. In other words, there is no existing study to our knowledge that uses consumption data and estimates ρ . Therefore, this paper will leave ρ unrestricted (even in the RIP version) and exploit consumption and income data jointly to pin down its value. Since many incomplete markets models are still calibrated using versions of the RIP process, the results of this exercise should be useful for calibrating those models. The third contribution of this paper will be in the method used for estimation—indirect inference—which is much less restrictive than, and has several important advantages over, the GMM approach used in previous work.

1.1 Why Look at Consumption-Savings Choice?

Even if one is only interested in the first two questions raised above, using information revealed by intertemporal choices has important advantages. This is because one difficulty of using income data alone is that identification between HIP and RIP models partly depends on the behavior of the *higher-order* autocovariances of income.

To see this clearly, consider the case where the panel data set contains income observations on a single cohort over time. In this case, the second moments of the cross-sectional distribution for this cohort are given by:

$$\begin{aligned} \text{var}(y_t^i) &= [\sigma_\alpha^2 + 2\sigma_{\alpha\beta}t + \sigma_\beta^2 t^2] + \text{var}(z_t^i) + \sigma_\varepsilon^2 \\ \text{cov}(y_t^i, y_{t+n}^i) &= [\sigma_\alpha^2 + \sigma_{\alpha\beta}(2t+n) + \sigma_\beta^2 t(t+n)] + \rho^n \text{var}(z_t^i), \end{aligned} \quad (2)$$

where $t = 1, \dots, T$, and $n = 1, \dots, T - t$. There are two sources of identification between the RIP and HIP processes, 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 (i.e., the diagonal elements of the variance-covariance matrix), which is shown on the first line of (2). 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 t^2 is σ_β^2), whereas a non-convex shape would be captured by the presence of AR(1) shocks.

The second source of identification is provided by the autocovariances displayed in the second line. The covariance between ages t and $t + n$ 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 t and n , which allows covariances to be decreasing, increasing or non-monotonic in n at each t . The second term captures the effect of the AR(1) shock, and notice that for a given t , 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 age t , 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 .

Notice that for a cohort with 40 years of working life, there are only 40 variance terms, but many more—780 ($= (40 \times 41) / 2 - 40$) to be precise—autocovariances, which provide

crucial information for distinguishing between HIP and RIP processes. The main difficulty is that because of sample attrition, fewer and fewer individuals contribute to these higher autocovariances, raising important concerns about potential selectivity bias. To give a rough idea, if one uses labor income data from the Panel Study of Income Dynamics (PSID), and selects all individuals who are observed in the sample for 3 years or more (which is a typical sample selection criterion), the number of individuals contributing to the 20th autocovariance will be about 1/5 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 models.

In contrast, because of its forward-looking nature, even short-run movements in consumption, and the immediate response of consumption to income innovations contain information about the perceived *long-run* behavior of the income process. Therefore even lower-order covariances of consumption would help in distinguishing HIP from RIP. (Notice that the dynamic aspect of the consumption-savings choice also distinguishes it from other decisions, such as labor supply, which are static in nature, unless one models intertemporally non-separable preferences in leisure.)

2 Bayesian Learning about Income Profiles

Embedding the HIP process into a life-cycle model requires one take a stand on what individuals know about their own β^i . We follow Guvenen (2007) and assume that individuals enter the labor market with some prior belief about their β^i and then update their beliefs over time in a Bayesian fashion. Notice that the prior variance of this belief (denote by $\hat{\sigma}_{\beta|0}^2$) measures how uncertain individuals are about their own β^i at time zero, addressing question 3 above.

We now cast the learning process as a Kalman filtering problem which allows us to obtain recursive updating formulas for beliefs. Individuals (know α^i), observe y_t^i , and must learn about $\mathbf{S}_t^i \equiv (\beta^i, z_t^i)$.⁵ It is convenient to express the learning process as a Kalman filtering problem using the state-space representation. In this framework, the “state equation” describes the evolution of the vector of state variables that is unobserved by the decision

⁵Guvenen (2007) also allows for learning about α^i and shows that it has a minimal effect on the behavior of the model. Therefore, we abstract from this feature which eliminates one state variable and simplifies the problem.

maker:

$$\underbrace{\begin{bmatrix} \beta^i \\ z_{t+1}^i \end{bmatrix}}_{\mathbf{S}_{t+1}^i} = \underbrace{\begin{bmatrix} 1 & 0 \\ 0 & \rho \end{bmatrix}}_{\mathbf{F}} \underbrace{\begin{bmatrix} \beta^i \\ z_t^i \end{bmatrix}}_{\mathbf{S}_t^i} + \underbrace{\begin{bmatrix} 0 \\ \eta_{t+1}^i \end{bmatrix}}_{\boldsymbol{\nu}_{t+1}^i}$$

Even though the parameters of the income profile have no dynamics, including them into the state vector yields recursive updating formulas for beliefs using the Kalman filter. A second (observation) equation expresses the observable variable(s) in the model—in this case, log income—as a linear function of the underlying hidden state and a transitory shock:

$$y_t^i = \alpha^i + \begin{bmatrix} t & 1 \end{bmatrix} \begin{bmatrix} \beta^i \\ z_t^i \end{bmatrix} + \varepsilon_t^i = \alpha^i + \mathbf{H}'_t \mathbf{S}_t^i + \varepsilon_t^i$$

We assume that both shocks have *i.i.d* Normal distributions and are independent of each other, with \mathbf{Q} and R denoting the covariance matrix of $\boldsymbol{\nu}_t^i$ and the variance of ε_t^i respectively. To capture an individual's initial uncertainty, we model his prior belief over (β^i, z_1^i) by a multivariate Normal distribution with mean

$$\widehat{\mathbf{S}}_{1|0} \equiv (\widehat{\beta}_{1|0}^i, \widehat{z}_{1|0}^i)$$

and variance-covariance matrix:

$$\mathbf{P}_{1|0} = \begin{bmatrix} \sigma_{\beta,0}^2 & 0 \\ 0 & \sigma_{z,0}^2 \end{bmatrix}$$

where we use the short-hand notation $\sigma_{\cdot,t}^2$ to denote $\sigma_{\cdot,t+1|t}^2$. After observing $(y_t^i, y_{t-1}^i, \dots, y_1^i)$, the posterior belief about \mathbf{S}^i is Normally distributed with a mean vector $\widehat{\mathbf{S}}_t^i$, and covariance matrix \mathbf{P}_t . Similarly, let $\widehat{\mathbf{S}}_{t+1|t}^i$ and $\mathbf{P}_{t+1|t}$ denote the one-period-ahead *forecasts* of these two variables respectively. These two variables play central roles in the rest of our analysis. Finally, log income has a Normal distribution conditional on an individual's beliefs:

$$y_{t+1}^i | \widehat{\mathbf{S}}_t^i \sim N \left(\mathbf{H}'_{t+1} \widehat{\mathbf{S}}_{t+1|t}^i, \mathbf{H}'_{t+1} \mathbf{P}_{t+1|t} \mathbf{H}_{t+1} + R \right). \quad (3)$$

In this particular problem, the standard Kalman filtering equations can be manipulated to obtain some simple expressions that will become useful later. To this end, define:

$$\begin{aligned} A_t &\equiv t\sigma_{\beta,t|t-1}^2 + \sigma_{\beta z,t|t-1}, \\ B_t &\equiv t\sigma_{\beta z,t|t-1} + \sigma_{z,t|t-1}^2, \\ X_t &\equiv \text{var}_{t-1}(y_t^i) = A_t t + B_t + R \end{aligned}$$

Using the Kalman recursion formulas:

$$\begin{bmatrix} \widehat{\beta}_{t+1|t}^i \\ \widehat{z}_{t+1|t}^i \end{bmatrix} = \begin{bmatrix} \widehat{\beta}_{t|t-1}^i \\ \rho \widehat{z}_{t|t-1}^i \end{bmatrix} + \begin{bmatrix} A_t/X_t \\ \rho B_t/X_t \end{bmatrix} \left(y_t^i - \left(\widehat{\beta}_{t|t-1}^i t + \widehat{z}_{t|t-1}^i \right) \right)$$

Define the innovation to beliefs:

$$\widehat{\xi}_t = y_t^i - \left(\widehat{\beta}_{t|t-1}^i t + \widehat{z}_{t|t-1}^i \right)$$

Then we can rewrite:

$$\widehat{\beta}_{t+1|t}^i - \widehat{\beta}_{t|t-1}^i = (A_t/X_t) \widehat{\xi}_t \quad (4)$$

$$\widehat{z}_{t+1|t}^i - \rho \widehat{z}_{t|t-1}^i = (\rho B_t/X_t) \widehat{\xi}_t \quad (5)$$

An important point to note is that $\widehat{\xi}_t$ and (the true innovation to income) η_t^i do not need to have the same sign, a point that will play a crucial role below. Finally, the posterior variances evolve:

$$\sigma_{\beta,t+1|t}^2 = \sigma_{\beta,t|t-1}^2 - \frac{A_t^2}{X_t} \quad (6)$$

$$\sigma_{z,t+1|t}^2 = \rho^2 \left[\sigma_{z,t|t-1}^2 - \frac{B_t^2}{X_t} \right] + R \quad (7)$$

For a range of parameterizations A/X has an inverse U-shape over the life-cycle. Therefore, beliefs about β^i changes (and precision rises) slowly early on but become faster over time. In contrast, B/X declines monotonically. As shown in Guvenen (2007), optimal learning in this model has some interesting features. In particular, learning is very slow and the speed of learning has a non-monotonic pattern over the life-cycle (which is due to the fact that A/X has an inverse U-shape). If instead the prior uncertainty were to resolve quickly, consumption behavior after the first few years would not be informative about the prior uncertainty faced by individuals ($\widehat{\sigma}_{\beta|0}^2$).

Finally we discuss how an individual's prior belief about β^i is determined. Suppose that the distribution of income growth rates in the population is generated as $\beta^i = \beta_k^i + \beta_u^i$, where β_k^i and β_u^i are two random variables, independent of each other, with zero mean and variances of $\sigma_{\beta_k}^2$ and $\sigma_{\beta_u}^2$. Clearly then, $\sigma_{\beta}^2 = \sigma_{\beta_k}^2 + \sigma_{\beta_u}^2$. The key assumption we make is that individual i observes the realization of β_k^i , but not of β_u^i (hence the subscripts indicate *known* and *unknown*, respectively). Under this assumption, the prior mean of individual i is $\widehat{\beta}_{1|0}^i = \beta_k^i$, and the prior variance is $\sigma_{\beta,0}^2 = \sigma_{\beta_u}^2 = (1 - \lambda) \sigma_{\beta}^2$, where we define $\lambda = 1 - \sigma_{\beta_u}^2 / \sigma_{\beta}^2$,

as the fraction of variance known by individuals. Two polar cases deserve special attention. If $\lambda = 0$, individuals do not have any *private prior information* about their income growth rate (i.e., $\sigma_{\beta,0}^2 = \sigma_{\beta}^2$ and $\widehat{\beta}_{1|0}^i = \bar{\beta}$ for all i , where $\bar{\beta}$ is the population average). On the other hand if $\lambda = 1$, each individual observes β^i completely and faces no prior uncertainty about its value.

2.1 The HIP Model

Consider an environment where each individual lives for T years and works for the first R ($< T$) years of his life, after which he retires. Individuals do not derive utility from leisure and hence supply labor inelastically.⁶ During the working life, the income process is given by the HIP process specified in equation (1). During retirement, the individual receives a pension which is given by a fixed fraction of the individual's income in period R .⁷ There is a risk-free bond that sells at price P^b (with a corresponding net interest rate $r^f \equiv 1/P^b - 1$). Individuals can also borrow at the same interest rate up to an age-specific borrowing constraint \underline{W}_{t+1} , specified below.

The relevant state variables for this dynamic problem are the asset level, ω_t^i , and his current forecast of the true state in the current period, $\widehat{\mathbf{S}}_t$. The dynamic programming problem of the individual can be written as:

$$V_t^i(\omega_t^i, \widehat{\mathbf{S}}_t^i) = \max_{C_t^i, \omega_{t+1}^i} \left\{ U(C_t^i) + \delta E_t \left[V_{t+1}^i(\omega_{t+1}^i, \widehat{\mathbf{S}}_{t+1}^i) \right] \right\}$$

s.t.

$$C_t^i + a_{t+1}^i = \omega_t \tag{8}$$

$$\omega_t = (1 + r) a_t^i + Y_t^i \tag{9}$$

$$a_{t+1}^i \geq \underline{W}_{t+1}, \quad \text{and}$$

and Kalman recursions

for $t = 1, \dots, R - 1$, where $Y_t^i \equiv e^{y_t^i}$ is the *level* of income, and V_t^i is the value function of a t year-old individual. The evolutions of the vector of beliefs and its covariance matrix are governed by the Kalman recursions given in equations (4, 5, 6, 7). Finally, the expectation is

⁶The labor supply choices of both the husband and wife appear to be important for drawing robust inference about the nature of income risk. Therefore, we intend to introduce labor supply choice for both spouses in future versions of this paper. Such extensions are conceptually feasible with indirect inference, although it increases computational costs.

⁷A more realistic Social Security system will be introduced in a later version of the paper.

taken with respect to the conditional distribution of y_{t+1}^i given by equation (3), since this is the source of all uncertainty in the model.

During retirement, pension income is constant and since there is no other source of uncertainty or learning, the problem simplifies significantly:

$$\begin{aligned} V_t^i(\omega_t^i, Y^i) &= \max_{c_t^i, \omega_{t+1}^i} [U(C_t^i) + \delta V_{t+1}^i(\omega_{t+1}^i, Y^i)] \\ \text{s.t.} \quad Y^i &= \Phi(Y_R^i), \text{ and eq. (8, 9)} \end{aligned} \tag{10}$$

for $t = R, \dots, T$, and $V_{T+1} \equiv 0$.

2.2 The RIP Model

The second model is essentially the same as the first one, with the exception that the income process is now given by a RIP process. Because with a RIP process all individuals share the same life-cycle income profile (α, β) , there is no learning about individual profiles, the problem simplifies significantly. Specifically, the dynamic programming problem of a typical worker is:

$$\begin{aligned} J_t^i(\omega_t^i, z_t^i) &= \max_{c_t^i, \omega_{t+1}^i} \{U(c_t^i) + \delta E [J_{t+1}^i(\omega_{t+1}^i, z_{t+1}^i) | z_t^i]\} \\ \text{s.t.} \quad &\text{equations (8, 9)} \end{aligned}$$

for $t = 1, \dots, R - 1$, where J_t^i is the value function of a t year-old individual. Notice that we assume the worker observes the persistent component of the income process, z_t^i , separately from y_t^i . This is the standard assumption in the existing consumption literature which uses the RIP process, and we follow them for comparability. Finally, because there is no income risk after retirement, the problem of a retiree is the same as in (10) above.

Notice that the HIP model does *not* nest the RIP model described here, although it comes quite close. In particular, when $\sigma_\beta^2 \equiv 0$ the HIP process does reduce to the RIP process, but now in the consumption-savings model individuals are assumed not to observe the AR(1) shock and the i.i.d shock separately (whereas in the RIP model described here, they do). We choose the RIP model not nested in the HIP model because it corresponds more closely to the framework studied in the consumption literature.

2.3 Modeling Partial Insurance

In this section we describe how we introduce partial insurance into the framework described above. Making explicit assumptions on the nature of partial insurance allows us to distinguish between the persistence of income shocks and the consumption smoothing due to partial insurance available to individuals over and above the self-insurance inherent in the life-cycle permanent income model.

[To be added]

2.4 Introducing Endogenous Labor Supply Choice

[To be added]

3 An Indirect Inference Approach

Indirect inference is a simulation-based method for estimating, or making inferences about, the parameters of economic models. It is most useful in estimating models for which the likelihood function (or any other criterion function that might form the basis of estimation) is analytically intractable or too difficult to evaluate, as is the case here: neither one of the consumption-savings models described above yields simple estimable equations that would allow a maximum likelihood or GMM estimation. Previous studies (which focused only on the RIP model) made a number of simplifying assumptions, such as the absence of binding borrowing constraints, separability between consumption and leisure in the utility function, a simplified retirement structure, and so on, and employed several approximations to the true structural equations in order to make GMM feasible.

Instead, the hallmark of indirect inference is the use of an “auxiliary model” to capture aspects of the data upon which to base the estimation. One key advantage of indirect inference over GMM is that this auxiliary model does *not* need to correspond to any valid moment condition of the structural model for the estimates of the structural parameters to be consistent. This allows significant flexibility in choosing an auxiliary model: it can be any sufficiently rich statistical model relating the model variables to each other as long as each structural parameter of the economic model has an independent effect on the likelihood of the auxiliary model.⁸ This also allows one to incorporate many realistic features into the

⁸In addition to some regularity conditions that the auxiliary model has to satisfy the precise specification

structural model without having to worry about whether or not one can directly derive the likelihood (or moment conditions for GMM) in the presence of these features.

While indirect inference shares a basic similarity to MSM (Method of Simulated Moments), it differs from MSM in its use of an auxiliary model to form “moment conditions”. In particular, indirect inference allows one to think in terms of structural and dynamic relationships of economic models that are difficult to express as simple unconditional moments as is often done with MSM. We illustrate this in the description of the auxiliary model below.

3.1 Towards an Auxiliary Model

To understand the auxiliary model that will be used, it is useful to elaborate on the dependence of consumption choice on income shocks. As noted above, the key idea behind an auxiliary model is that it should be an econometric model that is easy to estimate, yet one that captures the key statistical relations between the variables of interest in the model. Good candidates for an auxiliary model are provided by structural relationships that hold in models that are similar to the HIP and RIP models described above, and yet simple enough to allow the derivation of such relationships.

To this end, consider a simplified version of the HIP model, where we assume: (i) quadratic utility; (ii) $\delta(1+r^f) = 1$, and (iii) no retirement. Further consider a simpler form of the HIP process:

$$Y_t^i = \alpha^i + \beta^i t + z_t^i, \quad (11)$$

where income, instead of its logarithm, is linear in the underlying components, and we set $\varepsilon_t^i \equiv 0$. Under these assumptions, optimal consumption choice implies

$$\Delta C_t = \frac{1}{\varphi_t} \left[(1 - \gamma) \sum_{s=0}^{T-t} \gamma^s (E_t - E_{t-1}) Y_{t+s}^i \right], \quad (12)$$

where $\gamma = 1/(1+r^f)$ and $\varphi_t = (1 - \gamma^{T-t+1})$ is the annuitization factor. Substituting the simple HIP process in (11), we have:

$$\begin{aligned} E_t(Y_{t+s}^i) &= \alpha^i + \widehat{\beta}_t^i(t+s) + \rho^s \widehat{z}_t^i \\ (E_t - E_{t-1})Y_{t+s}^i &= (\widehat{\beta}_t^i - \widehat{\beta}_{t-1}^i)(t+s) + \rho^s \widehat{\eta}_t^i \end{aligned}$$

of the auxiliary model will also matter for the efficiency of the estimator.

Substituting this into (12), one can show:

$$\Delta C_t = \Phi_{t,T}^r \left(\widehat{\beta}_t^i - \widehat{\beta}_{t-1}^i \right) + \Psi_{T-t}^{r,\rho} \widehat{\eta}_t^i \quad (13)$$

where:

$$\begin{aligned} \Phi_{t,T}^r &\equiv \left[\left(\frac{\gamma}{1-\gamma} \right) + \frac{t - (T+1)\gamma^{T-t+1}}{1-\gamma^{T-t+1}} \right] \\ \Psi_{T-t}^{r,\rho} &\equiv \frac{1-\gamma}{1-\gamma\rho} \left[\frac{\left(1 - (\gamma\rho)^{T-t+1} \right)}{\left(1 - \gamma^{T-t+1} \right)} \right] \end{aligned}$$

Note that $\Phi_{t,T}^r$ is a (known) slightly convex increasing function of t , and $\Psi_{T-t}^{r,\rho}$ is constant and equal to 1 when $\rho = 1$. Recall that the Kalman filtering formulas above implied:

$$\begin{aligned} \widehat{\beta}_t^i - \widehat{\beta}_{t-1}^i &= (A_t/X_t) \widehat{\xi}_t \\ \widehat{z}_t^i - \rho \widehat{z}_{t-1}^i &= (B_t/X_t) \widehat{\xi}_t \end{aligned} \quad (14)$$

which is obtained easily from equations (4), but now $\widehat{\xi}_t$ has to be reinterpreted as the level deviation: $Y_t^i - \left(\widehat{\beta}_{t|t-1}^i t + \widehat{z}_{t|t-1}^i \right)$. Plugging this, we get in the HIP model:

$$\Delta C_t = \left[\Phi_{t,T}^r (A_t/X_t) + \Psi_{T-t}^{r,\rho} (\rho^s B_t/X_t) \right] \times \widehat{\xi}_t \quad (15)$$

Instead in the RIP model we have:

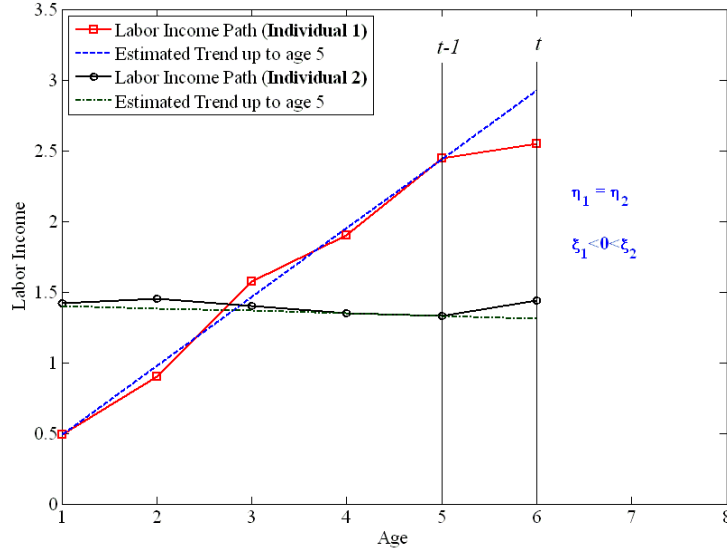
$$\Delta C_t = \Psi_{T-t}^{r,\rho} \times \eta_t^i \quad (16)$$

The last two equations underscore the key difference between the two frameworks: in the RIP model only current η_t^i matters for consumption response, whereas in the HIP model the entire history of shocks matters.⁹ As a result, two individuals hit by the same η_t^i may react differently depending on their history. Specifically, in the HIP model η_t^i and $\widehat{\xi}_t$ may have different signs. Therefore, an increase in income ($\Delta Y_t^i > 0$) may cause a fall in consumption ($\Delta C_t^i < 0$). In the RIP model, this will never happen.

An example of this case is shown in figure 1. This graph plots the income paths of two individuals, where we continue to assume $\varepsilon_t^i \equiv 0$ for simplicity. Individual 1 experiences a

⁹It is true that if individuals could not separately observe z_t and ε_t in the RIP model but were solving a signal extraction problem instead, the history of shocks would also matter in the RIP model. However, the specific predictions implied by the HIP model described below would still not hold in such a model.

Figure 1: Distinguishing HIP from RIP (from Consumption Changes)




faster average income growth rate in the first five periods than individual 2, but observes the same rise in income between periods five and six. If these income paths are generated by a RIP process (and individuals correctly perceives them as such), then both individuals will adjust their consumption growth by exactly the same amount between periods five and six. Instead, if the truth is as in the HIP model, individual 1 will have formed a belief that his income growth rate is higher than that of individual 2, and was expecting his income to be closer to the trend line (shown by the dashed blue line). Therefore, even though his income increases, it is significantly below the trend ($\hat{\xi}_t < 0$), which causes him to revise down his beliefs about his true β^i , and consequently his consumption level from equation (15). Specifically, we have:

Prediction 1: The HIP model with Bayesian learning predicts that controlling for current income growth, consumption growth will be a *decreasing* function of average past income *growth rate*. Instead, the RIP model predicts no dependence on past income growth rate of this kind.

It is also possible to obtain a closed-form expression for the consumption (level) in the simplified version of the HIP model described above (Here we simply give an intuitive description of the information contained in the level of consumption, rather than going through the derivation). One can easily see that the *level* of consumption contains information about

Figure 2: Distinguishing HIP from RIP (from Consumption Levels)




whether individuals perceive their income process as HIP or RIP. An example of this is shown in figure 2. This example is most easily explained when income shocks are permanent ($\rho = 1$), which we assume for the moment. As before, individuals realize different income growth rates up to period 3. Under the RIP model, both individuals' forecast of their future income is the same as their current income (shown with the horizontal dashed lines). In contrast with a HIP process, individual 1 will expect a higher income growth rate and therefore a much higher lifetime income than individual 2. Therefore, the first individual will have a higher consumption level than individual 2 at the same age, despite the fact that their current income levels are very similar. Therefore, we have:

Prediction 2: The HIP model predicts that controlling for the current level of income *and* past average income level, an individual's current consumption level will be an increasing function of his past income growth *rate*.

Finally, it is also easy to see that the level of consumption is also informative about how much prior information individuals have about their own β^i within the HIP framework (question 3 raised in the introduction). To see this, consider the next figure (3) which is a slight variation of the previous one. Here both individuals are assumed to have observed the same path of income growth up to period 3 even though their true β^i are different. (This is possible since there are many stochastic shocks to the income process over time (coming

Figure 3: Determining the Amount of Prior Knowledge in HIP



from η_t), and the contribution of β^i to income is quite small). In this case, under the HIP model, if individuals have no private prior information about their own true β^i (which will be the case when $\lambda = 0$) then both individuals should have the same consumption level. The more prior information each individual has about his true β^i the higher will be the consumption of the first individual compared to the second. Therefore, an auxiliary model can capture this relationship by focusing on the following dynamic relationship:

Prediction 3: if $\lambda > 0$, then controlling for *past* income growth (as well as the current income level and past average consumption level) the consumption level of an individual will be increasing in his *future* income growth as well. This is because in this case the individual has more information about his true β^i than is known to the econometrician and what is revealed by his past income growth.

These three examples illustrate how one can use the structural relationships that hold true exactly in a somewhat simplified version of the economic model of interest in order to come up with an auxiliary model. Indirect inference allows one to think in terms of these rich dynamic relationships instead of a set of moments (covariances, etc.). Below we are going to write a parsimonious auxiliary model that will capture these dynamic relationships to identify HIP from RIP and will also determine the degree of prior information (or equivalently, uncertainty) individuals face upon entering the labor market in the case of the HIP process.

3.1.1 A Parsimonious and Feasible Auxiliary Model

As shown above, the HIP model implies:

$$\Delta C_t = \Pi_{t,T}^{r,\rho} \left(Y_t^i - \left(\widehat{\beta}_{t|t-1}^i t + \widehat{z}_{t|t-1}^i \right) \right), \quad (17)$$

where $\Pi_{t,T}^{r,\rho} \equiv \Phi_{t,T}^r(A_t/X_t) + \Psi_{T-t}^{r,\rho}(B_t/X_t)$. However, since $\widehat{\beta}_{t|t-1}^i$ and $\widehat{z}_{t|t-1}^i$ are unobserved by the econometrician, this regression is not feasible as an auxiliary model. Moreover, we derived this relationship assuming a simplified HIP income process, quadratic utility, no borrowing constraints, and no retirement period, none of which is true in the life-cycle model we would like to estimate. Fortunately, as mentioned earlier, none of these issues represent a problem for the consistency of the estimates of the structural parameters that we are interested in.¹⁰

We approximate the relationship in (17) with the following regression:

$$c_t = a_{10} + a_{11}y_{t-1} + a_{12}c_{t-1} + a_{13}\bar{y}_{1,t-2} + a_{14}\bar{y}_{t+1,T} + a_{15}\overline{\Delta y}_{1,t-1} + a_{16}\overline{\Delta y}_{t+1,T-t} + a_{17}t + \epsilon_t \quad (18)$$

where c_t is the logarithm of consumption; y denotes the logarithm of labor income; $\bar{y}_{a,b}$ denotes the average of log income from time a to b ; and similarly $\overline{\Delta y}_{a,b}$ denotes the average growth rate of log income from time a to b . Notice that we use the logarithm of variables rather than the level; since the utility function is CRRA and income is log-Normal this seems to be a more natural specification.¹¹ This regression captures the three predictions made by the HIP and RIP models discussed above by adding the past and future income growth rate as well as past and future income levels. We also write the regression without restricting the coefficient on lagged consumption (a_{12}) which allows for both a difference specification (when $a_{12} = -1$) as well as a level specification ($a_{12} = 0$) as discussed above. Finally we also add a time trend.

To complete the auxiliary model we also add a second equation with the current labor income as the dependent variable, but without changing the right hand side variables. The final model is:

¹⁰Two important differences of the present paper from Guvenen (2007) is that that paper (i) only estimated $\widehat{\sigma}_{\beta|0}^2$ from consumption data, taking all other parameters as estimated from income data, and (ii) only used the rise in within-cohort consumption inequality as a moment condition. The present paper instead (i) brings consumption data to bear on the estimation of the *entire* vector of structural parameters, and (ii) does this by systematically focusing on the dynamic relationship between consumption and income movements.

¹¹Although, the auxiliary model would correspond to the structural equation in (19) more closely if the coefficients were time-varying, this would increase the number of parameters in the auxiliary model substantially. Our experience is that the small sample performance of the estimator is better when the auxiliary model is more parsimonious, and therefore we opt for the specification here.

$$\begin{bmatrix} c_t \\ y_t \end{bmatrix} = \begin{bmatrix} a_{10} \\ a_{20} \end{bmatrix} + \begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} & a_{15} & a_{16} & a_{17} \\ a_{21} & a_{22} & a_{23} & a_{24} & a_{25} & a_{26} & a_{27} \end{bmatrix} \begin{bmatrix} y_{t-1} \\ c_{t-1} \\ \bar{y}_{1,t-2} \\ \bar{y}_{t+1,T} \\ \overline{\Delta y}_{1,t-1} \\ \overline{\Delta y}_{t+1,T-t} \\ t \end{bmatrix} + \begin{bmatrix} b_{11} & 0 \\ b_{12} & b_{22} \end{bmatrix} \begin{bmatrix} \epsilon_{1t} \\ \epsilon_{2t} \end{bmatrix} \tag{19}$$

To implement the indirect inference estimator, we choose the values of the structural parameters so that the (approximate) likelihood of the observed data (as defined by the auxiliary model) is as large as possible. That is, given a set of structural parameters, we simulate data from the model, use this data to estimate the auxiliary model parameters, and evaluate the likelihood defined by the auxiliary model at these parameters. We then vary the structural parameters so as to maximize this likelihood. Viewed from another perspective, we are simply minimizing the difference between the (log) likelihood evaluated at two sets of auxiliary model parameters: the estimates in the observed data and the estimates in the simulated data (given a set of structural parameters). The advantage of this approach over other approaches to indirect inference (such as efficient method of moments or minimizing a quadratic form in the difference between the observed and simulated auxiliary model parameters) is that it does not require the estimation of an optimal weighting matrix. It is, however, less efficient asymptotically than the other two approaches, though this inefficiency is small when the auxiliary model is close to being correctly specified (and vanishes in the case of correct specification).

3.2 The Data

3.2.1 Constructing an Imputed Panel of Consumption

An important impediment to the previous efforts to use consumption data has been the lack of a sufficiently long panel on consumption expenditures. The Panel Study of Income Dynamics (PSID) has a long panel dimension but covers limited categories of consumption whereas the Consumer Expenditure Survey (CE) has detailed expenditures over a short period of time (four quarters). As a result, most previous work has either used food expenditures as a measure of non-durables consumption (available in PSID), or resorted to using repeated cross-sections from CEX under additional assumptions.

In a recent paper Blundell, Pistaferri and Preston (2006b, BPP) develop a structural imputation method which imputes consumption expenditures in PSID using information from CE. The basic approach involves estimating a demand system for food consumption as a function of nondurable expenditures, a wide set of demographic variables, and relative prices as well as the interaction of nondurable expenditures with all these variables. The key condition is that all the variables in the demand system must be available in the CE data set, and all but non-durable expenditures must be available in PSID. One then estimates this demand system from CEX. As long as this demand system is monotonic in nondurable expenditures, one can invert it and obtain a panel of imputed consumption in the PSID.

BPP show that several statistics of the imputed measure of consumption in PSID compare remarkably well to their counterparts from CEX. In this paper we use an extension of the method proposed by these authors. In particular, these authors include time dummies interacted with nondurable expenditures in the demand system to allow for the elasticity to change over time, which they find to be important for the accuracy of the imputation procedure. However, the CEX is only available after 1980, whereas we would like to use the entire length of PSID from 1968 to 1992, which makes the use of time dummies impossible. Therefore, instead of time dummies, we include additional terms, specifically the interaction of nondurable expenditures with food and fuel inflation rate, and we allow for the price elasticity of food to be non-linear in the prices of food and fuel. This modification delivers an imputed consumption measure that has a rather good fit (comparable to that in BPP) to the statistics from CEX.¹²

Since food and non-food consumption are jointly determined, some of the right hand side variables are endogenous. In addition, nondurable expenditures are likely to suffer from measurement error (as is the case in most survey data sets), which necessitates an instrumental variables approach. We instrument log nondurable expenditures (as well as its interaction with demographics and prices) with the cohort-year-education specific average of the log of the husband's hourly wage and the cohort-year-education specific average of the log of the wife's hourly wage (as well as their interaction with the demographics and prices).

Table 1 reports the estimate of the demand system using the CEX data. Several terms with the log of nondurable expenditures are significant as well as several of the demographic and price variables. We invert this equation as described above to obtain the imputed measure of consumption. Notice that the unit of analysis is a household.

¹²Another difference from BPP is that we also use the 1972-73 CEX data in some versions of the IV regression to better capture the demand system in the 1970's. Those results are not reported here yet and will be included in the next version of the paper.

Table 1: ESTIMATING A DEMAND SYSTEM FOR FOOD IN CEX USING INSTRUMENTAL VARIABLES REGRESSION

Dependent variable is: <i>log total food consumption</i>			
$\ln(c)$	0.6204 ** [0.0964]	<i>age</i>	0.0023 [0.0140]
$\ln(c) \times age$	0.0030* [0.0015]	age^2	-0.0003 ** [0.0000]
$\ln(c) \times year$	-0.0011 [0.0037]	Northeast	0.0556 ** [0.0066]
$\ln(c) \times$ High school drop	-0.0884 [0.0514]	Midwest	0.0405 ** [0.0063]
$\ln(c) \times$ High school grad	-0.0178 [0.0404]	South	-0.0016 [0.0056]
$\ln(c) \times$ One child	0.0103 [0.0281]	Family size	0.0440 ** [0.0035]
$\ln(c) \times$ Two children	-0.0113 [0.0326]	$\ln p_{food}$	1.1902 [1.1222]
$\ln(c) \times$ Three children	0.0228 [0.0340]	$\ln p_{food} \times 1 \{5\% < \Delta \ln p_{food} < 8\%\}$	0.0134 [0.0071]
$\ln(c) \times \Delta \ln p_{food}$	-0.0034 [0.0023]	$\ln p_{food} \times 1 \{8\% < \Delta \ln p_{food} < 13\%\}$	-0.0122 [0.0076]
$\ln(c) \times \Delta \ln p_{fuel}$	0.0010 ** [0.0003]	$\ln p_{fuel}$	-1.8229* [0.8332]
One child	-0.0664 [0.2703]	$\ln p_{fuel} \times 1 \{5\% < \Delta \ln p_{fuel} < 8\%\}$	-0.0254 [0.0147]
Two children	0.1709 [0.3140]	$\ln p_{fuel} \times 1 \{8\% < \Delta \ln p_{fuel} < 13\%\}$	-0.0397* [0.0192]
Three children	-0.1602 [0.3282]	$\ln p_{fuel} \times 1 \{13\% < \Delta \ln p_{fuel}\}$	-0.0597* [0.0255]
High school dropout	0.8226 [0.4861]	White	0.1006 ** [0.0073]
High school graduate	0.1781 [0.3850]	Constant	4.4553* [2.1667]
Observations			14742

Note: Standard errors are in Parenthesis

In PSID, households report their total taxable income which includes labor income, transfers and financial income of all the members in the household. The measure of labor income we use subtracts financial income from this measure, and therefore, includes the labor income of the head and spouse as well as several categories of transfer income (unemployment benefits, social security income, pension income, worker’s compensation, welfare payments, child support, and financial help from relatives are the main components).

4 Results

In this section we report a variety of results. First, we begin with the RIP process and illustrate the performance of the estimator by reporting some Monte Carlo simulations. We then report some preliminary estimates of the RIP process using actual data. We then turn to the HIP process and report some Monte Carlo simulations, although using a different set up for parameter values, etc. to illustrate the performance of the estimator in significantly different environments.

4.1 Estimates of the RIP Process

4.1.1 Monte Carlo Experiments

We first report some Monte Carlo simulation results to evaluate the performance of the indirect inference estimation of the RIP process. We set the length of life to 10 periods and the number of individuals, N , to 4,000. We also allow for 75 percent of observations to be missing in a way that is randomly distributed across T and N (therefore, the effective sample size is 10,000). We set the borrowing constraints to zero: $\underline{W}_t = 0.0$ for all t . With this calibration constraints are binding for 22 percent of the population in a given period. Notice that the existence of frequently binding borrowing constraints is a challenge to the estimation procedure as it introduces a kink in the decision rule. As noted above, binding constraints have been ignored in most of the previous work mentioned above to make a GMM estimation feasible. We assume logarithmic preferences and set the interest rate, r , to 5 percent. We also set $\delta(1+r) = 1$ where δ is the discount rate. We simulate 100 data sets and the results below

As can be seen in Table 3 the results are quite encouraging. All the parameters of interest are estimated tightly and with a very small bias. We add classical measurement error to

Table 2: BASELINE PARAMETERIZATION FOR MONTE CARLO SIMULATION

Annual model		
Parameter		Value
δ	Time discount rate	1/1.05
r	Risk-free rate	0.05
ϕ	Relative risk aversion	1
R	Retirement age	10
N	Number of observations	10,000

Table 3: MONTE CARLO RESULTS FROM RIP PROCESS, NO MEASUREMENT ERROR

	ρ	σ_η	σ_ε	σ_α
True Value	.900	.200	.100	.300
Starting value	.750	.240	.080	.330
Mean Estimate	.902	.201	.099	.299
Standard deviation	.024	.009	.004	.011

both consumption and income:

$$y_t^* = y_t + u_t^y,$$

$$c_t^* = c_t + u_t^c$$

where y_t^* and c_t^* are measured income and consumption respectively. The results are presented in Table 4. Again, there is very little bias and the parameters are estimated tightly.

Table 4: MONTE CARLO RESULTS FROM RIP PROCESS, WITH MEASUREMENT ERROR

	ρ	σ_η	σ_ε	σ_{u^y}	σ_{u^c}	σ_α
True Value	.900	.200	.100	.130	.160	.300
Start. val.	.800	.240	.080	.200	.090	.330
Mean Estimate	.899	.199	.099	.129	.160	.302
Std. dev.	.018	.005	.012	.014	.001	.016

Table 5: RESULTS FROM RIP PROCESS USING REAL DATA

Data	ρ	σ_η	σ_ε	σ_{u^y}	σ_{u^c}	σ_α
Y only	.965	.153	.171	—		.34
C only	.996	.169	.010	—	.261	.40
Y, C	.975	.151	.179		.268	.36
Y, C	.971	.152	.158	.088	.267	.39

4.1.2 Results with Real Data

We now present the preliminary estimates obtained by using PSID data. The estimates in the first row are obtained by using only the second line of the auxiliary model, that is, without using any consumption data. The estimated parameter values are in line with those reported in the previous literature (c.f., Heathcote, Storesletten, Violante (2006), Guvenen (2005), Topel (1991) among others). The second row repeats the estimation, but now only using the consumption equation in the auxiliary model. The estimate of the persistence parameter is somewhat higher, but the results are broadly similar to the previous one. The next two rows use the full auxiliary model and allows for measurement error in consumption and income. The estimates are similar to the previous rows and overall appear quite plausible.

The fact that the estimates of the RIP process using income data are very similar to those obtained using consumption data should not be interpreted to mean that the latter does not contain any more information than the former. This is because as shown in Guvenen (2005), if the true data generating process is HIP, the RIP specification creates a potentially substantial upward bias in the estimates of ρ . In other words, if in reality income profiles have heterogeneous growth rates but the econometrician ignores this fact (as we do when we estimate a RIP process) the estimated persistence will be biased towards 1, leaving little room for other parameters to vary. Therefore, the information contained in consumption data is likely to play a more important role when we estimate the HIP model below.

4.2 Estimates of the HIP Process

4.2.1 Monte Carlo Experiments

In this section, we apply the proposed methodology to the estimation of the full HIP model with learning. To demonstrate the ability of this estimation method to uncover the true

structural parameter vector in spite of (very) incomplete individual histories and substantial measurement error, we conduct a small Monte Carlo study using 50 “observed” data sets drawn from the true data generating process. The set up here is quite different than in the RIP case before and mimics the actual estimation to be conducted later below more closely.

The missing observations in the Monte Carlo study are chosen to be exactly the same as in the observed data (we include only individuals with at least five observations, for a total of about 2,500 individuals and 30,000 observations). In the model, the discount factor and the interest rate are both set to 4%, the number of years in the working life is set to 41, and the number of years in the retirement period is set to 30. The model incorporates a simplified Social Security system in which individuals receive, in each of the retirement periods, 30% of their income at age 65 (the last year of the working life). Individuals have isoelastic utility with coefficient of relative risk aversion equal to 2. The borrowing constraint is set close to the loosest possible value consistent with almost sure repayment of debts at the end of life, so that few individuals hit the borrowing constraint during their lifetimes. This setup is nearly identical to the one that will be used in the estimation using “real” data, so it is a good test of the performance of the proposed estimation methodology.

We use the auxiliary model given in equation (19). When constructing the regressors, we use consumption and income data corrupted by measurement error. These explanatory variables explain a substantial amount of the variation in measured consumption and income (the R^2 s in the two regressions are about 0.7) and t -statistics on the coefficients are all statistically significant at (at least) the 5% level.

Incomplete histories are handled by “filling in” missing values in a reasonable way: in particular, for each individual we estimate a linear time-trend model of either log consumption or log income using the observed data for that individual, and then use predicted values from this regression for time periods in which data is not observed. We use exactly the same procedure in both the simulated and observed data. The filled-in data is then treated as observed data when constructing the right-hand side variables in the auxiliary model. The (approximate) likelihood, however, includes contributions only from those time periods in which the left-hand side variables are observed (i.e., not missing).

The results are contained in Table 4. The “true values” for the parameters are set to the estimates reported in Guvenen (2007). The initial values of the parameters are set randomly to $\pm 10\%$ of the true values. Each Monte Carlo run takes about 15 minutes on a state-of-the-art workstation. Clearly, the estimation method works well: bias is virtually absent and standard deviations are small. Although it is difficult, if not impossible, to prove identification in this setup, the results suggest strongly that local identification near the true

Table 6: MONTE CARLO RESULTS FROM HIP PROCESS, WITH MEASUREMENT ERROR

	σ_α	ρ	σ_η	σ_ε	$\sigma_\beta \times 100$	λ	σ_{u^y}	σ_{u^c}
True Value	.148	.821	.170	.216	1.959	.600	.200	.200
Mean Estimate	.140	.823	.170	.218	1.942	.602	.200	.200
Std dev	.033	.021	.008	.012	.054	.052	.002	.010
Max estimate	.186	.884	.190	.238	2.047	.742	.203	.220
Min estimate	.050	.773	.144	.190	1.813	.509	.196	.175

parameter vector does indeed hold. These results are very encouraging and suggest strongly that the proposed methodology is a feasible and practical method for estimating structural consumption-saving models with missing data and multiple sources of heterogeneity.

4.2.2 Results with Real Data

To be added.

5 Conclusions

The joint dynamics of consumption and labor income contains potentially rich information that can distinguish between the RIP and HIP models. Monte Carlo results from the RIP model suggests that the indirect inference method works very well, even in the presence of frequently binding borrowing constraints, missing observations, retirement income, and so on, that make the auxiliary model a poor approximation to the structural relationships that need to hold in the model. We plan to introduce time and cohort effects in variances and non-separable leisure into the utility function. We also aim to conduct formal statistical tests to assess the extent to which the RIP and HIP models find support in the data.

References

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

- [2] Baker, Michael (1997): “Growth-rate Heterogeneity and the Covariance Structure of Life-Cycle Earnings,” *Journal of Labor Economics*, Vol. 15, pp. 338-375.
- [3] Becker, Gary S. (1965): “Human Capital : A Theoretical and Empirical Analysis, with Special Reference to Education,” reprinted 1994, University of Chicago Press.
- [4] Ben-Porath, Yoram (1967): “The Production of Human Capital and the Life Cycle of Earnings,” *Journal of Political Economy*, Vol. 75, pp 352-365.
- [5] Blundell, Richard, and Ian Preston (1998): “Consumption inequality and income uncertainty,” *Quarterly Journal of Economics*, Vol. 113, pp. 603-640.
- [6] Blundell, Richard, Luigi Pistaferri, and Bruce Preston (2006a): “Consumption Inequality and Partial Insurance,” Stanford University Working Paper.
- [7] Blundell, Richard, Luigi Pistaferri, and Bruce Preston (2006b): “Imputing Consumption in the PSID Using Food Demand Estimates in the CEX,” Stanford University Working Paper.
- [8] Cunha, Flavio, James Heckman, and Salvador Navarro (2005): “Separating uncertainty from heterogeneity in life cycle earnings,” *Oxford Economic Papers*, forthcoming.
- [9] Deaton, Angus, and Christina Paxson (1994): “Intertemporal Choice and Inequality,” *Journal of Political Economy*, Vol. 102, pp. 437-467.
- [10] Gourinchas, Pierre-Olivier, and Jonathan Parker (2002): “Consumption Over the Life Cycle,” *Econometrica*, Vol 70, pp. 47-89.
- [11] Guvenen, Fatih (2005): “An Empirical Investigation of Labor Income Processes,” Working Paper, University of Texas at Austin.
- [12] Guvenen, Fatih (2007): “Learning Your Earning: Are Labor Income Shocks Really Very Persistent?” *American Economic Review*, forthcoming.
- [13] Guvenen, Fatih and Burhanettin Kuruscu (2006): “A Quantitative Analysis of the Evolution of the U.S. Wage Distribution: 1970—2000,” Working Paper, University of Texas at Austin.
- [14] Haider, Steve J. (2001): “Earnings Instability and Earnings Inequality of Males in the United States: 1967-1991,” *Journal of Labor Economics*, Vol. 19, pp. 799-836.
- [15] Hause, John C. (1980): “The Fine Structure of Earnings and the On-the-Job Training Hypothesis,” *Econometrica*, Vol. 48, pp. 1013-29.

- [16] Heathcote, Jonathan, Kjetil Storesletten, Gianluca Violante (2006): “The Macroeconomic Implications of Rising Wage Inequality in the United States,” Working Paper, NYU.
- [17] Huggett, Mark, Gustavo Ventura, and Amir Yaron (2004): “Risky Human Capital,” Working Paper, Wharton School, University of Pennsylvania.
- [18] 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.
- [19] 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.
- [20] Meghir Costas and Luigi Pistaferri (2004): “Income Variance Dynamics and Heterogeneity,” *Econometrica*, Vol. 72, pp. 1-32.
- [21] Moffitt, Robert A., and Peter Gottschalk (1994): “Trends in the Autocovariance Structure of Earnings in the US, 1969-1987, Working Paper, Brown University.
- [22] Pischke, Jörn-Steffen (1995): “Individual Income, Incomplete Information, and Aggregate Consumption,” *Econometrica*, Vol. 63, pp. 805-840
- [23] Topel, Robert H. (1990): “Specific Capital, Mobility and Wages: Wages rise with Job Seniority,” NBER Working Paper No. 3294.