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 11-030

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

by

Fatih Karahan and Serdar Ozkan

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

Fatih Karahan† Serdar Ozkan‡

September 2, 2011

*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; E-mail: karahan@econ.upenn.edu, http://www.econ.upenn.edu/~karahan

‡Corresponding author. Federal Reserve Board, 20th and C Streets, NW, Washington, DC, 20551, USA. Tel.: +1202 721-4558. E-mail: Serdar.Ozkan@frb.gov, www.serdarozkan.me

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 an age 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 find that shocks to earnings are only moderately persistent (around 0.75) for young workers. Persistence rises with age up to unity until midway in life. The variance of persistent shocks exhibits a U-shaped profile over the life cycle (with a minimum of 0.01 and a maximum of 0.05). These results suggest that the standard specification in the literature (with constant persistence and variances) cannot capture the earnings dynamics of young workers. We also argue that a calibrated job turnover model can account for these non-flat profiles. The key idea is that workers sort into better jobs and settle down as they age; in turn, magnitudes of wage growth rates decline, thereby decreasing variance of shocks. Furthermore the decline in job mobility results in higher persistence. Finally, we investigate the implications of age profiles for consumption-savings behavior. The welfare cost of idiosyncratic risk implied by the age-dependent income process is 34 percent lower compared with its age-invariant counterpart. This difference is mostly due to a higher degree of consumption insurance for young workers, for whom persistence is moderate. These results suggest that age profiles of persistence and variances should be taken into account when calibrating life-cycle models.

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

JEL: C33, D31, D91, E21, J31

1 Introduction

Two important determinants of labor income risk are the persistence and variance of shocks. 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 and estimate a novel specification for idiosyncratic earnings that allows for a life-cycle profile in the persistence and variance of earnings shocks.

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 novel feature of our specification is that both the persistence parameter of the $AR(1)$ component and the variance of innovations to transitory and persistent components are allowed to vary by age. This paper, to the best of our knowledge, is the first study that estimates a lifetime profile of earnings persistence and variance together.[1]

We show that this specification is identified and estimate it using earnings data from the Panel Study of Income Dynamics (PSID). Our results reveal that persistence increases at early stages in the working life: starting from 0.70 it rises to unity. These differences are sizable: 70 percent of a shock received during the early years in the labor market dies out over the next five years, whereas if the shock is received at age 40, 85 percent of it would still remain after five years. However, we find a U-shaped profile for the variance of persistent shocks: A shock of one standard deviation implies a 26 percent change in annual earnings for a 24-year old. The corresponding number for a 40-year old is only 12 percent. As for the variance of transitory shocks, we find a sizable increase early on but a flat profile for the remaining working life. These results suggest that the standard specification in the literature (with constant persistence and variances) cannot capture the earnings dynamics of young workers.

We then ask the question of whether these life-cycle profiles are statistically significant. To tackle this question, we estimate life-cycle profiles by partitioning the working life into three stages. Here, we assume that persistence and variances are constant within a stage but might differ from one to the other. We test whether the persistence and the variance of

[1] There are several other studies that take into account variation in persistence and variance of shocks. Baker and Solon (2003) and Gottschalk and Moffitt (2011) allow for age-specific variances in transitory shocks, and Sabelhaus and Song (2010) also let both the permanent and the transitory shocks vary with age and cohort. Hause (1980) estimates a process that has an $AR(1)$ component with time-specific persistence and variance of shocks. Alvarez, Browning, and Ejrnæs (2010) investigate the heterogeneity in the persistence of shocks across individuals. Feigenbaum and Li (2008) find a U-shaped earnings uncertainty profile over the life cycle.

shocks differ significantly across the three age intervals. We strongly reject the hypothesis of a flat profile for persistence and the variance of persistent shocks but not for variance of transitory shocks.

The estimates of persistence in the literature are close to unity. The age-dependent estimate of persistence lies substantially below one for most of the lifetime. We argue that the high persistence in the literature is driven by targeting the linear, if not convex, increase in residual earnings inequality over the working life. Namely, estimation avoids lower levels of persistence, which would imply a concave rise in inequality. The age-dependent income process matches the inequality profile without high levels of persistence; thanks to the inverse relationship between persistence and the variance of labor income shocks that our estimates reveal.

To explore one possible mechanism behind the rise in the increase in persistence and decrease in the variance of persistent shocks early in life, we study the implications of the job turnover model by Jovanovic (1979). In this model, unemployed workers match with firms and draw a match-specific productivity, unobservable to both the firm and the worker. Output is the sum of match productivity and a white noise. Firms pay workers their expected productivity. After observing the output, both the worker and the firm update their beliefs about the match productivity in a Bayesian fashion. In the end of the period, workers decide whether to quit and meet another firm or stay on the same job based on their beliefs.

In a simple calibration exercise, we show that the model is quite successful in generating the age profiles in the data, i.e., the increase in persistence and the decrease in variance of persistent shocks early in working life. The mechanism behind this result can be summarized as follows: The model implies that wages of stayers follow a random walk, whereas the autocorrelation of wages is very small for quitters. The overall persistence is a combination of these two. As workers age, they sort into jobs with better match productivities and settle down, which results in an increase in the number of stayers, thereby resulting in an increase in persistence. Similarly, as match productivity is being revealed, the magnitude of changes in beliefs and thus wages decrease and in turn the variance of persistent shocks declines. This mechanism is known to have empirical relevance (Flinn (1986)). Therefore, we also view these results as complementary to our econometric analysis, providing justification for the age profiles.

We then investigate the economic implications 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. For this purpose, 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 its age-invariant counterpart.

We find that, in an economy with natural borrowing constraints (NBC), the age-dependent income process implies a much higher consumption insurance against persistent shocks: Around 56 percent of persistent shocks translate into consumption growth under the age-dependent income process compared to 38 percent for the age-invariant specification. Most of this difference comes from young workers for whom the degree of insurance is as high as 70 percent under the age-dependent process as opposed to 30 percent for the age-invariant specification. This difference is due to the level of persistence, which is particularly low for young workers under the age-dependent process. In the presence of highly persistent shocks, agents refrain from borrowing against the possibility of a long sequence of low income realizations. Insurance against such shocks is, therefore, mostly through assets. This type of self-insurance is not possible for young agents, since they don't have enough wealth.

In an economy with zero borrowing constraints (ZBC), consumption insurance is lower for both specifications compared with the NBC economy. Now, the gap in consumption insurance between the age-dependent and the age-invariant processes is smaller: 38 percent for the age-dependent vs. 30 percent for the age-invariant. The decrease in the gap is due to young workers who lack the borrowing option to insure against moderately persistent shocks.

We also compare the welfare costs of idiosyncratic risk implied by the age-dependent process with the age-invariant one. We find substantial differences: In the NBC (ZBC) economy, welfare costs of lifetime shocks is 3.13 percent (5.80 percent) under the age-dependent income process, whereas this number is 4.76 percent (6.80 percent) for the age-invariant specification.

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 3 presents the structural job turnover model. Section 4 presents the life-cycle model that is used to study the consumption-savings implications of the age-dependent process. Finally, Section 5 concludes.

2 Empirical Analysis

In this section we describe the statistical model for earnings, and discuss the data and our benchmark sample. The empirical findings are presented at the end of this section.

2.1 An Age-Dependent Income Process

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

$$y^i_{h,t} = f\left(X^i_{h,t};\theta_t\right) + \tilde{y}^i_{h,t}. \tag{1}$$

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. $X^i_{h,t}$ is a vector of observables that includes a cubic polynomial in age and education dummies for less than a high school diploma, high school diploma, and 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.

Residual income is decomposed into a fixed effect, an $AR(1)$ component, and a transitory component. This representation is parsimonious, yet it captures the salient features of the data well. Therefore, it is widely used in the literature. This paper extends the standard specification to allow for a lifetime profile in the persistence parameter, the variance of persistent and transitory shocks:

$$
\begin{aligned}
\tilde{y}_{h,t}^{i} &= \alpha_i + z_{h,t}^{i} + \phi_t \varepsilon_h^i, \qquad (2)\\
z_{h,t}^{i} &= \rho_{h-1} z_{h-1,t-1}^{i} + \pi_t \eta_h^i,\ z_{1,t}^{i} \sim N\left(0, \pi_t^2 \sigma_{z_1}^2\right),\\
\text{and} \qquad & \eta_h^i \sim N\left(0, \sigma_{\eta,h}^2\right)\ \varepsilon_h^i \sim N\left(0, \sigma_{\varepsilon,h}^2\right).
\end{aligned}
$$

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.[2] 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. $z_{1,t}^i$ captures the initial variation in the persistent component.[3] 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.

A number of studies document the evolution of residual inequality for the United States in the last three decades (e.g., Gottschalk and Moffitt (2011); Heathcote, Perri, and Violante (2010); and Panousi, Vidangos, Heim, and DeBacker (2011)). We follow Gottschalk and Moffitt (1995) and control for the change in residual inequality over time with ϕ_t and π_t,

[2] These changes are potentially correlated with future promotions. However, we follow the literature and assume that these shocks are i.i.d. in nature (see Lillard and Willis (1978); Lillard and Weiss (1979); MaCurdy (1982); Abowd and Card (1989); and Baker (1997)). A notable exception is Hryshko (2011).

[3] Our benchmark sample is composed of workers who are at least 24 years old. Therefore, it is reasonable to think that they already have some labor market experience, in turn, they have accumulated some persistent shocks by age 24.

representing the time loading factors for transitory and permanent shocks, respectively.[4]

Having introduced the age-dependent income process, an immediate concern is identification. Can the variance-covariance structure of earnings data tell us how changes in earnings differ in variance and persistence over age and time together? The identification discussion allows us to connect the statistical model to the moments in the data and makes the estimation procedure meaningful.[5] The next proposition establishes that the income process (2) is identified and provides a proof:

Proposition 1: Specification (2) is identified in levels up to the normalizations that $\rho_1 = \rho_2$, $\pi_1 = \phi_1 = 1$, $\phi_H = \phi_{H-1}$, and $\sigma^2_{\eta,H} = \sigma^2_{\eta,H-1}$.

Proof: See Appendix A.

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 feature 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 guides our sample selection process.

2.2 Sample Selection and Estimation Method

This section briefly describes the data and the variable definitions used in the empirical analysis. We use 30 waves of the Panel Study of Income Dynamics (PSID) between 1968-

[4]A related approach would be to control for cohort effects. Heathcote, Storesletten, and Violante (2005) provide some evidence that time effects are more pronounced than cohort effects. Thus, we choose to control for time effects.

[5]Heathcote, Storesletten, and Violante (2010) also explain the identification of time loading factors for a standard income process.

1997. We estimate our model using both annual earnings and the average hourly wage of male heads of households.[6] Here, we present the results for earnings data. Estimation results for wage data are reported in Appendix B.2; the results are qualitatively the same. In order to have a large number of cohorts, we include an individual in our benchmark sample if he satisfies the following criteria for three, not necessarily consecutive, years:[7] (i) the individual has an average hourly wage between \$2 and \$400 in 1993 dollars, (ii) his age is between 24 and 60, and (iii) he worked between 520 and 5,110 hours during the calendar year. We also exclude people from the Survey of Economic Opportunity (SEO) sub-sample in 1968. These criteria are fairly standard in the literature and leave us with 4,324 individuals and 56,156 observations.

We employ an equally weighted minimum distance estimator. We minimize the distance between the moments of the $(T \times T)$ and $(H \times H)$ empirical variance-covariance structure of residual earnings and their theoretical counterparts implied by income process 2. In particular, we target all the variance and covariance terms over age, $cov\left(\tilde{y}_h^i, \tilde{y}_{h+n}^i\right)$, and over time, $cov\left(\tilde{y}_t^i, \tilde{y}_{t+n}^i\right)$, to which at least 150 individuals contribute. This leaves us with 1,067 moments.[8] To obtain the theoretical counterpart of $cov\left(\tilde{y}_h^i, \tilde{y}_{h+n}^i\right)$, we average $cov\left(\tilde{y}_{h,t}^i, \tilde{y}_{h+n,t+n}^i\right)$ over t. Similarly, we compute the theoretical counterpart of $cov\left(\tilde{y}_t^i, \tilde{y}_{t+n}^i\right)$ by averaging $cov\left(\tilde{y}_{h,t}^i, \tilde{y}_{h+n,t+n}^i\right)$ over h. Due to small sample considerations explained in Al-

[6]Earnings in the PSID are composed of wages, bonuses, commissions, and the labor portion of self-employment. Hourly wage is defined as earnings divided by annual hours.

[7]In one robustness check, we restrict our sample to people with three consecutive income spells. The results are reported in Appendix B.6. In an additional check we require people to have at least 10 (not necessarily consecutive) observations. The results for this sample are reported in Appendix B.7.

[8]If we require that there are at least 30 observations in a moment to be targeted in the estimation, we end up with 60 more moments and this does not have any substantial effect on our results.

tonji and Segal (1996), our minimum distance estimator employs the identity matrix as the weighting matrix.

2.3 Estimation Results

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

We estimate the lifetime profile of shocks and persistence in two ways. First, we estimate a nonparametric specification, that is, we don't impose any functional form on the lifetime profiles. Then, we assume the life-cycle profiles follow a cubic function of age and estimate its parameters. Figure 1 shows the results for persistence. The point estimates for the nonparametric estimation are shown in dots along with the 95 percent bootstrap confidence interval in dashed lines and the point estimates are shown in Table 6 and 7. We employ a block bootstrap with 150 repetitions.[9] The results of the cubic specification are shown in the solid blue line. The parameter estimates as well as bootstrap standard errors are reported on the left panel of Table 1.

Figure 1 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 around age 40. The differences also appear to be economically large (although a quantitative evaluation needs to await the consumption model in Section 4). For example, more than 70 percent of a change in a 24-year-old's earnings dies out in five years. This number is only around 15 percent for a 40-year-old individual.

[9]Increasing the number of repetitions does not change the standard errors.

Figure 1: Persistence Profile

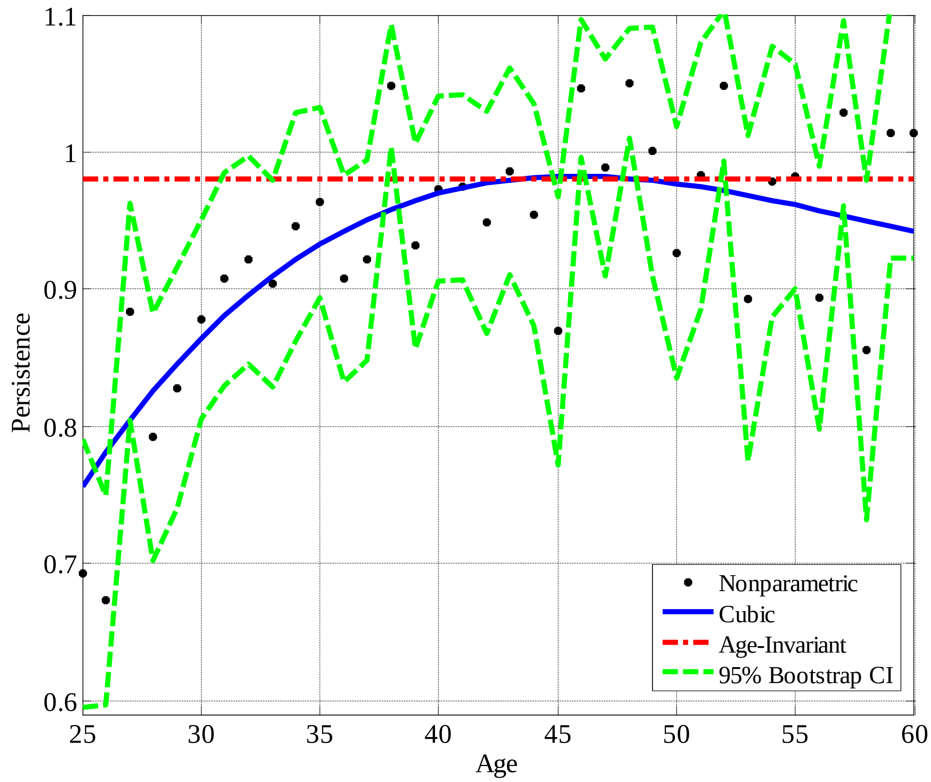

Figure 2: Variance Profile of Persistent Shocks

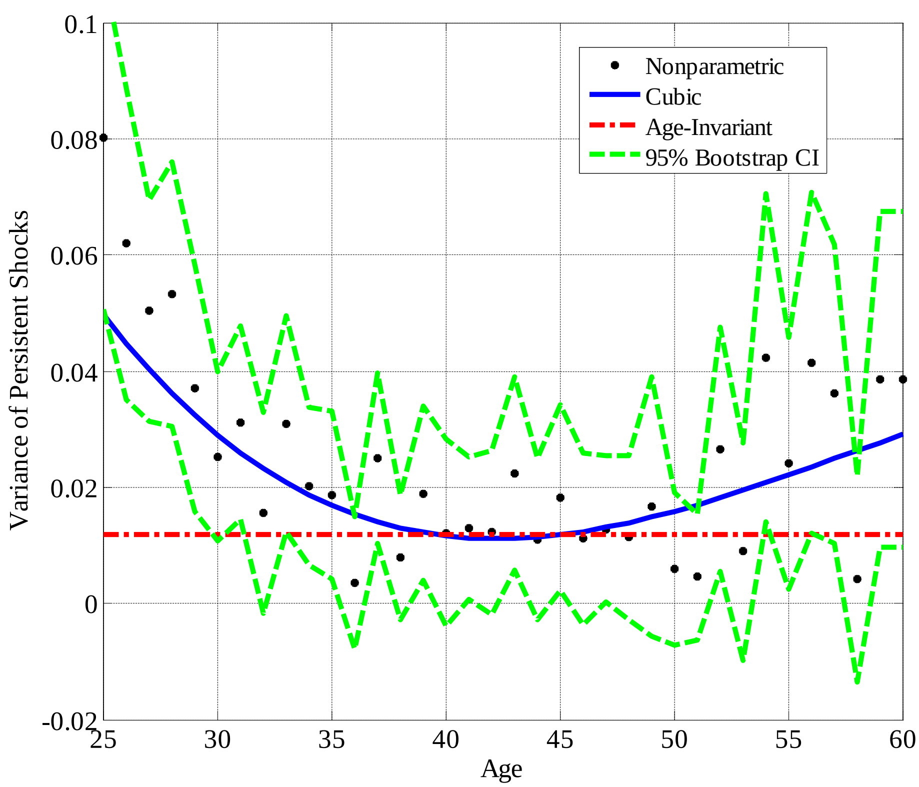

The variance of persistent shocks, shown in Figure 2, follows a different pattern. It exhibits a U-shaped profile over the lifetime. Early in life, shocks are larger compared with those in the 40s. The variance starts around 0.06, decreases to around 0.01 by age 35, and remains roughly flat for 10 years. Shocks toward the end of the life cycle are larger, which manifest in a variance of around 0.035. These differences again appear to be economically large; a one-standard-deviation persistent shock implies a 26 percent change in earnings at age 24, whereas a one-standard-deviation shock implies only a 12 percent change for a 40-year old.

Figure 3: Variance Profile of Transitory Shocks

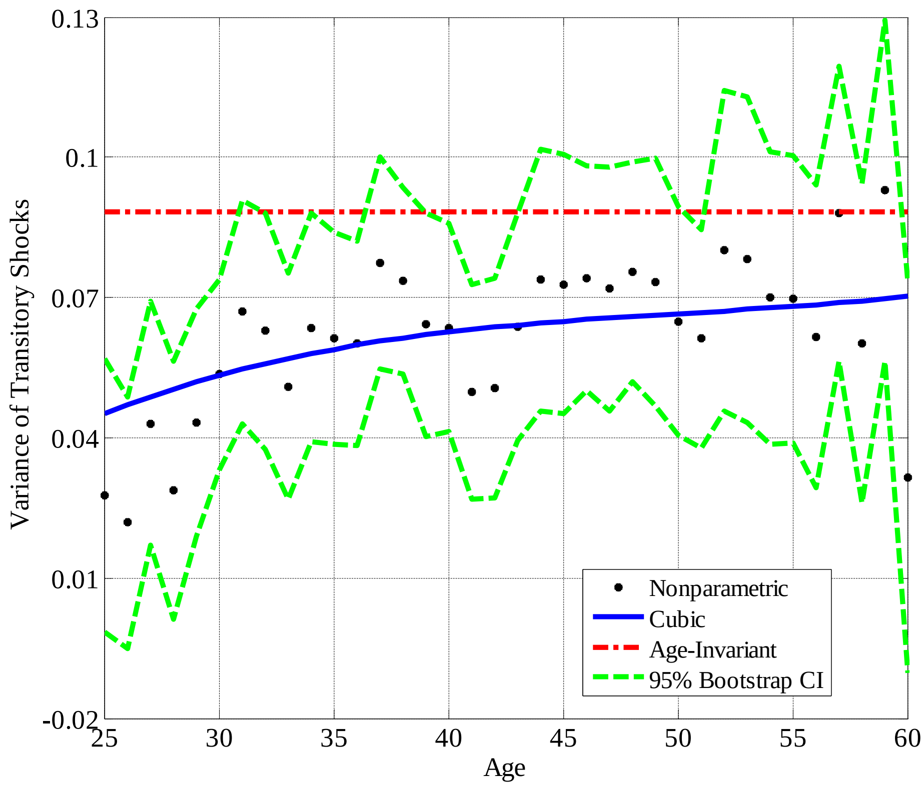

Figure 3 plots the variance of transitory shocks. There is a sizable increase early on; it increases from around 0.03 to 0.07 by age 35. The profile is flat after age 35. Note that

Table 1: Estimation Results for Cubic Specification

	Age-Dependent σ^2_ϵ				Age-Invariant σ^2_ϵ			
x	$\gamma_{x,0}$	$\gamma_{x,1}$	$\gamma_{x,2}$	$\gamma_{x,3}$	$\gamma_{x,0}$	$\gamma_{x,1}$	$\gamma_{x,2}$	$\gamma_{x,3}$
σ^2_α	0.0803				0.0768			
	(0.0159)				(0.0167)			
$\sigma^2_{z_1}$	0.0849				0.0803			
	(0.0180)				(0.0186)			
ρ	0.7003	0.2974	-0.0978	0.0095	0.7596	0.2039	-0.0535	0.0028
	(0.0604)	(0.1035)	(0.0607)	(0.0107)	(0.0524)	(0.1059)	(0.0670)	(0.0120)
σ^2_η	0.0607	-0.0593	0.0215	-0.0021	0.0518	-0.0405	0.0105	-0.0002
	(0.0129)	(0.0237)	(0.0135)	(0.0023)	(0.0100)	(0.0219)	(0.0135)	(0.0024)
σ^2_ϵ	0.0410	0.0221	-0.0069	0.0008	0.0564			
	(0.0177)	(0.0385)	(0.0233)	(0.0040)	(0.0049)			

Note: The numbers in brackets are bootstrap standard errors. γ's are the coefficients of a cubic polynomial. Specifically, for $x = \rho, \sigma^2_\eta, \sigma^2_\epsilon : x_h = \gamma_{x,0} + \gamma_{x,1} * h/10 + \gamma_{x,2} * (h/10)^2 + \gamma_{x,3} * (h/10)^3$

in our specification, transitory component soaks up the measurement error.[10] As we will discuss in Section 2.4, this non-flat profile is not statistically significant.

What features of the data give rise to the increase in persistence early in the life cycle? For this, we refer to the identification argument presented in Appendix A, where we argue that the ratio of two-period ahead covariance to one-period ahead covariance at age $h+1$, corrected for fixed effects, (henceforth, Φ^{21}_h) yields a consistent estimate for the persistence parameter.[11] The need to correct for the fixed effect arises because both of these covariance terms contain the variance of the fixed effects. In correcting for fixed effects, we use our baseline estimate ($\sigma^2_\alpha = 0.075$), which is in line with the estimates in the literature. Figure 4

[10] 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 sizes of transitory shocks is not distinguishable from the measurement error once we assume fully transitory errors. In Appendix B.4, we model the transitory component as the sum of an MA(1) and an i.i.d. component. The latter is assumed to be measurement error, and its estimate is taken from PSID validation studies.

[11] Note that, abstracting from time effects, (6) implies $\xi^{21}_h = \left[cov\left(\tilde{y}^i_h, \tilde{y}^i_{h+2}\right) - \sigma^2_\alpha\right] / \left[cov\left(\tilde{y}^i_h, \tilde{y}^i_{h+1}\right) - \sigma^2_\alpha\right] = \rho_{h+1}$ for $h = 1, \ldots, H-2$.

plots the empirical counterpart of Φ_h^{21} in the solid line along with the estimated persistence profile in dots. The age profile of Φ^{21} closely resembles the estimate of the persistence profile: It increases from below 0.7 to above 0.9.[12]

In general, the age profile of ξ_h^{21} depends on the level of fixed effects. To check for robustness, we plot Φ_h^{21} for the case where there are no fixed effects ($\sigma_\alpha^2 = 0$), shown in dashed lines in Figure 4. We see that the increase in persistence is robust to the variance of fixed effects, though the steepness depends on it. Note that the estimation of an upward sloping persistence profile is a result of targeting a fairly complicated variance-covariance structure. Figure 4 confirms this increase over the lifetime from a much simpler look at the data.

Some of the changes in persistence and variance that we observe 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. To control for the effect of part-time workers, we

---

[12] The increase in this ratio is also consistent with an age-invariant income process as in Meghir and Pistaferri (2004), which consists of a permanent component and an MA(1) transitory component: $\tilde{y}_{iht} = \alpha_i + p_{iht} + e_{iht}$, where $p_{iht} = p_{ih-1t-1} + \zeta_{iht}$, $e_{iht} = \epsilon_{iht} + \theta\epsilon_{ih-1t-1}$ and the shocks ζ_{iht} and ϵ_{iht} are uncorrelated at all leads and lags. For this income process, the ratio equals:

$$\frac{cov(\tilde{y}_h, \tilde{y}_{h+2}) - \sigma_\alpha^2}{cov(\tilde{y}_h, \tilde{y}_{h+1}) - \sigma_\alpha^2} = \frac{var(p_h)}{var(p_h) + \theta var(\epsilon_h)} = \frac{1}{1 + \theta\frac{var(\epsilon_h)}{var(p_h)}}.$$

Since $var(p_h)$ increases with age, $\frac{1}{1+\theta\frac{var(\epsilon_h)}{var(p_h)}}$ is increasing over the life cycle, which is consistent with the data presented in Figure 4.

On the one hand, under this specification, the ratio of three-period ahead covariance to two-period ahead covariance, Φ_h^{32}, is constant and equal to one:

$$\Phi_h^{32} = \frac{cov(\tilde{y}_h, \tilde{y}_{h+3}) - \sigma_\alpha^2}{cov(\tilde{y}_h, \tilde{y}_{h+2}) - \sigma_\alpha^2} = \frac{var(p_h)}{var(p_h)} = 1.$$

On the other hand, under the age-dependent income process, the same ratio equals ρ_{h+2}. The empirical counterpart of $\frac{cov(\tilde{y}_h,\tilde{y}_{h+3})-\sigma_\alpha^2}{cov(\tilde{y}_h,\tilde{y}_{h+2})-\sigma_\alpha^2}$ (see Appendix B.10) exhibits an increasing profile over the life cycle. This favors the age-variant persistence profile over the age-invariant one.

restrict our sample to only full-time workers in Appendix B.5. The conclusion for persistence is similar: persistence is increasing. Though, the conclusion for the variance of persistent shocks is somewhat different. It is still the case that variance of persistent shocks for younger workers is significantly larger than the variance of shocks for middle-aged workers. However, the difference between the middle age and old age is not significant. This suggests that some of the increase in the variance of persistent shocks for older workers is driven by early/partial retirement.

Figure 4: Ratio of Covariances: $\Phi_h^{21} = \frac{cov(y_h, y_{h+2}) - \sigma_\alpha^2}{cov(y_h, y_{h+1}) - \sigma_\alpha^2}$, $h = 23, \cdots, 50$

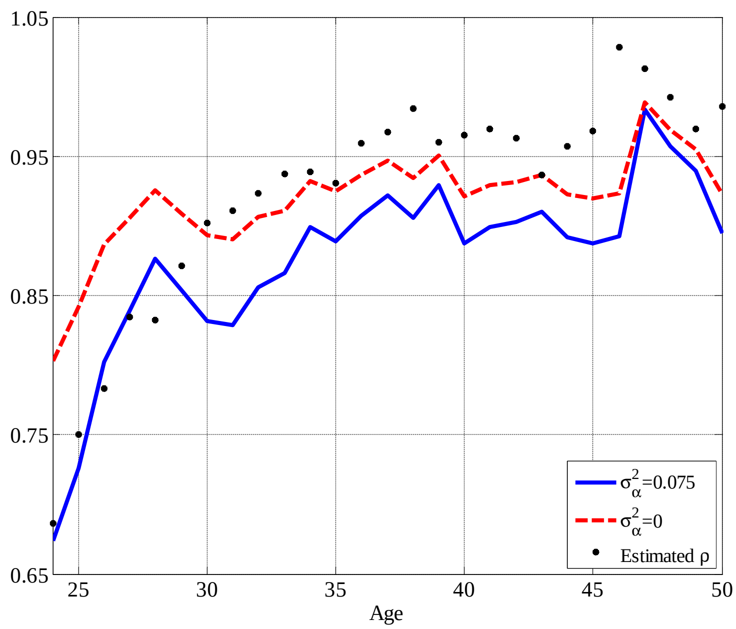

Note: This figure plots the ratio of two-year ahead covariance to one-year ahead covariance, Φ^{21}, corrected for the variance of fixed effects, along with the estimated persistence profile. All three series are smoothed by a moving average method with a three-year span.

To explore the differences in age profiles between workers with and without college degrees, we estimate the age-dependent income process on a sample of workers with a college degree and on a sample of workers without one. The results are reported in Appendix B.9

in Tables 17 and 18. We find that persistence for college workers is increasing over the working life, as opposed to being hump shaped for the non-college sample. The variance of persistent shocks is decreasing for college graduates and U-shaped for those without a college degree.

2.4 Significance Tests

We now turn to the question of statistical significance, that is, we want to see whether the non-flat pattern is statistically significant. For this purpose we consider a model in which working life is divided into three stages (age intervals); young, middle, and old ages. This model restricts the persistence and variances to be constant within an interval but allows them to differ from one to the other. The age bins correspond to ages 24-33 (young), 34-52 (middle), and 53-60 (old).[13] These intervals allows us to identify systematic differences across age intervals.

Point estimates are shown in the first three columns of Table 2 along with bootstrap standard errors in parenthesis. The results, once again, point to the same life cycle profiles of persistence and variance of shocks. We test whether the persistence profile exhibits a hump-shaped pattern. Similarly we investigate if variance of persistent shocks follows a U-shape. Finally we test whether the increase in transitory shocks is statistically significant. Formally the null hypotheses are: $H_0 : \rho_1 \geq \rho_2$, $H_0 : \rho_2 \leq \rho_3$, $H_0 : \sigma^2_{\eta,1} \leq \sigma^2_{\eta,2}$, $H_0 : \sigma^2_{\eta,2} \geq \sigma^2_{\eta,3}$, $H_0 : \sigma^2_{\epsilon,2} \leq \sigma^2_{\epsilon,1}$, and $H_0 : \sigma^2_{\epsilon,2} \leq \sigma^2_{\epsilon,3}$. The results are summarized in the last two columns of

[13] In choosing these intervals, we are motivated by the results of the nonparametric estimates shown on Figures 1-3. Most of the changes in parameters occur in the first 10 and last 8 years of the working life. As we argue in the next section, changes in the parameters are driven by job mobility of workers, which is high in the first 10 years. This also guides us in choosing the initial interval.

Table 2.

Table 2: Estimation and Test Results for Age Bins

	σ^2_α	$\sigma^2_{z_1}$
	0.0707	0.0767
	(0.0268)	(0.0255)

	$[24,33]$	$[34,52]$	$[53,60]$	Test 1 p-value	Test 2 p-value
				$H_0: \rho_1 \geq \rho_2$	$H_0: \rho_2 \leq \rho_3$
ρ	0.8783	0.9712	0.9608	0.00	0.63
	(0.0283)	(0.0141)	(0.0223)		
				$H_0: \sigma^2_{\eta,1} \leq \sigma^2_{\eta,2}$	$H_0: \sigma^2_{\eta,2} \geq \sigma^2_{\eta,3}$
σ^2_η	0.0273	0.0130	0.0258	0.00	0.04
	(0.0045)	(0.0026)	(0.0072)		
				$H_0: \sigma^2_{\epsilon,2} \leq \sigma^2_{\epsilon,1}$	$H_0: \sigma^2_{\epsilon,2} \geq \sigma^2_{\epsilon,3}$
σ^2_ϵ	0.0558	0.0588	0.0675	0.35	0.79
	(0.0063)	(0.0066)	(0.0109)		

Notes: [24,33], [34,52], [53,60] are age intervals in which the persistence and variances are assumed to be constant. Point estimates are shown in the first three columns along with bootstrap standard errors in parenthesis. The last two columns report p-values for bootstrap significance tests.

We find that the persistence for young workers is statistically smaller than that of middle-aged workers. However, we cannot reject the null hypothesis that the persistence in the last age bin is different than that in the second. As for the variance of persistent shocks, parameter of the second age interval is significantly lower than that of the first and third (at 5 percent significance level). However, the non-flat profile in the variance of transitory shocks is insignificant with p-values of 0.347 and 0.213.

Note that the standard errors of linear and higher-order terms in the cubic specification are large such that one might suspect non-flat profiles implied by this specification are insignificant (see Table 1). However, unlike quadratic polynomials, cubic polynomials can generate hump-shaped (or U-shaped) profiles for different combinations of signs of coefficients. Indeed, the correlation structure between parameters of the cubic polynomial is such that for almost all bootstrap runs, the implied persistence profile is increasing and the variance profile for persistent shocks is U-shaped. However, for the lifetime profile of the transitory variance, a significant number of bootstrap repetitions do not imply an increase over the first 10-15 years. In a previous version of this paper (Karahan and Ozkan (2009)), we impose a quadratic polynomial on lifetime profiles and find that both the linear and quadratic terms are significant for persistence and for the variance of persistent shocks.

Overall these results suggest that persistence and variance of persistent shocks have non-flat profiles over the life cycle but not the variance of transitory shocks. Thus, from now on in our analysis, we assume that variance of transitory shocks is age-invariant and use estimates of the cubic specification with a constant variance of transitory shocks. The results are reported in the right panel of Table 1.

2.5 Comparison with the Literature

We now compare the age-dependent process with the age-invariant version of this specification, that is, a specification consisting of a fixed effect, an $AR(1)$ component, and an i.i.d. transitory component, where the persistence and variance of shocks are age-invariant. The age-invariant specification is widely used in quantitative models featuring income risk. In order for these cases to be comparable, we estimate this model on the benchmark sample. The estimates are shown in dashed lines on Figures 1-3 as well as in Table 3. Our estimate of persistence, 0.98, 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.

Table 3: Estimates of the Age-Invariant Specification

	σ_α^2	ρ	$\sigma_{z_1}^2$	σ_η^2	σ_ϵ^2
Point Estimates	0.0146	0.9802	0.0774	0.0113	0.0831
Standard Errors	(0.0265)	(0.0114)	(0.0157)	(0.0016)	(0.0087)

Notes: Table reports estimates of the age-invariant process. $\sigma_{z_1}^2$ is the initial variance of the persistent component. Note that the estimate of σ_α^2 is smaller than the estimates in the literature (see Kaplan (2010)). This is because the specification also allows for a nonzero initial condition in the persistent component.

In what follows, we will argue that targeting the lifetime profile of residual inequality in the data results in an upward bias in persistence if one does not allow for age-specific persistence and variance. For the age profile of residual inequality, we first compute $\widehat{var(\tilde{y}_{h,t})}$ for every year t and age h. An individual in year t contributes to $\widehat{var(\tilde{y}_{h,t})}$ if he is between ages $h-2$ and $h+2$.[14] We then regress these variances on a full set of age and year dummies and report the age dummies. The resulting profile is shown in Figure 5. The rise in residual inequality over the lifetime is almost linear, if not convex. The increase is particularly steep after age 35.

For the age-invariant process, the corresponding theoretical variances are given by

$$var(\tilde{y}_h) = \sigma_\alpha^2 + \sigma_\eta^2 \sum_{j=1}^{h-1} \rho^{2j} + \sigma_{z_1}^2 \rho^{2h} + \sigma_\epsilon^2 \,,$$

where $\sigma_{z_1}^2$ represents the initial variance of the persistent component. So long as $\rho < 1$,

[14]We do this in order not to have too few individuals in each (h, t)-cell (similar to Guvenen (2009)).

residual inequality has a well-defined limit, say, $var^*(\tilde{y})$. It can easily be shown that $var(\tilde{y}_h)$ will converge to $var^*(\tilde{y})$ from below in a concave fashion. 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 linear. The empirical variance profile on Figure 5 implies that the fit would be poor if ρ is too far away from 1. Targeting these moments puts an upward pressure on ρ and drives it close to 1.

Figure 5: Lifetime Profile of Residual Inequality

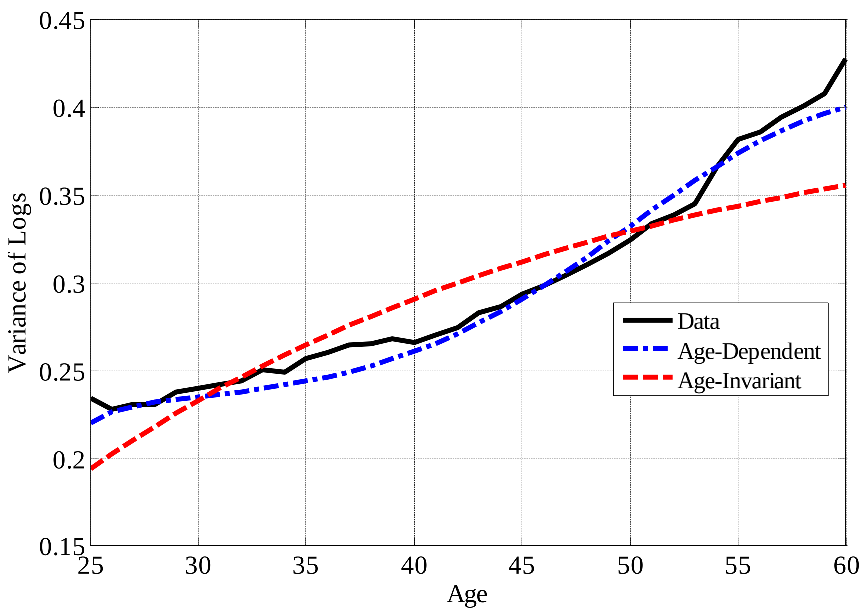

Note: This figure compares the lifetime profile of residual inequality implied by the age-dependent, age-invariant specifications, and its empirical counterpart. For the age-dependent specification, we use the estimates of the cubic specification with constant variance of transitory shocks. As for the empirical counterpart, we control for time effects.

At this point, it is worth stressing that the age-dependent income process does not need to contain a unit root or a highly persistent component to match the inequality profile. On Figure 5 we also plot the inequality profile implied by the cubic specification of the age-dependent process in the dash-dotted line. The model captures the increase in lifetime inequality even if persistence for young individuals is very low. This is by means of 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.

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 introducing this type of heterogeneity results in a lower estimate of persistence. 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. In fact, the age profile of residual inequality implied by the age-dependent process is convex for most of the life cycle.

Meghir and Pistaferri (2004) also allow for age effects while modeling conditional variances of transitory and permanent shocks, which are found to be insignificant. Their point estimates reveal a U-shape for the variance of permanent and transitory shocks. However, these are found to be statistically insignificant. In Appendix B.8.1, we provide the estimates of the age-dependent specification with constant persistence; i.e. we only let the variance of transitory and permanent shocks vary by age. We find that the transitory shocks are then U-shaped and that the variance of persistent shocks are increasing at the end of the working life but are flat for most of the life cycle.

A process containing a random walk component and an AR(1) component with age dependence in the variance of innovations (Baker and Solon (2003); Gottschalk and Moffitt (2011)) can generate most of the age dependence in the variance-covariance structure that we use to identify the age profile of persistence and variance of shocks (see Figure 4). The advantage of the age-dependent specification over this is that it is more suitable for use in quantitative life cycle macro models, since it requires one less state variable.

2.6 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 of log earnings. This is expected since the estimation targeted the moments in levels with a larger number of parameters. How about the fit for the variance-covariance 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 for the variance of income growth rates as well as one-lag covariances.

The theoretical moments for the age-dependent process (abstracting from time effects) are given by:

$$\begin{aligned} var\left(\Delta y_{i,h}\right) &= \left(\rho_{h-1}-1\right)^2 var\left(z_{i,h-1}\right)+\sigma^2_{\eta,h}+\sigma^2_{\epsilon,h}+\sigma^2_{\epsilon,h-1} \\ \text{and } cov\left(\Delta y_{i,h},\Delta y_{i,h+1}\right) &= \rho_{h-1}\left(\rho_{h-1}-1\right)\left(\rho_h-1\right) var\left(z^i_{h-1}\right)+(\rho_h-1)\sigma^2_{\eta,h}-\sigma^2_{\epsilon,h}. \end{aligned}$$

To compute the empirical counterparts, we compute these moments for all (h,t) cells and regress them on age and year dummies. The dots in the left panel of Figure 6 plot the resulting variance profile. This reveals a U-shaped profile. Theoretical moments implied by the cubic specification, shown in the solid line, show that the age-dependent process does a good job of capturing the U-shape. However, the same moments implied by the age-invariant process, shown in the dashed line, cannot match this profile.

To assess how the age-dependent process fits the covariance profile of differences at one lag, we plot the empirical and theoretical counterparts of this profile on the right panel of Figure 6. In the data, the covariance profile is almost flat over the lifetime, which is very similar to what is implied by the age-dependent specification. However, the empirical covariance is closer to zero. As for the age-invariant process, covariance profile is also flat but much further away from zero compared with the age-dependent specification.

Overall, we conclude that the age-dependent income process achieves a better fit for the moments in levels without worsening the fit for the moment structure in differences. If anything, it fits the empirical variance and one-lag covariance profiles better than the age-invariant specification.

Figure 6: Variance Profile of Income Growth Rates

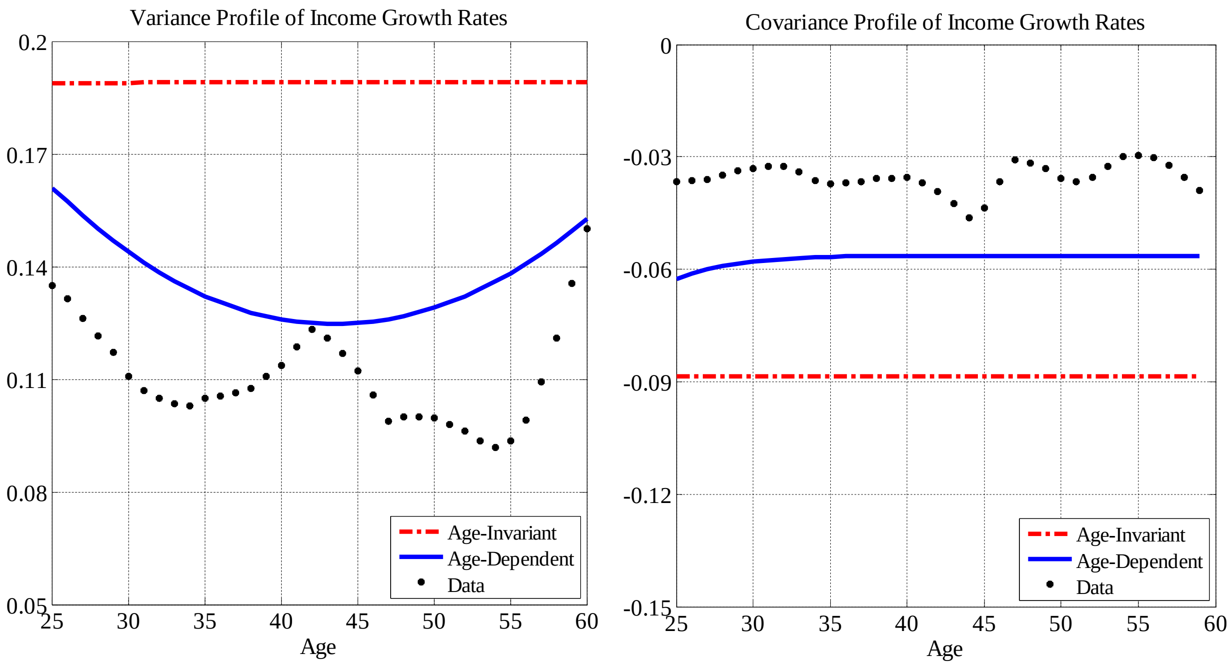

3 An Economic Rationale for the Age-Dependent Specification

Through a series of econometric analyses, we have shown that the 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 profiles. 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.

Note that the increase in persistence and decrease in variance of persistent shocks take place in the first 10 years of the working life, which coincides with the period where job turnover of workers is high (see Topel and Ward (1992)). Thus, another mechanism, again related to mobility, would be 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 the National Longitudinal Survey of Youth (NLSY/66) in favor of this theory. We now study the wage dynamics implied by this model.

3.1 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 are a continuum 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, ξ, from a normal distribution with mean μ_ξ and variance σ_ξ^2. The match-specific productivity is not known by the firm or the worker. Employed workers with tenure, t, receive their compensation, w_t, before production takes place. Output of the match, y_t, is given by $y_t = \xi + \nu_t$, where ν_t is an i.i.d. normal random variable with mean 0 and variance σ_ν^2. After observing the output, workers and employers update their beliefs about the match productivity in a Bayesian fashion. Since the information set of the worker and the firm are the same, their beliefs are identical. By means of normality assumptions, this belief is normally distributed as well.

Let $m_{t|t-1}$ denote the mean of the belief about ξ in period t conditional on all of the information up to period $t-1$, and let $1/p_t$ denote the variance, thereby p_t denoting the precision. Similarly, $p_\mu = 1/\sigma_\mu^2$ and $p_\nu = 1/\sigma_\nu^2$ denote the precision of the distribution of ξ and ν_t, respectively. Finally, $\omega_t \sim N(0, 1/p_t)$ represents the deviation of the belief from the true productivity. The law of motion for these are governed by:

$$\begin{aligned} m_{t+1|t} &= m_{t|t-1}\frac{p_t}{p_t + p_\nu} + y_t\frac{p_\nu}{p_t + p_\nu}, \\ p_t &= p_\mu + (t-1)p_\nu, \\ \text{and } y_t &= \underbrace{m_{t|t-1} + \omega_t}_{\xi} + \nu_t. \end{aligned} \tag{3}$$

For simplicity, we assume that firms pay workers their expected productivity before production takes place $\left(w_t = m_{t|t-1}\right)$. After updating the beliefs, a worker decides whether to break the match. We assume that upon breaking the match, she immediately meets another employer with a new match productivity.[15]

[15]The initial beliefs are given by the unconditional mean of the distribution for match productivity, thus they are the same for every quitter ($m_{1|0} = \mu_\xi$).

3.2 A Quantitative Evaluation of the Model

In order to evaluate the performance of this model on earnings dynamics, we calibrate the model, simulate it, and then estimate the age-dependent income process using residual wages from simulated data. Our exercise shows that the model has the potential to replicate our empirical findings for non-flat age profiles.

3.2.1 Calibration

This is a fairly stylized model with only five parameters: r, δ, μ_ξ, σ_ξ^2, and σ_ν^2. The model period is one year. The interest rate, r, is set to an annual rate of 3 percent. We set δ to 1/37 to match an average working life of 37 years, motivated by our dataset. The model allows the normalization of the mean of match productivity; we set μ_ξ to a computationally convenient value.

We calibrate the remaining two parameters; the variance of match productivity, σ_ξ^2, and the variance of the i.i.d. shock, σ_ν^2, by targeting two moments from our empirical findings. σ_ξ^2 has a pronounced effect on the level of the variance of persistent shocks. In the data, changes in the variance of persistence shocks at older ages are due to reasons not captured by this model (for example, health shocks). Thus, we target the average of the first 10 years' variance of persistent shocks.

In the model, an increase in σ_ν^2 increases the time it takes for the match quality to be revealed. This increase in turn increases the time to settle down into jobs, which can be approximated by average persistence over the last 25 years.[16] Our second target is therefore

[16] As we discuss in the next section, persistence increases as workers settle down into jobs.

this average. At no point in the calibration do we target the profile of persistence and the variance of shocks. Table 4 summarizes our calibration exercise.

Table 4: Calibrating Model Parameters

Parameter	Value
r, interest rate	3%
δ, death probability	1/37
μ_ξ, mean of match productivity	10
σ_ξ^2, variance of match productivity	0.50
σ_ν^2, variance of iid productivity shock	0.50

Empirical Moments Used in Calibration

Moment	Data	Model
Average variance of persistent shocks in first 10 years	0.0495	0.0490
Average persistence profile in last 25 years	0.968	0.965

3.2.2 Simulation Results

We simulate 10,000 individuals, run the first stage regressions to obtain the residuals, and estimate the nonparametric specification of the age-dependent process. Figure 7 shows the results.

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 a random walk. Namely, $w_t = m_{t|t-1}$. Equation (3) implies that $w_{t+1} = w_t\frac{p_t}{p_t+p_\nu} + y_t\frac{p_\nu}{p_t+p_\nu} = w_t\left(\frac{p_t}{p_t+p_\nu} + \frac{p_\nu}{p_t+p_\nu}\right) + \frac{p_\nu}{p_t+p_\nu}(\omega_t + \nu_t) = w_t + \chi_t$, where $\chi_t \sim N(0, \frac{p_\nu}{p_t(p_t+p_\nu)})$. On the other hand, job switchers always get the unconditional mean of the match-specific component μ_ξ, implying a low correlation between current and future wages. Therefore, persistence is lower for them. The persistence of income changes in the overall sample is a combination of the persistence of these two subsamples. Over the lifetime, the fraction of switchers declines with age because workers settle into more productive jobs as they age.[17] Thus, they are less likely to switch to other jobs. This implies a rising persistence profile. Furthermore, the bottom panel of Figure 7 shows a decreasing variance profile for persistent shocks. This decrease is because both the number of stayers increases and the variance of innovations to wages declines with tenure for stayers. Namely, the variance of χ_t is decreasing, since p_t is increasing in t.

Figure 7: Simulation Results for the Learning Model

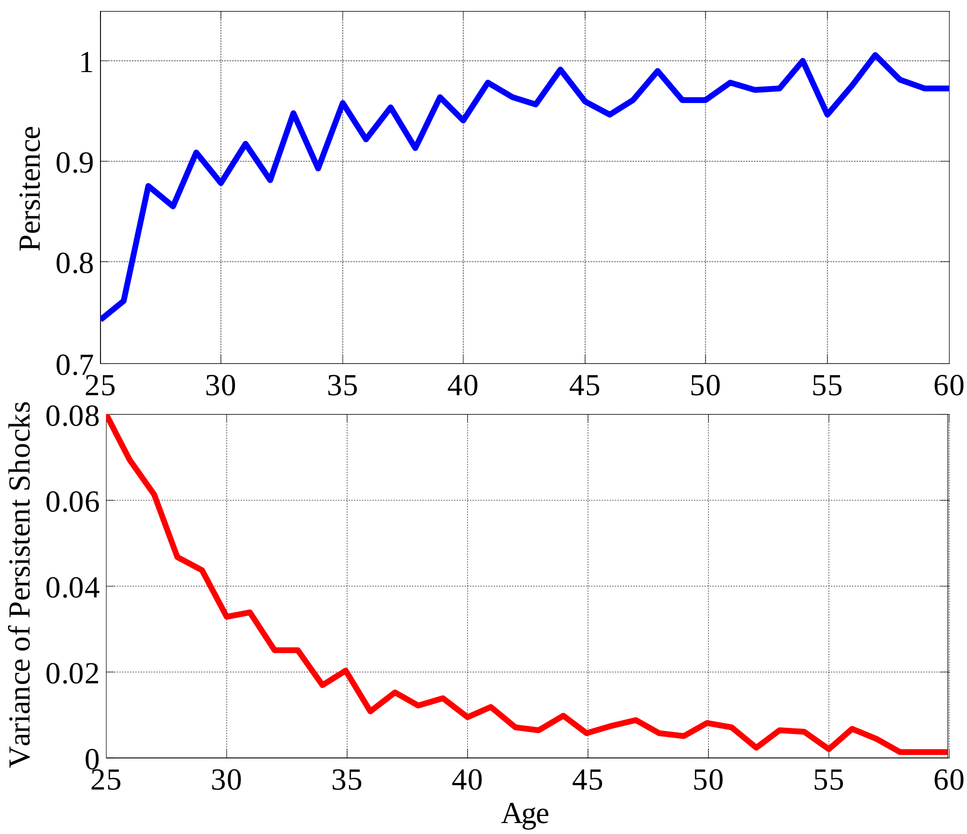

This section presented a theoretical background for our empirical findings. We have

[17]This is consistent with the empirical findings on worker turnover (see Topel and Ward (1992)).

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.

4 Consumption-Savings Implications

There is a large literature that rejects full consumption insurance for the U.S. economy (Cochrane (1991); Mace (1991); and Attanasio and Davis (1996)) making the nature of labor income risk an important object for economic research. This paper so far has established the existence of a non-flat lifetime profile in persistence and variance of shocks. We now investigate its consumption-savings 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 the age-invariant process.

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

$$E \sum_{h=1}^{H} \beta^h u\left(c_h^i\right),$$

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 only buy and sell 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

$$c_h^i + \frac{a_{h+1}^i}{1+r} \;=\; a_h^i + y_h^i\,, \tag{4}$$

where y_h^i is the labor earnings at age h. Agents face an age-dependent borrowing constraint, $\bar{A}_h$. We study welfare costs in two economies: a natural borrowing constraint economy (NBC) and a zero borrowing constraint economy (ZBC).[18] 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 is composed of both a deterministic part, which is common to everyone, and an idiosyncratic component, which captures individuals' earnings risk. We consider two specifications for the idiosyncratic component: i) the age-dependent income process and ii) the age-invariant process as we discussed in Section 2.5. The first is calibrated according to the cubic specification with constant variance of

[18] The natural borrowing limit is the maximum amount that an agent can pay back for sure out of future earnings.

transitory shocks reported in Table 1. The parameters of the latter come from the estimates reported in Table 3. The deterministic component of earnings is estimated using the PSID data.

There is a social security system that pays a pension after retirement. 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 Guvenen, Kuruscu, and Ozkan (2009) and is set to mimic the properties of the US social security system.

One period in our model corresponds to a calendar year. Agents enter the economy at age 24, retire at 60, and die with probability 1 at age 84. We assume CRRA preferences and set the parameter of relative risk aversion to 2. We take the risk-free interest rate to be 3 percent. We pin down the discount factor β by targeting an aggregate wealth to income ratio of 3. The Bellman equations of the model, and further detail of its calibration are explained in Appendix C.

4.1 Consumption Insurance against Labor Income Shocks

We now turn to the differences in consumption insurance induced by the age-dependent and the age-invariant processes. For each specification, we calibrate the discounting factor, β, to match an aggregate wealth to income ratio of 3. We compute the degree of consumption insurance at age h as:

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

where η_h^i is the persistent shock faced by worker i at age h. This measures the amount of change in the persistent component that does not translate into consumption growth.

Figure 8 plots ϕ_h over the life cycle for both processes in the NBC and ZBC economies. It is clear that persistent shocks from the age-dependent process are better insured throughout the lifetime. In the NBC economy, on average, 56 percent of persistent shocks are insured under the age-dependent process, whereas the corresponding number for the age-invariant process is only 40 percent.

Strikingly, most of this difference comes from younger adults. In the age-invariant process, the profile of insurance tracks the profile of assets. This difference is because persistence is constant and high throughout the working life, and agents abstain from borrowing in response to a highly persistent bad income shock. Therefore, insurance against such shocks is mainly through assets, which young households have very little of. The increase in assets over the lifetime allows them to fare better against highly persistent shocks, which results in an increasing consumption insurance profile.

However, for the age-dependent income process, insurance for young households is much larger, precisely because the $AR(1)$ component is moderately persistent for them. Consumption insurance first decreases until middle age and then increases until the end of working life. This U-shape happens because of the combination of two effects: i) households get richer and can better insure themselves against persistent shocks and ii) persistence of shocks becomes larger, making them harder to insure. In the initial phase of the life cycle (24-40) the latter dominates the former, and insurance decreases with age. Later on, assets are large enough that they compensate the increase in persistence. Thus, insurance increases with age.

Similar comments apply for the ZBC economy. Insurance decreases for both specifications once we impose no borrowing, but it is still larger under the age-dependent income process, though by a smaller margin. The decrease in the gap is due to young households, who can insure against moderately persistent shocks via borrowing in an NBC economy.

Figure 8: Insurance Against Persistent Shocks

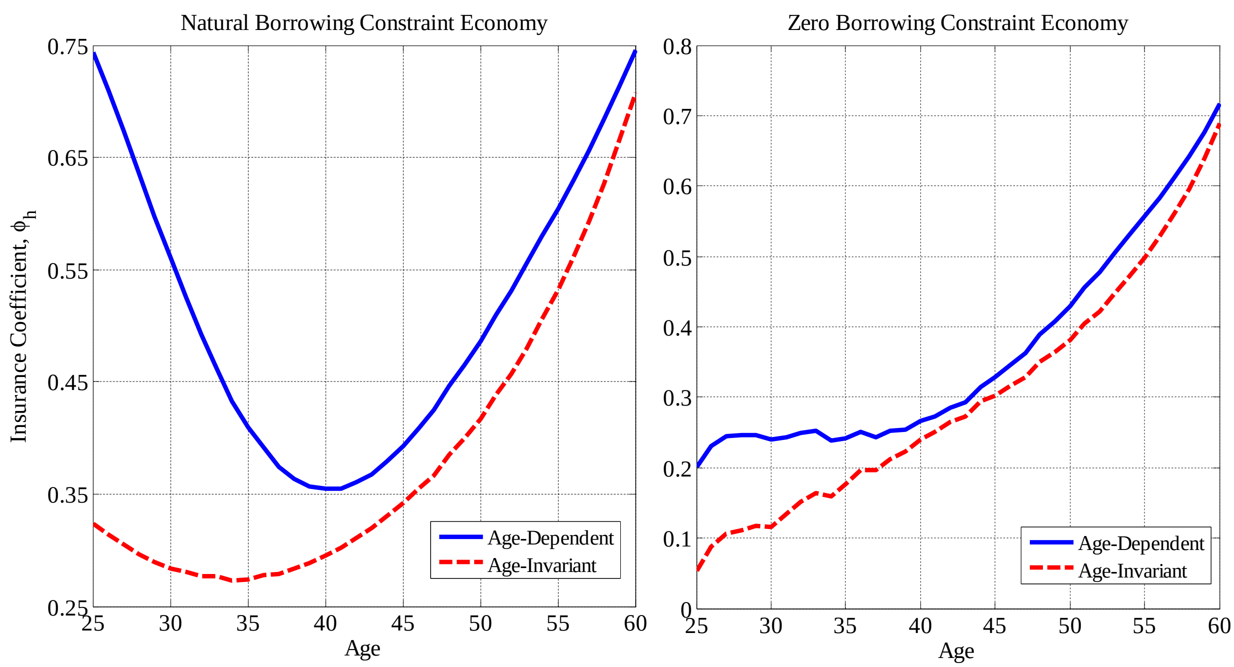

4.2 Welfare Costs of Earnings Risk

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 1 and 2). 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.[19]

Table 5: Welfare Costs under Different Income Processes

	Welfare Costs (1)	(2)	Insurance (3)
	NBC Economy		
Age-Dependent	15.22%	3.13%	0.56
Age-Invariant	14.73%	4.76%	0.38
	ZBC Economy		
Age-Dependent	18.20%	5.80%	0.38
Age-Invariant	16.8%	6.80%	0.30

Note: Column (1) shows the welfare cost of total idiosyncratic risk including risk due to fixed effects, initial variation in persistent component (z_1^i) as well as life-cycle shocks. Column (2) presents welfare costs of life-cycle shocks, i.e., shocks accumulated after workers enter labor force. Column (3) shows the insurance coefficient against persistent shocks.

The upper panel of Table 5 shows the results for the NBC economy. Column (1) shows the total welfare cost of idiosyncratic risk, that is, welfare costs of income risk for a person who has not entered the labor force yet. The first two rows correspond to the age-dependent and age-invariant processes, respectively. The age-dependent income process delivers higher welfare costs. This is due to the fact that the level of inequality in the beginning and at the end of the life cycle is lower for the age-invariant specification (see Figure 5). At this point it is not clear how much of these differences is driven by shocks and how much is driven by differences in initial conditions. In column (2) we report the welfare costs of shocks over

[19]The 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 ($\gamma = 2$).

the life cycle for a person with the average fixed effect ($\alpha^i = 0$) and the average initial persistent component ($z_1^i = 0$). The differences are significant in the NBC economy: An agent in a world with age-invariant income process is willing to give up 4.76 percent of her consumption every period in order to have perfect insurance. The same number is only 3.13 percent for an agent in the age-dependent world.[20]

The bottom panel of Table 5 presents the results for the ZBC economy.[21] As expected, welfare costs have increased compared with the NBC economy for both specifications. Note that the increase is larger for the age-dependent process, and thus, the difference between the two processes decreases. However, welfare costs are still lower for the age-dependent process (5.8 percent vs. 6.8 percent). 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.

These results 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 presents 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

[20]There is a caveat in this analysis. The increase in inequality implied by age-dependent specification is slightly higher than the age-invariant process. Moreover, inequality both at the beginning and at the end of working life are higher for the age-dependent specification. These might bias the welfare costs for the age-dependent process. In an earlier version of this paper (Karahan and Ozkan (2009)), we conducted two experiments, where we changed the parameters of the age-invariant process to match the level and the rise of inequality implied by the age-invariant process. The results are qualitatively similar.

[21]For 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 against the increase in earnings.

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.

5 Conclusion

Most of the existing literature on income processes has assumed constant persistence and variance of income shocks over the life cycle. As a result, macroeconomists have calibrated life-cycle models using these flat profiles. In this paper, we have estimated a novel specification for labor income risk that allows the persistence and variance of shocks to change over the lifetime. Our results reveal that persistence is only moderate for young workers and increases up to unity by age 40. The variance of persistent shocks exhibits a U-shaped profile. These results suggest that the standard specification in the quantitative macro literature (with constant persistence and variances) cannot capture the earnings dynamics of young workers. We have also argued that these non-flat profiles have significant implications for consumption insurance. The welfare costs of idiosyncratic risk implied by the age-dependent income process is significantly lower compared with the age-invariant process. This difference has important implications for the evaluation of redistributive policies.

There is a large literature that has mostly focused on statistical representations of idiosyncratic income risk. However, there is less work connecting wage generating structural models to these income processes.[22] Using a structural model of worker turnover, this paper argues that the high job mobility of young workers can explain the earnings dynamics

[22] Notable exceptions include Huggett, Ventura, and Yaron (2006), and Postel-Vinay and Turon (2010).

implied by the age-dependent process. As a future work, we plan to investigate whether these non-flat profiles can help us tell different theories of wages apart.

References

Abowd, J. M., and D. Card (1989): "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, 57(2), 411–45. 9

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. 11

Alvarez, J., M. Browning, and M. Ejrnæs (2010): "Modelling Income Processes with Lots of Heterogeneity," *Review of Economic Studies*, 771-4, 1353–1381. 4

Attanasio, O., and S. J. Davis (1996): "Relative Wage Movements and the Distribution of Consumption," *Journal of Political Economy*, 104(6), 1227–62. 32

Baker, M. (1997): "Growth-Rate Heterogeneity and the Covariance Structure of Life-Cycle Earnings," *Journal of Labor Economics*, 15(2), 338–375. 9

Baker, M., and G. Solon (2003): "Earnings Dynamics and Inequality among Canadian Men, 1976-1992: Evidence from Longitudinal Income Tax Records," *Journal of Labor Economics*, 21, 289–321. 4, 24

Bound, J., C. Brown, G. J. Duncan, and W. L. Rodgers (1994): "Evidence on the Validity of Cross-Sectional and Longitudinal Labor Market Data," *Journal of Labor Economics*, 12, 345–368. 53

Cagetti, M. (2003): "Wealth Accumulation over the Life Cycle and Precautionary Savings," *Journal of Business and Economic Statistics*, 21, No. 3, 339–353. 63

Cochrane, J. H. (1991): "A Simple Test of Consumption Insurance," *Journal of Political Economy*, 99(5), 957–76. 32

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. 4

Flinn, C. J. (1986): "Wages and Job Mobility of Young Workers," *Journal of Political Economy*, Vol. 94, No. 3, 88–110. 6, 27, 32

Gottschalk, P., and R. A. Moffitt (1995): "Trends in the Transitory Variance of Earnings in the United States," *Working Paper*. 9

——— (2011): "Trends in the Transitory Variance of Male Earnings in the U.S., 1970-2004," *Working Paper*. 4, 9, 24

Gourinchas, P.-O., and J. A. Parker (2002): "Consumption over the Life Cycle," *Econometrica*, 70(1), 47–89. 63

Guvenen, F. (2009): "An Empirical Investigation of Labor Income Processes," *Review of Economic Dynamics*, 12, No.1, 58–79. 21, 23

Guvenen, F., B. Kuruscu, and S. Ozkan (2009): "Taxation of Human Capital and Wage Inequality: A Cross-Country Analysis," *NBER Working Paper*. 34, 64

Hause, J. C. (1980): "The Fine Structure of Earnings and the On-the-Job Training Hypothesis," *Econometrica*, 48(4), 1013–1029. 4

Heathcote, J., F. Perri, and G. L. Violante (2010): "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006," *Review of Economic Dynamics*, 13(1), 15–51. 9

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. 10

Heathcote, J., K. Storesletten, and G. L. Violante (2010): "The Macroeconomic Implications of Rising Wage Inequality in the United States," *Journal of Political Economy*, 118(4), 681–722. 10

Hryshko, D. (2011): "Excess Smoothness of Consumption in an Estimated Life Cycle Model," *University of Alberta Working Paper*. 9

Huggett, M., G. Ventura, and A. Yaron (2006): "Human capital and earnings distribution dynamics," *Journal of Monetary Economics*, 53, 265–290. 39

Jovanovic, B. (1979): "Job Matching and the Theory of Turnover," *Journal of Political Economy*, Vol. 87, No. 5, Part 1, 972–990. 5, 27, 32

Kambourov, G., and I. Manovskii (2008): "Rising Occupational and Industry Mobility in the United States: 1968-1997," *International Economic Review*, 49, 41–79. 3

Kaplan, G. (2010): "Inequality and the Lifecycle," *Working Paper*. 21

Karahan, F., and S. Ozkan (2009): "On the Persistence of Income Shocks over the Life Cycle: Evidence and Implications," *PIER Working Paper*. 20, 38

Krueger, D., F. Perri, L. Pistaferri, and G. L. Violante (2010): "Cross-sectional Facts for Macroeconomists," *Review of Economic Dynamics*, 13, 1–14. 24

Lillard, L. A., and Y. Weiss (1979): "Components of Variation in Panel Earnings Data: American Scientists 1960-70," *Econometrica*, 47(2), 437–454. 9

Lillard, L. A., and R. J. Willis (1978): "Dynamic Aspects of Earning Mobility," *Econometrica*, Vol. 46, No.5, 985–1012. 9

Mace, B. J. (1991): "Full Insurance in the Presence of Aggregate Uncertainty," *Journal of Political Economy*, 99, 928–956. 32

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. 9

Meghir, C., and L. Pistaferri (2004): "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72(1), 1–32. 16, 23

Panousi, V., I. Vidangos, B. T. Heim, and J. M. DeBacker (2011): "Rising Inequality: Transitory or Permanent? New Evidence from a Panel of U.S. Tax Returns 1987-2006," *Indiana University-Bloomington: School of Public & Environmental Affairs Research Paper Series*, No. 2011-01-01. 9

Postel-Vinay, F., and H. Turon (2010): "On-the-Job Search, Productivity Shocks, and the Individual Earnings Process," *International Economic Review*, 51(3), 599–629. 39

Sabelhaus, J., and J. Song (2010): "The Great Moderation in Micro Labor Earnings," *Journal of Monetary Economics*, 57, 391–403. 4

Storesletten, K., C. I. Telmer, and A. Yaron (2004): "Consumption and Risk Sharing over the Life Cycle," *Journal of Monetary Economics*, 51, 609–633. 63

Topel, R. H., and M. P. Ward (1992): "Job Mobility and the Careers of Young Men," *The Quarterly Journal of Economics*, 107(2), 439–79. 27, 31

APPENDICES

A Identification

Here, we provide the proof of identification for the age-dependent specification in (2). The variance-covariance structure of this specification is given by:

$$var(\tilde{y}^i_{h,t}) = \sigma^2_\alpha + var(z^i_{h,t}) + \phi^2_t \sigma^2_{\epsilon,h} \quad (5)$$

$$cov(y^i_{h,t}, y^i_{h+n,t+n}) = \sigma^2_\alpha + \rho_h \rho_{h+1} \cdots \rho_{h+n-1} var\left(z^i_{h,t}\right) \quad (6)$$

$$var(z^i_{h,t}) = \rho^2_{h-1} var(z^i_{h-1,t-1}) + \pi^2_t \sigma^2_{\eta,h} \quad (7)$$

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

Proof of Proposition 1: We start by assuming that we know the variance of the fixed effect, σ^2_α, 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^2_α.

Note that since we assume that σ^2_α is known, we can construct $cov\left(\tilde{y}^i_{h,t}, \tilde{y}^i_{h+n,t+n}\right) - \sigma^2_\alpha$. (6) implies $\left[cov\left(\tilde{y}^i_{h,t}, \tilde{y}^i_{h+2,t+2}\right) - \sigma^2_\alpha\right] / \left[cov\left(\tilde{y}^i_{h,t}, \tilde{y}^i_{h+1,t+1}\right) - \sigma^2_\alpha\right] = \rho_{h+1}$ for $h = 1, \ldots, H-2$. This pins down the whole profile of ρ_h for $h = 2, 3, \ldots, H-1$.[23] Note also that by normalization $\rho_1 = \rho_2$.

Now, our goal is to recover the schedule of $var\left(z^i_{h,t}\right)$. Once we recover these, we can use (7) to identify the loading factors and variances of persistent shocks, $\{\pi_t\}_{t=1}^{t=T}$ and

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

$\left\{\sigma_{\eta,h}^2\right\}_{h=2}^{h=H-1}$. Note that

$$\frac{cov\left(\tilde{y}_{h,t}^i, \tilde{y}_{h+1,t+1}^i\right) - \sigma_\alpha^2}{\rho_h} = var\left(z_{h,t}^i\right) \tag{8}$$

Since ρ_h is pinned down for $h \geq 1$, (8) recovers $var\left(z_{h,t}^i\right)$ for $h = 1, \ldots, H-1$, $t = 1, \ldots, T-1$. Please note that $var(z_{H,t}^i)$ for $t = 1, .., T$ and $var(z_{h,T}^i)$ for $h = 1, .., H$ are not identified yet.

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 $cov\left(\tilde{y}_{2,1}^i, \tilde{y}_{5,4}^i\right)$suffices to recover σ_α^2 uniquely:

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

$$\Rightarrow \frac{cov\left(\tilde{y}_{2,1}^i, \tilde{y}_{5,4}^i\right) - \sigma_\alpha^2}{cov\left(\tilde{y}_{2,1}^i, \tilde{y}_{4,3}^i\right) - \sigma_\alpha^2} = \frac{cov\left(\tilde{y}_{3,1}^i, \tilde{y}_{5,3}^i\right) - \sigma_\alpha^2}{cov\left(\tilde{y}_{3,1}^i, \tilde{y}_{4,2}^i\right) - \sigma_\alpha^2}$$

$$\Rightarrow \sigma_\alpha^2 = \frac{cov\left(\tilde{y}_{2,1}^i, \tilde{y}_{4,3}^i\right) cov\left(\tilde{y}_{3,1}^i, \tilde{y}_{5,3}^i\right) - cov\left(\tilde{y}_{2,1}^i, \tilde{y}_{5,4}^i\right) cov\left(\tilde{y}_{3,1}^i, \tilde{y}_{4,2}^i\right)}{cov\left(\tilde{y}_{2,1}^i, \tilde{y}_{4,3}^i\right) + cov\left(\tilde{y}_{3,1}^i, \tilde{y}_{5,3}^i\right) - cov\left(\tilde{y}_{2,1}^i, \tilde{y}_{5,4}^i\right) - cov\left(\tilde{y}_{3,1}^i, \tilde{y}_{4,2}^i\right)}$$

Now, we are ready to identify the loading factors and variances of persistent shocks. Using the normalization that $\pi_1 = 1$, we get $\sigma_{z,1}^2$. Tracking $var\left(z_{1,t}^i\right)$ along t identifies π_t for $t = 2, \ldots, T-1$. Consequently, tracing (7) along the age dimension identifies $\sigma_{\eta,h}^2$ for $h = 2, \ldots, H-1$. By assumption $\sigma_{\eta,H}^2 = \sigma_{\eta,H-1}^2$ which gives us $var(z_{H,1}^i)$.

Now, our goal is to recover π_T. First, we identify $\sigma^2_{\epsilon,1}$ using equation 5 for $h = 1$ and $t = 1$. Then again using equation 5 for $h = 1$, $t = T$, we can get $var(z^i_{1,T})$. Equation 7 for $h = 1$ and $t = T$ pins down π_T. We now 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 5 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 5 is

$$var(\tilde{y}^i_{h,1}) - \sigma^2_\alpha - var(z^i_{h,1}) = \sigma^2_{\epsilon,h} \quad \text{for } h = 1, \ldots, H$$

identifying $\sigma^2_{\epsilon,h}$ over the life cycle (except for $h = H - 1$). Fixing h, tracking 5 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 $t = 2, \ldots, T-1$.

B Robustness Checks and Other Estimation Results

B.1 Parameter Estimates for the Nonparametric Specification of the Age-Dependent Income Process

Here, we present the point estimates of the nonparametric specification as well as their bootstrap standard errors. These results are plotted in Figures 1-3.

Table 6: Estimation Results for the Nonparametric Specification: Age Profiles

Age	ρ_{age}	$\sigma^2_{\eta,age}$	$\sigma^2_{\varepsilon,age}$	Age	ρ_{age}	$\sigma^2_{\eta,age}$	$\sigma^2_{\varepsilon,age}$	Age	ρ_{age}	$\sigma^2_{\eta,age}$	$\sigma^2_{\varepsilon,age}$
24	0.6929	0.1100	0.0586	37	0.9213	0.0250	0.0774	50	0.9266	0.0060	0.0649
	(0.0498)	(0.0802)	(0.0149)		(0.0372)	(0.0074)	(0.0116)		(0.0468)	(0.0067)	(0.0125)
25	0.6929	0.0802	0.0277	38	1.0485	0.0079	0.0736	51	0.9828	0.0045	0.0613
	(0.0498)	(0.0152)	(0.0149)		(0.0229)	(0.0054)	(0.0102)		(0.0494)	(0.0055)	(0.0119)
26	0.6730	0.0620	0.0218	39	0.9316	0.0190	0.0642	52	1.0485	0.0266	0.0801
	(0.0389)	(0.0137)	(0.0137)		(0.0384)	(0.0077)	(0.0122)		(0.0280)	(0.0108)	(0.0175)
27	0.8837	0.0504	0.0430	40	0.9732	0.0122	0.0635	53	0.8934	0.0090	0.0782
	(0.0403)	(0.0098)	(0.0133)		(0.0344)	(0.0082)	(0.0113)		(0.0607)	(0.0095)	(0.0178)
28	0.7924	0.0533	0.0287	41	0.9744	0.0129	0.0498	54	0.9782	0.0422	0.0700
	(0.0458)	(0.0116)	(0.0141)		(0.0343)	(0.0063)	(0.0118)		(0.0504)	(0.0144)	(0.0160)
29	0.8279	0.0371	0.0433	42	0.9488	0.0122	0.0506	55	0.9821	0.0241	0.0696
	(0.0448)	(0.0108)	(0.0124)		(0.0413)	(0.0072)	(0.0120)		(0.0418)	(0.0110)	(0.0157)
30	0.8779	0.0253	0.0536	43	0.9860	0.0224	0.0638	56	0.8938	0.0414	0.0617
	(0.0369)	(0.0074)	(0.0104)		(0.0383)	(0.0085)	(0.0124)		(0.0487)	(0.0150)	(0.0165)
31	0.9077	0.0311	0.0669	44	0.9543	0.0110	0.0737	57	1.0284	0.0361	0.0879
	(0.0396)	(0.0084)	(0.0122)		(0.0413)	(0.0071)	(0.0143)		(0.0343)	(0.0131)	(0.0161)
32	0.9216	0.0155	0.0628	45	0.8697	0.0182	0.0729	58	0.8560	0.0042	0.0601
	(0.0387)	(0.0088)	(0.0129)		(0.0500)	(0.0081)	(0.0141)		(0.0632)	(0.0091)	(0.0174)
33	0.9040	0.0308	0.0510	46	1.0465	0.0111	0.0742	59	1.0141	0.0386	0.0928
	(0.0386)	(0.0095)	(0.0124)		(0.0257)	(0.0075)	(0.0122)		(0.0468)	(0.0148)	(0.0186)
34	0.9457	0.0202	0.0635	47	0.9887	0.0128	0.0719	60		0.0386	0.0315
	(0.0426)	(0.0095)	(0.0125)		(0.0403)	(0.0064)	(0.0133)			(0.0148)	(0.0214)
35	0.9633	0.0186	0.0613	48	1.0500	0.0113	0.0756				
	(0.0355)	(0.0074)	(0.0115)		(0.0203)	(0.0072)	(0.0120)		σ^2_α		
36	0.9078	0.0035	0.0601	49	1.0005	0.0167	0.0734		0.0752		
	(0.0386)	(0.0058)	(0.0112)		(0.0461)	(0.0114)	(0.0136)		(0.0115)		

Note: Table shows the point estimates of the nonparametric specification of the age profiles as well as the variance of fixed effects. The numbers in parenthesis are bootstrap standard errors.

Table 7: Estimation Results for the Nonparametric Specification: Time Loading Factors

Year, t	π_t	ϕ_t	Year, t	π_t	ϕ_t	Year, t	π_t	ϕ_t
1968	0.7181	0.8497	1978	0.8495	0.8798	1988	1.2158	1.2287
	(0.0704)	(0.0476)		(0.0442)	(0.0412)		(0.0612)	(0.0465)
1969	1.1811	0.8299	1979	1.1338	0.8064	1989	1.2907	1.2133
	(0.0588)	(0.0433)		(0.0572)	(0.0416)		(0.0636)	(0.0560)
1970	1.0494	0.7575	1980	1.0000	1.0000	1990	1.1645	1.1522
	(0.0611)	(0.0455)		(0.0000)	(0.0000)		(0.0623)	(0.0452)
1971	0.9618	1.0008	1981	1.1897	0.8752	1991	1.2185	1.1295
	(0.0518)	(0.0420)		(0.0536)	(0.0431)		(0.0551)	(0.0487)
1972	1.0447	0.8159	1982	1.5450	0.8848	1992	1.2285	1.2901
	(0.0562)	(0.0401)		(0.0671)	(0.0431)		(0.0529)	(0.0514)
1973	0.8620	0.9024	1983	1.2007	1.1962	1993	0.6559	1.4131
	(0.0467)	(0.0443)		(0.0527)	(0.0509)		(0.0485)	(0.0563)
1974	0.9404	0.9295	1984	1.5082	1.2010	1994	1.3502	1.2017
	(0.0512)	(0.0357)		(0.0619)	(0.0496)		(0.0580)	(0.0516)
1975	1.0397	0.9441	1985	1.2881	1.1673	1995	1.2234	1.1426
	(0.0550)	(0.0405)		(0.0597)	(0.0462)		(0.0523)	(0.0475)
1976	1.0176	0.9566	1986	1.2068	1.1726	1996	1.1253	1.1115
	(0.0528)	(0.0437)		(0.0577)	(0.0477)		(0.0493)	(0.0431)
1977	1.2315	0.8932	1987	1.1426	1.2078	1997	0.3000	1.1115
	(0.0610)	(0.0402)		(0.0604)	(0.0501)		(0.0262)	(0.0431)

Note: Table shows the point estimates of the nonparametric specification of the time loading factors. The numbers in parenthesis are bootstrap standard errors. π_t and ϕ_t are the loading factors of the persistent and transitory shocks at time t, respectively.

B.2 Results with Wage Data

Table 8: 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^2_η	0.0495 (0.0089)	-0.0033 (0.0009)	0.0001 (0.0000)	$H_0 : \gamma_{\sigma^2_\eta,1} \geq 0$ 0.0000	$H_0 : \gamma_{\sigma^2_\eta,2} \leq 0$ 0.0000
σ^2_α	0.0695 (0.0236)				
σ^2_ϵ	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 9: 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^2_η	0.0280 (0.0073)	0.0133 (0.0038)	0.0243 (0.0069)	$H_0 : \sigma^2_{\eta,1} \leq \sigma^2_{\eta,2}$ 0.0000	$H_0 : \sigma^2_{\eta,2} \geq \sigma^2_{\eta,3}$ 0.0480
σ^2_α	0.0699 (0.0102)				
σ^2_ϵ	0.0522 (0.0171)				

* The numbers in brackets are standard errors.
** The last three columns report the p-values of the corresponding tests.

B.3 Transitory Component Modeled as MA(1)

Table 10: Estimation and Test Results with an MA(1) Component

	σ^2_α	$\sigma^2_{z_1}$	θ
	0.0575	0.0474	0.174
	(0.0308)	(0.0169)	(0.0208)

	$[24,33]$	$[34,52]$	$[53,60]$	Test 1 p-value	Test 2 p-value
				$H_0: \rho_1 \geq \rho_2$	$H_0: \rho_2 \leq \rho_3$
ρ	0.909	0.979	0.966	0.00	0.21
	(0.0245)	(0.0141)	(0.0190)		
				$H_0: \sigma^2_{\eta,2} \leq \sigma^2_{\eta,1}$	$H_0: \sigma^2_{\eta,2} \geq \sigma^2_{\eta,3}$
σ^2_η	0.0131	0.0068	0.0142	0.01	0.047
	(0.0022)	(0.0016)	(0.0035)		
				$H_0: \sigma^2_{\epsilon,2} \leq \sigma^2_{\epsilon,1}$	$H_0: \sigma^2_{\epsilon,2} \geq \sigma^2_{\epsilon,3}$
σ^2_ϵ	0.0639	0.0627	0.0715	0.59	0.22
	(0.0072)	(0.0065)	(0.011)		

Notes:Estimated process: $\tilde{y}^i_{h,t} = \alpha^i + z^i_{h,t} + \varepsilon^i_{h,t} + \nu^i_{h,t}$, where $z^i_{h,t} = \rho_{h-1} z^i_{h-1,t-1} + \eta^i_{h,t}$, $\varepsilon^i_{h,t} = \epsilon^i_{h,t} + \theta\epsilon^i_{h-1,t-1}$, $\eta^i_{h,t} \sim N(0, \sigma^2_{\eta,h})$ and $\epsilon^i_{h,t} \sim N(0, \sigma^2_{\epsilon,h})$. [24,33], [34,52], [53,60] are age intervals in which the persistence and variances are assumed to be constant. Point estimates are shown in the first three columns along with bootstrap standard errors in parenthesis. The last two columns report p-values for bootstrap significance tests.

B.4 External Estimate of Measurement Error

Table 11: Estimation and Test Results with an MA(1) Component and External Measurement Error

	σ^2_α	$\sigma^2_{z_1}$	θ
	0.0588	0.0450	0.218
	(0.0285)	(0.0143)	(0.0268)

	[24, 33]	[34, 52]	[53, 60]	Test 1 p-value	Test 2 p-value
				$H_0 : \rho_1 \geq \rho_2$	$H_0 : \rho_2 \leq \rho_3$
ρ	0.908	0.979	0.965	0.00	0.21
	(0.0248)	(0.0136)	(0.0186)		
				$H_0 : \sigma^2_{\eta,2} \leq \sigma^2_{\eta,1}$	$H_0 : \sigma^2_{\eta,2} \geq \sigma^2_{\eta,3}$
σ^2_η	0.0126	0.0065	0.0137	0.00	0.035
	(0.0027)	(0.0018)	(0.0039)		
				$H_0 : \sigma^2_{\epsilon,2} \leq \sigma^2_{\epsilon,1}$	$H_0 : \sigma^2_{\epsilon,2} \geq \sigma^2_{\epsilon,3}$
σ^2_ϵ	0.0491	0.0479	0.0563	0.59	0.30
	(0.0066)	(0.0061)	(0.0112)		

Notes:Estimated process: $\tilde{y}^i_{h,t} = \alpha^i + z^i_{h,t} + \varepsilon^i_{h,t} + \nu^i_{h,t}$, where $z^i_{h,t} = \rho_{h-1} z^i_{h-1,t-1} + \eta^i_{h,t}$, $\varepsilon^i_{h,t} = \epsilon^i_{h,t} + \theta \epsilon^i_{h-1,t-1}$, $\eta^i_{h,t} \sim N(0, \sigma^2_{\eta,h})$, $\epsilon^i_{h,t} \sim N(0, \sigma^2_{\epsilon,h})$, and $\nu^i_{h,t}$ is classical measurement error with $\sigma^2_\nu = 0.015$ (see Bound, Brown, Duncan, and Rodgers (1994)). [24,33], [34,52], [53,60] are age intervals in which the persistence and variances are assumed to be constant. Point estimates are shown in the first three columns along with bootstrap standard errors in parenthesis. The last two columns report p-values for bootstrap significance tests.

B.5 Full-Time Sample

Table 12: Estimation and Test Results on a Sample with Full-Time Workers

	σ^2_α	$\sigma^2_{z_1}$
	0.0613	0.0703
	(0.0255)	(0.022)

	$[24, 33]$	$[34, 52]$	$[53, 60]$	Test 1 p-value	Test 2 p-value
				$H_0 : \rho_1 \geq \rho_2$	$H_0 : \rho_2 \leq \rho_3$
ρ	0.896	0.976	0.977	0.00	0.57
	(0.0253)	(0.0117)	(0.0210)		
				$H_0 : \sigma^2_{\eta,2} \leq \sigma^2_{\eta,1}$	$H_0 : \sigma^2_{\eta,2} \geq \sigma^2_{\eta,3}$
σ^2_η	0.0239	0.0118	0.0195	0.01	0.067
	(0.0038)	(0.0023)	(0.0062)		
				$H_0 : \sigma^2_{\epsilon,2} \leq \sigma^2_{\epsilon,1}$	$H_0 : \sigma^2_{\epsilon,2} \geq \sigma^2_{\epsilon,3}$
σ^2_ϵ	0.0517	0.0547	0.0590	0.32	0.38
	(0.0065)	(0.0060)	(0.0088)		

Notes:Estimated process: $\tilde{y}^i_{h,t} = \alpha^i + z^i_{h,t} + \epsilon^i_{h,t}$, where $z^i_{h,t} = \rho_{h-1} z^i_{h-1,t-1} + \eta^i_{h,t}$, $\eta^i_{h,t} \sim N(0, \sigma^2_{\eta,h})$ and $\epsilon^i_{h,t} \sim N(0, \sigma^2_{\epsilon,h})$. [24,33], [34,52], [53,60] are age intervals in which the persistence and variances are assumed to be constant. Point estimates are shown in the first three columns along with bootstrap standard errors in parenthesis. The last two columns report p-values for bootstrap significance tests.

B.6 Consecutive

Table 13: Estimation and Test Results on Workers with Consecutive Observations

	σ^2_α	$\sigma^2_{z_1}$
	0.0719	0.0755
	(0.0291)	(0.0278)

	[24, 33]	[34, 52]	[53, 60]	Test 1 p-value	Test 2 p-value
				$H_0 : \rho_1 \geq \rho_2$	$H_0 : \rho_2 \leq \rho_3$
ρ	0.87	0.971	0.959	0.00	0.31
	(0.032)	(0.0158)	(0.0205)		
				$H_0 : \sigma^2_{\eta,2} \leq \sigma^2_{\eta,1}$	$H_0 : \sigma^2_{\eta,2} \geq \sigma^2_{\eta,3}$
σ^2_η	0.0274	0.0132	0.0259	0.00	0.0533
	(0.0044)	(0.0029)	(0.0062)		
				$H_0 : \sigma^2_{\epsilon,2} \leq \sigma^2_{\epsilon,1}$	$H_0 : \sigma^2_{\epsilon,2} \geq \sigma^2_{\epsilon,3}$
σ^2_ϵ	0.0553	0.0597	0.0675	0.37	0.36
	(0.0070)	(0.0065)	(0.0097)		

Notes:Estimated process: $\tilde{y}^i_{h,t} = \alpha^i + z^i_{h,t} + \epsilon^i_{h,t}$, where $z^i_{h,t} = \rho_{h-1} z^i_{h-1,t-1} + \eta^i_{h,t}$, $\eta^i_{h,t} \sim N(0, \sigma^2_{\eta,h})$ and $\epsilon^i_{h,t} \sim N(0, \sigma^2_{\epsilon,h})$. [24,33], [34,52], [53,60] are age intervals in which the persistence and variances are assumed to be constant. Point estimates are shown in the first three columns along with bootstrap standard errors in parenthesis. The last two columns report p-values for bootstrap significance tests.

B.7 Sample with 10 Years of Observations

Table 14: Estimation and Test Results on a Sample with 10 Years of Observations

	σ^2_α	$\sigma^2_{z_1}$
	0.0787	0.0678
	(0.026)	(0.0244)

	$[24,33]$	$[34,52]$	$[53,60]$	Test 1 p-value	Test 2 p-value
				$H_0 : \rho_1 \geq \rho_2$	$H_0 : \rho_2 \leq \rho_3$
ρ	0.86	0.971	0.958	0.00	0.33
	(0.037)	(0.0143)	(0.0216)		
				$H_0 : \sigma^2_{\eta,2} \leq \sigma^2_{\eta,1}$	$H_0 : \sigma^2_{\eta,2} \geq \sigma^2_{\eta,3}$
σ^2_η	0.0277	0.0135	0.0207	0.00	0.13
	(0.0054)	(0.0030)	(0.0055)		
				$H_0 : \sigma^2_{\epsilon,2} \leq \sigma^2_{\epsilon,1}$	$H_0 : \sigma^2_{\epsilon,2} \geq \sigma^2_{\epsilon,3}$
σ^2_ϵ	0.0474	0.0574	0.0663	0.11	0.15
	(0.0078)	(0.0071)	(0.01057)		

Notes:Estimated process: $\tilde{y}^i_{h,t} = \alpha^i + z^i_{h,t} + \epsilon^i_{h,t}$, where $z^i_{h,t} = \rho_{h-1} z^i_{h-1,t-1} + \eta^i_{h,t}$, $\eta^i_{h,t} \sim N(0, \sigma^2_{\eta,h})$ and $\epsilon^i_{h,t} \sim N(0, \sigma^2_{\epsilon,h})$. [24,33], [34,52], [53,60] are age intervals in which the persistence and variances are assumed to be constant. Point estimates are shown in the first three columns along with bootstrap standard errors in parenthesis. The last two columns report p-values for bootstrap significance tests.

B.8 Estimates of the Age-Dependent Process with some Age-Invariant Parameters

B.8.1 Estimates of the Age-Dependent Process with Constant Persistence

Table 15: Estimation and Test Results with Constant Persistence

	σ^2_α	$\sigma^2_{z_1}$
	0.0705	0.0248
	(0.0299)	(0.0155)

	$[24, 33]$	$[34, 52]$	$[53, 60]$	Test 1 p-value	Test 2 p-value
ρ	0.962	0.962	0.962		
	(0.014)	(0.014)	(0.014)		
				$H_0 : \sigma^2_{\eta,2} \leq \sigma^2_{\eta,1}$	$H_0 : \sigma^2_{\eta,2} \geq \sigma^2_{\eta,3}$
σ^2_η	0.00773	0.00913	0.0152	0.20	0.00
	(0.0015)	(0.0016)	(0.0032)		
				$H_0 : \sigma^2_{\epsilon,2} \leq \sigma^2_{\epsilon,1}$	$H_0 : \sigma^2_{\epsilon,2} \geq \sigma^2_{\epsilon,3}$
σ^2_ϵ	0.0842	0.0548	0.0688	0.00	0.06
	(0.0091)	(0.0062)	(0.0105)		

Notes:Estimated process: $\tilde{y}^i_{h,t} = \alpha^i + z^i_{h,t} + \epsilon^i_{h,t}$, where $z^i_{h,t} = \rho_{h-1} z^i_{h-1,t-1} + \eta^i_{h,t}$, $\eta^i_{h,t} \sim N(0, \sigma^2_{\eta,h})$ and $\epsilon^i_{h,t} \sim N(0, \sigma^2_{\epsilon,h})$. [24,33], [34,52], [53,60] are age intervals in which the persistence and variances are assumed to be constant. Point estimates are shown in the first three columns along with bootstrap standard errors in parenthesis. The last two columns report p-values for bootstrap significance tests.

B.8.2 Estimates of the Age-Dependent Process with Constant Variance of Shocks

Table 16: Estimation and Test Results with Constant Variance of Shocks

	σ^2_α	$\sigma^2_{z_1}$
	0.0843	0.0271
	(0.0172)	(0.0106)

	$[24,33]$	$[34,52]$	$[53,60]$	Test 1 p-value	Test 2 p-value
				$H_0 : \rho_1 \geq \rho_2$	$H_0 : \rho_2 \leq \rho_3$
ρ	0.916	0.950	0.980	0.007	0.033
	(0.024)	(0.014)	(0.014)		
σ^2_η	0.0104	0.0104	0.0104		
	(0.0018)	(0.0018)	(0.0018)		
σ^2_ϵ	0.0667	0.0667	0.0667		
	(0.0070)	(0.0070)	(0.070)		

Notes:Estimated process: $\tilde{y}^i_{h,t} = \alpha^i + z^i_{h,t} + \epsilon^i_{h,t}$, where $z^i_{h,t} = \rho_{h-1} z^i_{h-1,t-1} + \eta^i_{h,t}$, $\eta^i_{h,t} \sim N(0, \sigma^2_{\eta,h})$ and $\epsilon^i_{h,t} \sim N(0, \sigma^2_{\epsilon,h})$. [24,33], [34,52], [53,60] are age intervals in which the persistence and variances are assumed to be constant. Point estimates are shown in the first three columns along with bootstrap standard errors in parenthesis. The last two columns report p-values for bootstrap significance tests.

B.9 Results for College and Non-College Samples

In this section, we investigate how the profiles of persistence and variance of shocks look like for households with a college degree and those without one. Tables XX and XX report the results for these two samples.

Table 17: Estimation and Test Results for College Sample

	σ^2_α	$\sigma^2_{z_1}$
	0.0698	0.0865
	(0.0267)	(0.0255)

	$[24, 33]$	$[34, 52]$	$[53, 60]$	Test 1 p-value	Test 2 p-value
				$H_0 : \rho_1 \geq \rho_2$	$H_0 : \rho_2 \leq \rho_3$
ρ	0.8741	0.9738	1.0160	0.00	0.16
	(0.0283)	(0.0141)	(0.0223)		
				$H_0 : \sigma^2_{\eta,2} \leq \sigma^2_{\eta,1}$	$H_0 : \sigma^2_{\eta,2} \geq \sigma^2_{\eta,3}$
σ^2_η	0.0376	0.0163	0.0207	0.00	0.37
	(0.0045)	(0.0026)	(0.0073)		
				$H_0 : \sigma^2_{\epsilon,2} \leq \sigma^2_{\epsilon,1}$	$H_0 : \sigma^2_{\epsilon,2} \geq \sigma^2_{\epsilon,3}$
σ^2_ϵ	0.0512	0.0560	0.0689	0.39	0.21
	(0.0064)	(0.0066)	(0.0108)		

Notes:Estimated process: $\tilde{y}^i_{h,t} = \alpha^i + z^i_{h,t} + \epsilon^i_{h,t}$, where $z^i_{h,t} = \rho_{h-1} z^i_{h-1,t-1} + \eta^i_{h,t}$, $\eta^i_{h,t} \sim N(0, \sigma^2_{\eta,h})$ and $\epsilon^i_{h,t} \sim N(0, \sigma^2_{\epsilon,h})$. [24,33], [34,52], [53,60] are age intervals in which the persistence and variances are assumed to be constant. Point estimates are shown in the first three columns along with bootstrap standard errors in parenthesis. The last two columns report p-values for bootstrap significance tests.

Table 18: Estimation and Test Results for Non-College Sample

	σ^2_α	$\sigma^2_{z_1}$
	0.0833	0.0574
	(0.0268)	(0.0255)

	$[24,33]$	$[34,52]$	$[53,60]$	Test 1 p-value	Test 2 p-value
				$H_0:\rho_1 \geq \rho_2$	$H_0:\rho_2 \leq \rho_3$
ρ	0.8757	0.9481	0.9342	0.007	0.67
	(0.0283)	(0.0141)	(0.0223)		
				$H_0:\sigma^2_{\eta,2} \leq \sigma^2_{\eta,1}$	$H_0:\sigma^2_{\eta,2} \geq \sigma^2_{\eta,3}$
σ^2_η	0.0228	0.0160	0.0310	0.06	0.067
	(0.0045)	(0.0026)	(0.0073)		
				$H_0:\sigma^2_{\epsilon,2} \leq \sigma^2_{\epsilon,1}$	$H_0:\sigma^2_{\epsilon,2} \geq \sigma^2_{\epsilon,3}$
σ^2_ϵ	0.0626	0.0585	0.0701	0.72	0.20
	(0.0063)	(0.0063)	(0.0109)		

Notes:Estimated process: $\tilde{y}^i_{h,t} = \alpha^i + z^i_{h,t} + \epsilon^i_{h,t}$, where $z^i_{h,t} = \rho_{h-1} z^i_{h-1,t-1} + \eta^i_{h,t}$, $\eta^i_{h,t} \sim N(0, \sigma^2_{\eta,h})$ and $\epsilon^i_{h,t} \sim N(0, \sigma^2_{\epsilon,h})$. [24,33], [34,52], [53,60] are age intervals in which the persistence and variances are assumed to be constant. Point estimates are shown in the first three columns along with bootstrap standard errors in parenthesis. The last two columns report p-values for bootstrap significance tests.

B.10 Identification of Persistence

Figure 9: Ratio of Covariances: $\Phi_h^{32} = \frac{cov(y_{i,h}, y_{i,h+3}) - \sigma_\alpha^2}{cov(y_{i,h}, y_{i,h+2}) - \sigma_\alpha^2}$, $h = 22, \cdots, 50$

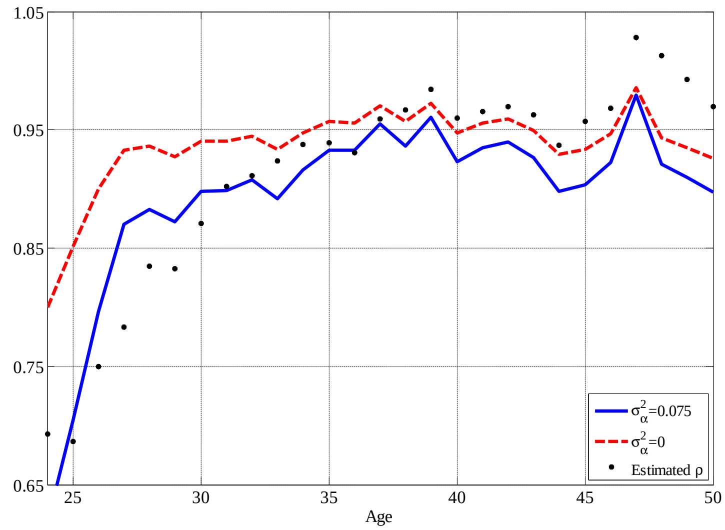

Note: This figure plots the ratio of 3-year ahead covariance to 2-year ahead covariance, Φ^{32}, corrected for the variance of fixed effects, along with the estimated persistence profile. All 3 series are smoothed by a moving average method with a 3-year span.

C Consumption-Savings Model and Its Calibration

C.1 Value Functions

Let $V_h\left(a_h^i,\alpha^i,z_h^i,\varepsilon_h^i\right)$ 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

$$\begin{aligned}
V_h^i\left(a_h^i,\alpha^i,z_h^i,\varepsilon_h^i\right) &= \max_{a_{h+1}^i,c_h^i} u\left(c_h^i\right)+\beta E V_{h+1}\left(a_{h+1}^i,\alpha^i,z_{h+1}^i,\varepsilon_{h+1}^i\right) \\
s.t. & \quad (4) \text{ and} \\
\log\left(y_h^i\right) &= \beta_0+\beta_1 h+\beta_2 h^2+\beta_3 h^3+\alpha^i+z_h^i+\varepsilon_h^i \\
z_{h+1}^i &= \rho_h z_h^i+\eta_h^i \\
a_{h+1}^i &\geq -\bar{A}_{h+1}
\end{aligned}$$

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

$$\begin{aligned}
V_h^i\left(a_h^i,\alpha^i,z_R^i\right) &= \max_{a_{h+1}^i,c_h^i} u\left(c_h^i\right)+\beta V_{h+1}\left(a_{h+1}^i,\alpha^i,z_R^i\right) \\
s.t. & \quad (4) \\
& \quad ln y_h^i = \Phi(\alpha^i, z_R^i) \\
& \quad a_{h+1}^i \geq -\bar{A}_{h+1}
\end{aligned}$$

C.2 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.[24] 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 1-3.

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

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

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 Guvenen, 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.[25]

[25]The 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.