

Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
<http://economics.sas.upenn.edu/pier>

PIER Working Paper 10-012

“On the Persistence of Income Shocks over the Life Cycle:
Evidence and Implications”, Second Version

by

Fatih Karahan and Serdar Ozkan

<http://ssrn.com/abstract=1585884>

On the Persistence of Income Shocks over the Life Cycle: Evidence and Implications*

Fatih Karahan[†] Serdar Ozkan[‡]

April 5, 2010

Abstract

How does the persistence of earnings change over the life cycle? Do workers at different ages face the same variance of idiosyncratic shocks? This paper proposes a novel specification for residual earnings that allows for a lifetime profile in the persistence and variance of labor income shocks. We show that the statistical model is identified and estimate it using PSID data. We strongly reject the hypothesis of a flat life-cycle profile for persistence and variance of persistent shocks, but not for the variance of transitory shocks. Shocks to earnings are only moderately persistent (around 0.75) for young individuals. Persistence rises with age up to unity until midway in life. On the other hand, the variance of persistent shocks exhibits a U-shaped profile over the life cycle (with a minimum of 0.01 and a maximum of 0.045). Our estimate of persistence, for most of the working life, is substantially lower than typical estimates in the literature. We investigate the implications of these profiles for consumption-savings behavior with a standard life-cycle model. The welfare cost of idiosyncratic risk implied by the age-dependent income process is 32% lower compared to an $AR(1)$ process without age profiles. This is mostly due to a higher degree of consumption insurance for young workers, for whom persistence is moderate. We conclude that the welfare cost of idiosyncratic risk will be overstated if one does not account for the age profiles in the persistence and variance of shocks.

Keywords: Idiosyncratic income risk, Incomplete markets models, Earnings persistence, Consumption insurance

JEL: C33, D31, D91, E21, J31

*We would like to thank Dirk Krueger, Iourii Manovskii, Flavio Cunha, Jesus Fernandez-Villaverde, Fatih Guvenen, Greg Kaplan, Kjetil Storesletten, Michael Palumbo, Geng Li, Ivan Vidangos, Kurt Mitman, seminar participants at the Penn Macro club, and the Board of Governors of the Federal Reserve System for their valuable comments and encouragement. We also thank Fatih Guvenen for kindly sharing his dataset. Of course, all errors are our own.

[†]University of Pennsylvania; karahan@econ.upenn.edu; <http://www.econ.upenn.edu/~karahan>

[‡]University of Pennsylvania; ozkan@econ.upenn.edu; <http://www.econ.upenn.edu/~ozkan>

1 Introduction

How does the persistence of earnings change over the life cycle? Do workers at different ages face the same variance of idiosyncratic shocks? Answers to these questions are central to many economic decisions in the presence of incomplete financial markets. Uninsured idiosyncratic risk affects the dynamics of wealth accumulation, consumption inequality, and the effectiveness of self-insurance through asset accumulation. Thus, income risk is an important object of study for quantitative macroeconomics. Moreover, the age profile of persistence can be informative about the economic mechanisms underlying earnings risk. For these purposes, we propose a novel process for idiosyncratic earnings that allows for a life-cycle profile in the persistence and variance of earnings shocks.

Two important determinants of labor income risk are the persistence and variance of shocks. The persistence governs how long the effect of a shock lasts. For example, in the case of an unexpected health problem, this represents the time to full recovery. The variance, on the other hand, captures the magnitude by which shocks affect earnings. The goal of this paper is to estimate the lifetime profiles of these two components.

We are motivated by the observation that changes in earnings occur for different reasons over the life span. For young workers, mobility because of a mismatch or demand shocks to occupations might play an important role ([Kambourov and Manovskii \(2008\)](#)). Midway through a career, settling down into senior positions as well as bonuses, promotions or demotions may account for earnings dynamics. Older people are more likely to develop health problems that reduce their productivity. These changes differ in nature, and more specifically, in persistence and magnitude. Thus, we suspect that variance and persistence of shocks are constant throughout a lifetime.

In our analysis, we decompose residual earnings into an individual-specific fixed effect, a persistent component and a transitory component. The fixed effect captures permanent differences among individuals. The persistent component captures lasting changes in earnings and it is modeled as an $AR(1)$ process. The transitory component encompasses both measurement error and temporary changes in earnings and is i.i.d. The novel feature of our specification is that both the persistence parameter of the $AR(1)$ process and the variance of innovations to transitory and persistent components are allowed to vary by age. Besides allowing for age profiles, we also account for changes in variances over time. This paper, to our best knowledge, is the first study that estimates a lifetime profile of earnings persistence

and variance.¹

We next turn to identification. Particularly, which features of the data tell us how changes in earnings vary in persistence and variance over the lifetime? We show that these profiles can be identified using the variance covariance structure of levels of earnings. Intuitively, we identify the age profile of persistence by tracking the covariance structure over lags for a given age. The variance of persistent shocks is obtained by exploiting the variation in the covariance structure over age for a given lag. Finally, the variance of transitory shocks is recovered from the variance structure. The proof is rigorously discussed in Appendix A.

Using earnings data from the PSID, we first estimate a nonparametric specification, i.e., without imposing any functional form on the lifetime profiles of persistence and variance. Our results reveal that persistence is increasing at early stages in the working life. Young agents face only moderately persistent shocks: 70 percent of a shock received during the early years in the labor market dies out over the next 5 years. Shocks for workers midway through their careers are more enduring. If the shock was received at age 40, 85 percent of it would still remain after 5 years. On the other hand, we find a U-shaped profile for the variance of persistent shocks: A shock of one standard deviation implies a 26% change in annual earnings for a 24 year old. The corresponding number for a 40 year old is only 12%. These are sizable differences. As for the variance of transitory shocks, we find a slight increase early on but a flat profile for the remaining working life.²

We then ask the question of whether these life-cycle profiles are statistically significant. To tackle this question, we proceed in two ways. First, we estimate a quadratic function for the age profiles and test whether the coefficients on the linear and quadratic terms are zero. Then, in order to complement this approach, we also estimate life-cycle profiles by partitioning the working life into 3 stages. Here, we assume that persistence and variance are constant within a stage but might differ from one to the other. Again, we test whether the profile of persistence and the variance of persistent shocks are flat over the lifetime. Both of these tests strongly reject the hypothesis of a flat profile for persistence and the variance of persistent shocks.

¹Meghir and Pistaferri (2004) allow for an age profile in the variance of permanent and transitory shocks. They don't find evidence for a nontrivial profile.

²The 95% bootstrap confidence intervals point to a flat profile. In our specification, transitory shocks also capture classical measurement error. Therefore, it is not surprising to find a flat profile for transitory shocks.

The estimates of persistence in the literature are close to unity.³ Our age-specific estimate of persistence lies substantially below 1 for most of the lifetime. We argue that the high persistence in the literature is driven by targeting the almost linear increase in lifetime earnings inequality. Namely, estimation avoids lower levels of persistence, which would imply a concave rise in inequality. The age-dependent income process can capture the linear shape without high levels of persistence. This is possible because of the inverse relationship between persistence and the variance of labor income shocks that our estimates reveal. When persistence goes up with age, the additional increase it induces in inequality is compensated by a decrease in the variance.

We then investigate the economic implications and economic significance of the age-dependent income process. In particular, we ask how much the presence of age profiles matters for the insurability of labor income shocks and the welfare costs of idiosyncratic risk. To address these issues, we study a standard life-cycle model featuring incomplete financial markets and a social security system. We compare the consumption-savings implications of the age-dependent income process with a standard $AR(1)$ process (with constant persistence and variance).

We start with an economy with natural borrowing constraints (NBC). We find that both of the processes imply very similar consumption and asset profiles. However, they differ significantly in the degree of consumption insurance against persistent shocks. We measure the level of insurance as the fraction of shocks to earnings that do not lead to consumption changes (Blundell, Pistaferri, and Preston (2008)). Around 44% of persistent shocks translate into consumption growth under the age-dependent process compared to 60% under the standard $AR(1)$ specification. Most of this difference comes from young workers for whom the degree of insurance is as high as 70% under the age-dependent process as opposed to 30% under the $AR(1)$ process. This is due to the level of persistence, which is particularly low for young workers under the age-dependent process. It is well known that persistence is an important determinant of insurance; transitory shocks are easily insured by borrowing (e.g., Kaplan and Violante (2008), Gourinchas and Parker (2002)). In the presence of very persistent shocks, agents refrain from borrowing against the possibility of a long sequence of low income states. Insurance against such shocks is, therefore, mostly through assets. This is not possible for young agents, since they don't have enough wealth.

³Estimates of specifications that account for the heterogeneity in income growth rates find lower levels of persistence. In particular, Guvenen (2009) estimates persistence at around 0.82.

Persistence is fairly moderate for young workers under the age-dependent income process, which explains the higher insurance coefficients early in careers.

Note that the low levels of persistence under the age-dependent process are compensated by the larger variances of shocks. On the one hand, lower persistence implies better insurability. On the other hand, larger variance implies more instability. In order to evaluate this tradeoff quantitatively, we compare the welfare costs of idiosyncratic risk implied by the age-dependent process with a standard $AR(1)$ process. We find sizable differences: An agent living in the $AR(1)$ economy is willing to give up around 14.85% of her consumption permanently in return for perfect insurance as opposed to only 9.97% for an agent under the age-dependent income process.

As discussed above, the differences in welfare costs are mostly due to higher insurability. The fact that the age-dependent income process results in larger insurance coefficients relies crucially on the extent of borrowing limits. In order to quantify the effect of borrowing limits, we study an economy in which borrowing is ruled out altogether (ZBC economy). The degree of consumption insurance goes down by a significant amount under the age-dependent specification, especially for young workers, for whom insurance falls from around 70% to 26%. This shows the importance of borrowing constraints for young workers.

The decrease in the degree of insurance does have welfare consequences: Welfare costs increase compared to the NBC economy for both of the specifications. The increase is larger for the age-dependent process, lowering the differences between the two processes. However, welfare costs are still significantly lower under the age-dependent income process (12.5% and 16.37%, respectively for the age-dependent and the $AR(1)$ processes).

1.1 Related Literature

Our contribution is twofold. First, we contribute to the literature that models idiosyncratic earnings risk. The estimates of statistical models are used as an input in macroeconomic models with heterogeneous agents. Different specifications will induce different economic decisions; therefore, one needs a good measure of labor income risk. A partial list of such papers includes [Lillard and Willis \(1978\)](#), [Lillard and Weiss \(1979\)](#), [MaCurdy \(1982\)](#), [Abowd and Card \(1989\)](#) and [Baker \(1997\)](#), although none of the papers above have investigated the lifetime profiles of persistence and variances. Our paper fills that void.

A notable exception is [Meghir and Pistaferri \(2004\)](#), which estimates a process with a fully permanent component, an $MA(q)$ component where q is estimated, and a fully

transitory component. Their focus is on conditional heteroskedasticity in permanent and transitory shocks. Similar to our paper, they also allow for age profiles in the variance of permanent and transitory shocks. However, unlike our paper, they do not allow persistence to change over the life cycle. They find no evidence in favor of changing variance over the lifecycle. In this paper, we argue that it is crucial to allow persistence to change with age.

Another paper related to ours is [Hause \(1980\)](#). Using data on Swedish white collar workers, he estimates a process that has an $AR(1)$ component with time-specific persistence and variance of shocks. Since his data set contains only workers born in 1943, it is not clear whether these profiles are age or time-specific. Our paper takes advantage of the rich panel structure of the PSID and separates changes over time from changes over the life cycle.

Recently, [Guvenen \(2009\)](#) argues for the existence of growth rate heterogeneity and finds evidence against unit roots. The evidence he brings forward is twofold. First, he points to the convexity in the variance profile of earnings. Second, he exploits the increase in higher order covariances. He argues that these can be captured through growth rate heterogeneity but not by highly persistent shocks. The age-dependent income process can inherently capture these features of the data without growth rate heterogeneity. [Alvarez, Browning, and Ejrnaes \(2006\)](#) investigates the role of heterogeneity in income dynamics of individuals and find significant heterogeneity among seemingly homogeneous individuals. Our paper can be thought as complementary to theirs in that we focus on *observed* heterogeneity, that is, heterogeneity across age.

Another approach to infer the nature of earnings risk is to make use of economic choices. [Guvenen \(2007\)](#), [Storesletten, Telmer, and Yaron \(2004\)](#) and [Guvenen and Smith \(2009\)](#) are papers that bring consumption data into the picture to make inference about the nature of income risk. [Cunha, Heckman, and Navarro \(2004\)](#) use schooling decisions and decompose residual earnings into a component that is foreseen and acted upon (heterogeneity) and a component that is unanticipated (shocks). [Feigenbaum and Li \(2008\)](#) also make this distinction and measure income uncertainty as the variance of income forecasting errors at different ages. They find a U-shaped uncertainty profile over the life cycle. [Altonji, Smith, and Vidangos \(2009\)](#) consider a structural approach to estimate a joint model of earnings, employment, job changes, wage rates, and work hours.

We also contribute to the literature on consumption insurance. [Blundell, Pistaferri, and Preston \(2008\)](#) develop and apply a methodology to measure the degree of consumption insurance against permanent and transitory shocks. [Kaplan and Violante \(2008\)](#) argue

that the lifetime profile of insurance coefficients in the data is not consistent with a life-cycle model that features a standard $AR(1)$ process, since this implies that the insurance profile follows the profile of assets, which is roughly increasing over the life cycle. However, [Blundell, Pistaferri, and Preston \(2008\)](#) find a roughly flat insurance profile in the data. We show that under the age-dependent income process proposed in this paper, the profile of insurance need not be increasing.

The rest of the paper is organized as follows: In [Section 2](#) we describe the statistical model that we estimate, discuss its identification and present our results. [Section C](#) discusses the implications of the learning model for residual wages. [Section 3](#) presents the life-cycle model that is used to study the consumption-savings implications of the age-dependent process and compares its welfare consequences to a standard $AR(1)$ process. Finally, [Section 4](#) concludes.

2 Empirical Analysis

In this section we describe the statistical model for earnings. We start with a simple age-dependent income process and discuss its identification. We then introduce the full-blown model, but the proof of identification is left to the appendix. Empirical results are discussed at the end of this section.

2.1 An Age-Dependent Income Process

Let \tilde{y}_h^i be the residual component of earnings of individual i at age h , which is obtained by running cross-sectional regressions of earnings on observables.⁴ The details of this first-stage regression are presented later. Residual income is decomposed into a fixed effect, an $AR(1)$ component, and a transitory component. This representation is simple, yet it captures the salient features of the data well. Therefore, it is widely used in the literature.⁵ This paper

⁴Some papers, such as [Guvenen \(2009\)](#), use potential experience as the explanatory variable instead of age which is defined as $age - \max(\text{schooling}, 12) - 6$. This is used as a proxy for actual experience in order to avoid endogeneity issues. We use age since potential experience is collinear with it. We carried out the same analysis with potential experience, and the results are similar (see [Appendix B.2](#)).

⁵Some papers, including [Meghir and Pistaferri \(2004\)](#) and [Hryshko \(2008\)](#), allow for a fixed effect, a permanent component (unit root), a fully transitory component and a persistent component that is modeled either as an $MA(q)$ or $AR(1)$.

extends the standard specification to allow for a lifetime profile in the persistence parameter, the variance of persistent and transitory shocks:

$$\begin{aligned}
\tilde{y}_h^i &= \alpha^i + z_h^i + \varepsilon_h^i & (1) \\
z_h^i &= \rho_{h-1} z_{h-1}^i + \eta_h^i \\
\eta_h^i &\sim iid(0, \sigma_{\eta,h}^2) \quad \varepsilon_h^i \sim iid(0, \sigma_{\varepsilon,h}^2)
\end{aligned}$$

Here, α^i is an individual-specific fixed effect that captures the variation in initial conditions such as innate ability. ε_h^i is a fully transitory component that encompasses both measurement error and temporary changes in earnings such as bonuses and overtime pay.⁶ z_h^i is the persistent component of idiosyncratic income at age h that captures lasting changes in earnings such as promotions and health status. Each period the individual is hit by a persistent shock of size η_h^i . The magnitude of this shock is governed by the variance $\sigma_{\eta,h}^2$ and the extent to which it lasts is determined by the persistence parameter ρ . The key innovation of our paper is to allow for an age profile in the variance of shocks, $\sigma_{\eta,h}^2$ and $\sigma_{\varepsilon,h}^2$, as well as in the durability of the persistent shocks, ρ_h .

The age profiles capture the idea that changes in earnings occur for different reasons throughout the life span. For example, young households experience high mobility because of a mismatch or demand shocks to occupations. On the other hand, middle-aged workers settle down into senior positions and experience promotions or demotions that lead to changes in earnings. As for older people, the causes of earnings instability are more likely to be health problems. These sources of earnings dynamics differ in nature, and more specifically, in persistence and magnitude. Thus, we suspect that the variance and the persistence of shocks are not flat throughout the lifetime. Rather than imposing constant parameters throughout the lifetime, we let the data speak for itself.

Having introduced the age-dependent income process, an immediate concern is identification. Which features of the data tell us how changes in earnings vary in variance and persistence over the lifetime? The identification discussion allows us to connect the statistical model to the moments in the data and makes the estimation procedure meaningful. Intuitively, we identify the profile of persistence by tracking the covariance structure over

⁶These changes are potentially correlated with future promotions. However, we follow the literature and assume that these shocks are i.i.d. in nature.

lags for a given age. The variance of persistent shocks is obtained by exploiting the variation in the covariance structure over ages for a given lag. Finally, the variance of transitory shocks is recovered from the variance structure.

The next proposition establishes that the income process (1) is identified and provides a formal proof:

Proposition 1: Specification (1) is identified in levels up to the normalization that $\rho_1 = \rho_2$.

Proof: We use the variance-covariance structure in levels that is implied by specification (1) and outline a strategy to identify the parameters of the statistical model. Below we present this variance-covariance structure.

$$\text{var}(\tilde{y}_h^i) = \sigma_\alpha^2 + \text{var}(z_h^i) + \sigma_{\varepsilon,h}^2 \quad h = 1, \dots, H \quad (2)$$

$$\text{cov}(\tilde{y}_h^i, \tilde{y}_{h+n}^i) = \sigma_\alpha^2 + \left(\prod_{j=h}^{h+n-1} \rho_j \right) \text{var}(z_h^i), \quad (3)$$

$$h = 1, \dots, H-1 \quad n = 1, \dots, H-h$$

$$\text{var}(z_h^i) = \rho_{h-1}^2 \text{var}(z_{h-1}^i) + \sigma_{\eta,h}^2 \quad h = 1, \dots, H \quad (4)$$

Let's first assume that we know the variance of the fixed effect, σ_α^2 , and show that we can identify all the remaining parameters. Then we come back to argue that the unused moment conditions are enough to pin down σ_α^2 *uniquely*.

Note that since we assume that σ_α^2 is known, we can construct $\text{cov}(\tilde{y}_h^i, \tilde{y}_{h+n}^i) - \sigma_\alpha^2$. (3) implies $[\text{cov}(\tilde{y}_h^i, \tilde{y}_{h+2}^i) - \sigma_\alpha^2] / [\text{cov}(\tilde{y}_h^i, \tilde{y}_{h+1}^i) - \sigma_\alpha^2] = \rho_{h+1}$ for $h = 1, \dots, H-2$. This pins down the whole profile of ρ_h for $h = 2, 3, \dots, H-1$ except for ρ_H .⁷ Since ρ_h is already pinned down for $h > 1$, $\text{cov}(\tilde{y}_h^i, \tilde{y}_{h+1}^i) - \sigma_\alpha^2 = \rho_h \text{var}(z_h^i)$ recovers $\text{var}(z_h^i)$ for $1 < h < H$. Note that it is not possible to identify ρ_1 and $\text{var}(z_1^i)$ separately. We make the identifying assumption that $\rho_1 = \rho_2$. This then pins down $\text{var}(z_1^i)$. Using the information contained in (2), we recover $\sigma_{\varepsilon,h}^2 \forall h$.

Note that all of the parameters recovered so far depend on σ_α^2 . It remains to be shown that the unused covariances uniquely pin this down. We now show that $\text{cov}(\tilde{y}_2^i, \tilde{y}_5^i)$ suffices

⁷Note that ρ_H does not enter the variance-covariance profile at all, so it is, in fact, not a parameter of the model.

to recover σ_α^2 uniquely:

$$\begin{aligned}
\text{cov}(\tilde{y}_2^i, \tilde{y}_5^i) &= \sigma_\alpha^2 + \rho_4 \rho_3 \rho_2 \text{var}(z_2^i) \\
&= \sigma_\alpha^2 + \rho_4 \rho_3 \rho_2 \left[\frac{\text{cov}(\tilde{y}_2^i, \tilde{y}_3^i) - \sigma_\alpha^2}{\rho_2} \right] \\
&= \sigma_\alpha^2 + \left[\frac{\text{cov}(\tilde{y}_3^i, \tilde{y}_5^i) - \sigma_\alpha^2}{\text{cov}(\tilde{y}_3^i, \tilde{y}_4^i) - \sigma_\alpha^2} \right] \left[\frac{\text{cov}(\tilde{y}_2^i, \tilde{y}_4^i) - \sigma_\alpha^2}{\text{cov}(\tilde{y}_2^i, \tilde{y}_3^i) - \sigma_\alpha^2} \right] [\text{cov}(\tilde{y}_2^i, \tilde{y}_3^i) - \sigma_\alpha^2] \\
&\Rightarrow \frac{\text{cov}(\tilde{y}_2^i, \tilde{y}_5^i) - \sigma_\alpha^2}{\text{cov}(\tilde{y}_2^i, \tilde{y}_4^i) - \sigma_\alpha^2} = \frac{\text{cov}(\tilde{y}_3^i, \tilde{y}_5^i) - \sigma_\alpha^2}{\text{cov}(\tilde{y}_3^i, \tilde{y}_4^i) - \sigma_\alpha^2} \\
\Rightarrow \sigma_\alpha^2 &= \frac{\text{cov}(\tilde{y}_2^i, \tilde{y}_4^i) \text{cov}(\tilde{y}_3^i, \tilde{y}_5^i) - \text{cov}(\tilde{y}_2^i, \tilde{y}_5^i) \text{cov}(\tilde{y}_3^i, \tilde{y}_4^i)}{\text{cov}(\tilde{y}_2^i, \tilde{y}_4^i) + \text{cov}(\tilde{y}_3^i, \tilde{y}_5^i) - \text{cov}(\tilde{y}_2^i, \tilde{y}_5^i) - \text{cov}(\tilde{y}_3^i, \tilde{y}_4^i)}
\end{aligned}$$

Finally, we use (4) to identify $\sigma_{\eta,h}^2 \forall h$.⁸ This completes the proof. Notice that there are still unused moments meaning that the process is overidentified.

2.2 Full Model

In order to better account for earnings dynamics, we extend the basic specification introduced in the previous section by incorporating time effects in variances.

Let $y_{h,t}^i$ denote the log of annual earnings of individual i of age h at time t . To obtain the residual income $\tilde{y}_{h,t}^i$, we run cross-sectional first-stage regressions of earnings on observables. More specifically,


$$y_{h,t}^i = f(X_{h,t}^i; \theta_t) + \tilde{y}_{h,t}^i \quad (5)$$

The first component in this specification, f is a function of age and schooling and captures the life-cycle component of earnings that is common to everyone. $X_{h,t}^i$ is a vector of observables that includes a cubic polynomial in age and an education dummy, indicating whether the individual has a college degree. The parameter θ is indexed by t to allow the coefficients on age and schooling to change over time and captures changes in returns to age and schooling that took place over time.

⁸The result in proposition 1 tells us that $\sigma_{\varepsilon,H}^2$ and $\sigma_{\eta,H}^2$ are unidentified. This is to be anticipated, since distinguishing between persistent changes and transitory changes requires us to observe the individual for several periods (at least one) after the change and see how long the change affects the wage. Obviously, for the last age this is not possible.

Figure 1 plots the evolution of residual inequality for the U.S. during our sample period of 1967-1995. It is obvious that there is a significant change in residual inequality starting in the late 1970s. Ignoring the changes that took place over time might bias our estimates of the age profile of shocks. In particular, changes that occur over time can be misinterpreted as changes during the life cycle. The rich panel structure of the PSID helps us to distinguish life-cycle effects from time effects: We observe individuals with a given age at different points in time, and thus at a given year, we observe individuals of different ages. This allows us to separate what is due to calendar time from a life-cycle phenomenon. For this particular reason, it is important to have a large number of cohorts in order to accurately separate these effects. This observation will guide our sample selection process, as we will explain in 2.3.

Figure 1: Residual Inequality over Time



Here we follow [Gottschalk and Moffitt \(1995\)](#), who argue that significant changes took place in the variance of transitory shocks as well as persistent shocks and modify (1) as:

$$\begin{aligned}
\tilde{y}_{h,t}^i &= \alpha_i + z_{h,t}^i + \phi_t \varepsilon_h^i & (6) \\
z_{h,t}^i &= \rho_{h-1,t-1} z_{h-1,t-1}^i + \pi_t \eta_h^i \\
\eta_h^i &\sim N(0, \sigma_{\eta,h}^2) \quad \varepsilon_h^i \sim N(0, \sigma_{\varepsilon,h}^2),
\end{aligned}$$

where ϕ_t and π_t represent time loading factors for transitory and permanent shocks, respectively.⁹

We leave the formal identification proof for the generalized version to Appendix A, since it doesn't provide any further insight. Here is a heuristic argument. The loading factors on persistent shocks, π_t , will be identified through the changes in the covariances over time. The difference in the covariances between age 1 and age 2 at different points in time must have come from the change in the respective loading factors. Once we have pinned down the profile of π_t 's we then look at the variance profile over time for a given age h . Changes in this variance can be due to a change in the variance of the transitory component or the persistent component. Since we have already identified the profile of π , whatever remains unexplained will be picked up by ϕ , the time profile of transitory shocks. Once we control for the time effects in the variance and covariance structure, the identification of the parameters governing the age profile follows from the previous result.

A related approach would be to control for cohort effects. It is reasonable to think that different cohorts face different economic environments; thus the changes in the residual variance structure may be due to the fact that there are different cohorts at different points in time. It would be better to allow for cohort effects and time effects in variances at the same time but this is not possible because age, time and cohort are perfectly collinear. [Heathcote, Storesletten, and Violante \(2005\)](#) provide some evidence that time effects are more pronounced than cohort effects.¹⁰

⁹This assumes that changes over time have affected everyone at the same age in the same way.

¹⁰Another issue regarding our econometric analysis is measurement error. It has been widely documented that earnings in the PSID contain substantial measurement error. In this paper, we assume that transitory changes also capture the measurement error. The true size of transitory shocks is not distinguishable from the measurement error once we assume fully transitory errors. [Meghir and Pistaferri \(2004\)](#) decompose residual income into a completely permanent component, a transitory component that is modeled as $MA(q)$ and an i.i.d. component that they assume to be measurement error. [Bound and Krueger \(1991\)](#) provide evidence in favor of somewhat persistent measurement errors.

2.3 Data and Sample Selection

This section briefly describes the data and the variable definitions used in the empirical analysis. We use the first 29 waves of the Panel Study of Income Dynamics (PSID). We estimate our model using both annual earnings and the average hourly wage of male heads of households as the measure of labor income. Here, we present the results for earnings data. Estimation results for wage data are reported in Appendix B.1; the results are qualitatively the same. We include an individual in our baseline sample if he satisfies the following criteria for 3 not necessarily consecutive years: (i) the individual has reported positive labor earnings and hours, (ii) his age is between 24 and 60, (iii) he worked between 520 and 5110 hours during the calendar year, and (iv) had an average hourly wage between \$2 and \$400 in 1993 dollars. We also exclude people from the poverty sub-sample in 1968 (SEO). These criteria are fairly standard in the literature and leave us with 4380 individuals and 53,864 observations. Sample statistics are reported in Appendix D.

We exclude individuals younger than 24 to abstract from young part-time workers. Adults older than 60 are also left out to avoid issues related to early retirement. The early retirement of the elderly increases the variance of residual earnings by a substantial amount, since some people quit their jobs for low-paying, less intensive jobs. We did our analysis for a sample between ages 20 and 65; our results are even stronger for this sample. Some of the changes in persistence and variance that we observe for that sample the phenomenon known as might be driven by young individuals who move from part-time to full-time employment or by older individuals who are heterogeneous in retirement age. Therefore, in our baseline case, we present the conservative results. We report the results for the larger sample in Appendix B.2.

Another issue with our sample selection criteria is the minimum number of years. Our choice is guided by the identification argument presented in 2.2. Recall that we need to observe people of the same age at different points in time (and vice versa). Requiring individuals to stay longer in the sample decreases the number of cohorts that we have in the data, since it gets rid of the early cohorts.¹¹

¹¹Of course, another source of concern is the sample size; if we were to require individuals to remain in the sample longer, we would end up with fewer observations. This is important for us, since we are increasing the number of parameters of the specification along the life-cycle dimension.

2.4 Estimation Results

In this section, we present our estimation results. The emphasis is on the existence of a nontrivial lifetime profile.

We employ an equally weighted minimum distance estimator. We minimize the distance between the moments from the theoretical variance-covariance structure and the corresponding moments in the data. In particular, we target all the variance and covariance terms over age, $cov(\tilde{y}_h^i, \tilde{y}_{h+n}^i)$, and over time $cov(\tilde{y}_t^i, \tilde{y}_{t+n}^i)$, but we use only those moments to which at least 150 individuals contribute. To obtain the theoretical counterpart of $cov(\tilde{y}_h^i, \tilde{y}_{h+n}^i)$, we average $cov(\tilde{y}_{h,t}^i, \tilde{y}_{h+n,t+n}^i)$ over t . Similarly, we compute the theoretical counterpart of $cov(\tilde{y}_t^i, \tilde{y}_{t+n}^i)$ by averaging $cov(\tilde{y}_{h,t}^i, \tilde{y}_{h+n,t+n}^i)$ over h . This leaves us with more than 1000 moments. Due to small sample considerations explained in [Altonji and Segal \(1996\)](#), our minimum distance estimator employs the identity matrix as the weighting matrix.

We start by estimating the lifetime profile of shocks and persistence nonparametrically, i.e., without imposing any functional form on the lifetime profiles. [Figure 2](#) shows the results for persistence. The point estimates are plotted with dots and the 95% bootstrap confidence interval is shown with dashed lines.

[Figure 2](#) reveals an interesting fact: Early in life, shocks are moderately persistent. Persistence starts around 0.70 for young individuals and increases with age up to unity by the age of 45. The differences also appear to be economically large (although a more precise evaluation needs to await the consumption model in [Section 3](#)). For example, more than 70% of a change in a 24-year-old's earnings dies out in 5 years. This number is only around 15% for a 40-year-old individual.

The variance of persistent shocks (see [Figure 3](#)) follows the opposite pattern. Early in life, shocks are larger compared to in the 40s. The variance starts around 0.05, decreases to around 0.01 by age 35 and remains roughly flat for 10 years. Shocks toward the end of the life cycle are larger, which manifests itself in a variance of around 0.035. These differences again appear to be economically large; a one-standard-deviation persistent shock implies a 26% change in earnings at age 24, whereas a one standard deviation shock implies only a 12% rise for a 40 year old.

Figure 2: Persistence Profile




Figure 3: Variance Profile of Persistent Shocks

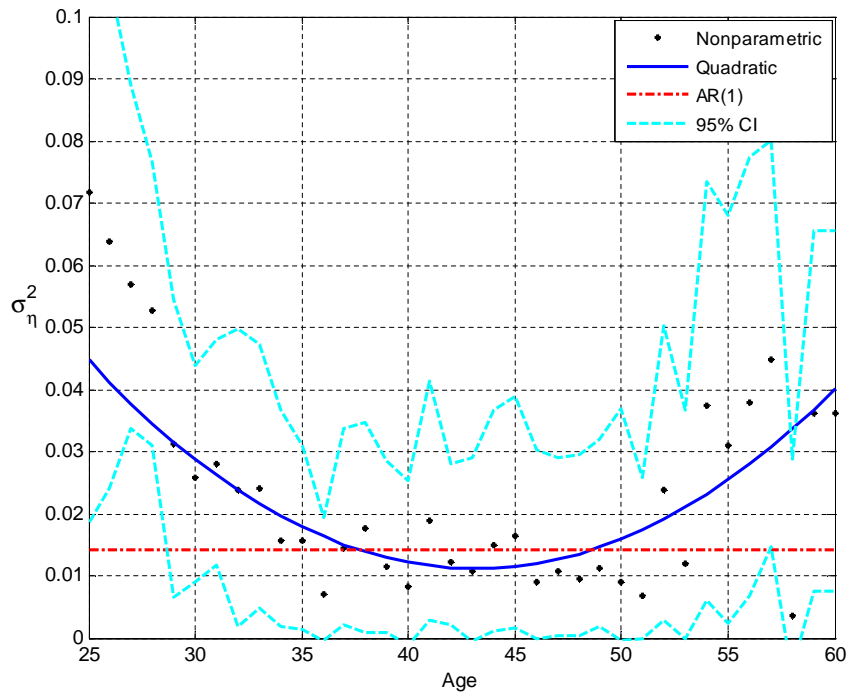
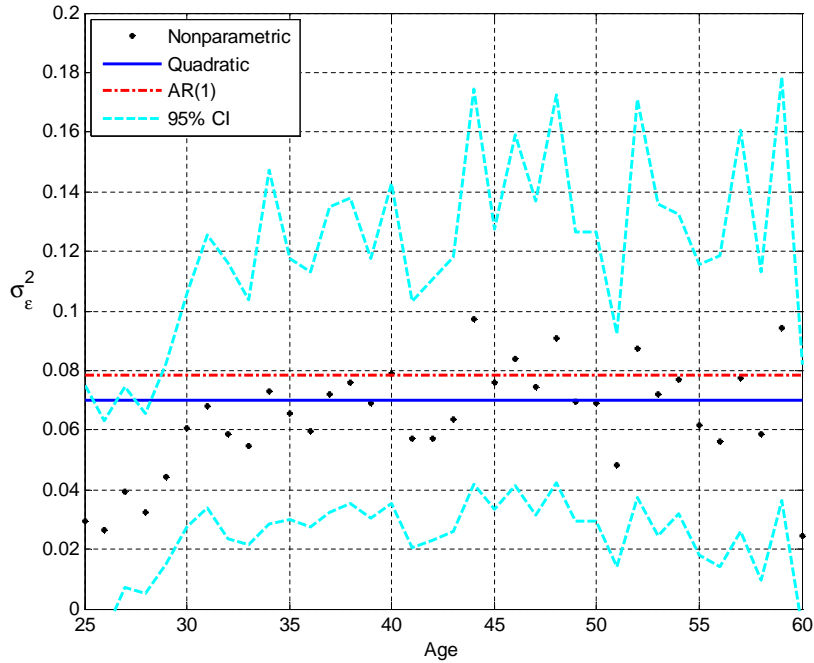


Figure 4 plots the variance of transitory shocks. Note that although there is a slight increase early on, it is not statistically significant. This is not very surprising, since the transitory component absorbs the classical measurement error, which we would expect to be flat. In what follows, we take the variance of the transitory component to be constant over the life cycle.

Figure 4: Variance Profile of Transitory Shocks




What features of the data give rise to this profile of persistence? In other words, we want to learn what moments in the data identify the increase in persistence early in the lifecycle. For this, we refer to the identification argument presented in Section 2.1, where we argued that the ratio of 2-period ahead covariance to 1-period ahead covariance, corrected for fixed effects, yields a consistent estimate for the persistence parameter.¹² The need to correct for the fixed effect arises because both of these covariance terms contain the variance of the fixed effects. We now plot the empirical counterpart of this ratio in Figure 5.

In correcting for fixed effects, we use our baseline estimate ($\sigma_\alpha^2 = 0.08$), which is in line with the estimates in the literature. The solid line plots the moving average of the

¹²Recall that (3) implies $[cov(\tilde{y}_h^i, \tilde{y}_{h+2}^i) - \sigma_\alpha^2] / [cov(\tilde{y}_h^i, \tilde{y}_{h+1}^i) - \sigma_\alpha^2] = \rho_{h+1}$ for $h = 1, \dots, H - 2$.

ratio over the lifetime. The shape of the ratio closely resembles our estimate of persistence profile (shown in dots on the same figure): It increases from 0.78 to 0.94, paralleling our estimation results in Figure 2. In general, the shape of this ratio depends on the level of fixed effects. To check for the robustness of this, we plot the ratio for the case where there are no fixed effects ($\sigma_\alpha^2 = 0$), which is shown in dashed lines. We see that the increase in persistence is robust to the variance of fixed effects, though the steepness depends on it. Note that our estimation of an upward sloping persistence profile is a result of targeting a fairly complicated variance-covariance structure. The finding in Figure 5 confirms this increase over the lifetime from a much simpler look at the data.

Figure 5: Identification of Persistence



2.5 Comparison with the Literature

We now compare the age-dependent process with the benchmark specification, i.e., a specification consisting of a fixed effect, an $AR(1)$ component where the persistence and variance of shocks are constant throughout life, and an i.i.d. transitory component with constant variance. In order for these cases to be comparable, we estimate this model using our data. The dashed lines on Figures 2-4 show the point estimates for persistence, and variance of persistent and transitory shocks. Our estimate of persistence, 0.978, is in line with the

estimates in the literature, which range from 0.96-1.0. It is surprising to see that for most of the life cycle, persistence in the age dependent process is significantly lower than the estimate of persistence for the benchmark case. As the examples above have shown, these differences can be economically significant. We will make this point clear in Section 3.

In what follows, we will argue that targeting the lifetime profile of inequality in the data results in an upward bias in persistence if one does not allow for age-specific persistence and variance. To do so, we compute the lifetime profile of inequality from the data. To control for time effects in variances, we compute the variance of residuals for each age-year bin, $\widehat{var}(\tilde{y}_{h,t})$. We then regress these on a full set of age and year dummies and report age dummies.¹³ The resulting profile is shown in Figure 2.5.¹⁴

This figure shows a steady rise in inequality of around 20 log points. The increase is particularly steep after age 35. For the benchmark process, the corresponding theoretical moments are given by

$$var(\tilde{y}_h) = \sigma_\alpha^2 + \sigma_\eta^2 \sum_{j=0}^{h-1} \rho^{2j} + \sigma_{z_0}^2 \rho^{2h} + \sigma_\epsilon^2,$$

where $\sigma_{z_0}^2$ represents the initial variance of the persistent component. So long as $\rho < 1$, residual inequality has a well-defined limit, say, $var^*(\tilde{y}_h)$. It can easily be shown that $var(\tilde{y}_h)$ will converge to $var^*(\tilde{y}_h)$ from below in a concave fashion.¹⁵ The degree of concavity is more pronounced the farther away ρ is from unity. In the case of a unit root, the variance profile will be linearly increasing, regardless of $var^*(\tilde{y}_h)$. Figure 2.5 obviously implies that the fit would be poor if ρ is far away from 1. Targeting these moments results in an upward bias and drives ρ close to 1 because the statistical model is misspecified.

At this point, it is worth stressing that the age-dependent income process does not need to contain unit roots or very high levels of persistence to match the inequality profile. Figure 2.5 also plots the smoothed inequality profile implied by our estimates. The model captures the increase in lifetime inequality even if persistence for young individuals is very low. The


¹³In order not to have too few individuals contributing to these variances, we include an individual in an age-year bin if he is within 2 years of that age.

¹⁴Some papers choose to control for cohort effects rather than time effects when reporting lifetime profile of inequality. We have decided to control for time effects for the sake of consistency, since the estimation controls for time effects.

¹⁵Here we implicitly assume that $var(\tilde{y}_0) < var^*(\tilde{y}_h)$, which is necessary to have an increasing lifetime profile.

mechanism is due to the inverse relationship between persistence and the variance of labor income shocks. When persistence goes up with age, the additional increase it induces in inequality is compensated by a decrease in the variance and vice versa. In this manner, the model is able to replicate the increase in the empirical variance profile with lower levels of persistence.

Figure 6: Lifetime Profile of Residual Inequality



[Guvenen \(2009\)](#) estimates a process that allows growth rates of earnings to differ across individuals. He finds support for significant heterogeneity in income growth rates and shows that ignoring this heterogeneity introduces an upward bias for the estimate of persistence. This paper shows that even if one takes the alternative view that agents are subject to similar income profiles, accounting for age-specific persistence and variances reduces the estimates of persistence significantly.

The evidence he brings forward for growth rate heterogeneity is twofold: First, he points to the convexity in the variance profile of earnings and argues that this feature of the data indicates the presence of growth rate heterogeneity. Second, he exploits the shape of higher order covariances, which features an increase in higher lags. This, he argues, can be captured through growth rate heterogeneity but not by highly persistent shocks. It is worthwhile to note that the age-dependent income process can naturally capture these features of the data without growth rate heterogeneity.

As we mentioned in [1.1](#), [Meghir and Pistaferri \(2004\)](#) also allow for age effects while modeling conditional variances of transitory and persistent shocks, which are found to be

insignificant. Since their specification assumes fully permanent shocks, i.e., persistence is constant at unity, it rules out the inverse relationship between variance and persistence that is crucial in our results. A flat profile in persistence suppresses the nontrivial lifetime profile in the variance of persistent shocks.

2.6 Significance Tests

We now turn to the question of statistical significance. Rather than making age-by-age comparisons using our nonparametric estimates, we want to see whether there is a significant pattern that is not flat. For this purpose, we proceed in two ways. First, we conjecture a quadratic function for the age profiles of the persistence and variance of persistent shocks and estimate its parameters from the data. This assumes that life-cycle effects are smooth in age. Yet, time effects are modeled nonparametrically; i.e., there are separate loading factors for each year. More specifically, we estimate

$$x_h = \gamma_{x,0} + \gamma_{x,1}h + \gamma_{x,2}h^2,$$

where x is the variable of interest, such as ρ and σ_η^2 . The quadratic polynomial is flexible enough to capture the profiles shown in Figures 2 and 3. We then test the hypothesis that the age pattern is flat. For each test, we compute the p-value as the fraction of the bootstrap runs for which the null hypothesis is violated. The results of the estimation and the test are presented in Table 1. The implied age profiles of the persistence and variance of shocks are plotted in Figures 2 and 3. Note that these line up well with the nonparametric estimates.¹⁶

The first row of Table 1 shows the results for persistence. The first three columns report point estimates along with bootstrap standard errors in parentheses. In the last two columns, we investigate the statistical significance of these coefficients. Column 5 (Column 6) tests if the linear (quadratic) term is significantly positive (negative). We find that in all bootstrap runs the linear (quadratic) term is positive (negative). The same analysis for the variance of persistent shocks, reported in the second row, shows that the linear and quadratic terms are significant at a 99% confidence level as well. Thus, based on the polynomial estimation, we reject that these profiles are constant over the lifecycle.

¹⁶As explained before, we assume a constant profile of variance for transitory innovations.

Table 1: Estimation and Test Results for Quadratic Specification

x	$\gamma_{x,0}$	$\gamma_{x,1}$	$\gamma_{x,2}$	Test 1	Test 2
ρ	0.7638 (0.0523)	0.0190 (0.0045)	-0.0004 (0.0001)	$H_0 : \gamma_{\rho,1} \leq 0$ 0.0000	$H_0 : \gamma_{\rho,2} \geq 0$ 0.0000
σ_η^2	0.0525 (0.0096)	-0.0041 (0.0009)	0.0001 (0.0000)	$H_0 : \gamma_{\sigma_\eta^2,1} \geq 0$ 0.0000	$H_0 : \gamma_{\sigma_\eta^2,2} \leq 0$ 0.0000
σ_α^2	0.0833 (0.0234)				
σ_ϵ^2	0.0702 (0.0233)				

* The numbers in brackets are bootstrap standard errors.

** The last three columns report the P-values for the corresponding test.

In order to complete the picture, we choose a specification that is in between the polynomial and the nonparametric specifications. We consider a model in which working life is divided into 3 stages. This model restricts the persistence and variance to be constant within an age interval but allows them to differ from one to the other. The bins correspond to ages 24-33, 34-52 and 53-60. More specifically, for $x = \rho, \sigma_\eta^2$:

$$x_h = \begin{cases} \delta_{x,1} & \text{if } h \in [24, 33] \\ \delta_{x,2} & \text{if } h \in [34, 52] \\ \delta_{x,3} & \text{if } h \in [53, 60] \end{cases}$$

These intervals give flexibility to the model in capturing arbitrary changes in parameters over the life cycle without disrupting the parsimonious structure. Furthermore, we do not want to bias the results by imposing a misspecified functional form. Time effects are still modeled nonparametrically. Figure 7 provides estimation results for this case along with 95% confidence intervals. The results, once again, point to the same profile over the lifecycle. The variance of persistent shocks follows a U-shape and the persistence is hump shaped. Confidence intervals show that persistence in the second age bin is significantly larger than in the first one. The difference in persistence between the second and third bins is, however, not significant. As for the variance, the second bin has a significantly lower variance than the other two bins. To be more formal, we test the hypotheses $H_0 : \rho_1 \geq \rho_2$, $H_0 : \rho_2 \leq \rho_3$, $H_0 : \sigma_{\eta,1}^2 \leq \sigma_{\eta,2}^2$ and $H_0 : \sigma_{\eta,2}^2 \geq \sigma_{\eta,3}^2$. The results are summarized in Table 2.

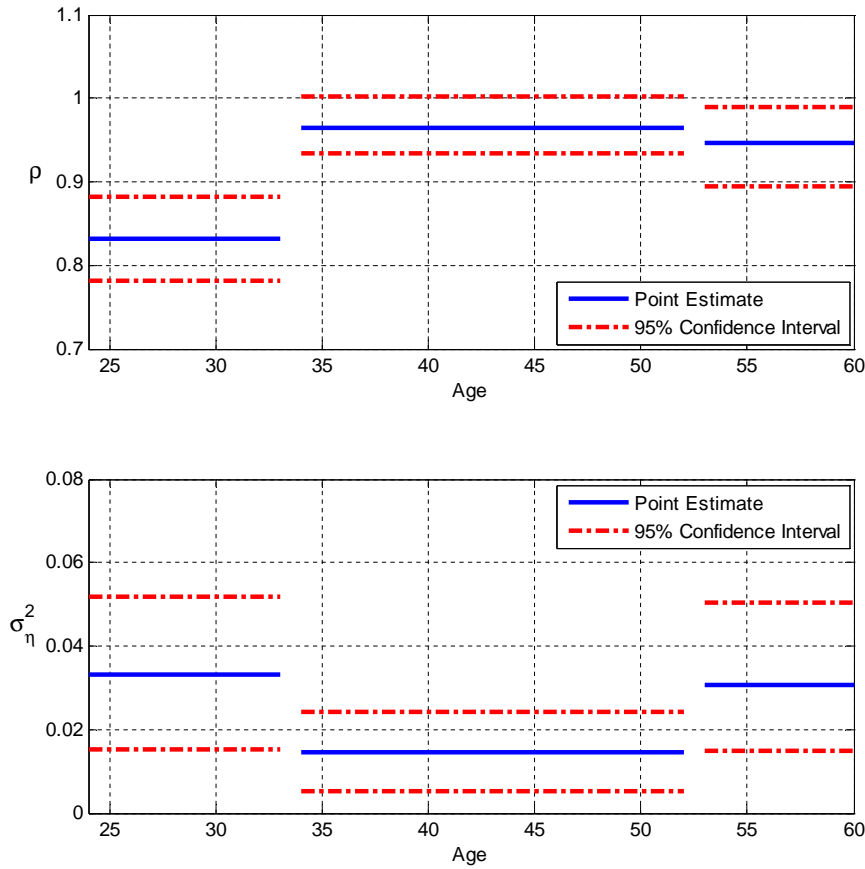
Table 2: Estimation and Test Results for Age Bins

	$\delta_{x,1}$	$\delta_{x,2}$	$\delta_{x,3}$	Test 1	Test 2
ρ	0.8326 (0.0266)	0.9648 (0.0170)	0.9458 (0.0265)	$H_0 : \rho_1 \geq \rho_2$ 0.0000	$H_0 : \rho_2 \leq \rho_3$ 0.2800
σ_η^2	0.0295 (0.0075)	0.0128 (0.0039)	0.0273 (0.0071)	$H_0 : \sigma_{\eta,1}^2 \leq \sigma_{\eta,2}^2$ 0.0000	$H_0 : \sigma_{\eta,2}^2 \geq \sigma_{\eta,3}^2$ 0.0120
σ_α^2	0.0956 (0.0102)				
σ_ϵ^2	0.0732 (0.0194)				

* The numbers in brackets are standard errors.

** The last three columns report the p-values of the corresponding tests.

Figure 7: Results for Age Bins



We find that the persistence in the first stage is statistically smaller than the second stage. However, we cannot reject the hypothesis that the persistence in the last stage is different than that in the second. For variance, the second bin is significantly lower than the first and third.

Both the analysis of the polynomial and the age-bin specifications provides strong evidence that these profiles are significant.

2.7 The Fit for Income Growth Rates

The previous sections have illustrated how the age-dependent process does a better job in fitting the variance-covariance structure using log earnings (levels). This is expected since the estimation targeted the moments in levels with a larger number of parameters. How about the fit for the variance structure of income *growth rates* (differences)? Is the fit for levels better at the expense of a worse fit for income growth rates? It is well known in the literature that the estimates of canonical income processes using levels are strikingly different than the estimates using income growth rates, suggesting misspecification of the model (Krueger, Perri, Pistaferri, and Violante (2010)). This section investigates this aspect of the age-dependent process.


The variance structure of growth rates (abstracting from time effects) is as follows:

$$\text{var}(\Delta y_{i,h}) = (\rho_{h-1} - 1)^2 \text{var}(z_{i,h-1}) + \sigma_{\eta,h}^2 + \sigma_{\epsilon,h}^2 + \sigma_{\epsilon,h-1}^2$$

The *-marked series in Figure 8 plots the variance profile of income growth rates in the data. This reveals a U-shaped profile. In order to evaluate the performance of the age-specific income process, we calculate and plot $\text{var}(\Delta y_{i,h})$ using our estimates from the polynomial specification (dashed line). Similarly, we plot the corresponding series implied by a standard $AR(1)$ process (solid line). This figure clearly replicates the misspecification we discussed above: Both the age-dependent and the $AR(1)$ processes are far from matching the *level* of variances. However, it is worthwhile to note that the age-dependent process can capture the U-shaped profile. This is facilitated by the U-shape in the variance profile of persistent shocks, $\sigma_{\eta,h}^2$.

The age-dependent income process achieves a better fit for the moments in levels without worsening the fit for the moment structure in differences. Although it cannot match the magnitude of the variance of income growth rates in the data, it can replicate the age profile.

Figure 8: Variance Profile of Income Growth Rates



Further Remarks

A natural follow-up question is how to explain these profiles economically. Which economic forces give rise to these? To speculate about one mechanism, these profiles could be due to differences in insurance opportunities against earnings shocks between young and old workers. For example, in case of an adverse demand shock to individual’s occupation, one might switch to a different one if she is young. For an old worker, though, switching is costlier (e.g. because of occupation-specific human capital). Therefore, shocks of the same nature can translate into innovations with different persistence over the working life.

Another mechanism, again related to mobility, is learning about the match quality, first studied by [Jovanovic \(1979\)](#). In his setup, neither the worker nor the firm know the productivity of the match before employment. After observing the output, match productivity is revealed to both parties in a Bayesian fashion. This generates endogenous movements in wages and job turnover. In [Appendix C](#) we show that a very stylized version of this model implies an increasing persistence profile and a decreasing variance over the working life. Since this type of models are shown to have empirical support ([Flinn \(1986\)](#)), we view this as an additional evidence for our findings.

3 Consumption-Savings Implications

There is a large literature that rejects full insurance for the US economy (Cochrane (1991), Mace (1991), Attanasio and Davis (1996)) making the nature of labor income risk an important object to study. This paper so far has established the existence of a nonflat lifetime profile in persistence and variance of shocks. We now investigate its economic implications. In particular, we are interested in the insurability of labor income shocks and the welfare costs of idiosyncratic risk under different specifications for earnings. To address these issues, we consider a standard life-cycle model that features incomplete financial markets and a social security system and compare the implications of the age-dependent income process with a standard $AR(1)$ process. There are several reasons to expect different consequences for welfare costs and consumption insurance. First, as we have discussed above, the age-dependent income process implies lower persistence but larger shocks for young agents. Kaplan and Violante (2008) show that for reasonably calibrated versions of a Bewley model, the insurability of shocks is decreasing in persistence. Therefore, one might expect a higher level of insurance for young agents under the age-dependent income process than under the standard process. This will imply lower welfare costs of risk compared to the benchmark case. On the other hand, shocks to earnings are larger for young agents, which in turn results in larger welfare costs. Ultimately, whether welfare costs are larger or smaller becomes a quantitative question.

We now describe the model that we use to study the question. The economy is populated by a continuum of agents that have preferences over consumption that are ordered according to

$$E \sum_{h=1}^H \beta^h u(c_h^i) \quad (7)$$

where c_h^i denotes the consumption of agent i at age h . They engage in labor market activities for the first R years of their life and retire afterward. After retirement, they live up to a maximum age of H .

Financial markets are incomplete in that agents can buy and sell only a risk-free bond. Letting r denote the risk-free interest rate and a_h^i denote the asset level of individual i of age h , the budget constraint is given by

$$c_h^i + \frac{a_{h+1}^i}{1+r} = a_h^i + y_h^i, \quad (8)$$

where y_h^i is the labor earnings at age h . Agents are allowed to borrow up to an age-dependent level, denoted by \bar{A}_h . We assume that everyone of the same age faces the same borrowing limit and we experiment with two extreme cases: a natural borrowing limit and a zero borrowing limit.¹⁷ It is important to investigate these two cases for the question we have in mind, because the evaluation of the tradeoff between persistence and variance of shocks depends crucially on the extent of the borrowing limit. Namely, if borrowing limits are loose, the not-so-persistent but large shocks to young agents can be well insured by borrowing. On the other hand, in case of tight borrowing limits, the magnitude of shocks matters more.

While in the labor market, agents' earnings have two components. The deterministic part is common to everyone and follows a quadratic polynomial in age. The idiosyncratic component captures individual earnings risk and is modeled as discussed in 2.1:

$$\begin{aligned} \ln y_h^i &= \gamma_0 + \gamma_1 h + \gamma_2 h^2 + \tilde{y}_h^i \\ \tilde{y}_h^i &\sim (1) \end{aligned} \tag{9}$$

We consider the implications of two specifications for the income process: i) the age-dependent income process and ii) an $AR(1)$ process with constant persistence and variance of shocks over the lifetime: $\rho_h = \rho$, $\sigma_h^2 = \sigma^2 \forall h$.

There is a social security system that pays a pension after retirement.¹⁸ We model the retirement salary as a function of the fixed effect and the persistent component of income in the last period, $\ln y_h^i = \Phi(\alpha^i, z_R^i)$. This function is modeled as in [Guvener, Kuruscu, and Ozkan \(2009\)](#) and is set to mimic the properties of the US social security system. Its details are discussed in 3.1.

Let $V_h(a_h^i, \alpha^i, z_h^i, \varepsilon_h^i)$ denote the value function of an agent at age $h \leq R$, with asset holdings a_h^i , fixed effect α^i , persistent component of labor income z_h^i and transitory component of income ε_h^i . The agent's programming problem can be written recursively as

¹⁷The natural borrowing limit is the maximum amount that an agent can pay back with future earnings for sure.

¹⁸Since this is a partial equilibrium framework, we do not model social security taxes and do not consider the government's budget.

$$\begin{aligned}
V_h^i(a_h^i, \alpha^i, z_h^i, \varepsilon_h^i) &= \max_{a_{h+1}^i, c_h^i} u(c_h^i) + \beta EV_{h+1}(a_{h+1}^i, \alpha^i, z_{h+1}^i, \varepsilon_{h+1}^i) \\
s.t. & \quad (8) \text{ and } (9) \\
a_{h+1}^i &\geq -\bar{A}_{t+1}
\end{aligned}$$

Upon retirement, the agent has a constant stream of income from social security and faces no risk. His problem is given by:

$$\begin{aligned}
V_h^i(a_h^i, \alpha^i, z_R^i) &= \max_{a_{h+1}^i, c_h^i} u(c_h^i) + \beta V_{h+1}(a_{h+1}^i, \alpha^i, z_R^i) \\
s.t. & \quad (8) \\
& \quad \ln y_h^i = \Phi(\alpha^i, z_R^i) \\
& \quad a_{h+1}^i \geq -A_{h+1}^-
\end{aligned}$$

3.1 Calibration

One period in our model corresponds to a calendar year. Agents enter the economy at age 24, retire at 60 and are dead by age 84. We assume CRRA preferences and set the parameter of relative risk aversion to 2.¹⁹ We take the risk-free interest rate to be 3%.

As suggested by [Storesletten, Telmer, and Yaron \(2004\)](#), among others, the crucial part of our calibration is to pin down the discount factor β . We set this parameter to match an aggregate wealth to income ratio of 3. This is important, since the amount of wealth held by individuals affects the insurability and welfare costs of labor income shocks. We define aggregate wealth as the sum of positive asset holdings. Aggregate income is the sum of labor earnings (excluding retirement pension).

The deterministic component of earnings is estimated using the PSID data. It has a hump-shaped profile where earnings grow by 60% during the first 25 years and then decrease by 18% until the end of the working life. For the residual component of earnings, we consider two specifications: the age-dependent and the $AR(1)$ processes. The first is calibrated according to the quadratic specification reported in [Table 1](#). The parameters of the latter come from our estimates in [Figures 2-4](#).

¹⁹This is within the range of estimates in the literature ([Gourinchas and Parker \(2002\)](#), [Cagetti \(2003\)](#)).

In a realistic model of the retirement system, a pension would be a function of lifetime average earnings, but this would introduce one more continuous state variable to the problem of the household. We refrain from doing so, since this would complicate the model without adding any further insight for our purposes. In our model, the retirement pension is a function of predicted average lifetime earnings. We first regress average lifetime earnings on last period’s earnings net of the transitory component and use the coefficients to predict an individual’s average lifetime earnings, denoted by $\hat{y}_{LT}(\alpha^i, z_R^i)$. Following [Güvenen, Kuruscu, and Ozkan \(2009\)](#) we use the following pension schedule:

$$\Phi(\alpha^i, z_R^i) = a * AE + b * \hat{y}_{LT}(\alpha^i, z_R^i),$$

where AE is the average earnings in the population. The first term is the same for everyone and captures the insurance aspect of the system. The second term is proportional to \hat{y}_{LT} and governs the private returns to lifetime earnings. We set $a = 16.78\%$, and $b = 35.46\%$.

We discretize all three components of earnings using 61, 11, and 11 grid points for the persistent component, transitory component, and fixed effect, respectively. The value function and policy rules are solved using standard techniques on an exponentially spaced grid for assets of size 100. The economy is simulated with 50,000 individuals.²⁰

3.2 Simulation Results

In this section, we report the differences in consumption behavior induced by the age-dependent and the $AR(1)$ processes. For every specification, we calibrate the discounting factor, β , to match an aggregate wealth to income ratio of 3. We first start by showing the results for the economy with natural borrowing constraints (NBC). The resulting discount factors for the age-dependent and $AR(1)$ processes are $1/(1 + 0.041)$ and $1/(1 + 0.042)$, respectively (see [Table 3](#)). [Figure 9](#) shows mean asset and consumption profiles. Note that the asset and consumption profiles are very similar for both specifications.²¹ However, even

²⁰The number of grids for the income process is sufficient, since simulated earnings are very close to theoretical earnings. We find that increasing the grid for assets does not change Euler errors significantly. Also, increasing the number of people we simulate does not change the model statistics. We conclude that the current precision is sufficient.

²¹The model is able to generate a hump-shaped profile for consumption, as reported in [Krueger and Fernandez-Villaverde \(2009\)](#), but the timing of the hump is later. This fit can be improved by adding mortality risk or health shocks in older ages ([Palumbo \(1999\)](#)).

though agents are more impatient in the $AR(1)$ economy, consumption growth of young individuals is steeper. This points to the differences in precautionary motives (Carroll (1997)).

Figure 9: Mean Asset and Consumption Profiles for NBC




Figure 10: Consumption Inequality for NBC




Figure 10 shows the inequality profiles of consumption implied by the two income processes. Recall from Figure 2.5 that the initial level of earnings inequality is lower for the $AR(1)$ process, but that the increases over the lifetime are roughly equal.²² Thus, we focus on the increase in consumption inequality rather than levels: The increase implied by the $AR(1)$ process is 21 log points, whereas the age-dependent income process implies a rise of only 17 log points. This shows that the shocks in the age-dependent process economy are more insurable.

To make this point clearer, we provide a measure of insurance against persistent shocks and investigate the differences between the two processes. Following Kaplan and Violante (2008) and Blundell, Pistaferri, and Preston (2008), we compute the degree of consumption insurance at age h as:

$$\phi_h = 1 - \frac{\text{cov}(\Delta c_h^i, \eta_h^i)}{\text{var}(\eta_h^i)},$$


where η_h^i is the persistent shock faced by worker i at age h . This measures the amount of change in earnings that does not translate into consumption growth. Figure 11 plots ϕ_h over the life cycle for both processes. It is obvious that persistent shocks from the age-dependent process are better insured throughout the lifetime. On average, 56% of persistent shocks are insured under the age-dependent process, whereas the corresponding number for the $AR(1)$ process is only 40%. Strikingly, most of this difference comes from younger adults. Recall that for them the level of persistence is particularly low under the age-dependent process. It is well known in the literature that persistence is an important determinant of insurance. Transitory shocks are easily self-insured by using the risk-free bond (Kaplan and Violante (2008)). On the other hand, in the presence of a very persistent component, agents abstain from borrowing because of the possibility of a long series of bad income states. Insurance against such shocks is, therefore, mostly through asset accumulation. This is not possible for young agents since, on average, they are poor. Under the age-dependent income process, persistence is fairly moderate for young workers, implying insurance coefficients as large as 70%.

Another striking difference between the two processes is the profiles of insurance coefficients. In the $AR(1)$ process, the profile of insurance tracks the profile of assets. This is consistent with the previous explanation, since persistence is constant and high throughout the working life and insurance mainly depends on the amount of assets. Blundell, Pistaferri,

²²The increase in the $AR(1)$ process is only 0.01 higher.

and Preston (2008) approximate insurance coefficients against permanent shocks in the data and find that this is roughly flat over the life cycle.²³ Thus, the profile of insurance implied by an $AR(1)$ process is not consistent with the data (Kaplan and Violante (2008)). The left panel of Figure 11 shows, however, that the age profile of insurance in a Bewley model need not track the profile of assets. Note that the profile of assets under the age-dependent process is very similar to the one under $AR(1)$, but the insurance profiles are drastically different. This is solely due to the profile of persistence. Young agents, as explained above, have access to better insurance since shocks are not very persistent. Insurance decreases with age in the early part of the working life, since persistence increases. After age 40, on the other hand, agents have enough assets so that the change in persistence has virtually no effect on the profile of insurance and thus insurance increases due to the increase in assets.

Figure 11: Insurance against Persistent Shocks



²³They develop an approximation to insurance coefficients in a life-cycle model assuming that residual earnings consist of a completely permanent and a fully transitory component and that there are no borrowing constraints.

3.3 Welfare Costs of Earnings Risk

We now turn to welfare costs of idiosyncratic risk under the two processes. Recall that the low levels of persistence under the age-dependent process is compensated by the larger variance of shocks (Figures 2 and 3). On the one hand, lower persistence implies better insurability. On the other hand, larger variance implies more instability. In order to evaluate this tradeoff quantitatively, we compute the fraction of lifetime consumption that an individual would be willing to give up in order to live in an economy without earnings risk.²⁴ The results are reported in Table 3.

Table 3: Welfare Costs under Different Income Processes

Natural Borrowing Limit					
	(1)	(2)	(3)	(4)	(5)
	β	$\frac{wealth}{income}$	Shocks+Fixed	Shocks	Insurance
Age-Dependent	$1/(1 + 0.0410)$	2.9994	15.73%	9.97%	0.56
AR(1)	$1/(1 + 0.0420)$	3.0001	16.71%	14.85%	0.40
Experiment 1	$1/(1 + 0.0414)$	2.9995	19.06%	13.51%	0.39
Experiment 2	$1/(1 + 0.0418)$	2.9994	19.08%	13.55%	0.41

Tight Borrowing Limit					
	(1)	(2)	(3)	(4)	(5)
	β	$\frac{wealth}{income}$	Shocks+Fixed	Shocks	Insurance
Age-Dependent	$1/(1 + 0.0561)$	3.0009	18.84%	12.53%	0.39
AR(1)	$1/(1 + 0.0562)$	3.0008	18.51%	16.37%	0.31
Experiment 1	$1/(1 + 0.0549)$	3.0013	20.83%	14.72%	0.30
Experiment 2	$1/(1 + 0.0558)$	3.0009	21.01%	14.90%	0.31

Column 3 shows the welfare costs of not being able to insure against idiosyncratic risk as well as fixed effects. The first two rows correspond to the age-dependent and $AR(1)$

²⁴The formula for welfare costs, χ , is given by

$$\chi = 1 - \left(\frac{V}{V_{Complete}} \right)^{1/(1-\gamma)},$$

where V is the expected lifetime utility in the economy for which welfare costs are calculated, $V_{Complete}$ is the expected lifetime utility in the complete markets economy and γ is the coefficient of relative risk aversion in the $CRRA$ utility function ($\gamma = 2$).

processes, respectively. The age-dependent income process delivers lower welfare costs, even though the level of inequality at the end of the life cycle is lower for the $AR(1)$ process (see Figure 2.5). At this point it is not clear how much of these differences is driven by shocks and how much is driven by permanent differences. In order to properly account for the costs of shocks, we compute the welfare cost of idiosyncratic shocks only.²⁵ These are reported in Column 4. The differences between welfare costs are now even larger: An agent living in the $AR(1)$ world is willing to give up 15% of her consumption every period in order to have perfect insurance. The same number is only 10% for an agent in the age-dependent world. We conclude that the effect of lower persistence dominates the effect of larger instability. These are sizable differences.

There is a caveat in this analysis: The increase in earnings inequality over the working life is slightly higher in the $AR(1)$ process (0.1997 vs. 0.1863). Also, the level of inequality at the beginning of life is lower for the $AR(1)$ process. In order to correct for these, we modify the parameters of the $AR(1)$ process such that the inequality at the beginning and the end of the lifetime is the same for both processes. More specifically, we adjust the variance of the fixed effect in order to match the inequality at the beginning of the lifecycle. To match the increase we do the following two experiments: First, we keep the persistence the same but decrease the variance of persistent shocks from 0.0143 to 0.0129. Second, we keep the variance the same but decrease persistence from 0.978 to 0.9747. The last two rows in the top panel report the results for these experiments, respectively.

Note that the results for both experiments are very close. Since we increased the variance of fixed effects, the overall costs of inequality increased compared to the second row (from 16.7% to 19.1%). In addition, since the increase in inequality over the lifetime is now lower, the welfare costs of shocks are lower, too. However, they are still substantially larger than the welfare costs under the age-dependent specification. The difference in welfare costs almost corresponds to 4% of lifetime consumption.

As explained above, the driving force for welfare differences is the insurability of earnings shocks. The fact that the age-dependent income process results in larger insurance coefficients relies crucially on the extent of borrowing limits. Young agents would have little ability to insure even against moderately persistent shocks if they cannot borrow freely.

²⁵We follow [Storesletten, Telmer, and Yaron \(2004\)](#) and ask how much an agent with the average fixed effect would be willing to give up in order to live in the economy with complete financial markets.

In other words, the evaluation of the tradeoff between durability and magnitude might reverse. In order to quantify how much it matters, we take it to the extreme and redo the same analysis for an economy where there is no borrowing at all.²⁶ The bottom panel of Table 3 presents the results.

The last column reveals that, as expected, insurance goes down by a significant amount. The right panel of Figure 11 plots the lifetime profile of insurance coefficients for the ZBC economy. Note that the difference between the age-dependent and $AR(1)$ processes is significantly smaller compared to the NBC economy. The difference between the NBC and ZBC economies is substantial for young individuals, for whom insurance falls from around 70% to 26%. The main mechanism of insurance for young agents under the age-dependent process is borrowing. Since this is not allowed in the ZBC economy, insurance goes down significantly.

The decrease in the degree of insurance will have welfare consequences. Column 4 on the bottom panel of Table 3 shows the welfare costs of idiosyncratic risk for the ZBC economy. As expected, welfare costs have increased compared to the NBC economy for both of the specifications. Note that the increase is larger for the age-dependent process, and thus, the differences between the two processes are now lower. However, it is still the case that welfare costs are lower for the age-dependent process. These results hold also with the experiments explained above. We conclude that the evaluation of welfare costs is substantially different for the two processes; however, the margin depends on the amount of borrowing allowed.²⁷

4 Conclusion

In the presence of incomplete financial markets, the nature of labor income risk becomes an important determinant of individual decision making. In this paper, we have proposed a novel specification for labor income risk that allows the persistence and variance of shocks

²⁶For the case with tight borrowing constraints, the complete markets economy in the welfare calculations is the one with full insurance against income risk but with no borrowing.

²⁷Our findings have implications for the Credit CARD Act of 2009. One of the provisions of this act restricts individuals under the age of 21 from obtaining credit cards without the consent of their parents. If shocks were completely permanent, then access to credit would be less crucial since they would not use the option of borrowing. This paper presented evidence that young agents face very large variances of income shocks that are moderately persistent. As discussed above, the borrowing limit for young individuals have significant welfare consequences under such an income process. Thus using credit lines in this environment can go a long way as an insurance mechanism making access to credit crucial for young individuals.

to change over the lifetime and estimated it using data from the PSID. We have found evidence for a nonflat profile in the persistence and variance of persistent shocks, but not in transitory shocks. Our results reveal that persistence follows a hump shape over the working life: It starts at 0.75, increases up to unity by age 40 and then slightly decreases to around 0.95. On the other hand, the variance of persistent shocks exhibits a U-shaped profile (with a minimum of 0.01 and a maximum of 0.045).

We have investigated the implications of these profiles for consumption and savings behavior with a life-cycle model. We have found that under natural borrowing constraints, the welfare costs of idiosyncratic risk implied by the age-dependent income process is significantly lower compared to a standard $AR(1)$ process. This is mostly due to a higher degree of consumption insurance for young workers, for whom persistence is low. Namely, the low level of persistence allows agents to insure themselves against persistent shocks by borrowing. This mechanism relies crucially on the extent of borrowing limits. In order to quantify the effect of borrowing limits, we have studied an economy with no borrowing. The results are qualitatively the same, although the difference between specifications in the ZBC economy is smaller. We conclude that the welfare cost of idiosyncratic risk will be overstated if one does not account for the age profiles in the persistence and variance of shocks.

Our findings have implications for the Credit CARD Act of 2009. One of the provisions of this act restricts young individuals from obtaining credit cards. According to this paper, young agents face very large variances of income shocks that are moderately persistent. This makes access to credit crucial for them.

The benefits of public insurance policies are commonly based on the gains from redistribution, which can be proxied by the welfare costs of inequality. This paper presented evidence that once the researcher accounts for the age-dependent nature of labor income risk, welfare costs are much smaller.

References

- ABOWD, J. M., AND D. CARD (1989): “On the Covariance Structure of Earnings and Hours Changes,” *Econometrica*, 57(2), 411–45. 5
- ALTONJI, J. G., AND L. M. SEGAL (1996): “Small-Sample Bias in GMM Estimation of Covariance Structures,” *Journal of Business & Economic Statistics*, 14(3), 353–66. 14
- ALTONJI, J. G., A. SMITH, AND I. VIDANGOS (2009): “Modeling Earnings Dynamics,” *NBER Working Paper No: 14743*. 6
- ALVAREZ, J., M. BROWNING, AND M. EJRNÆS (2006): “Modelling Income Processes with Lots of Heterogeneity,” *Working Paper*. 6
- ATTANASIO, O., AND S. J. DAVIS (1996): “Relative Wage Movements and the Distribution of Consumption,” *Journal of Political Economy*, 104(6), 1227–62. 25
- BAKER, M. (1997): “Growth-Rate Heterogeneity and the Covariance Structure of Life-Cycle Earnings,” *Journal of Labor Economics*, 15, No. 2, 338–375. 5
- BLUNDELL, R., L. PISTAFERRI, AND I. PRESTON (2008): “Consumption Inequality and Partial Insurance,” *American Economic Review*, 98(5), 1887–1921. 4, 6, 7, 30
- 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*, 9(1), 1–24. 12
- CAGETTI, M. (2003): “Wealth Accumulation over the Life Cycle and Precautionary Savings,” *Journal of Business and Economic Statistics*, 21, No. 3, 339–353. 27
- CARROLL, C. D. (1997): “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *The Quarterly Journal of Economics*, 112(1), 1–55. 29
- COCHRANE, J. H. (1991): “A Simple Test of Consumption Insurance,” *Journal of Political Economy*, 99(5), 957–76. 25
- CUNHA, F., J. HECKMAN, AND S. NAVARRO (2004): “Separating Uncertainty from Heterogeneity in Life Cycle Earnings,” IZA Discussion Papers 1437, Institute for the Study of Labor (IZA). 6

- FEIGENBAUM, J., AND G. LI (2008): “Lifecycle Dynamics of Income Uncertainty and Consumption,” *Finance and Economics Discussion Series, Washington: Board of Governors of the Federal Reserve System*, 27. [6](#)
- FLINN, C. J. (1986): “Wages and Job Mobility of Young Workers,” *Journal of Political Economy*, Vol. 94, No. 3, 88–110. [24](#), [48](#), [51](#)
- GOTTSCHALK, P., AND R. A. MOFFITT (1995): “Trends in the Transitory Variance of Earnings in the United States,” *Working Paper*. [11](#)
- GOURINCHAS, P.-O., AND J. A. PARKER (2002): “Consumption over the Life Cycle,” *Econometrica*, 70(1), 47–89. [4](#), [27](#)
- GUVENEN, F. (2007): “Learning Your Earnings: Are Labor Income Shocks Really Very Persistent?,” *American Economic Review*, 97, No.3, 687–712. [6](#)
- (2009): “An Empirical Investigation of Labor Income Processes,” *Review of Economic Dynamics*, 12, No.1, 58–79. [4](#), [6](#), [7](#), [19](#)
- GUVENEN, F., B. KURUSCU, AND S. OZKAN (2009): “Taxation of Human Capital and Wage Inequality: A Cross-Country Analysis,” *Working Paper*. [26](#), [28](#)
- GUVENEN, F., AND A. SMITH (2009): “Inferring Labor Income Risk from Economic Choices: An Indirect Inference Approach,” *Working Paper*. [6](#)
- HAUSE, J. C. (1980): “The Fine Structure of Earnings and the On-the-Job Training Hypothesis,” *Econometrica*, 48(4), 1013–1029. [6](#)
- HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2005): “Two Views of Inequality Over the Life Cycle,” *Journal of the European Economic Association*, 3(2-3), 765–775. [12](#)
- HRYSHKO, D. (2008): “RIP to HIP: The Data Reject Heterogeneous Labor Income Profiles,” Discussion paper, University of Alberta. [7](#)
- JOVANOVIC, B. (1979): “Job Matching and the Theory of Turnover,” *Journal of Political Economy*, Vol. 87, No. 5, Part 1, 972–990. [24](#), [48](#), [51](#)

- KAMBOUROV, G., AND I. MANOVSKII (2008): “Rising Occupational and Industry Mobility in the United States: 1968-1997,” *International Economic Review*, 49, 41–79. [2](#)
- KAPLAN, G., AND G. L. VIOLANTE (2008): “How Much Insurance in Bewley Models?,” Discussion paper, New York University. [4](#), [6](#), [25](#), [30](#), [31](#)
- KRUEGER, D., AND J. FERNANDEZ-VILLAYERDE (2009): “Consumption and Saving over the Life Cycle: How Important are Consumer Durables?,” *Working Paper, University of Pennsylvania*. [28](#)
- KRUEGER, D., F. PERRI, L. PISTAFERRI, AND G. L. VIOLANTE (2010): “Cross-sectional Facts for Macroeconomists,” *Review of Economic Dynamics*, 13, 1–14. [23](#)
- LILLARD, L. A., AND Y. WEISS (1979): “Components of Variation in Panel Earnings Data: American Scientists 1960-70,” *Econometrica*, 47(2), 437–454. [5](#)
- LILLARD, L. A., AND R. J. WILLIS (1978): “Dynamic Aspects of Earning Mobility,” *Econometrica*, Vol. 46, No.5, 985–1012. [5](#)
- MACE, B. J. (1991): “Full Insurance in the Presence of Aggregate Uncertainty,” *Journal of Political Economy*, 99, 928–956. [25](#)
- MACURDY, T. E. (1982): “The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis,” *Journal of Econometrics*, 18(1), 83–114. [5](#)
- MEGHIR, C., AND L. PISTAFERRI (2004): “Income Variance Dynamics and Heterogeneity,” *Econometrica*, 72(1), 1–32. [3](#), [5](#), [7](#), [12](#), [19](#)
- PALUMBO, M. G. (1999): “Uncertain Medical Expenses and Precautionary Saving Near the End of the Life Cycle,” *The Review of Economic Studies*, 66, No. 2, 395–421. [28](#)
- STORESLETTEN, K., C. I. TELMER, AND A. YARON (2004): “Consumption and Risk Sharing over the Life Cycle,” *Journal of Monetary Economics*, 51, 609–633. [6](#), [27](#), [33](#)

APPENDICES

A Identification

Here, we provide the proof of identification for the full model (6). Again, we will make use of the variance-covariance structure implied by this model. This structure is given by:

$$\text{var}(\tilde{y}_{h,t}^i) = \sigma_\alpha^2 + \text{var}(z_{h,t}^i) + \phi_t^2 \sigma_{\epsilon,h}^2 \quad (10)$$

$$\text{cov}(y_{h,t}^i, y_{h+n,t+n}^i) = \sigma_\alpha^2 + \rho_h \rho_{h+1} \cdots \rho_{h+n-1} \text{var}(z_{h,t}^i) \quad (11)$$

$$\text{var}(z_{h,t}^i) = \rho_{h-1}^2 \text{var}(z_{h-1,t-1}^i) + \pi_t^2 \sigma_{\eta,h}^2 \quad (12)$$

Proposition: The process in (6) is identified up to the normalizations that $\rho_1 = \rho_2$, $\pi_1 = \phi_1 = \phi_H = 1$ and $\sigma_{\eta,H}^2 = \sigma_{\eta,H-1}^2$.

Proof: The proof is very similar to the one for the simpler specification. We start by assuming that we know the variance of the fixed effect, σ_α^2 , and show that we can identify all the remaining parameters. Then we come back to argue that the unused moment conditions are enough to pin down σ_α^2 .

Note that since we assume that σ_α^2 is known, we can construct $\text{cov}(\tilde{y}_{h,t}^i, \tilde{y}_{h+n,t+n}^i) - \sigma_\alpha^2$. (11) implies $[\text{cov}(\tilde{y}_{h,t}^i, \tilde{y}_{h+2,t+2}^i) - \sigma_\alpha^2] / [\text{cov}(\tilde{y}_{h,t}^i, \tilde{y}_{h+1,t+1}^i) - \sigma_\alpha^2] = \rho_{h+1}$ for $h = 1, \dots, H - 2$. This pins down the whole profile of ρ_h for $h = 2, 3, \dots, H - 1$ except for ρ_H .²⁸ Note also that by normalization $\rho_1 = \rho_2$.

Now, our goal is to recover the schedule of $\text{var}(z_{h,t}^i)$. Once we recover these, we can use (12) to identify the loading factors and variances of persistent shocks, $\{\pi_t\}_{t=1}^{t=T}$ and $\{\sigma_{\eta,h}^2\}_{h=1}^{h=H-1}$. Note that

$$\frac{\text{cov}(\tilde{y}_{h,t}^i, \tilde{y}_{h+1,t+1}^i) - \sigma_\alpha^2}{\rho_h} = \text{var}(z_{h,t}^i) \quad (13)$$

Since ρ_h is pinned down for $h \geq 1$, (13) recovers $\text{var}(z_{h,t}^i)$ for $h = 1, \dots, H - 1$, $t = 1, \dots, T - 1$. Please note that $\text{var}(z_{H,t}^i)$ for $t = 1, \dots, T$ and $\text{var}(z_{h,T}^i)$ for $h = 1, \dots, H$ are not identified yet.

²⁸Note that ρ_H does not enter the variance-covariance profile at all, so it is, in fact, not a parameter of the model.

Note that all of the parameters recovered so far depend on σ_α^2 . It remains to be shown that the unused covariances uniquely pin this down. We now show that $\text{cov}(\tilde{y}_{2,1}^i, \tilde{y}_{5,4}^i)$ suffices to recover σ_α^2 uniquely:

$$\begin{aligned}
\text{cov}(\tilde{y}_{2,1}^i, \tilde{y}_{5,4}^i) &= \sigma_\alpha^2 + \rho_4 \rho_3 \rho_2 \text{var}(z_{2,1}^i) \\
&= \sigma_\alpha^2 + \rho_4 \rho_3 \rho_2 \left[\frac{\text{cov}(\tilde{y}_{2,1}^i, \tilde{y}_{3,2}^i) - \sigma_\alpha^2}{\rho_2} \right] \\
&= \sigma_\alpha^2 + \left[\frac{\text{cov}(\tilde{y}_{3,1}^i, \tilde{y}_{5,3}^i) - \sigma_\alpha^2}{\text{cov}(\tilde{y}_{3,1}^i, \tilde{y}_{4,2}^i) - \sigma_\alpha^2} \right] \left[\frac{\text{cov}(\tilde{y}_{2,1}^i, \tilde{y}_{4,3}^i) - \sigma_\alpha^2}{\text{cov}(\tilde{y}_{2,1}^i, \tilde{y}_{3,2}^i) - \sigma_\alpha^2} \right] [\text{cov}(\tilde{y}_{2,1}^i, \tilde{y}_{3,2}^i) - \sigma_\alpha^2] \\
&\Rightarrow \frac{\text{cov}(\tilde{y}_{2,1}^i, \tilde{y}_{5,4}^i) - \sigma_\alpha^2}{\text{cov}(\tilde{y}_{2,1}^i, \tilde{y}_{4,3}^i) - \sigma_\alpha^2} = \frac{\text{cov}(\tilde{y}_{3,1}^i, \tilde{y}_{5,3}^i) - \sigma_\alpha^2}{\text{cov}(\tilde{y}_{3,1}^i, \tilde{y}_{4,2}^i) - \sigma_\alpha^2} \\
\Rightarrow \sigma_\alpha^2 &= \frac{\text{cov}(\tilde{y}_{2,1}^i, \tilde{y}_{4,3}^i) \text{cov}(\tilde{y}_{3,1}^i, \tilde{y}_{5,3}^i) - \text{cov}(\tilde{y}_{2,1}^i, \tilde{y}_{5,4}^i) \text{cov}(\tilde{y}_{3,1}^i, \tilde{y}_{4,2}^i)}{\text{cov}(\tilde{y}_{2,1}^i, \tilde{y}_{4,3}^i) + \text{cov}(\tilde{y}_{3,1}^i, \tilde{y}_{5,3}^i) - \text{cov}(\tilde{y}_{2,1}^i, \tilde{y}_{5,4}^i) - \text{cov}(\tilde{y}_{3,1}^i, \tilde{y}_{4,2}^i)}
\end{aligned}$$

Now, we are ready to identify the loading factors and variances of persistent shocks. Since $\text{var}(z_{0,t}^i) = 0$, $\text{var}(z_{1,t}^i) = \pi_t^2 \sigma_{\eta,1}^2$. Using the normalization that $\pi_1 = 1$, we get $\sigma_{\eta,1}^2$. Tracking $\text{var}(z_{1,t}^i)$ along t identifies π_t for $t = 2, \dots, T-1$. Consequently, tracing (12) along the age dimension identifies $\sigma_{\eta,h}^2$ for $h = 2, \dots, H-1$. By assumption $\sigma_{\eta,H}^2 = \sigma_{\eta,H-1}^2$ which gives us $\text{var}(z_{H,1}^i)$.

Now let's identify $\sigma_{\epsilon,1}^2$ using equation 10 for $h = 1$ and $t = 1$. Then again using equation 10 for $h = 1$, $t = T$ we can get $\text{var}(z_{1,T}^i)$. Equation 12 for $h = 1$ and $t = T$ pins down π_T . Now we have recovered the entire π_t profile.

The unidentified parameters so far are the lifetime profile of transitory variances and their respective loading factors over time. We will show that the information contained in 10 is sufficient to identify both of these parameters, thanks to our identifying assumptions of $\phi_1 = 1$ and $\phi_T = 1$. An immediate consequence of 10 is

$$\text{var}(\tilde{y}_{h,1}^i) - \sigma_\alpha^2 - \text{var}(z_{h,1}^i) = \sigma_{\epsilon,h}^2 \quad \text{for } h = 1, \dots, H$$

identifying $\sigma_{\epsilon,h}^2$ over the life cycle (except for $H-1$). Fixing h , tracking 10 over t , and using the fact that we already identified all the parameters except the profile of loading factors on transitory variances, it is easy to see that ϕ_t can be recovered for $h = 2, \dots, H-1$.

B Robustness

B.1 Results with Wage Data

Recall that the paper presented results using earnings data. One concern with earnings is that dynamics that are in reality due to changes in hours can be interpreted as shocks. This requires us to check the robustness of our results using data on wages. Wage in our data set is defined as the ratio of annual earnings to hours worked during that year. Figures 12-15 show the results for wage data.

Figure 12: Persistence Profile




Figure 13: Variance Profile of Persistent Shocks




Figure 14: Variance Profile of Transitory Shocks


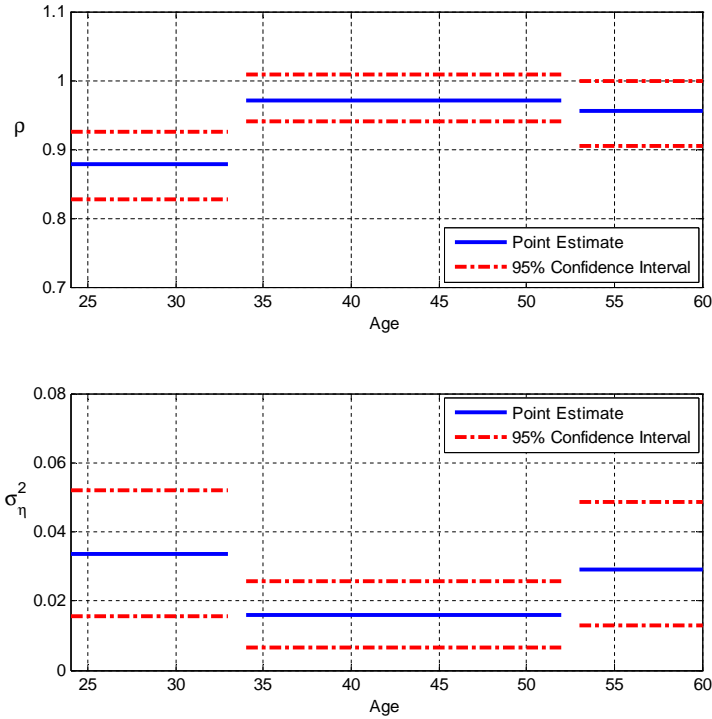


Figure 15: Results for Wages with Age Bins



The following tables present point estimates as well as the results of significance tests.

Table 4: Estimation and Test Results for Quadratic Specification (Wage Data)

x	$\gamma_{x,0}$	$\gamma_{x,1}$	$\gamma_{x,2}$	Test 1	Test 2
ρ	0.7862 (0.0534)	0.0163 (0.0048)	-0.0003 (0.0001)	$H_0 : \gamma_{\rho,1} \leq 0$ 0.0000	$H_0 : \gamma_{\rho,2} \geq 0$ 0.0000
σ_η^2	0.0495 (0.0089)	-0.0033 (0.0009)	0.0001 (0.0000)	$H_0 : \gamma_{\sigma_\eta^2,1} \geq 0$ 0.0000	$H_0 : \gamma_{\sigma_\eta^2,2} \leq 0$ 0.0000
σ_α^2	0.0695 (0.0236)				
σ_ϵ^2	0.0528 (0.0179)				

* The numbers in brackets are bootstrap standard errors.

** The last three columns report the P-values for the corresponding test.

Table 5: Estimation and Test Results for Age Bins (Wage Data)

	$\delta_{x,1}$	$\delta_{x,2}$	$\delta_{x,3}$	Test 1	Test 2
ρ	0.8774 (0.0266)	0.9706 (0.0170)	0.9558 (0.0265)	$H_0 : \rho_1 \geq \rho_2$ 0.0040	$H_0 : \rho_2 \leq \rho_3$ 0.3480
σ_η^2	0.0280 (0.0073)	0.0133 (0.0038)	0.0243 (0.0069)	$H_0 : \sigma_{\eta,1}^2 \leq \sigma_{\eta,2}^2$ 0.0000	$H_0 : \sigma_{\eta,2}^2 \geq \sigma_{\eta,3}^2$ 0.0480
σ_α^2	0.0699 (0.0102)				
σ_ϵ^2	0.0522 (0.0171)				

* The numbers in brackets are standard errors.

** The last three columns report the p-values of the corresponding tests.

B.2 Results with Potential Experience for Ages 20-64

Now, we check the robustness of our findings with respect to age criteria. Recall that we required an individual to be between the ages of 24 and 60. In Figures 16-19, we present the results for the sample with individuals between 20 and 64. Recall, also, that we used age as the variable that defines the life cycle. Here, we use potential experience as an alternative.²⁹

²⁹This also means that we use potential experience instead of age in our first-stage regressions.

Figure 16: Persistence Profile

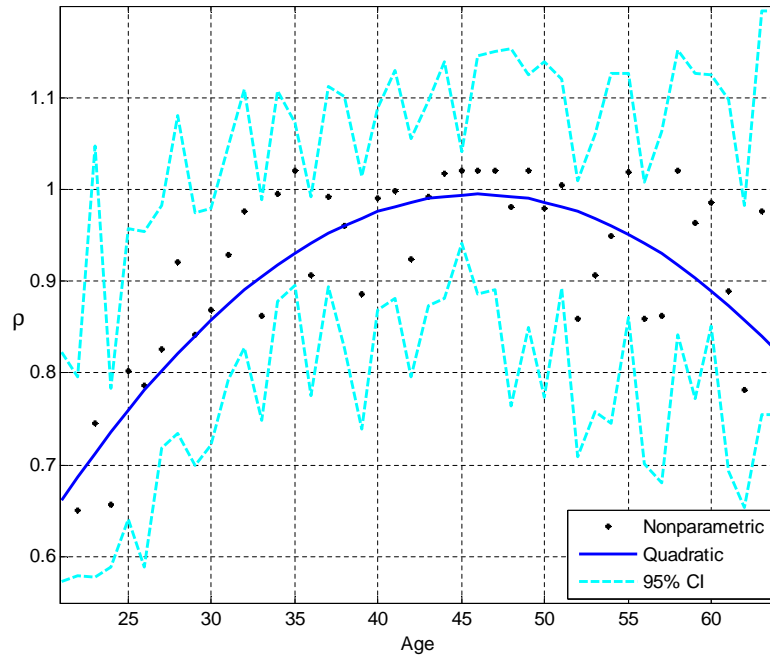


Figure 17: Variance Profile of Persistent Shocks




Figure 18: Variance Profile of Transitory Shocks


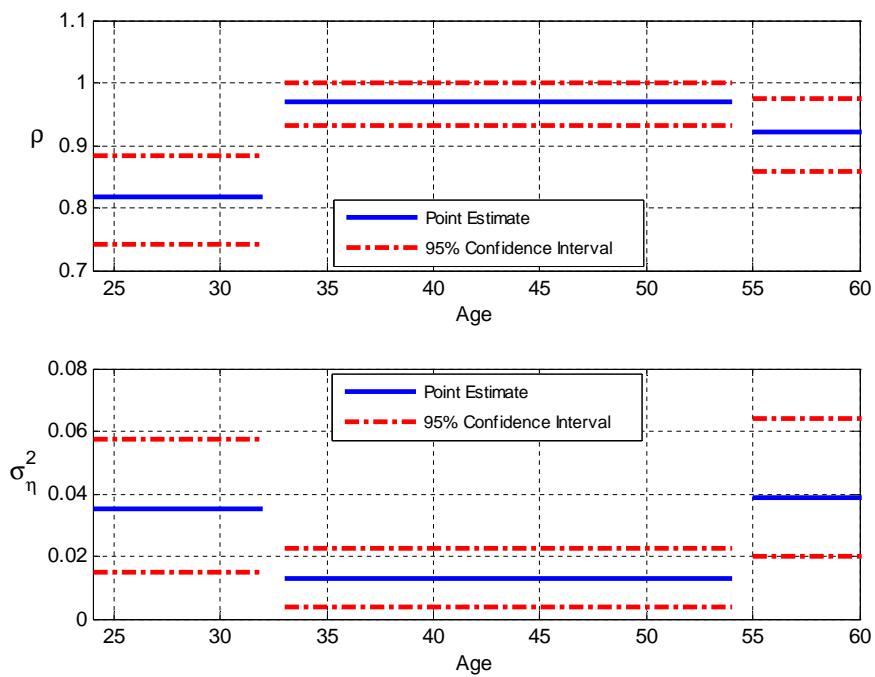


Figure 19: Results for Potential Experience with Age Bins



The following tables present point estimates as well as the results of significance tests.

Table 6: Estimation and Test Results for Quadratic Specification: Potential Experience

x	$\gamma_{x,0}$	$\gamma_{x,1}$	$\gamma_{x,2}$	Test 1	Test 2
ρ	0.6052 (0.0505)	0.0289 (0.0030)	-0.0005 (0.0001)	$H_0 : \gamma_{\rho,1} \leq 0$ 0.0000	$H_0 : \gamma_{\rho,2} \geq 0$ 0.0000
σ_η^2	0.0943 (0.0117)	-0.0071 (0.0009)	0.0001 (0.0000)	$H_0 : \gamma_{\sigma_\eta^2,1} \geq 0$ 0.0000	$H_0 : \gamma_{\sigma_\eta^2,2} \leq 0$ 0.0000
σ_α^2	0.0940 (0.0214)				
σ_ϵ^2	0.0755 (0.0200)				

* The numbers in brackets are bootstrap standard errors.

** The last three columns report the P-values for the corresponding test.

Table 7: Estimation and Test Results for Age Bins: Potential Experience

	$\delta_{x,1}$	$\delta_{x,2}$	$\delta_{x,3}$	Test 1	Test 2
ρ	0.8184 (0.0359)	0.9693 (0.0170)	0.9218 (0.0278)	$H_0 : \rho_1 \geq \rho_2$ 0.0000	$H_0 : \rho_2 \leq \rho_3$ 0.0920
σ_η^2	0.0351 (0.0091)	0.0129 (0.0044)	0.0386 (0.0101)	$H_0 : \sigma_{\eta,1}^2 \leq \sigma_{\eta,2}^2$ 0.0000	$H_0 : \sigma_{\eta,2}^2 \geq \sigma_{\eta,3}^2$ 0.0000
σ_α^2	0.0983 (0.0133)				
σ_ϵ^2	0.0996 (0.0221)				

* The numbers in brackets are standard errors.

** The last three columns report the p-values of the corresponding tests.

C An Economic Rationale for the Age-Dependent Specification

Through a series of econometric analyses, we have shown that persistence and variance of innovations to earnings exhibit non-trivial age profiles. A natural follow-up question would be which economic forces may give rise to these. In this section, we elaborate on the economic rationale behind having an age-dependent income process.

To speculate about one mechanism, these profiles could be due to differences in insurance opportunities against earnings shocks between young and old workers. For example, in case of an adverse demand shock to an individual's occupation, one might switch to a different one if she is young. For an old worker, though, switching is costlier (e.g. because of occupation-specific human capital). Therefore, shocks of the same nature can translate into innovations with different persistence over the working life.

Another mechanism, again related to mobility, is learning about the match quality, first studied by [Jovanovic \(1979\)](#). In his setup, neither the worker nor the firm know the productivity of the match before employment. After observing the output, match productivity is revealed to both parties in a Bayesian fashion. This generates endogenous movements in wages and job turnover. [Flinn \(1986\)](#) presents evidence from NLSY/66 in favor of this theory. We now study the wage dynamics implied by a simple version of [Jovanovic \(1979\)](#).

C.1 A Model of Job Mobility

Our economy consists of a continuum of workers endowed with one unit of time per period. Workers maximize the present value of their lifetime earnings and discount future earnings at a constant interest rate of r . They are subject to death with constant probability, δ . There is measure one of firms that have access to a constant returns to scale production technology. Labor is the only input to the production.

At the beginning of a period, unemployed workers meet with firms, form a match and draw a productivity specific to the match, $\hat{\mu}$, from a normal distribution with mean μ and variance σ_μ^2 . The match-specific productivity is not known to the firm and the worker. Employed workers with tenure t receive their compensation, w_t , before production takes place. Output of the match, y_t , is the sum of the match-specific productivity $\hat{\mu}$, and an

i.i.d. shock, ν_t . The latter is normally distributed with mean 0 and variance σ_ν^2 . After observing the output, beliefs are updated in a Bayesian fashion. Because of normality assumptions, they are characterized by the mean and the precision of the point estimate about $\hat{\mu}$. Let $\hat{m}_{t|t-1}$ denote the mean about $\hat{\mu}$ in period t conditional on all the information up to period $t-1$ and let p_t denote the precision.³⁰ The law of motion for these are governed by,

$$\begin{aligned}\hat{m}_{t+1|t} &= \hat{m}_{t|t-1} \frac{p_t}{p_t + p_\nu} + y_t \frac{p_\nu}{p_t + p_\nu} \\ p_t &= p_\mu + (t-1)p_\nu \\ y_t &= \underbrace{\hat{m}_{t|t-1}}_{\hat{\mu}} + \omega_t + \nu_t\end{aligned}\tag{14}$$

where $\omega_t \sim N(0, 1/p_t)$ represents the deviation of the belief from the true productivity $\hat{\mu}$, $p_\mu = 1/\sigma_\mu^2$, and $p_\nu = 1/\sigma_\nu^2$.

For simplicity, we assume that firms pay workers their expected productivity before production takes place (i.e. $w_t = \hat{m}_{t|t-1}$). After updating the beliefs, a worker decides whether to break the match. If she decides to break the match, she has to pay a fixed cost, C , which represents the direct and foregone earnings costs of changing a job.³¹

The value function of the worker is

$$\begin{aligned}W_t(\hat{m}_{t|t-1}) &= w_t + \beta E \max \{W_{t+1}(\hat{m}_{t+1|t}), W_1(\hat{m}_{1|0}) - C\} \\ s.t. &\quad (14)\end{aligned}$$

where $\hat{m}_{1|0} = \mu$ and $\beta = \frac{\delta}{1+r}$.³²

C.2 Simulation Results

In order to evaluate this model, we simulate data from the model and estimate the age-dependent income process. We should note that we do not calibrate the model to match


³⁰Since the information set of the worker and the firm are the same, their beliefs are identical.

³¹We do not model unemployment in the sense that workers meet new firms and start working immediately in the next period.

³²The initial beliefs are given by the unconditional mean of the distribution for match productivity, thus they are the same for every quitter.

any targets in the data. Ours is an exercise of showing that the model has the potential to generate age profiles and replicate our empirical findings.³³ Figure 20 shows the results.

Figure 20: Simulation Results for the Learning Model



The top panel shows that persistence profile is increasing with age. The mechanism behind this increase can be summarized as follows. First, let's consider a worker who stays in the same job. Her wage can be expressed as the sum of her previous wage and a mean-zero innovation, implying random walk.³⁴ On the other hand, job switchers always get

³³To be more precise, we set $\mu = 2$, $\sigma_\mu^2 = 1$, $\sigma_\nu^2 = 4$; $\beta = 1/(1+r) = 0.95$, and $C = 0$. We simulate 10000 individuals, run the first stage regressions to obtain the residuals and estimate the nonparametric specification of the age-dependent process.

³⁴Recall that $w_t = \hat{m}_{t|t-1}$. Equation (14) implies that $w_{t+1} = w_t \frac{p_t}{p_t+p_\epsilon} + y_t \frac{p_\epsilon}{p_t+p_\epsilon} = w_t \left(\frac{p_t}{p_t+p_\epsilon} + \frac{p_\epsilon}{p_t+p_\epsilon} \right) + \frac{p_\epsilon}{p_t+p_\epsilon} (\omega_t + \nu_t) = w_t + \xi_t$, where $\xi_t \sim N(0, \frac{p_\epsilon}{p_t(p_t+p_\epsilon)})$.

the unconditional mean of the match-specific component μ , implying 0 covariance between current and future wages. Therefore, persistence is lower for them. The persistence of the overall sample is a combination of the persistence of these two subsamples. Over the lifetime, the fraction of switchers is declining with age due to a selection argument, implying a rising persistence profile. Furthermore, the bottom panel of Figure 20 shows a decreasing variance profile for persistent shocks.³⁵ This is because the variance of innovations to wages declines with tenure for stayers.³⁶

This section presented a theoretical background for our empirical findings. We have illustrated that a very stylized model of learning (à la Jovanovic (1979)) implies an increasing persistence profile and a decreasing variance over the working life. The mechanism discussed here is known to have empirical relevance (see Flinn (1986)). Therefore, we also view these results as complementary to our econometric analysis in Section 2, providing independent evidence for the age profiles.

³⁵Note that the variance of persistent shocks are very low. This is because we did not calibrate the model to match the data.

³⁶According to the previous footnote, the variance of ξ_t is decreasing, since p_t is increasing in t .

D Data

We use the first 29 waves of the Panel Study of Income Dynamics (PSID). We include an individual in our baseline sample if he satisfies the following criteria for 3 not necessarily consecutive years: (i) the individual has reported positive labor earnings and hours, (ii) his age is between 24 and 60, (iii) he worked between 520 and 5110 hours during the calendar year, and (iv) had an average hourly real wage between a minimum of \$2 and a maximum of \$400 in 1993. We also exclude people from the poverty sub-sample in 1968 (SEO). These criteria are fairly standard in the literature and leave us with 4380 individuals and 53,864 observations. Tables 8 and 9 present some summary statistics.

Table 8: Summary of the Data by Year

Year	Obs	Earnings	Hours	Wages	Year	Obs	Earnings	Hours	Wages
1967	1289	36194.52	2366.20	15.80	1982	1951	38137.18	2187.37	17.52
1968	1367	36819.11	2322.09	16.25	1983	1983	40026.40	2244.20	17.95
1969	1400	36480.51	2273.38	16.58	1984	2015	40691.30	2246.78	18.19
1970	1437	36489.48	2285.34	16.51	1985	2086	40642.82	2250.73	18.15
1971	1497	37413.72	2316.87	16.69	1986	2106	42052.72	2267.04	18.33
1972	1580	38466.86	2314.00	17.12	1987	2121	42159.23	2289.65	18.41
1973	1614	37860.49	2262.98	17.18	1988	2124	42223.60	2296.47	18.22
1974	1642	36042.35	2223.94	16.66	1989	2125	42285.32	2294.95	18.30
1975	1702	37614.29	2266.62	17.04	1990	2122	42223.05	2257.22	18.62
1976	1741	38438.27	2271.08	17.32	1991	2050	45065.64	2236.87	20.25
1977	1795	39065.12	2271.54	17.73	1992	1994	41896.25	2238.78	18.87
1978	1854	38916.62	2259.41	17.64	1993	2221	42608.05	2265.50	19.02
1979	1881	38562.79	2228.56	17.71	1994	2217	44090.57	2286.83	19.37
1980	1928	37832.31	2196.97	17.68	1995	2098	46904.14	2278.47	20.81
1981	1924	37909.72	2165.78	17.60					

Table 9: Summary of the Data by Age

Age	Obs	Earnings	Hours	Wage	Age	Obs	Earnings	Hours	Wage
24	1508	25536.80	2172.93	12.04	43	1441	47467.94	2301.82	21.04
25	1867	27463.69	2191.75	12.77	44	1388	47676.97	2299.21	21.18
26	2025	28897.09	2230.11	13.22	45	1307	47257.60	2294.15	20.92
27	2100	30853.31	2280.07	13.82	46	1274	46160.94	2275.78	20.63
28	2133	32438.98	2269.40	14.63	47	1262	46182.96	2261.90	20.78
29	2117	33688.80	2260.75	15.11	48	1188	46978.38	2261.67	21.18
30	2127	35048.59	2283.71	15.58	49	1114	44870.02	2232.10	20.36
31	2100	36676.33	2294.13	16.34	50	1050	45727.01	2227.46	20.85
32	2071	37277.30	2299.87	16.46	51	997	44800.87	2214.57	20.44
33	2014	39220.43	2306.97	17.26	52	960	44676.26	2204.17	20.34
34	1947	39891.28	2285.42	17.59	53	923	45083.90	2223.70	20.03
35	1932	40876.66	2300.92	17.93	54	870	43177.46	2176.21	19.85
36	1874	42673.70	2321.84	18.64	55	832	44436.93	2166.60	20.36
37	1807	42933.08	2294.45	18.92	56	803	43051.21	2155.09	20.11
38	1712	43780.25	2307.00	19.29	57	757	43936.76	2169.68	20.01
39	1660	45151.08	2320.45	19.42	58	731	44140.43	2136.73	20.56
40	1657	44865.85	2301.67	19.85	59	670	41813.15	2095.75	19.97
41	1577	46068.74	2296.17	20.17	60	558	39711.87	2072.45	19.25
42	1511	46267.70	2306.93	20.50					