

Federal Reserve Bank of New York
Staff Reports

What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk?

Fatih Guvenen
Fatih Karahan
Serdar Ozkan
Jae Song

Staff Report No. 710
February 2015



This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk?

Fatih Guvenen, Fatih Karahan, Serdar Ozkan, and Jae Song

Federal Reserve Bank of New York Staff Reports, no. 710

February 2015

JEL classification: E24, J24, J31

Abstract

We study the evolution of individual labor earnings over the life cycle, using a large panel data set of earnings histories drawn from U.S. administrative records. Using fully nonparametric methods, our analysis reaches two broad conclusions. First, earnings shocks display substantial deviations from lognormality—the standard assumption in the literature on incomplete markets. In particular, earnings shocks display strong negative skewness and extremely high kurtosis—as high as 30 compared with 3 for a Gaussian distribution. The high kurtosis implies that, in a given year, most individuals experience very small earnings shocks, and a small but non-negligible number experience very large shocks. Second, these statistical properties vary significantly both over the life cycle and with the earnings level of individuals. We also estimate impulse response functions of earnings shocks and find important asymmetries: Positive shocks to high-income individuals are quite transitory, whereas negative shocks are very persistent; the opposite is true for low-income individuals. Finally, we use these rich sets of moments to estimate econometric processes with increasing generality to capture these salient features of earnings dynamics.

Key words: earnings dynamics, life-cycle earnings risk, nonparametric estimation, kurtosis, skewness, non-Gaussian shocks, normal mixture

Guvenen: University of Minnesota, Federal Reserve Bank of Minneapolis, and NBER (e-mail: guvenen@umn.edu). Karahan: Federal Reserve Bank of New York (e-mail: fatih.karahan@ny.frb.org). Ozkan: University of Toronto (e-mail: serdar.ozkan@utoronto.ca). Song: Social Security Administration (e-mail: jae.song@ssa.gov). For helpful critiques and comments, the authors thank Joe Altonji, Andy Atkeson, Richard Blundell, Michael Keane, Giuseppe Moscarini, Fabien Postel-Vinay, Kjetil Storesletten, Anthony Smith, and seminar and conference participants at various universities and research institutions. The views expressed in this paper are those of the authors and do not necessarily reflect the positions of the Social Security Administration, the Federal Reserve Banks of Minneapolis and New York, or the Federal Reserve System.

1 Introduction

This year about 2 million young American men will enter the labor market for the first time. Over the next 40 years, each of these men will go through his unique adventure in the labor market, involving a series of surprises—finding an attractive career, being offered a dream job, getting promotions and salary raises, and so on—as well as disappointments—experiencing unemployment, failing in one career and moving on to another one, suffering health shocks, and so on. These events will vary not only in their initial significance (upon impact) but also in how durable their effects turn out to be in the long run.¹

An enduring question for economists is whether these wide-ranging labor market histories, experienced by a diverse set of individuals, display sufficiently simple regularities that would allow researchers to characterize some general properties of earnings dynamics over the life cycle. Despite a vast body of research since the 1970s, it is fair to say that many aspects of this question remain open. For example, what does the probability distribution of earnings shocks look like? Is it more or less symmetric, or does it display important signs of skewness? More generally, how well is it approximated by a lognormal distribution, an assumption often made out of convenience? And, perhaps more important, how do these properties differ across low- and high-income workers or change over the life cycle? A host of questions also pertain to the dynamics of earnings. For example, how sensible is it to think of a single persistence parameter to characterize the durability of earnings shocks? Do positive shocks exhibit persistence that is different from negative shocks? Clearly, we can add many more questions to this list, but we have to stop at some point. If so, which of these many properties of earnings shocks are the most critical in terms of their economic importance and therefore should be included in this short list, and which are of second-order importance?

One major reason why many of these questions remain open has been the heretofore unavailability of sufficiently rich panel data on individual earnings histories.² Against this backdrop, the goal of this paper is to characterize the most salient aspects of life-

¹In this paper, we focus on the earnings dynamics of men so as to abstract away from the complexities of the female nonparticipation decision. We intend to undertake a similar study that focuses on the earnings dynamics of women.

²With few exceptions, most of the empirical work in this area has been conducted using the Panel Study of Income Dynamics (including the previous work of the authors of this paper), which contains between 500 to 2,000 households per year depending on the selection criteria and suffers from shortcomings that are typical of survey data, such as survey response error, attrition, and so on.

cycle earnings dynamics using a large and confidential panel data set from the U.S. Social Security Administration. The substantial sample size—of more than 200 million observations from 1978 to 2010—allows us to employ a fully nonparametric approach and take what amounts to high-resolution pictures of individual earnings histories.

In deciding what aspects of the earnings data to focus on, we were motivated in this paper by a growing body of theoretical work (reviewed in the next section), which attributes a central role to skewness and kurtosis of economic variables for questions ranging from the effects of monetary policy to optimal taxation, and from the determinants of wealth inequality to asset prices. Therefore, we focus on the first four moments of earnings changes over the life cycle. This analysis reaches two broad conclusions. First, the distribution of individual earnings shocks displays important deviations from lognormality. Second, the magnitude of these deviations (as well as a host of other statistical properties of earnings shocks) varies greatly both over the life cycle and with the earnings level of individuals. Under this broad umbrella of “non-normality and life-cycle variation,” we establish four sets of empirical results.

First, starting with the first moment, we find that average earnings growth over the life cycle varies strongly with the *level* of lifetime earnings: the median individual by lifetime earnings experiences an earnings growth of 38% from ages 25 to 55, whereas for individuals in the 95th percentile, this figure is 230%; for those in the 99th percentile, this figure is almost 1500%.³

Second, turning to the third moment (postponing the second moment for now), we see that earnings shocks are negatively skewed, and this skewness becomes more severe as individuals get older or their earnings increase (or both). Furthermore, this increasing negativity is due entirely to upside earnings moves becoming smaller from ages 25 to 45, and to increasing “disaster” risk (the risk of a sharp fall in earnings) after age 45. Although these implications may appear quite plausible, they are not captured by a lognormal specification, which implies zero skewness.

Third, studying the fourth (standardized) moment, we find that earnings changes display very high kurtosis. What kurtosis measures is most easily understood by looking at the histogram of log earnings changes, shown in Figure 1 (left panel: annual change; right panel: five-year change). Notice the sharpness in the peak of the empirical density,

³A positive relationship between lifetime earnings and life-cycle earnings growth is to be expected (since, all else equal, fast earnings growth will lead to higher lifetime earnings). What is surprising is the magnitudes involved, which turn out to be hard to match standard income processes.

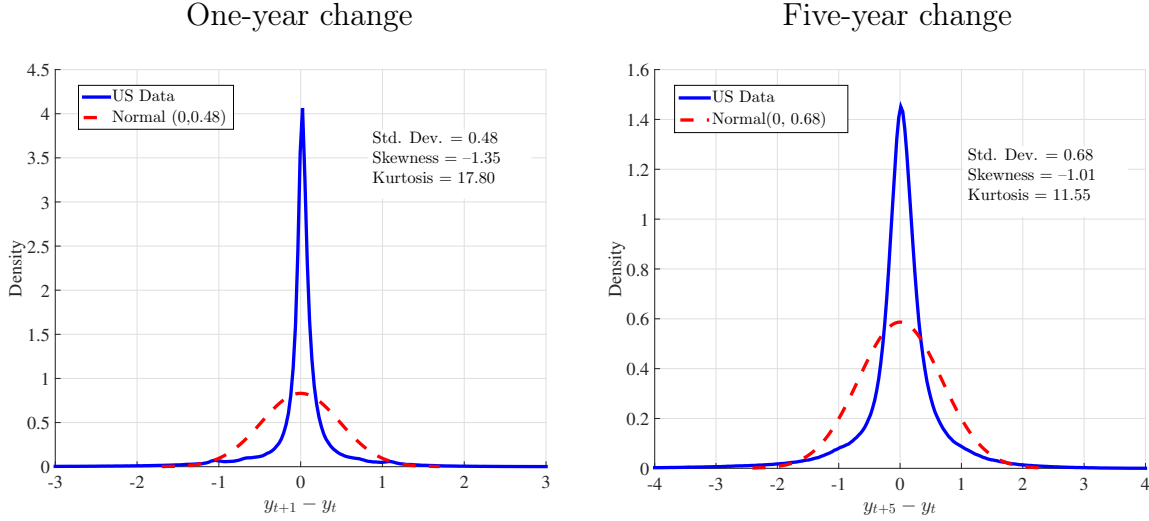


FIGURE 1 – Histogram of Log Earnings Changes. Note: The first year t is 1995, and the data are for all workers in the base sample defined in Section 2.

how little mass there is on the “shoulders” (i.e., the region around $\pm\sigma$), and how long the tails are compared with a normal density chosen to have the same standard deviation as in the data. Thus, there are far more people with very small earnings changes in the data compared with what would be predicted by a normal density. Furthermore, this *average* kurtosis masks significant heterogeneity across individuals by age and earnings: prime-age males with recent earnings of \$100,000 (in 2005 dollars) face earnings shocks with a kurtosis as high as 35, whereas young workers with recent earnings of \$10,000 face a kurtosis of only 5. This life-cycle variation in the nature of earnings shocks is one of the key focuses of the present paper.

What do these statistics mean for economic analyses of risk? Although a complete answer is beyond the scope of this paper, in Section 7 we provide some illustrative calculations. They suggest that the risk premium that will be demanded to bear the measured earnings fluctuations can be anywhere from *four to twenty times* larger than the one calculated with a Gaussian distribution with the same standard deviation. Although these figures are suggestive, and a complete answer requires a fuller investigation, these back-of-the-envelope calculations provide a glimpse into the potential of these documented higher-order moments for economic analyses.

Fourth, we characterize the dynamics of earnings shocks by estimating non-parametric impulse response functions conditional on the recent earnings of individuals and on the size of the shock that hits them. We find two types of asymmetries. One, fixing the shock

size, positive shocks to high-earnings individuals are quite transitory, whereas negative shocks are very persistent; the opposite is true for low-earnings individuals. Two, fixing the earnings level of individuals, the strength of mean reversion differs by the size of the shock: large shocks tend to be much more transitory than small shocks. To our knowledge, both of these findings are new in the literature. These kinds of asymmetries are hard to detect via the standard approach in the literature, which relies on the autocovariance matrix of earnings—the *second* cross-moments of the panel data.⁴ In this regard, our approach is in the spirit of the recent macroeconomics literature that views impulse responses as key to understanding time-series dynamics in aggregate data (e.g., [Christiano *et al.* \(2005\)](#), [Borovicka *et al.* \(2014\)](#)).

While this nonparametric approach allows us to establish key features of earnings dynamics in a robust fashion—which we view as the main contribution of this paper—a tractable parametric process is indispensable for conducting quantitative economic analyses. The standard approach in the earnings dynamics literature is to estimate the parameters of linear time-series models by matching the variance-covariance matrix of log earnings residuals. This approach has two difficulties. First, the strong deviations from lognormality documented in this paper call into question the wisdom of focusing exclusively on covariances at the expense of the rich variation in higher-order moments, which will miss key features of earnings risk faced by workers. Second, the covariance matrix approach makes it difficult to select among alternative econometric processes, because it is difficult to judge the relative importance—from an economic standpoint—of the covariances that a given model matches well and those that it does not. This is an important shortcoming given that virtually every econometric process used to calibrate economic models is statistically rejected by the data.

With these considerations in mind, in [Section 5](#), we follow a different route and target the four sets of empirical moments—broadly corresponding to the first four moments of earnings changes—described above, employing a method of simulated moments (MSM) estimator. We believe this is a more transparent approach: economists can more easily judge whether or not each of these moments is relevant for the economic questions they have in hand. Therefore, they can decide whether the inability of a particular stochastic process to match a given moment is a catastrophic failure or a tolerable shortcoming.

⁴These asymmetries are difficult to detect because a covariance lumps together all sorts of earning changes—large, small, positive, and negative—to produce a single statistic. This approach, although economical in its use of scarce data, masks lots of interesting heterogeneity, as revealed by our analysis.

Specifically, we estimate a set of stochastic processes with increasing generality to provide a reliable “user’s guide” for applied economists.⁵ Two main findings stand out. First, allowing for a rich mixture of AR(1) processes seems essential for matching the salient features of the data, especially the large deviations from normality. Second, a heterogenous income profiles (HIP) component also plays a key role in explaining data features, but only when considered together with the mixture structure. A corollary to these findings is that the workhorse model in the literature—a persistent AR(1) (or random walk) process plus a transitory shock with normal innovations—fails to match most of the prominent features of the earnings data documented in this paper.

The paper is organized as follows. In Section 2, we describe the data and the empirical approach. Section 3 presents the findings on the cross-sectional moments of earnings growth and Section 4 presents the impulse response analysis. Section 5 describes the parametric estimation and Section 6 presents its results. Section 7 concludes.

Related Literature

Since its inception in the late 1970s,⁶ the earnings dynamics literature has worked with the implicit or explicit assumption of a Gaussian framework, thereby making no use of higher-order moments beyond the variance-covariance matrix. One of the few exceptions is an important paper by Geweke and Keane (2000), who emphasize the non-Gaussian nature of earnings shocks and fit a normal mixture model to earnings innovations. More recently, Bonhomme and Robin (2009) analyze French earnings data over short panels and model the transitory component as a mixture of normals and the dependence patterns over time using a copula model. They find the distribution of this transitory component to be left skewed and leptokurtic. In this paper, we go beyond the overall distribution and find substantial variation in the degree of non-normality with age and earnings levels. Furthermore, the impulse response analysis shows the need for a different persistence parameter for large and small shocks, which is better captured as a mixing of AR(1) processes—a step beyond the normal mixture model.⁷

⁵The nonparametric analysis yields more than 10,000 empirical moments of individual earnings data. It is not feasible (or sensible) to estimate every conceivable stochastic process to match combinations of these moments. However, these moments are available for download as an Excel file (from the authors’ websites), so researchers can estimate their preferred specification(s).

⁶Earliest contributions include Lillard and Willis (1978), Lillard and Weiss (1979), Hause (1980), and MaCurdy (1982).

⁷Geweke and Keane (2007) study how regression models can be smoothly mixed, and our modeling approach shares some similarities with their framework.

Incorporating higher-order moments of earnings dynamics into economic models is still in its infancy. In an early attempt, [Mankiw \(1986\)](#) shows that if idiosyncratic earnings shocks become more negatively skewed during recessions, this could generate a large equity premium. Using nonparametric techniques and rich panel data, [Guvenen *et al.* \(2014\)](#) document that the skewness of individual income shocks becomes more negative in recessions, whereas the variance is acyclical. Building on this observation, [Constantinides and Ghosh \(2014\)](#) show that an incomplete markets asset pricing model with countercyclical (negative) skewness shocks generates plausible asset pricing implications, and [McKay \(2014\)](#) studies aggregate consumption dynamics in a business cycle model that is calibrated to match these skewness shocks. Turning to fiscal policy, [Golosov *et al.* \(2014\)](#) show that using an earnings process with negative skewness and excess kurtosis (targeting the empirical moments reported in this paper) implies a marginal tax rate on labor earnings for top earners that is substantially higher than under a traditional calibration with Gaussian shocks with the same variance.⁸

Methodologically, our work is most closely related to two important recent contributions. [Altonji *et al.* \(2013\)](#) estimate a joint process for earnings, wages, hours, and job changes, targeting a rich set of moments via indirect inference. [Browning *et al.* \(2010\)](#) also employ indirect inference to estimate an earnings process featuring “lots of heterogeneity” (as they call it). However, neither paper explicitly focuses on higher-order moments or their life-cycle evolution. The latter paper does model heterogeneity across individuals in innovation variances, as do we, and finds a lot of heterogeneity along that dimension in the data. In ongoing research, [Arellano *et al.* \(2014\)](#) also explore differences in the mean-reversion patterns of earnings shocks across households that differ in their earnings histories. Using data from the Panel Study of Income Dynamics, they find asymmetries in mean reversion that are consistent with those we document in Section 4.

Relatively little work has been done on the life-cycle evolution of earnings dynamics, which is the main focus of this paper. A few papers (including [Baker and Solon \(2003\)](#), [Meghir and Pistaferri \(2004\)](#), [Karahan and Ozkan \(2013\)](#), and [Blundell *et al.* \(2014\)](#)) allow age-dependent innovation variances but do not explore variation in higher-order moments. Our conclusion on the variance is consistent with this earlier work, indicating a decline in variance from ages 25 to 50, with a subsequent rise.

⁸Higher-order moments are gaining a more prominent place in recent work in monetary economics (e.g., [Midrigan \(2011\)](#) and [Berger and Vavra \(2011\)](#); see [Nakamura and Steinsson \(2013\)](#) for a survey) as well as in the firm dynamics literature (e.g., [Bloom *et al.* \(2011\)](#) and [Bachmann and Bayer \(2014\)](#)).

2 Empirical Analysis

2.1 The SSA Data

The data for this paper come from the Master Earnings File (MEF) of the U.S. Social Security Administration records. The MEF is the main source of earnings data for the SSA and contains information for every individual in the United States who was ever issued a Social Security number. Basic demographic variables, such as date of birth, place of birth, sex, and race, are available in the MEF along with several other variables. The earnings data in the MEF are derived from the employee's W-2 forms, which U.S. employers have been legally required to send to the SSA since 1978. The measure of labor earnings is annual and includes all wages and salaries, bonuses, and exercised stock options as reported on the W-2 form (Box 1). Furthermore, the data are uncapped (no top coding) since 1978. We convert nominal earnings records into real values using the personal consumption expenditure (PCE) deflator, taking 2005 as the base year. For background information and detailed documentation of the MEF, see [Panis *et al.* \(2000\)](#) and [Olsen and Hudson \(2009\)](#).

Constructing a nationally representative panel of males from the MEF is relatively straightforward. The last four digits of the SSN are randomly assigned, which allows us to pick a number for the last digit and select all individuals in 1978 whose SSN ends with that number.⁹ This process yields a 10% random sample of all SSNs issued in the United States in or before 1978. Using SSA death records, we drop individuals who are deceased in or before 1978 and further restrict the sample to those between ages 25 and 60. In 1979, we continue with this process of selecting the same last digit of the SSN. Individuals who survived from 1978 and who did not turn 61 continue to be present in the sample, whereas 10% of new individuals who just turn 25 are automatically added (because they will have the last digit we preselected), and those who died in or before 1979 are again dropped. Continuing with this process yields a 10% representative sample of U.S. males in every year from 1978 to 2010. Finally, the MEF has a small number of extremely high earnings observations. In each year, we cap (winsorize) observations above the 99.999th percentile in order to avoid potential problems with these outliers.

⁹In reality, each individual is assigned a transformation of their SSN number for privacy reasons, but the same method applies.




FIGURE 2 – Timeline For Rolling Panel Construction

Base Sample. Sample selection works in two steps. First, for each year we define a *base sample*, which includes all observations that satisfy three criteria, to be described in a moment. Second, to select the *final sample* for a given statistic that we analyze below, we select all observations that belong in the base sample in a collection of years, the details of which vary by the statistic and the year for which the statistic is constructed.

For a given year, the base sample is constructed as follows. First, we restrict attention to individuals between the ages of 25 and 60 to focus on working-age population. Second, we select workers whose annual wage/salary earnings exceeds a time-varying minimum threshold, denoted by $Y_{\min,t}$, defined as one-fourth of a full-year full-time (13 weeks at 40 hours per week) salary at half of the minimum wage, which amounts to annual earnings of approximately \$1,885 in 2010. This condition helps us avoid issues with taking the logarithm of small numbers and makes our analysis more comparable to the empirical earnings dynamics literature, where a condition of this sort is fairly standard (see, among others, [Abowd and Card \(1989\)](#), [Meghir and Pistaferri \(2004\)](#), and [Storesletten *et al.* \(2004\)](#)). Third, the base sample excludes individuals whose self-employment earnings exceed a threshold level, defined as the maximum of $Y_{\min,t}$ and 10% of the individual’s wage/salary earnings in that year. These steps complete the selection of the base sample. The selection of the final sample for a given statistic is described further below.

2.2 Empirical Approach

In the nonparametric analysis conducted in Sections 3 and 4, our main focus will be on individual-level log earnings changes (or growth) at one-year and five-year horizons. These earnings changes provide a simple and useful measure for discussing the dynamics of earnings without making strong parametric assumptions. In Sections 5 and 6, we will link these “changes” to underlying “shocks” or “innovations” to an earnings process by means of a parametric estimation.

To examine how the properties of earnings growth vary over the life cycle and in the

cross section, we proceed as follows. Let $\tilde{y}_{t,h}^i$ denote the log earnings of individual i who is h years old in year t . For each one- and five-year horizon starting in period t , we group individuals based on their age and *recent earnings* (hereafter, RE—to be defined precisely in a moment) as of time $t - 1$. If these groupings are done at a sufficiently fine level, we can think of all individuals within a given age/recent-earnings group to be ex ante identical (or at least very similar). Then, for each such group, we can compute the cross-sectional moments of earnings changes between t and $t + k$ ($k = 1, 2, \dots$), which can then be viewed as corresponding to the properties of shocks that individuals within each bin can expect to face (see Figure 2 for this rolling sample construction). This approach has the advantage that we can compute higher-order moments precisely, as each bin contains several hundred thousands of individuals.¹⁰ (Table I reports sample size statistics.)

Final Sample for Cross-Sectional Moments. We implement this approach by first grouping workers into five-year age bins based on their age in year $t - 1$: 25–29, 30–34, ..., 50–54, and 55–60. Then, within each age group, we select all individuals that were in the base sample in $t - 1$ and in at least two more years between $t - 5$ and $t - 2$.¹¹ For each one of these workers, we compute his average past earnings between years $t - 1$ and $t - 5$, denoted with $\tilde{Y}_{t-1}^i \equiv \sum_{s=1}^5 \exp(\tilde{y}_{t-s,h-s}^i)$. We set earnings observations below $Y_{\min,t}$ to the threshold for this computation. We also further control for age effects, because even within these narrowly defined age groups, age differences of a few years can systematically skew rankings in favor of older workers. To avoid this, we first estimate age dummies, denoted d_h , corresponding to the average of log earnings at each age,¹² and construct five-year average earnings from ages $h - 5$ to $h - 1$: $\sum_{s=1}^5 \exp(d_{h-s})$. We then normalize \tilde{Y}_{t-1}^i with this measure to clean age effects. Thus, our measure of *recent*

¹⁰A second possibility is that the properties of shocks depend more intimately on the characteristics individuals (such as the health, relationship skills, stamina, etc.), and different *types* of workers—for example, identified by their lifetime earnings—might face shocks with different properties. This suggests that perhaps we should group workers based on their lifetime earnings and study the properties of shocks for each group. We have an analogous set of results obtained by adopting this alternative perspective. It turns out that both approaches yield similar substantive conclusions, so we omit these results from the main text. These results are available upon request.

¹¹That is, in each of these years, the individual was in the qualifying age range with wage earnings exceeding $Y_{\min,t}$, and satisfied the no-self-employment condition.

¹²These are estimated from a pooled regression of log earnings on age and cohort dummies. Further details are given later in Section 3.1.

TABLE I – Sample Size Statistics for Cross-Sectional Moments

Age group	# Observations in Each RE Percentile Group			
	Median	Min	Max	Total ('000s)
25-29	381,606	337,603	674,986	40,871
30-34	640,596	566,235	712,307	63,835
35-39	642,226	441,425	721,966	61,891
40-44	606,344	356,700	707,502	57,151
45-49	524,177	320,935	686,753	50,699
50-54	383,887	291,117	619,987	42,554
55-59	240,273	215,842	407,997	27,634

Note: Each entry reports the statistics of the number of observations in each of the 100 RE percentile groups for each age. Cross-sectional moments are computed for each year and then averaged over all years, so sample sizes refer to the sum across all years of a given age by percentile group. The last column (“Total”) reports the sum of observations across all 100 RE percentile groups for the age group indicated.

earnings (hereafter, RE) is

$$\bar{Y}_{t-1}^i \equiv \frac{\tilde{Y}_{t-1}^i}{\sum_{s=1}^5 \exp(d_{h-s})}.$$

Our *final sample* for the cross-sectional moments is then obtained as follows. We rank individuals based on \bar{Y}_{t-1}^i , and divide them into (typically 100) age-specific RE percentile groups. Within each group, we drop those individuals who fail to qualify for the base sample in year t or $t + k$.¹³ Table I reports the summary statistics of the number of observations in each age/earnings cell (summed over all years). As seen here, the sample size is very large—the smallest cell size exceeds 200,000 observations and the average is close to 500,000—which allows us to compute all statistics very precisely.

3 Cross-sectional Moments of Earnings Growth

We begin our analysis by documenting empirical facts about the first four moments of earnings growth at short (one-year) and long (five-year) horizons. For computing moments of earnings growth, we work with the time difference of y_t^i , which is log earnings *net* of the age effect. Thus:

$$\Delta_k y_t^i \equiv (y_{t+k,h+k}^i - y_{t,h}^i) = (\tilde{y}_{t+k,h+k}^i - d_{h+k}) - (\tilde{y}_{t,h}^i - d_h).$$

¹³Therefore, the percentile bins are constructed using information only prior to t , whereas the number of observations within each bin also depends on being in the base sample in t and $t + k$.

We compute the cross-sectional moments of $\Delta_k y_t^i$ for each year, $t = 1980, 1981, \dots, 2009$ and then average these across all years.¹⁴

3.1 First Moment: Mean of Log Earnings Growth

We begin our analysis with the first moment—average earnings growth—and examine how it varies with age (i.e., over the life cycle) and, for reasons that will become clear in a moment, across groups of individuals that differ in their *lifetime earnings* (and not recent earnings). But first, to provide a benchmark, we follow the standard procedure in the literature, (e.g., [Deaton and Paxson \(1994\)](#)) to estimate the average life-cycle profile of log earnings. Although the procedure is well understood, its details matter for some of the discussions below, so we go over it in some detail.

The average life-cycle profile is obtained from panel data or repeated cross sections by regressing log individual earnings on a full set of age and (year-of-birth) cohort dummies. The estimated age dummies are plotted as circles in [Figure 3](#) and represent the average life-cycle profile of log earnings. It has the usual hump-shaped pattern that peaks around age 50. (On a side note, these age dummies turn out to be indistinguishable from a fourth order polynomial of age,¹⁵ a point also observed by [Murphy and Welch \(1990\)](#) in Current Population Survey data.)

One of the most important aspects of a life-cycle profile is the implied growth in average earnings over the life cycle (e.g., from ages 25 to 55). It is well understood that the magnitude of this rise matters greatly for many economic questions, because it is a strong determinant of borrowing and saving motives.¹⁶ In our data, this rise is about 80 log points, which is about 127%.¹⁷ Notice that feeding this life-cycle profile into a calibrated life-cycle model will imply that the median individual in the simulated sample experiences (on average) a rise of this magnitude from ages 25 to 55. One question we now address is whether this implication is consistent with what we see in the data. In other words, if we rank male workers in the U.S. data by their lifetime earnings, does the median worker experience an earnings growth of approximately 127%?


¹⁴We use $t = 1980$ as the first year of our analysis and therefore group individuals in 1979 based on their recent earnings computed over 1978 and 1979. Similarly, 2009 is the last feasible year for t , which allows us to construct the moments of one-year earnings changes between 2009 and 2010.

¹⁵Regressing the age dummies on a fourth order polynomial of age yields an average absolute deviation of only 0.3 log percent!

¹⁶See [Deaton \(1991\)](#), [Attanasio et al. \(1999\)](#), and [Gourinchas and Parker \(2002\)](#), among others.

¹⁷This figure lies on the high end of previous estimates from data sets such as the PSID, but not unseen before (cf. [Attanasio et al. \(1999\)](#)).

FIGURE 3 – Life-Cycle Profile of Average Log Earnings



This question can be answered directly with our data. First, we need to compute lifetime earnings for each individual. For this purpose, we select a subsample of individuals that have at least 33 years of data between the ages of 25 and 60. We further restrict our sample to individuals who (i) have earnings above $Y_{\min,t}$ for at least 15 years and (ii) are not self-employed for more than 8 years. We rank individuals based on their lifetime earnings, computed by summing their earnings from ages 25 through 60. Earnings observations lower than $Y_{\min,t}$ are set to this threshold. For individuals in a given lifetime earnings (hereafter, LE) percentile group, denoted LE_j , $j = 1, 2, \dots, 99, 100$, we compute growth in average earnings between any two ages h_1 and h_2 as $\log(\bar{Y}_{h_2,j}) - \log(\bar{Y}_{h_1,j})$, where $\bar{Y}_{h,j} \equiv \mathbb{E}(Y_h^i | i \in LE_j)$ and Y_h^i for a given individual may be zero.

Figure 4 plots the results for $h_1 = 25$ and $h_2 = 55$. Here, there are several takeaways. First, individuals in the median lifetime earnings group experience a growth rate of 38%, about one-third of what was predicted by the profile in Figure 3. Moreover, we have to look all the way above LE90 to find an average growth rate of 127%. However, earnings growth is very high for high-income individuals, with those in the 95th percentile experiencing a growth rate of 230% and those in the 99th percentile experiencing a growth rate of 1450%. Although some of this variation could be expected because individuals with high earnings growth are more likely to have high lifetime earnings, these magnitudes are too large to be accounted for by that channel, as we show below.


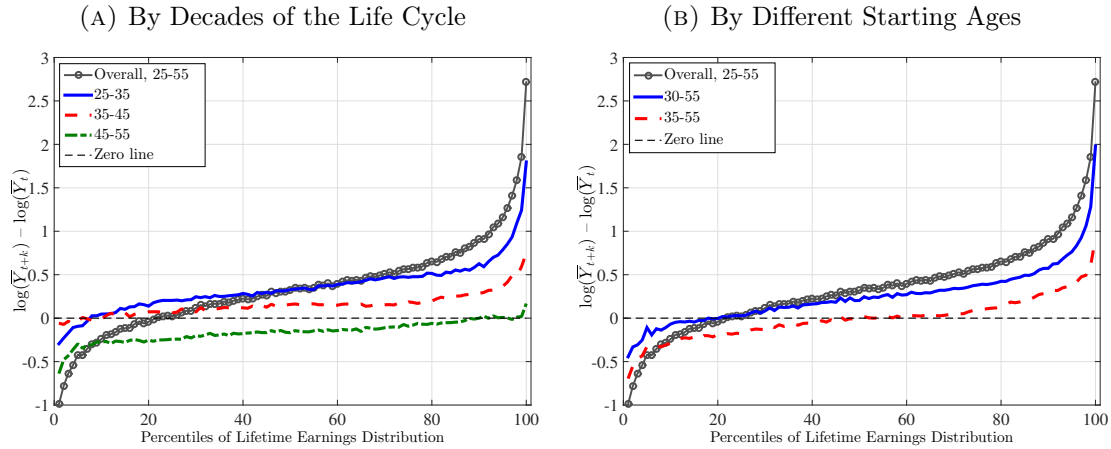



FIGURE 4 – Life-Cycle Earnings Growth Rates, by Lifetime Earnings Group

FIGURE 5 – Log Earnings Growth Over Sub-Periods of Life Cycle



Earnings Growth by Decades. How is earnings growth over the life cycle distributed over different decades of the life cycle? Figure 5a answers this question by plotting, separately, earnings growth from ages 25 to 35, 35 to 45, and 45 to 55. Across the board, the bulk of earnings growth happens during the first decade. In fact, for the median LE group, average earnings growth from ages 35 to 55 is zero (notice that the solid blue line and grey line with circles overlap at LE50). Second, with the exception of those in the top 10% of the LE distribution, all groups experience negative growth from ages 45 to 55. So, the peak year of earnings is strongly related to the lifetime earnings percentile.

FIGURE 6 – Standard Deviation of Earnings Growth




After age 45, the only groups that are experiencing growth on average are those who are in the top 2% of the LE distribution.

How do the results change if we consider a slightly later starting age? Figure 5b plots earnings growth starting at age 30 (solid blue line) and 35 (dashed red line). As can be anticipated from the previous discussion, from ages 35 to 55, average growth is zero for the median LE group and is very low for all workers below L70. Top earners still do very well though, experiencing a rise of 200 log points (or 640%) from ages 30 to 55 and a rise of 90 log points (or 146%) from ages 35 to 55. Those at the bottom of the LE distribution display the opposite pattern: average earnings drops by 70 log points (or 50%) from ages 35 to 55.

3.2 Second Moment: Variance

How does the dispersion of earnings shocks vary over the life cycle and by earnings groups? To answer this question, Figure 6 plots the standard deviation of one-year and five-year earnings growth by age and recent earnings (hereafter, RE) groups (as defined above, Section 2.2). The following patterns hold true for both short- and long-run growth rates. First, for every age group, there is a pronounced U-shaped pattern by RE levels, implying that earnings changes are less dispersed for individuals with higher RE up to about the 90th percentile (along the x -axis). This pattern reverts itself inside the top 10% as dispersion increases rapidly with recent earnings. Second, over the life cycle, the dispersion of shocks declines monotonically up to about age 50 (with the exception

FIGURE 7 – Skewness (Third Standardized Moment) of Earnings Growth



of very top earners) and then rises slightly for middle- to high-earning individuals from ages 50 to 55.

The life-cycle pattern is quite different for top earners who experience a monotonic *increase* in dispersion of shocks over the life cycle. In particular, for one-year changes, individuals at the 95th percentile of the RE distribution experience a slight increase from 0.45 in the youngest age group up to 0.51 in the oldest group (50–54). Those in the top 1% experience a larger increase from 0.62 in the first age group up to 0.75 in the oldest. Therefore, we conclude that the lower 95 percentiles and the top 5 percentiles display patterns with age and recent earnings that are the opposite of each other. The same theme will emerge again in our analysis of higher-order moments.

Standard Deviation of (Log) Earnings Levels. Although the main focus of this section is on earnings *growth*, the life-cycle evolution of the dispersion of earnings *levels* has been at the center of the incomplete markets literature since the seminal paper of [Deaton and Paxson \(1994\)](#). For completeness, and comparability with earlier work, we have estimated the within-cohort variance of log earnings over the life cycle and report it in [Figure A.2](#) in [Appendix A.1](#).

3.3 Third Moment: Skewness (or Asymmetry)

The lognormality assumption implies that the skewness of earnings shocks is zero. [Figure 7](#) plots the skewness, measured here as the third standardized moment,¹⁸ of one-

¹⁸More precisely, for random variable X , with mean μ and standard deviation σ , the third standardized moment is $\mathbb{E}[(X - \mu)^3] / \sigma^3$.

year (left) and five-year (right) earnings growth. The first point to observe is that every graph in both panels of Figure 7 lies below the zero line, indicating that earnings changes are negatively skewed at every stage of the life cycle and for all earnings groups. The second point, however, is that skewness is increasingly more negative for individuals with higher earnings and as individuals get older. Thus, it seems that the higher an individual’s current earnings, the more room he has to fall and the less room he has left to move up. And this is true for both short-run and long-run earnings changes. Curiously, and as was the case with the standard deviation, the life-cycle pattern in skewness becomes much weaker at the very top of the earnings distribution.

Another measure of asymmetry is provided by Kelly’s measure of skewness, which is defined as

$$\mathcal{S}_K = \frac{(P90 - P50) - (P50 - P10)}{P90 - P10}, \quad (1)$$

where Pxy refers to percentile xy of the distribution under study. Basically, \mathcal{S}_K measures the relative fractions of the overall dispersion (P90–P10) accounted for by the upper and lower tails. An appealing feature of Kelly’s skewness relative to the third standardized moment is that a particular value is easy to interpret. To see this, rearrange (1) to get


$$\frac{P90 - P50}{P90 - P10} = 0.5 + \frac{\mathcal{S}_K}{2}.$$

Thus, a negative value of \mathcal{S}_K implies that the lower tail (P50–P10) is longer than the upper tail (P90–P50), indicating negative skewness. Another property of Kelly’s measure is that it is less sensitive to extremes (above the 90th or below the 10th percentile of the shock distribution). Instead, it captures the shift in the weight distribution in the middling section of the shock distribution, whereas the third moment also puts a large weight on the relative lengths of each tail. (We examine the tails in more detail in the next subsection.)

In the left panel of Figure 8, we plot Kelly’s skewness, which is also negative throughout and becomes more negative with age, especially below RE60. However, it does not always get more negative with higher RE. This difference from the third standardized moment (Fig. 7a) indicates that as RE increases it is mostly the extreme negative shocks (captured by the third moment) that drive the negative skewness, rather than the more middling shocks—those between P10 and P90.

Figure 8b plots Kelly’s skewness for five-year changes, which reveals essentially the

FIGURE 8 – Kelly’s Skewness of Earnings Growth

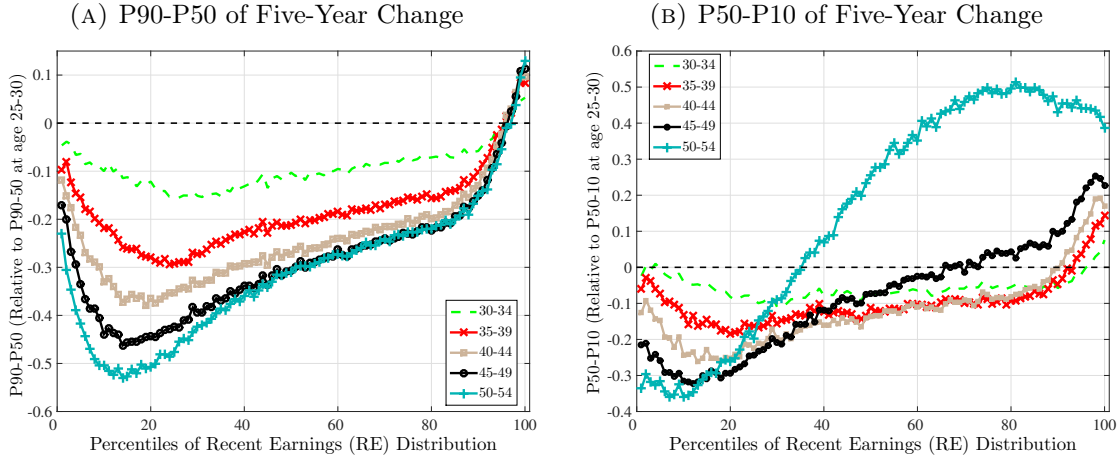


same pattern as with the third moment in Figure 7b: each measure shows a strong increase in left-skewness with both age and earnings (except for the very-high earners). Furthermore, the magnitude of skewness is substantial. For example, the Kelly’s skewness for five-year earnings change of -0.35 for individuals aged 45–49 and in the 80th percentile of the RE distribution implies that the P90-P50 accounts for 32% of P90-P10, whereas P50-P10 accounts for the remaining 68%. This is clearly different from a lognormal distribution, which is symmetric—both tails contribute 50% of the total.

While the preceding decomposition is useful, it does not answer a key question: is the increasingly more negative skewness over the life cycle primarily due to a compression of the upper tail (fewer opportunities to move up) or due to an expansion in the lower tail (increasing risk of falling a lot)? For the answer, we need to look at the *levels* of the P90-P50 and P50-P10 separately over the life cycle. The left panel of Figure 9 plots P90-P50 for different age groups minus the P90-P50 for 25- to 29-year-olds, which serves as a normalization. The right panel plots the same for P50-P10. One way to understand the link between these two graphs and skewness is that keeping P50-P10 fixed over the life cycle, if P90-P50 (left panel) declines with age, this causes Kelly’s skewness to become more negative. Similarly, keeping P90-P50 fixed, a rise in P50-P10 (right panel) has the same effect.

Turning to the data, up until age 45, both P90-P50 and P50-P10 decline with age (across most of the RE distribution). This leads to the declining dispersion that we have seen above. The shrinking P50-P10 would also lead to a *rising* skewness if it were not

FIGURE 9 – Kelly’s Skewness Decomposed: Change in P90-P50 and P50-P10 Relative to Age 25–30



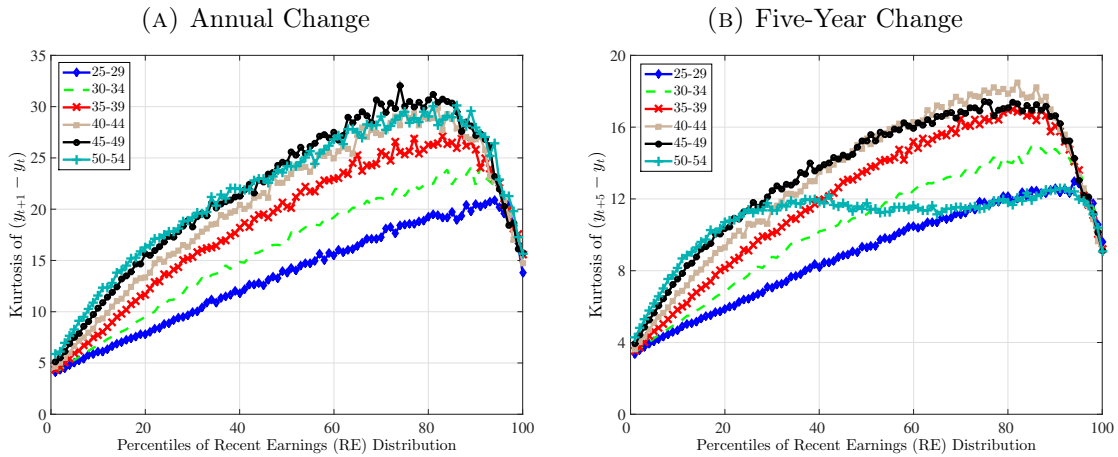
for the faster compression of P90-P50 during the same time. Therefore, from ages 25 to 45, the increasing negativity of skewness is *entirely* due to the fact that the upper end of the shock distribution compresses more rapidly than the compression of the lower end. After age 45, P50-P10 starts expanding rapidly (larger earnings drops becoming more likely), whereas P90-P50 stops compressing any further (stabilized upside). Thus, during this phase of the life cycle, the increasing negativity in Kelly’s skewness is due to increasing downward risks and not the disappearance of upward moves. The only exception to this pattern is, again, the top earners (RE95 and above) for whom P90-P50 actually never compresses over the life cycle, whereas the P50-P10 gradually rises as they get older. Therefore, as they climb the wage ladder, these individuals do not face a tightening ceiling, but do suffer from an increasing risk of falling a lot.

3.4 Fourth Moment: Kurtosis (Peakedness and Tailedness)

It is useful to begin by discussing what kurtosis measures. A useful interpretation has been suggested by Moors (1986), who described kurtosis as measuring how dispersed a probability distribution is away from $\mu \pm \sigma$.¹⁹ This is consistent with how a distribution with excess kurtosis often looks like: a sharp/pointy center, long tails, and little mass near $\mu \pm \sigma$. A corollary to this description is that for a distribution with high kurtosis, the usual way we think about standard deviation—as representing the size of the typical

¹⁹This can easily be seen by introducing a standardized variable $Z = (x - \mu)/\sigma$ and noting that kurtosis is $\kappa = \mathbb{E}(Z^4) = \text{var}(Z^2) + \mathbb{E}(Z^2)^2 = \text{var}(Z^2) + 1$. So κ can be thought of as the dispersion of Z^2 around its expectation, which is 1, or the dispersion of Z around +1 and -1.

FIGURE 10 – Kurtosis of Earnings Changes



shock—is not very useful. This is because very few realizations will be of a magnitude close to the standard deviation; instead, most will be either close to the median or in the tails.

With this definition in hand, let us now examine the earnings growth data. Figure 10a plots the kurtosis of annual earnings changes. First, notice that kurtosis increases monotonically with recent earnings up to the 80th to 90th percentiles for all age groups. That is, high-earnings individuals experience *even smaller* earnings changes of either sign, with few experiencing very large changes. Second, kurtosis increases over the life cycle, for all RE levels, except perhaps the top 5%. Furthermore, the peak levels of kurtosis range from a low of 20 for the youngest group, all the way up to 30 for the middle-age group (40–54).

To provide a more familiar interpretation of these kurtosis values, it is useful to calculate measures of concentration. The first three columns of Table II report the fraction of individuals experiencing a log earnings change (of either sign) of less than a threshold $x = 0.05, 0.10, 0.20, 0.50$, and 1.00, under alternative assumptions about the data-generating process. For the entire sample, the standard deviation of $y_{t+1}^i - y_t^i$ is 0.48. Assuming that the data-generating process is a Gaussian density with this standard deviation, only 8% of individuals would experience an annual earnings change of less than 5%. The true fraction in the data is 35%. Similarly, the Gaussian density predicts a fraction of 16% when the threshold is 0.10, whereas the true fraction is 54%. As an alternative calculation, we calculate the areas under the densities in three different

TABLE II – Fraction of Individuals with Selected Ranges of Log Earnings Change

$x :$	Prob($ y_{t+1}^i - y_t^i < x$)			$Range:\dagger$	Prob($(y_{t+1}^i - y_t^i) \in Range$)		
	Data*	$\mathcal{N}(0, 0.48^2)$	Ratio		Data	$\mathcal{N}(0, 0.48^2)$	Ratio
0.05	0.35	0.08	4.38	Center	0.653	0.250	2.61
0.10	0.54	0.16	3.38	Shoulders	0.311	0.737	0.42
0.20	0.71	0.32	2.23	Tails	0.036	0.013	2.77
0.50	0.86	0.70	1.22	$ x > 1.50$	0.023	0.002	11.5
1.00	0.94	0.96	0.98				

Notes: *The empirical distribution used in this calculation is for 1995-96, the same as in Figure 1.


\dagger The intervals are defined as follows: “Center” refers to the area inside the first intersection between the two densities in Figure 1: $[-0.122, 0.187]$. “Tails” refer to the areas outside the intersection point at the tails: $(-\infty, -1.226] \cup [1.237, \infty)$. “Shoulders” refer to the remaining areas of the densities.

ranges determined by the intersections of the two densities in the left panel of Figure 1. The center is the area inside the first set of intersections, and the Gaussian density has 25% of its mass in this area compared with 65% in the data. The shoulders are the second set of areas, marked again by the intersections, and the Gaussian density has almost three-quarters of its mass in this area, compared with only 31% in the data. Turning to the tails, the Gaussian density has only 1.3% of its mass in the tails compared with almost three times that amount in the data. Further, the last row of the right panel reports that a typical worker draws a shock larger than 150 log points (an almost five-fold increase or an 80% drop in earnings) once in a lifetime (or 2.3% annual chance), whereas this probability is 11.5 times less likely under a normal.

We now take a closer look at the tails of the earnings growth distribution compared with a normal density. Figure 11 plots the log density of the one-year change in the data versus the Gaussian density. This is essentially the same as the left panel of Figure 1 but with the y -axis now in logs. The lognormal density is an exact quadratic, whereas the data display a more complex pattern. Two points are worth noting. One, the data distribution has much thicker *and* longer tails compared with a normal distribution, and the tails decline almost linearly, implying a Pareto distribution at both ends, with significant weight at extremes.²⁰ Two, the tails are asymmetric, with the left tail declining much more slowly than the right, contributing the negative third standardized moment documented above. In fact, fitting linear regression lines to each tail yields a tail index of

²⁰A double-Pareto distribution is one where both tails are Pareto with possibly different tail indices.

FIGURE 11 – Tails of the Distributions: U.S. Data vs. Normal Density



2 for the right tail and 1.2 for the left tail—the latter showing especially high thickness.²¹

Overall, these findings show that earnings changes in the U.S. data exhibit important deviations from lognormality and raise serious concerns about the focus in the current literature on the covariances (second moments) alone. In particular, targeting the covariances alone can vastly overestimate the typical earnings shock received by the average worker and miss the substantial but infrequent jumps experienced by few.

Economic Models behind Skewness and Kurtosis. While the lognormal framework is often adopted for technical and empirical tractability, negative skewness and excess kurtosis are naturally generated by standard structural models of job search over the life cycle. For example, job ladder models in which workers do on-the-job search and move from job to job as they receive better offers and fall off the job ladder after unemployment not only will generate negative skewness but also will imply that skewness becomes more negative with age. This is because as the worker climbs the job ladder, the probability of receiving a wage offer much higher than the current wage will be declining. At the same time, as the worker moves higher up in the wage ladder, falling down to a flat unemployment surface (or disability) implies that there is more room to

²¹Notice that although the Pareto tail in the earnings distribution is well known, here the two Pareto tails emerge in earnings *changes*, for which much less empirical evidence exists.

fall. Furthermore, as the attractiveness of job offers declines with the current wage, more job offers will be rejected, and therefore the frequency of job-to-job transitions will also decline with age, implying that most wage changes will be small (within-job) changes. This increases the concentration of earnings changes near zero, which in turn raises the kurtosis of changes.

That said, the magnitudes of variation over the life cycle and by earnings levels that we documented in these moments are so large that it is an open question whether existing models of job search can be consistent with these magnitudes, and, if not, what kinds of modifications should be undertaken to make them consistent.²²

3.5 Robustness and Extensions

The results documented in the previous sections show important deviations from lognormality as well as clear patterns with age and past earnings. This raises the question of whether some of these findings are due to simple statistical artifacts—say, due perhaps to extreme shocks experienced by very few individuals—and whether the age and earnings patterns might be due to sample selection or other assumptions made in the construction of these statistics. In this section, we discuss five cases of interest and report all the relevant figures and analysis in Appendixes [A.2.1](#) to [A.2.2](#).

I. Decomposing Moments: Job-Stayers vs. Job-Switchers. Going back to the work of [Topel and Ward \(1992\)](#), economists have found important differences between the earnings changes that occur during an employment relationship and those that occur across jobs (see, more recently, [Low *et al.* \(2010\)](#), [Altonji *et al.* \(2013\)](#), and [Bagger *et al.* \(2014\)](#)). Therefore, it is of interest to ask how the empirical patterns we have documented so far relate to within- and between-job earnings changes. Our data set contains a unique employer identification number (EIN) for each job that a worker holds in a given year, which allows us to conduct such an analysis (see Appendix [A.2.1](#)).

II. Disentangling the Effects of Age and Recent Earnings. In the analysis so far, we have grouped workers first by age, and then within each age group, we have ranked and divided them into recent earnings percentiles. The implication is that RE

²²A high kurtosis in earnings changes could partly be due to heterogeneity across individuals in shock variances, as documented by [Chamberlain and Hirano \(1999\)](#) and [Browning *et al.* \(2010\)](#). We have explored this possibility by estimating earnings shock variances for each individual. While we do find significant heterogeneity in variances, consistent with these papers, the remaining kurtosis is still substantial. We do model and estimate individual-specific variances in Section [5](#).

percentiles in these figures are age group dependent. Therefore, when we fix an RE percentile and examine how statistics vary with age, we are simultaneously looking at changes with earnings, since earnings vary with age. The advantage of this approach was that it ensured each RE group contained a similar number of observations, whereas grouping workers based on the RE distribution in the overall sample will result in too many younger workers appearing in lower RE percentiles and vice versa for middle-age workers. In Appendix A.2.2 (Figures A.5 and A.6), we plot 3-D graphs of skewness and kurtosis, where we first group workers based on the RE distribution in the overall sample, and then within each RE group, we classify workers by age. Inspecting these graphs and comparing with their counterparts above shows that the main substantive conclusions described above are robust.

III. Averaging Earnings Over Neighboring Years. Recall that the statistics above were constructed by taking differences between two years, t and $t + k$. One concern is the robustness of this measure to small changes in the timing of earnings. For example, suppose that an individual’s income has been shifted from the last few months of year $t + k$ into the beginning of $t + k + 1$. If true, this would represent an earnings fluctuation that is easy to smooth, but could appear as a big negative shock between t and $t + k$. A similar comment applies to period t . To address this issue, we have constructed the same set of statistics for the second to fourth moments by using two-year average earnings. For the short-run and long-run variations, we use, respectively

$$\tilde{\Delta}y_t^i = \log(Y_{t+3}^i + Y_{t+2}^i) - \log(Y_t^i + Y_{t+1}^i) \quad \text{and} \quad \tilde{\Delta}_5y_t = \log(Y_{t+5}^i + Y_{t+6}^i) - \log(Y_t^i + Y_{t+1}^i).$$

The first measure becomes more like a two-year difference, whereas the second one is closer to a five-year difference as before. However, we are mostly interested in whether statistics are broadly robust and the qualitative patterns remain unchanged, so these are reasonable choices.

IV. Difference from Usual Earnings. Even though we condition on recent earnings over the past five years and require all individuals in the sample to be employed in year $t - 1$, it is conceivable that some individuals receive large positive shocks in period t , and the subsequent drop in earnings from t to $t + k$ is simply mean reversion—and not a new shock. The same argument applies for a large negative shock in t . To see if this might be important, we have constructed the same statistics using an alternative difference

measure, again for the short-run and long-run variations:

$$\Delta_{\text{short}}y^i = \log(Y_{t+1}^i) - \log(\bar{Y}_{t-1}^i) \quad \text{and} \quad \Delta_{\text{long}}y^i = \log(Y_{t+5}^i) - \log(\bar{Y}_{t-1}^i).$$

These are longer differences than before, since the base year is now centered around $t - 3$.

V. Trimming the Tails. As noted before, earnings growth displays very long tails, and even though measurement error is unlikely to explain it, it is still of interest to know, for example, how much of the very large kurtosis and negative skewness is due to the extreme observations and how much is due to the bulk of the distribution. Since the third and fourth moments are sensitive to tails, this is worth exploring (although Kelly’s skewness is already reassuring for the skewness). To this end, we have constructed the higher-order moments under alternative assumptions about the tails: (i) by dropping the top and bottom 1% of earnings growth observations, and (ii) by changing the lower threshold for sample exclusion from $Y_{\min,t}$ to be individual-specific and equal to 5% of each worker’s own recent earnings.²³


In Appendix A, we report the figures analogous to those above under these three robustness checks (III to V). Although the figures are quantitatively different, the difference is almost always small, and therefore the substantive conclusions of this analysis remain intact.

4 Dynamics of Earnings

A key dimension of life-cycle earnings risk is the persistence of earnings changes. Typically, this persistence is modeled as an AR(1) process or a low-order ARMA process (typically, ARMA(1,1)), and the persistence parameter is pinned down by the rate of decline of autocovariances with the lag order. The AR(1) structure, for example, predicts a geometric decline and the rate of decline is directly given by the mean reversion parameter. While this approach might be appropriate in survey data with small sample sizes, it imposes restrictions on the data that might be too strong, such as the uniformity of mean reversion for positive and negative shocks, for large and small shocks, and so on.

²³Because $Y_{\min,t}$ does not vary with an individual’s own earnings, high-income individuals can experience a larger fall in earnings *and* still remain in the sample, whereas low-income individuals would exit the sample with the same fall. This asymmetry might give the appearance of a more negative skewness for higher-income individuals.

FIGURE 12 – Impulse Responses, Prime-Age Workers with Median RE




Here, the substantial sample size allows us to characterize persistence without making parametric assumptions.

Final Sample for Impulse Response Analysis

The final sample for this analysis is slightly different from the one used in the previous section. In particular, our *final sample* includes all observations that are in the base sample in $t - 1$ and in at least two more years between $t - 5$ and $t - 2$, and furthermore satisfies the age (25–60) and no-self-employment condition in years t through $t + 3$, and in $t + 5$ and $t + 10$. To reduce the number of graphs to a manageable level, we aggregate individuals across demographic groups. First, we combine the first two age groups (ages 25 to 34) into “young workers,” and the last four groups (ages 35 to 55) into “prime-age workers.”

To this end, we rank and group individuals based on their average earnings from $t - 5$ to $t - 1$, then within each such group, we rank and group again by the size of the earnings change between $t - 1$ and t . Hence, all individuals within a given group obtained by crossing the two conditions have the same average earnings up to time $t - 1$ and experience the same earnings “shock” from $t - 1$ to t . For each such group of individuals, we then compute their average earnings change from t to $t+k$, for all values of $k = 1, 2, 3, 5, 10$. Specifically, we construct 21 groups based on their RE percentiles: 1–5, 6–10, 11–15, . . . , 86–90, 91–95, 96–99, 100. Then, we construct 20 equally-divided groups

FIGURE 13 – Impulse Responses (Rotated View), Prime-Age Workers with Median RE



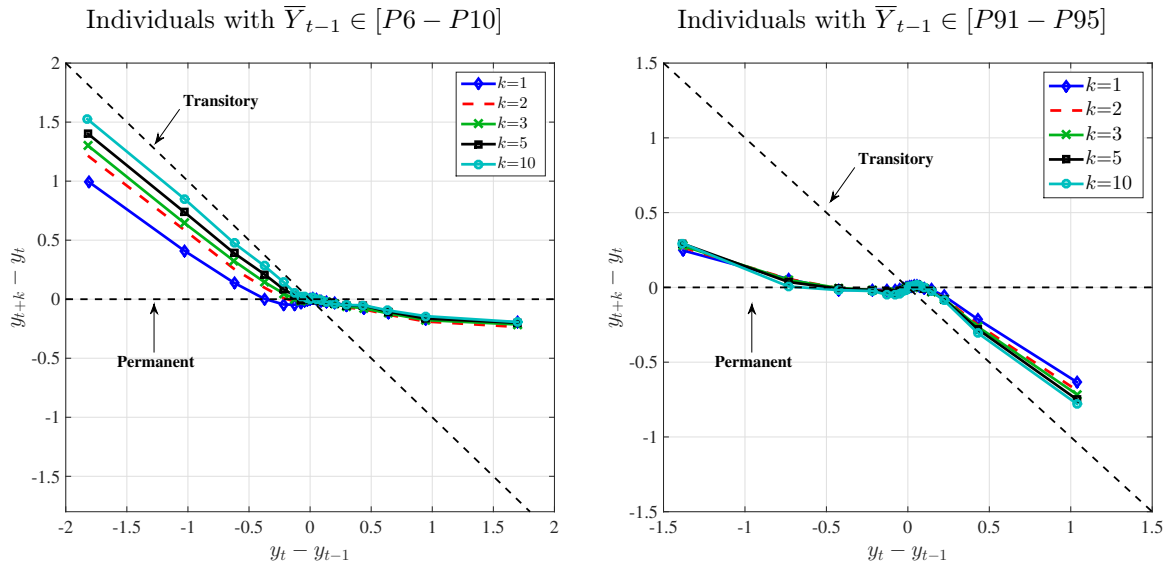
based on the percentiles of the shock, $y_t^i - y_{t-1}^i$: percentiles 1–5, 6–10, 11–15, \dots , 91–95, 96–100. Therefore, for every year t , we have 2 age groups, 21 RE groups and 20 groups for shock size (for a total of 840 groups). As before, we construct these groups separately for each t and assign workers based on these averages. Then, for workers in each group, we compute the average of log k -year earnings growth, $\mathbb{E} \left[y_{t+k}^i - y_t^i | \bar{Y}_{t-1}^i, y_t^i - y_{t-1}^i \right]$, for $k = 1, 2, 3, 5, 10$.

4.1 Impulse Response Functions

For prime-age males with median RE level (as of $t - 1$), Figure 12 plots 20 impulse response functions (one for each “shock” size, $y_t^i - y_{t-1}^i$). As seen here, the most positive 5% of shocks at time t are about 100 log points and the most negative are about -125 log points. Notice that the mean reversion pattern varies with the size of the shock, with much stronger mean reversion in $t + 1$ for large shocks and smaller reversion for smaller shocks. Furthermore, even at the 10-year horizon, a nonnegligible fraction of the shocks’ effect is still present, indicating a permanent component to these shocks.

To illustrate these patterns more clearly, Figure 13 plots the shock, $y_t^i - y_{t-1}^i$, on the x -axis and the fraction of each shock that has mean-reverted, $y_{t+k}^i - y_t^i$, on the y -axis for the median RE group. Thus, this figure contains the same information as in Figure 12

FIGURE 14 – Impulse Responses, Prime-Age Workers with Low or High RE




but is reported differently. Several remarks are in order. First, negative earnings changes tend to be less persistent than positive earnings changes. For example, a worker whose earnings rise by 100 log points between $t - 1$ and t loses about 50% of this gain in the following 10 years. It is also interesting to note that almost all of this mean reversion happens after one year, implying that whatever mean reversion there is happens very quickly. Turning to earnings losses: a worker whose earnings fall 100 log points recovers one-third of that loss by $t + 1$ and recovers more than two-thirds of the total within 10 years. Moreover, unlike with positive shocks, the recovery (hence mean reversion) is more gradual in response to negative shocks.

Second, the degree of mean reversion varies with the magnitude of earnings shocks. This is evident in Figure 13, where small shocks (i.e., those less than 10 log points in absolute value) look very persistent, whereas there is substantial mean reversion following larger earnings changes. A univariate autoregressive process with a single persistence parameter will fail to capture this behavior. In the next section, we will allow for multiple AR(1) processes to accommodate the variation in persistence by shock size.

In the next figure (14), we plot the same kind of impulse response functions but now for workers that are in the 10th percentile (left panel) and 90th percentile (right) of RE distribution. Notice that, for low-income individuals, negative shocks mean-revert much more quickly, whereas positive shocks are more persistent than before. The opposite is true for high-income individuals.

FIGURE 15 – Asymmetric Mean Reversion: Butterfly Pattern



We now extend the results of Figure 14 to the entire distribution of recent earnings. To make the comparison clear, we focus on a fixed horizon, 10 years, and plot the total mean reversion between t and $t + 10$ for the 6 RE groups in Figure 15. Starting from the lowest RE group (those individuals in the bottom 5% of the recent earnings distribution), notice that negative shocks are transitory, with an almost 80% mean reversion rate at the 10-year horizon. But positive shocks are quite persistent, with only about a 20% mean reversion at the same horizon. As we move up the RE distribution, the positive and negative branches of each graph start rotating in *opposite* directions, so that for the highest RE group, we have the opposite pattern: only 20 to 25% of negative shocks mean-revert at the 10-year horizon, whereas almost 75% of positive shocks mean-revert at the same horizon. We refer to this shape as the “butterfly pattern.”

5 Estimating Stochastic Processes for Earnings

With the few exceptions noted earlier, the bulk of the earnings dynamics literature relies on the (often implicit) assumption that earnings shocks can be approximated reasonably well with a lognormal distribution. This assumption, combined with linear time series models (e.g., an ARMA(p, q) process) to capture the accumulation of such shocks,

allowed researchers to focus their estimation to match the covariance matrix of log earnings either in levels or in first difference form, ignoring higher-order moments.²⁴

This approach has two problems. First, the broad range of evidence presented in the previous sections implies that this approach is likely to miss important aspects of the data and produces a picture of earnings risk that does not capture salient features of the risks faced by workers. Second, the covariance matrix estimation method makes it difficult to select among alternative models of earnings risk, because it is difficult to judge the relative importance—from an economic standpoint—of the covariances that a given model matches well and those that it does not. This is an especially important shortcoming given that virtually every econometric process used to calibrate economic models is statistically rejected by the data.

With these considerations in mind, we propose and implement a different approach that relies on matching the kinds of moments presented above. We believe that economists can more easily judge whether or not each one of these moments is relevant for the economic questions they have in hand. Therefore, they can decide whether the inability of a particular stochastic process to match a given moment is a catastrophic failure or an acceptable shortcoming. They can similarly judge the success of a given stochastic process in matching some moments and not others.

More concretely, in the first stage, we use each set of moments presented above as diagnostic tools to determine the basic components that should be included in the stochastic process that we will then fit to the data. Clearly, this stage requires extensive pre-testing and exploratory work. For example, to generate the life-cycle earnings growth patterns documented in Section 3.1, we considered three basic ingredients: (i) an AR(1) process + an i.i.d. shock, (ii) growth rate heterogeneity with no shocks, and (iii) a mixture of two AR(1) processes where each component receives a nonzero innovation with a certain probability. We picked the first two ingredients because of the widespread attention they garnered in the previous literature, and the third one based on our conjecture that it might perform well. We found that the first ingredient, on its own, could not generate the rich patterns of earnings growth revealed by the data, whereas the HIP process performed fairly well, and the AR(1) mixture process performed the best.

²⁴Exceptions include [Browning *et al.* \(2010\)](#), [Altonji *et al.* \(2013\)](#), and [Guvenen and Smith \(2014\)](#). Clearly, GMM or minimum distance estimation that is used to match such moments does not require the assumption of lognormality for consistency. But abstracting away from moments higher than covariances is a reflection of the belief that higher-order moments do not contain independent information, which relies on the lognormality assumption.

Therefore, we concluded that a stochastic process for earnings should include either (ii) or (iii) (or both) as one of its components. Although earnings growth data on their own could not determine which one of these pieces is more important, when we analyze these data along with the other moments, we will be able to obtain sharper identification of the parameters of these two components.

We conducted similar diagnostic analyses on the other cross-sectional moments (standard deviation, skewness, and kurtosis) as well as on impulse response moments. The variation in the second to fourth moments over the life cycle and with earnings levels seemed impossible to match without introducing explicit dependence of shocks to these two characteristics. Allowing the mixing probabilities to depend on age and earnings delivered much improved results, so we make this specification part of our benchmark.

5.1 A Flexible Stochastic Process

The most general econometric process we estimate has the following features: (i) a heterogeneous income profiles (hereafter, HIP) component of quadratic form, (ii) a mixture of three AR(1) processes, denoted by z , x , and ν , where each component receives a new innovation in a given year with probability $p_j \in [0, 1]$ for $j = z, x, \nu$; and (iii) an i.i.d. transitory shock. Here is the full specification:

$$\tilde{y}_t^i = (\alpha^i + \beta^i t + \gamma^i t^2) + z_t^i + x_t^i + v_t^i + \varepsilon_t^i \quad (2)$$

$$z_t^i = \rho_z z_{t-1}^i + \eta_{zt}^i \quad (3)$$

$$x_t^i = \rho_x x_{t-1}^i + \eta_{xt}^i \quad (4)$$

$$\nu_t^i = \rho_\nu \nu_{t-1}^i + \eta_{\nu t}^i, \quad (5)$$

where for $j = z, x, v$:

$$\eta_{jt}^{*i} \sim \mathcal{N}(\mu_j, \sigma_j^i) \quad \text{and} \quad \eta_{jt}^i = \eta_{jt}^{*i} \times \mathbb{I}\{s_{i,t} \in I_{p_j}\} \quad (6)$$

$$\log \sigma_j^i \sim \mathcal{N}(\bar{\sigma}_j - \frac{\sigma_{jj}^2}{2}, \sigma_{jj}^2), j = z, x, \quad \sigma_\nu^i \equiv \sigma_\nu. \quad (7)$$

Here, (α, β, γ) follows a multivariate normal distribution with zero mean and a covariance matrix to be estimated. The realizations of the three innovations $(\eta_{jt}, j = z, x, v)$ are mutually exclusive—only one of the three shocks is received per period. This is implemented by first drawing a standard uniform random variable, $s_{i,t}$, for a given individual

at age t and dividing the unit interval into three pieces: $I_{p_z} = [0, p_z]$, $I_{p_x} = (p_z, p_z + p_x]$, and $I_{p_v} = (p_z + p_x, 1]$, where $p_z + p_x \leq 1$ and $p_v = 1 - p_z - p_x$. Depending on which interval $s_{i,t}$ falls in, that innovation is set to its random draw, η_{jt}^{*i} , and the others are set equal to zero.

The specification in (7) implies that the innovation standard deviation for each AR(1) process has an individual-specific component that is lognormal, with mean $\bar{\sigma}_j$ and standard deviation proportional to σ_{jj} . To economize on parameters, we assume that the permanent innovation is identically distributed across individuals. Regarding the initial conditions of the persistent processes, z_0^i and x_0^i , we assume that they are drawn from a normal distribution with zero mean and standard deviation $\sigma_{j,0}$, $j = z, x$.²⁵ Finally, to avoid indeterminacy in the estimation, without loss of generality we impose $\rho_z < \rho_x < \rho_v$.

We allow the mixing probabilities to depend on age and on the lagged idiosyncratic component of earnings:

$$p_j(y_{t-1}^i, t) = a_j + b_j \times y_{t-1}^i + c_j \times t + d_j \times y_{t-1}^i \times t. \quad (8)$$

Recall that in Section 3, we defined y_t^i as log earnings net of age effects. It can alternatively be defined as $y_t^i = z_t^i + x_t^i + v_t^i + \varepsilon_t^i$. Finally, since p_j 's are probabilities, they must remain between $[0,1]$, so equation (8) is truncated outside of this range.²⁶

5.2 Details of Estimation Procedure

We estimate the parameters of the described stochastic process using a method of simulated moments (MSM) estimator. The empirical targets are: (i) moments from the life-cycle profile of average earnings, summarized in Figures 4 and 5; (ii) the standard deviation, skewness, and kurtosis of one-year and five-year earnings growth presented in Sections 3.2 to 3.4; and (iii) moments describing the impulse response functions presented in Section 4. In addition, and as noted above, the within-cohort variance of log earnings levels is a key dimension of the data that has been extensively studied in previous research. For both completeness and consistency with earlier work, we include these variances (summarized in Figure A.2) as a fourth set of moments.

²⁵The initial variance of the permanent component cannot be identified from the fixed effect variance and hence is normalized to zero.

²⁶We have also considered an alternative specification where the innovation variances are functions of earnings and age: $\sigma_j(y_{t-1}^i, t) = a_j + b_j \times y_{t-1}^i + c_j \times t + d_j \times y_{t-1}^i \times t$. After extensive experimentation with this formulation, we have found it to perform very poorly.

Accounting for Zeros. Recall that in order to construct the cross-sectional moments, we have dropped individuals who had very low earnings—below \bar{Y}_{\min} —in year t or $t+k$ so as to allow taking logarithms in a sensible manner.²⁷ Although this approach made sense for documenting empirical facts that are easy to interpret, for the estimation exercise, we would like to also capture the patterns of these “zeros” (or very low earnings observations), given that they clearly contain valuable information. This is feasible thanks to the flexibility of the MSM approach, and in Appendix B we describe how we modify the cross-sectional moments to allow these zeros in the estimation.

Aggregating Moments. If we were to match all data points in all these moments (i.e., for every RE percentile and every age group), it would yield more than 10,000 moments. Although this step is doable, not much is likely to be gained from such a level of detail, and it would make the diagnostics—that is, judging the performance of the estimation—quite difficult. To avoid this, we aggregate 100 RE percentiles into 10 to 15 groups and the 6 age groups into two (ages 25–34 and 35–55). Full details of how this aggregation is performed and which moments are targeted are included in Appendix B. After the aggregation procedure, we are left with 120 moments that capture average earnings over the life cycle (targeted ages are 25, 30, 35, 40, 45, 50, 55, and 60); 156 moments for cross-sectional statistics (standard deviation, skewness and kurtosis of one-year and five-year earnings growth); and 1,120 moments coming from the impulse response functions. Adding the 36 moments on the variance of log earnings, in sum, we target a total of $120 + 156 + 1,120 + 36 = 1,432$ moments.²⁸

Let m_n for $n = 1, \dots, N = 1,432$ denote a generic empirical moment, and let $d_n(\theta)$ be the corresponding model moment that is simulated for a given vector of earnings process parameters, θ . We simulate the entire earnings histories of 200,000 individuals who enter labor market at age 25 and work until age 60. When computing the model moments, we apply precisely the same sample selection criteria and employ the same methodology to the simulated data as we did with the actual data. To deal with potential issues that could arise from the large variation in the scales of the moments, we minimize the *scaled* deviation between each data target and the corresponding simulated model moment. For

²⁷We were able to include those below the threshold in sets (i) and (iii) because for those moments it made sense to first take averages, including zeros, and then take the logarithms of those averages.

²⁸The full set of moments targeted in the estimation are reported (in Excel format) as part of an online appendix available from the authors’ websites.

each moment n , define

$$F_n(\theta) = \frac{d_n(\theta) - m_n}{|m_n| + \gamma_n},$$

where $\gamma_n > 0$ is an adjustment factor. When $\gamma_n = 0$ and m_n is positive, F_n is simply the percentage deviation between data and model moments. This measure becomes problematic when the data moment is very close to zero, which is not unusual (e.g., impulse response of log earnings changes close to zero). To account for this, we choose γ_n to be equal to the 10th percentile of the distribution of the absolute value of the moments in a given set. The MSM estimator is

$$\hat{\theta} = \arg \min_{\theta} \mathbf{F}(\theta)'W\mathbf{F}(\theta), \tag{9}$$

where $\mathbf{F}(\theta)$ is a column vector in which all moment conditions are stacked, that is,

$$\mathbf{F}(\theta) = [F_1(\theta), \dots, F_N(\theta)]^T.$$

The weighting matrix, W , is chosen such that the life-cycle average earnings growth moments and impulse response moments are assigned a relative weight of 0.25 each, the cross-sectional moments of earnings growth receive a relative weight of 0.35, and the variance of log earnings is given a relative weight of 0.15.²⁹ The objective function is highly jagged in certain directions and highly nonlinear in general, owing to the fact that we target higher-order moments and percentiles of the distribution. Therefore, we employ a global optimization routine, described in further detail in [Güvenen \(2013\)](#), to perform the minimization in (9). Further details can be found in [Appendix B](#).

6 Results: Estimates of Stochastic Processes

Table [III](#) reports the parameter estimates. Before delving into the discussion of these estimates, we begin with an overview of what each of the eight columns aims to capture. Columns (1) to (3) take the general stochastic process (equations (2) to (8)) and restrict the persistent process ν_t to a unit root ($\rho_\nu \equiv 1$). The three columns only differ in the order of the HIP component: quadratic, linear, and none. Notice how reducing the order from quadratic to linear causes only a small rise in the objective value (from 16.65 to 17.17), whereas eliminating the HIP component altogether (going from column 1 to 2 to

²⁹More precisely, each life-cycle growth moment is weighed by 0.25/120, each cross-sectional moment by 0.35/156, each impulse response moment by 0.25/1120, and each variance moment by 0.15/36.

TABLE III – Estimates of Stochastic Process Parameters

<i>Specification:</i>	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Benchmark model		Best fit			Standard model		
HIP order	2	1	0	1	1	1	1	0
AR(1)	2	2	2	3	2	1	1	1
RW	1	1	1	0	0	0	0	0
Heterog. variances?	yes	yes	yes	yes	yes	yes	no	no
<i>Parameter</i>	<i>Values</i>							
σ_α	0.523	0.552	0.681	0.475	0.464	0.698	0.679	0.673
$\sigma_\beta \times 10$	0.149	0.130	—	0.137	0.191	0.067	0.252	—
$\sigma_\gamma \times 100$	0.043	—	—	—	—	—	—	—
$\text{corr}_{\alpha\beta}$	-0.48	-0.49	—	-0.04	0.07	-0.12	0.02	—
$\text{corr}_{\alpha\gamma}$	-0.06	—	—	—	—	—	—	—
$\text{corr}_{\beta\gamma}$	-0.44	—	—	—	—	—	—	—
ρ_z	0.249	0.259	0.182	0.083	0.226	1.004	0.962	1.00
ρ_x	0.368	0.425	0.315	0.512	0.639	—	—	—
ρ_ν	1.0*	1.0*	—	0.796	—	—	—	—
$\bar{\sigma}_z$	0.853	0.847	0.947	0.656	0.786	0.121	0.174	0.172
$\bar{\sigma}_x$	0.422	0.361	0.425	0.337	0.123	—	—	—
$\bar{\sigma}_\nu$	0.096	0.087	0.174	0.082	—	—	—	—
σ_ε	0.039	0.029	0.023	0.019	0.017	0.029	0.027	0.040
σ_{z0}	0.261	0.133	0.229	0.121	0.206	0.370	0.031	0.304
σ_{x0}	0.186	0.089	0.295	0.113	0.271	—	—	—
$\sigma_{\nu0}$	—	—	—	0.110	—	—	—	—
σ_{zz}	0.121	0.107	0.134	0.167	0.066	0.313	—	—
σ_{xx}	0.142	0.194	0.170	0.294	0.550	—	—	—
μ_z	-0.544	-0.426	-0.780	-0.518	-0.231	—	—	—
μ_x	0.020	0.021	0.197	0.237	0.175	—	—	—
μ_ν	0.068	0.060	0.064	0.303	—	—	—	—
<i>Mixture probabilities:</i>								
$a_z \times 1$	0.064	0.067	0.025	0.148	0.215	—	—	—
$b_z \times t$	0.036	0.038	0.001	0.005	0.027	—	—	—
$c_z \times y_{t-1}$	-0.207	-0.203	0.027	-0.244	-0.293	—	—	—
$d_z \times t \times y_{t-1}$	0.076	0.071	0.072	0.080	0.056	—	—	—
$a_x \times 1$	0.358	0.433	0.152	0.481	—	—	—	—
$b_x \times t$	-0.112	-0.222	-0.030	-0.019	—	—	—	—
$c_x \times y_{t-1}$	0.044	0.072	-0.066	-0.045	—	—	—	—
$d_x \times t \times y_{t-1}$	-0.328	-0.373	0.099	-0.156	—	—	—	—
<i>Objective value</i>	16.65	17.17	24.11	15.29	18.30	38.10	38.46	38.75
<i>Decomposition:</i>								
(i) Cross-section	5.527	5.901	9.097	5.448	6.130	19.730	19.944	19.918
(ii) Impulse resp.	9.923	10.134	14.033	8.812	11.108	17.550	17.465	17.795
(iii) Inc. growth	0.719	0.644	0.965	0.958	0.623	0.629	0.918	0.863
(iv) Inequality	0.484	0.492	0.020	0.076	0.435	0.194	0.131	0.179

Note: * $\rho_\nu = 1.0$ is imposed.

3) causes a larger jump, from 17.17 to 24.11. Consequently, in columns (4) to (6), we keep the linear HIP specification but instead vary the number of persistent components. That is, column (4) allows for three AR(1) components, achieved by relaxing the unit root constraint on ν_t and estimating its persistence ρ_ν ; column (5) reduces the number of AR(1) components to two, and column (6) reduces it to one, thereby eliminating the mixture structure. In other words, column (6) is the first case in which the process is Gaussian, because z receives a normal shock with 100% probability every period. Notice the substantial jump in the objective value when the mixture structure is eliminated going from column (5) to (6). As we will see in a moment, this result is a reflection of the very poor fit delivered by the Gaussian process.

Based on the overall goodness of fit and the relative economy of its parameters (relative to columns (1) and (4)), we will refer to the specification in column (2) as our *benchmark model*. However, it should be stressed that this is not a very strict preference, and depending on the application, either one of the specifications in columns (1) or (4) could serve as a more appropriate benchmark. For quick reference, we refer to column (4) as the *best fit model*, since it delivers the lowest objective value.

The process estimated in column (6) is quite simple relative to those that came before it, but it is still not as simple as what is used to calibrate economic models in the extant literature. Therefore, column (7) further slims down the process by shutting down the heterogeneity in innovation variances: $\sigma_{zz} \equiv 0$. Finally, column (8) shuts down the linear HIP component yet again to obtain a process that is essentially the canonical earnings process in the current literature: an AR(1) process plus an i.i.d. shock. We refer to this last and simplest process as the *standard model*.

With this overview, we are now ready to discuss the parameter estimates in more detail.

6.1 Parameter Estimates

Because the specifications in columns (1) and (2) deliver similar objective values and parameter estimates, we begin by discussing them together and in some detail. For the estimates in the remaining columns, we will mainly point out the changes from the benchmark model.

TABLE IV – Mixture Probabilities of Persistent Components

	Age groups				RE (Percentile) groups				
	25–29	35–39	45–49	55–60	1	10	50	90	100
<i>Benchmark (Col. 2)</i>									
p_z ($\rho_z = 0.26$)	0.09	0.07	0.11	0.26	0.25	0.19	0.12	0.09	0.08
p_x ($\rho_x = 0.43$)	0.32	0.12	0.06	0.04	0.60	0.29	0.09	0.04	0.03
p_ν ($\rho_\nu = 1.0^*$)	0.59	0.81	0.83	0.72	0.15	0.52	0.79	0.87	0.91
<i>Best fit (Col. 4)</i>									
p_z ($\rho_z = 0.08$)	0.11	0.06	0.09	0.19	0.30	0.20	0.09	0.05	0.05
p_x ($\rho_x = 0.51$)	0.44	0.22	0.13	0.14	0.66	0.51	0.18	0.07	0.00
p_ν ($\rho_\nu = 0.80$)	0.45	0.82	0.79	0.67	0.04	0.29	0.63	0.88	0.95

The Life-Cycle Profile. The estimate of σ_α is about 0.50, implying large permanent *level* differences in earnings across individuals.³⁰ More important, earnings growth rates display significant heterogeneity: σ_β is 0.0149 (or approximately 1.5%) in column (1), and the dispersion in the quadratic component is also sizable. The fixed effect is negatively correlated with the growth rate (about -0.50).³¹ The fact that the linear HIP specification in the benchmark model matches the moments nearly as well suggests that the additional flexibility afforded by the quadratic component is not critical for the empirical facts we study in this paper. But the quadratic specification introduces three more parameters (σ_γ , $\text{corr}_{\alpha\gamma}$, $\text{corr}_{\beta\gamma}$) relative to the linear one, which is why the latter is likely to be preferable in most applications.

The Mixture Component. We now turn to the stochastic component driving the dynamics: the most important element is the mixture component. The first AR(1) component has a low persistence of $\rho_z = 0.26$ but a very high innovation standard deviation that ranges from 0.65 for the lowest 1% of individuals to 1.08 for the top 1% (ranked by σ_z^i), with a mean of $\bar{\sigma}_x = 0.85$. The second AR(1) component is only slightly more persistent, $\rho_x = 0.43$ with about half the average innovation volatility of z : $\bar{\sigma}_x = 0.36$, and a dispersion for σ_x^i between 0.23 and 0.57 (for the bottom and top 1% respectively).³² The third and final persistent component is restricted to be a random walk ($\rho_\nu = 1.0$) and has $\sigma_\nu = 0.087$, assumed to be common to all individuals.

³⁰This figure is on the high end of earlier estimates; see, for example, Haider (2001), Storesletten *et al.* (2004), and Heathcote *et al.* (2010).

³¹These figures are comparable to earlier estimates. It is perhaps surprising that even though the moments targeted here include a much broader set of statistics than these previous studies, and the process considered here has many more components, the estimates of growth rate heterogeneity are not too different.

³²This sizable heterogeneity is consistent with that in the work of Browning *et al.* (2010).

We now turn to the mixture probabilities, which are nontrivial functions of age and past earnings. Rather than try to interpret the coefficients a_j to d_j ($j = z, x$) directly, it is more useful to calculate the values of p_j for different age and past earnings groups. The top panel of Table IV reports these probabilities for the benchmark model. As seen in the left panel, the very large innovations to z are realized with a low probability, ranging from 7% to 26%, and are flat until the middle ages and increase into old ages. The midsize innovations into x are realized with a probability of 32% for the youngest workers, declining to 4% for the oldest workers. The third process, ν_t (the most persistent component), receives innovation with the highest probability, ranging from 59% to 83% and displaying a hump-shaped pattern over the life cycle.³³ The top right panel of Table IV displays how these probabilities vary with the earnings level. The first two components with low persistence (z and x) show a monotonically declining probability with earnings, implying that the permanent component displays the opposite pattern.³⁴

To summarize the mixture structure, the component that receives the largest shocks displays the smallest persistence and the smallest probability of realization, whereas the component that receives the smallest shocks (about 10 times smaller than the first component) displays the highest persistence and the highest probability realization. Therefore, individuals typically receive small shocks that are quite persistent, and once in a while they receive much larger shocks (η_z or η_x) that are less persistent. It is this structure that is behind the success of the model in generating the very high kurtosis observed in the data (as well as its variation across groups, to be analyzed in a moment).


The mean value of each process is also worth commenting on. The volatile process z_t has a large negative mean, $\mu_z = -0.544$, which—together with the positive but quite small means of the other two components ($\mu_x = 0.02$, $\mu_\nu = 0.068$)—is largely responsible for the longer left tail of the earnings change distribution as well as the small hump on the left shoulder, in turn driving the negative skewness seen in the data.

Finally, the i.i.d. shock, ε , has a very small standard deviation of $\sigma_\varepsilon = 0.039$ in column (1) and even smaller in some of the other specifications in the table. It seems that relatively transitory shocks exist in the data and can be very large (such as η_z), but these are not shocks that are regularly realized every period.

³³Karahan and Ozkan (2013) find a similar pattern in the PSID: earnings shocks are only moderately persistent for young workers and persistence increases with age until during the first half of the life cycle.

³⁴Figure A.25 provides a more detailed look at these probabilities by plotting 3D graphs of p_j as a function of t and y_{t-1} .

FIGURE 16 – Histogram of Log Earnings Changes: US Data vs. Benchmark Model



Mixture Model with No HIP Component (Column 3). Shutting down the HIP component (setting $\beta^i \equiv 0$) leads to a considerably worse fit, almost equally due to a poorer fit to cross-sectional moments (objective rising from 5.901 to 9.097) and to impulse response moments (rising from 10.134 to 14.033). The most notable changes in parameter estimates happen in the volatilities: $\bar{\sigma}_z$ goes up from 0.85 to 0.95, and more important, the permanent innovation, $\bar{\sigma}_v$, nearly doubles, from 0.087 up to 0.174. However, as we elaborate below, the fit of this specification to some key moments are poor, so we do not discuss this case any further.

Best Fit (Column 4). So far, we have restricted ν_t to a random walk, both because this is a pervasive assumption in the earnings dynamics literature and because it eliminates one parameter (ρ_v), making it potentially more convenient for use in calibration. In column (4) we relax the unit root assumption and estimate a persistence parameter for this component of $\rho_v = 0.80$. This added flexibility leads to a non-negligible improvement in the fit of the model, attaining the lowest objective value (15.29), indicating that the unit root restriction is not supported by the data. In the bottom panel of Table IV, we report the innovation probabilities, p_j , for this specification. As seen here, the rough magnitudes of probabilities as well as their pattern of variation remain largely unchanged relative to the benchmark model in the top panel.

Simplified Benchmark (Column 5). Persistent components are expensive to include in quantitative economic models because each one introduces a new state variable. Therefore, in column 4, we take the benchmark model and eliminate the random walk component, reducing the number of parameters significantly (along with losing some flexibility). Not surprisingly, the fit worsens but not by a large amount (objective rising from 17.17 to 18.30). This is also borne out in the model-data comparisons of moments in the next section. In our view, this is the simplest specification that can reasonably be considered for calibrating economic models.

Models with No Mixture (Columns 6, 7, and 8). We now reduce the number of AR(1) components to 1. This also eliminates variation in moments with age and RE, since these were delivered by variations in mixing probabilities. As seen in Table III, the objective value more than doubles for all three cases. The estimated persistence is close to a unit root. And whether we include a HIP component (column 7) or allow for the innovation variance to be heterogeneous across individuals does little to improve this poor fit. Clearly, this critique applies to the standard model that is the workhorse specification in the incomplete markets literature. In light of these results, we find such overly simplified models to be unsuitable for serious quantitative economic work.

6.2 Model Fit to Key Moments


A large number of moments are targeted in the estimation, and discussing how the estimated models fit all of them is not feasible. Therefore, in this section, we discuss the model fit to some of the key moments and relegate the remaining moments to Appendix C. Because many of the targeted moments display patterns by age and RE levels that are visually easy to see, in Figures 16 to 18 we plot these moments as well as their counterparts simulated from each of the eight estimated models.³⁵

Figure 16 plots the histogram of $y_{t+1} - y_t$ generated by the benchmark model, superimposed on the data counterpart. The fit is quite good, except at the center where the model misses the size of the spike in the histogram. Next, Figure 17 plots the cross-sectional moments of the data for prime-age individuals. The top, middle, and bottom panels display, respectively, the standard deviation, skewness, and kurtosis of earnings


³⁵Eagle-eyed readers will notice that the statistics plotted in this section look slightly different from those reported before in Section 3. This is because here we are plotting the exact moments that are targeted in the estimation, which have been aggregated as described in the previous section and also include observations below the threshold level Y_{\min} .

FIGURE 17 – Fit of Estimated Model to Key Cross-Sectional Data Moments


(A) Std. dev. of $y_{t+1}^i - y_t^i$




(B) Std. dev. of $y_{t+5}^i - y_t^i$




(C) Skewness of $y_{t+1}^i - y_t^i$



(D) Skewness of $y_{t+5}^i - y_t^i$



(E) Kurtosis of $y_{t+1}^i - y_t^i$



(F) Kurtosis of $y_{t+5}^i - y_t^i$




TABLE V – Percentiles of Earnings Growth Distribution: Benchmark Model vs. Data

		Percentiles						
Percentile	P1	P5	P10	P25	P75	P90	P95	P99
Data	-1.84	-0.75	-0.35	-0.07	0.13	0.34	0.64	1.40
Model	-1.58	-0.73	-0.34	-0.09	0.12	0.34	0.65	1.33
		Percentile Differentials						
	P99 – P91		P95 – P5		P90 – P10		P75 – P25	
Data	3.24		1.39		0.69		0.20	
Model	2.91		1.38		0.68		0.22	

changes. The subfigures on the left are for one-year changes, and those on the right are for five-year changes. Although the graphs display data generated from each one of the eight specifications, we will focus on three important cases: (i) the benchmark model, (ii) the standard model, and (iii) column (3), which is the same as the benchmark model but without a HIP component. The first is shown with a solid black line with circles; the second one with a dashed brown line; and the last one with a thick gray solid line.


The first observation is that all specifications understate the standard deviation of shocks both at both short and long horizons. We have experimented with further (unreported) specifications but failed to match this aspect of the data. The reason seems to largely be that the earnings changes in the data display extremely long tails, and the mixture of three processes fails to completely capture this feature without compromising fit in other dimensions. This can be seen in Table V, which plots selected percentiles from the annual earnings growth distribution (top panel). As seen here, the benchmark model does an excellent job of matching the all percentiles between the 5th and 95th, but fails to generate the very low values of the lowest and highest percentiles. Consequently, the P99-P1 difference (bottom panel) is 3.24 in the data but only 2.91 in the model (a difference of 33 log percents).³⁶

That said, the benchmark model still performs better for this moment than other specifications (including the *best fit* model). It captures the declining pattern with RE percentiles, although fails to capture the sharp rise in volatility at the very top of the RE distribution. The no-HIP specification yields an increasing and then flat profile, whereas the standard model substantially underestimates the standard deviation (generating only

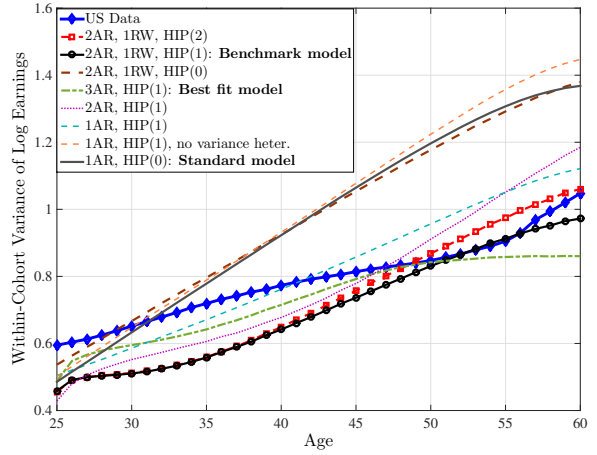
³⁶We suspect that mixing a fourth AR(1) process could improve fit along this dimension, but we have not tried that (demanding) specification. The current benchmark model takes more than one week to estimate on a cluster using 100+ state-of-the-art CPUs in parallel.

FIGURE 18 – Fit of Estimated Model, Continued

(A) Life-Cycle Earnings Growth: Ages 25 to 55



(B) Variance of Log Earnings (right)



one-quarter of the empirical value for the median worker).

Turning to skewness, we now see that the benchmark model performs considerably better, matching both the qualitative patterns as well as the overall level of skewness throughout the distribution for one-year changes. Furthermore, both the benchmark model and the best-fit specification perform nearly as well for skewness. However, the no-HIP model (column 3) generates less than half of the level of negative skewness at the one-year horizon and fails to generate *any* skewness at the five-year horizon. As could be expected, once the mixture component is shut down (columns 6, 7, and 8), the model generates zero skewness at all horizons, which is grossly inconsistent with the data.

The bottom panel shows the kurtosis patterns, and the benchmark model again does well. It manages to match both the qualitative pattern of rising kurtosis with RE levels up to the 90th percentile and the subsequent decline, as well as its levels throughout the picture. As with skewness, the other specifications that allow for both HIP and a mixture structure (columns 1, 4, and even 5) perform comparably well. However, if the HIP specification is eliminated (as in column 3), this reduces the level of kurtosis at its peak by almost half; further eliminating the mixture structure reduces the kurtosis to 4.5. Finally, the standard model is Gaussian with a kurtosis of 3.0.

Figure 18a plots the fit of each model to life-cycle earnings growth moments. The quadratic HIP specification in column (1) is the clear winner here, matching the patterns for all LE groups between the 5th and 90th percentiles. The benchmark model is a

close second. The best fit model does less well, overestimating growth for most income groups. However, all models miss the very high income growth for those with high lifetime earnings. This is a deficiency we have not been able to remedy. Future work should focus on understanding the earnings dynamics of top earners in the United States.

Finally, Figure 18b plots the cross-sectional variance of log earnings by age. The benchmark model and best fit model perform arguably the best, matching the total rise in inequality over the life cycle. The standard model and the models in columns (3) and (7) tie for the worst performance, overestimating the rise of inequality by nearly 100%.

7 Concluding Thoughts

Quantifying the Measured Earnings Risk

Consider the well-known thought experiment (Arrow (1965); Pratt (1964)), in which a decision maker chooses between (i) a static gamble that changes consumption, c , by a random proportion $(1 + \tilde{\delta})$ and (ii) a fixed payment π , called the risk premium, to avoid the gamble. An expected utility maximizer solves

$$U(c \times (1 - \pi)) = \mathbb{E} \left[U(c \times (1 + \tilde{\delta})) \right]. \quad (10)$$

Let us compare two scenarios. In the first one, $\tilde{\delta}^A$ is drawn from a Gaussian distribution with mean zero and standard deviation of 0.10. In the second, $\tilde{\delta}^B$ has the same first two moments as $\tilde{\delta}^A$, but it also has a skewness coefficient of -2 and a kurtosis of 30 (roughly corresponding to the one-year earnings change of a 45-year-old individual at P90 of the RE distribution). Further assume that U is a CRRA utility function with a curvature of $\theta = 10$. Equation (10) is an equation in one unknown. The solution for π under different assumptions is displayed in Table VI. As seen here, the willingness to pay to avoid $\tilde{\delta}^B$ is 22.15% compared to 4.88% for $\tilde{\delta}^A$, an amplification of 450%.

To provide a decomposition for the separate effects of skewness and kurtosis, the following equation is helpful. Taking a first-order Taylor-series approximation to the left hand side and fourth-order approximation to the right hand side, we get

$$\pi^* \approx \underbrace{\frac{\theta}{2} \times \sigma_\delta^2}_{\text{variance aversion}} - \underbrace{\left[\frac{(\theta + 1)\theta}{6} \times \sigma_\delta^3 \right]}_{\text{negative skewness aversion}} \times s_\delta + \underbrace{\left[\frac{(\theta + 2)(\theta + 1)\theta}{24} \times \sigma_\delta^4 \right]}_{\text{kurtosis aversion}} \times \kappa_\delta,$$

TABLE VI – Effect of Skewness and Kurtosis on Risk Premium

	Risk Premium (π)	
	3.0	30.0
<i>Kurtosis</i> →		
<i>Skewness</i> ↓		
0.0	4.88%	18.79%
-2.0	6.34%	22.15%

where s_δ and κ_δ are the skewness and kurtosis coefficients, respectively. This expression is a generalization of the well-known formula that only includes the first of these terms.³⁷ Among other things, it shows that the first term—what we often refer to as “risk aversion”—is perhaps more appropriately called “variance aversion,” because with higher-order moments, the individual also displays aversion to negative skewness and to excess kurtosis. These are components of the risk perceived by the individual and can significantly amplify the aversion. In particular, the coefficient in front of skewness is a quadratic in θ and a cubic in σ_δ . Therefore, keeping the skewness fixed, a higher curvature or dispersion can greatly amplify the risk premium. For kurtosis, the situation is even more extreme: the coefficient is a cubic in θ and a quartic in dispersion. As seen in the example, even a standard deviation of 0.10 and $\theta = 10$ delivers a substantially amplified risk premium. Increasing σ_δ from 0.10 to 0.25 makes the negative skewness aversion *twice* as large, and kurtosis aversion *twenty* times as large, as variance aversion.

Clearly, these calculations are meant to be suggestive rather than conclusive. A more thorough analysis in the future should start by measuring insurance and smoothing opportunities against these earnings fluctuations as well as distinguishing between an anticipated earnings change versus an unexpected shock.

Summary and Conclusions

Our analysis of the life-cycle earnings histories of millions of U.S. workers has reached two broad conclusions. First, the higher-order moments of individual earnings shocks display clear and important deviations from lognormality. In particular, earnings shocks display strong negative skewness (what can be viewed as individual disaster shocks) and extremely high kurtosis—as high as 35 compared with 3 for a Gaussian distribution. The high kurtosis implies that in a given year, most individuals experience very small earnings shocks, few experience middling shocks, and a small but non-negligible number

³⁷See Appendix B.3 for derivation.

experience extremely large shocks.

The second conclusion of this analysis is that these statistical properties of earnings shocks change substantially both over the life cycle and with the earnings level of individuals. For example, our estimates of the stochastic process show that low-income individuals experience very large earnings shocks with low persistence, whereas high-income individuals experience shocks that are very persistent but with much lower volatility. These properties of earnings shocks over the life cycle are likely to generate a range of interesting implications for economic behavior.

We have also estimated impulse response functions of earnings shocks and found significant asymmetries: positive shocks to high-earnings individuals are quite transitory, whereas negative shocks are very persistent; the opposite is true for low-earnings individuals. While these statistical properties are typically ignored in quantitative analyses of life-cycle models, they are fully consistent with search-theoretic models of careers over the life cycle. After establishing these empirical facts nonparametrically, we estimated what we think is the simplest earnings process that is broadly consistent with these salient features of the data.

A broader message of this paper is a call to reconsider the way researchers approach the study of earnings dynamics. The covariance matrix approach that dominates current work is too opaque and a bit mysterious: it is difficult to judge what it means to match or miss certain covariances in terms of their economic implications. With the current trend toward the increasing availability of very large panel data sets, we believe that researchers' priority in choosing methods should shift from efficiency concerns to transparency. The approach adopted here is an example of the latter, and we believe that economists can better judge what each moment implies for economic questions.

The nonparametric empirical facts documented in Sections 3 and 4 (some of which are reported in an online appendix to save space) add up to more than 10,000 empirical moments of individual earnings data. The richness of this information is far beyond what we are able to fully utilize in the estimation exercise in this paper. Furthermore, for different questions, it would make sense to focus on a different subset of these moments from what we have aimed for in this paper. With these considerations in mind, we make these detailed moments available online for download as an Excel file. We hope the empirical findings about the nature of earnings dynamics documented in this paper feed back into economic research and policy analyses.

References

- ABOWD, J. M. and CARD, D. (1989). On the covariance structure of earnings and hours changes. *Econometrica*, **57** (2), 411–45.
- ALTONJI, J., SMITH, A. A. and VIDANGOS, I. (2013). Modeling earnings dynamics. *Econometrica*, **81** (4), 1395–1454.
- ARELLANO, M., BLUNDELL, R. and BONHOMME, S. (2014). *Household Earnings and Consumption: A Nonlinear Framework*. mimeo, University College London.
- ARROW, K. (1965). *Aspects of the Theory of Risk Bearing*, Yrjö Jahnsson lectures, Yrjö Jahnssonin Saatio, Helsinki.
- ATTANASIO, O. P., BANKS, J., MEGHIR, C. and WEBER, G. (1999). Humps and bumps in lifetime consumption. *Journal of Business & Economic Statistics*, **17** (1), 22–35.
- BACHMANN, R. and BAYER, C. (2014). Investment dispersion and the business cycle. *American Economic Review*, **104** (4), 1392–1416.
- BAGGER, J., FONTAINE, F., POSTEL-VINAY, F. and ROBIN, J.-M. (2014). Tenure, experience, human capital and wages: A tractable equilibrium search model of wage dynamics. *American Economic Review*, **104** (6), 1551–1596.
- BAKER, M. and SOLON, G. (2003). Earnings dynamics and inequality among Canadian men, 1976–1992: Evidence from longitudinal income tax records. *Journal of Labor Economics*, **21** (3), 289–321.
- BERGER, D. and VAVRA, J. (2011). *Dynamics of the U.S. Price Distribution*. Working paper, Yale University.
- BLOOM, N., FLOETOTTO, M., JAIMOVICH, N., SAPORTA-EKSTEN, I. and TERRY, S. J. (2011). *Really Uncertain Business Cycles*. Working paper, Stanford University.
- BLUNDELL, R., GRABER, M. and MOGSTAD, M. (2014). Labor income dynamics and the insurance from taxes, transfers, and the family. *Journal of Public Economics*.
- BONHOMME, S. and ROBIN, J.-M. (2009). Assessing the equalizing force of mobility using short panels: France, 1990 a 2000. *Review of Economic Studies*, **76** (1), 63–92.

- BOROVICKA, J., HANSEN, L. P. and SHEINKMAN, J. A. (2014). *Shock Elasticities and Impulse Responses*. Working paper, University of Chicago.
- BROWNING, M., EJRNAES, M. and ALVAREZ, J. (2010). Modelling income processes with lots of heterogeneity. *Review of Economic Studies*, **77**, 1353–1381.
- CHAMBERLAIN, G. and HIRANO, K. (1999). Predictive distributions based on longitudinal earnings data. *Annales d'Économie et de Statistique*, **55-56**, 211–242.
- CHRISTIANO, L. J., EICHENBAUM, M. and EVANS, C. L. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, **113** (1), 1–45.
- CONSTANTINIDES, G. M. and GHOSH, A. (2014). *Asset Pricing with Countercyclical Household Consumption Risk*. Working paper, University of Chicago.
- DEATON, A. (1991). Saving and liquidity constraints. *Econometrica*, **59** (5), 1221–48.
- and PAXSON, C. (1994). Intertemporal choice and inequality. *Journal of Political Economy*, **102** (3), 437–67.
- GEWEKE, J. and KEANE, M. (2000). An empirical analysis of earnings dynamics among men in the PSID: 1968-1989. *Journal of Econometrics*, **96**, 293–356.
- and — (2007). Smoothly mixing regressions. *Journal of Econometrics*, **138**, 252–290.
- GOLOSOV, M., TROSHKIN, M. and TSYVINSKI, A. (2014). *Redistribution and Social Insurance*. Working paper, Princeton University.
- GOURINCHAS, P.-O. and PARKER, J. A. (2002). Consumption over the life cycle. *Econometrica*, **70** (1), 47–89.
- GUVENEN, F. (2013). Quantitative economics with heterogeneity. Book Manuscript, University of Minnesota.
- , OZKAN, S. and SONG, J. (2014). The nature of countercyclical income risk. *Journal of Political Economy*, **122** (3), 621–660.
- and SMITH, A. A. (2014). Inferring labor income risk and partial insurance from economic choices. *Econometrica*.

- HAIDER, S. J. (2001). Earnings instability and earnings inequality of males in the united states: 1967-1991. *Journal of Labor Economics*, **19** (4), 799–836.
- HAUSE, J. C. (1980). The fine structure of earnings and the on-the-job training hypothesis. *Econometrica*, **48** (4), 1013–1029.
- HEATHCOTE, J., STORESLETTEN, K. and VIOLANTE, G. L. (2010). The macroeconomic implications of rising wage inequality in the united states. *Journal of Political Economy*, **118** (4), 681–722.
- KARAHAN, F. and OZKAN, S. (2013). On the persistence of income shocks over the life cycle: Evidence, theory, and implications. *Review of Economic Dynamics*, **16** (3), 452–476.
- LILLARD, L. A. and WEISS, Y. (1979). Components of variation in panel earnings data: American scientists 1960-70. *Econometrica*, **47** (2), 437–454.
- and WILLIS, R. J. (1978). Dynamic aspects of earnings mobility. *Econometrica*, **46** (5), 985–1012.
- LOW, H., MEGHIR, C. and PISTAFERRI, L. (2010). Wage risk and employment risk over the life cycle. *American Economic Review*, **100** (4), 1432–1467.
- 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.
- MANKIW, N. G. (1986). The equity premium and the concentration of aggregate shocks. *Journal of Financial Economics*, **17** (1), 211–219.
- MCKAY, A. (2014). *Time-Varying Idiosyncratic Risk and Aggregate Consumption Dynamics*. Working paper, Boston University.
- MEGHIR, C. and PISTAFERRI, L. (2004). Income variance dynamics and heterogeneity. *Econometrica*, **72** (1), 1–32.
- MIDRIGAN, V. (2011). Menu costs, multiproduct firms, and aggregate fluctuations. *Econometrica*, **79** (4), 1139–1180.
- MOORS, J. J. A. (1986). The meaning of kurtosis: Darlington reexamined. *The American Statistician*, **40**, 283–284.

- MURPHY, K. M. and WELCH, F. (1990). Empirical age-earnings profiles. *Journal of Labor Economics*, **8** (2), 202–29.
- NAKAMURA, E. and STEINSSON, J. (2013). Price rigidity: Microeconomic evidence and macroeconomic implications. *Annual Review of Economics*, **5** (133-163).
- OLSEN, A. and HUDSON, R. (2009). Social security administration’s master earnings file: Background information. *Social Security Bulletin*, **69** (3), 29–46.
- PANIS, C., EULLER, R., GRANT, C., BRADLEY, M., PETERSON, C. E., HIRSCHER, R. and STINBERG, P. (2000). *SSA Program Data User’s Manual*. Baltimore, MD: Social Security Administration.
- PRATT, J. W. (1964). Risk aversion in the small and in the large. *Econometrica*, **32** (1/2), 122–136.
- STORESLETTEN, K., TELMER, C. I. and YARON, A. (2004). Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy*, **112** (3), 695–717.
- TOPEL, R. H. and WARD, M. P. (1992). Job mobility and the careers of young men. *The Quarterly Journal of Economics*, **107** (2), 439–79.
- ZHANG, H., CONN, A. R. and SCHEINBERG, K. (2010). A derivative-free algorithm for least-squares minimization. *SIAM Journal on Optimization*, **20** (6), 3555–3576.

Supplemental Online Appendix

NOT FOR PUBLICATION

A Appendix: Additional Figures

A.1 Further Figures

In this section, we report some additional figures of interest that are omitted from the main text due to space constraints. First, Figure A.1 plots the selected percentiles of the future earnings change distribution for every RE percentile. Second, Figure A.1 plots the cross-sectional variance of log earnings over the life cycle, constructed along the lines described in Deaton and Paxson (1994).

A.2 Robustness and Extensions

A.2.1 Decomposing Moments: Job-Stayers vs. Job-Switchers

In this section, we present the the cross-sectional statistics analyzed in Section 3 by first splitting the sample in each year depending on whether a worker switched employers or whether he stayed at the same job.

One challenge we face is that many workers hold multiple jobs in a given year, which requires us to be careful about how to think of job changes. We have explored several plausible definitions and found broadly very similar results. Here we describe one reasonable classification. A worker is said to be a “job-stayer” between years t and $t + 1$ if a given EIN (employer identification number) provides the largest share of his annual earnings (out of all his EINs in that year) in years $t - 1$ through $t + 2$, and if the same EIN provides at least 90% of his total annual earnings in years t and $t + 1$. A worker is defined as a “job-switcher” if he is not a job-stayer.³⁸



As seen in Figure A.4, the results are consistent with what we might expect. Job-stayers (i) face a dispersion of earnings changes that is less than half that of job-changers, (ii) face shocks that have zero or slightly positive skewness as opposed to job-switchers who face shocks that are very negatively skewed, and (iii) experience shocks with much higher kurtosis than job-switchers. In fact, kurtosis is as high as 43 for annual changes and 28.5 for five-year changes for job-stayers, but is less than 10 for job-switchers at both horizons.

Before concluding this section, we examine how the tails of the earnings change distribution affect the computed statistics and also examine how this effect varies by stayers and switchers. Unlike with survey-based data, here we are not too concerned that these tails might be dominated by measurement error, as most of these changes are likely to be genuine. Instead, we are simply interested in understanding what parts of the distribution are critical for the different moments that we estimated so far. Table A.1 reports the 2nd to 4th moments for the original sample used so far (left panel) as well as for a sample where we drop extreme observations, defined as those in the top or bottom 1% of the earnings change distribution. As expected, discarding the tails reduces all statistics, but does not change the ranking between stayers and switchers.

³⁸Clearly, this classification is quite stringent for classifying somebody as a job stayer, meaning that some individuals will be classified as job switchers even though they did not change a job. An alternative definition we have explored defines a job switcher directly as somebody who has an EIN that provides more than 50% of his annual earnings in year t and provides less than 10% of his annual earnings in year $t + 1$; and also has an EIN that provides less than 10% of his earnings in t and more than 50% in $t + 1$. The results were very similar to those reported here.

FIGURE A.1 – Quantiles of Earnings Changes

(A) Quantiles of One-Year Earnings Changes



(B) Quantiles of Five-Year Earnings Changes

FIGURE A.2 – Within-Cohort Variance of Log Earnings




FIGURE A.3 – Fraction Staying Jobs Between t and $t + 1$




FIGURE A.4 – Second to Fourth Moments of Annual Earnings Growth: Stayers vs Switchers

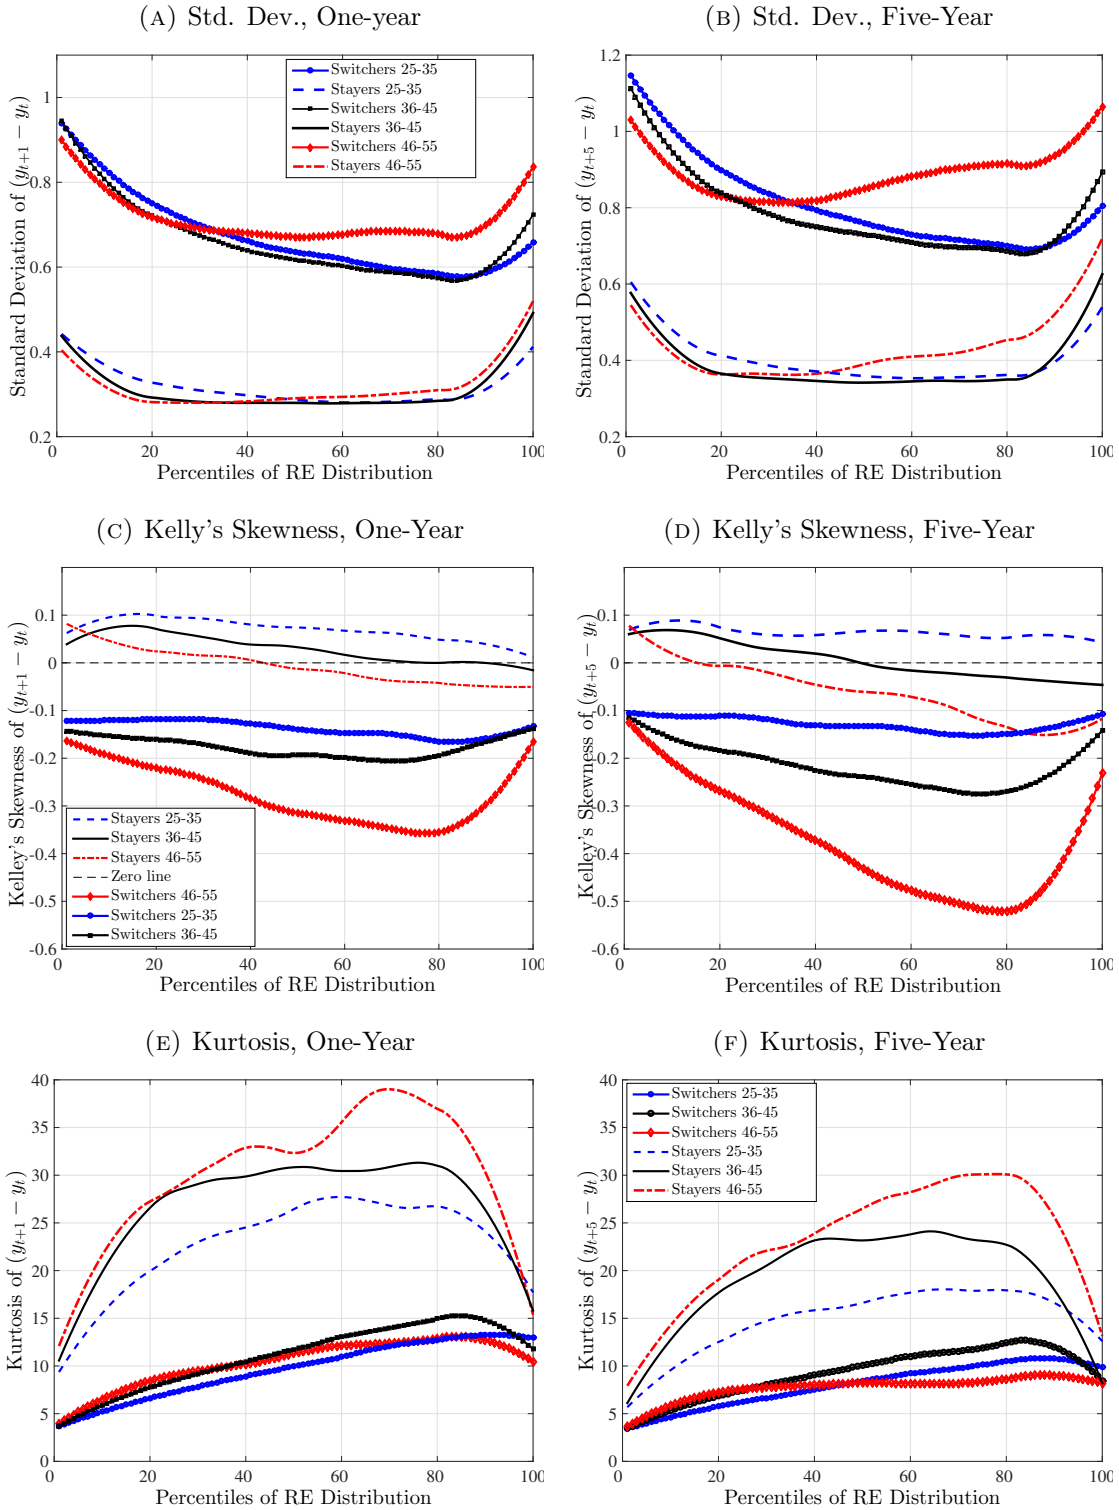



TABLE A.1 – Moments with Trimmed Data

	Original			Trim top/bot. 1%		
	Std. dev.	Skew.	Kurtosis	Std. dev.	Skew.	Kurtosis
One-Year						
All	0.49	-1.49	18.35	0.36	-0.89	8.91
Stayers	0.31	-1.51	42.81	0.20	-0.51	13.50
Switchers	0.67	-1.09	9.69	0.56	-0.85	6.17
Five-Year						
All	0.66	-1.06	12.16	0.53	-0.60	6.58
Stayers	0.41	-1.04	28.54	0.29	0.10	7.16
Switchers	0.78	-0.92	8.86	0.65	-0.67	5.61

A.2.2 Disentangling The Effects of Age and Recent Earnings (Case II)

Figures A.5 and A.6 plot the 3D graphs of skewness and kurtosis that jointly conditions on age and recent earnings as described in Section 3.5.


FIGURE A.5 – Skewness of One-year Earnings Change, Jointly Conditioning on Age and Recent Earnings



A.2.3 Cross-Sectional Facts

This section reports the analogs of figures in Section 3 by considering three additional robustness exercises. In particular, case III averages earnings over two consecutive years before

FIGURE A.6 – Kurtosis of One-year Earnings Change, Joint Conditioning on Age and Recent Earnings




computing statistics:

$$\tilde{\Delta}y_t^i = \log(Y_{t+3}^i + Y_{t+2}^i) - \log(Y_t^i + Y_{t+1}^i) \quad \text{and} \quad \tilde{\Delta}_5y_t = \log(Y_{t+5}^i + Y_{t+6}^i) - \log(Y_t^i + Y_{t+1}^i).$$


Case IV takes the difference between earnings in year $t + k$ and usual earnings:

$$\Delta_{\text{short}}y^i = \log(Y_{t+1}^i) - \log(\bar{Y}_{t-1}^i) \quad \text{and} \quad \Delta_{\text{long}}y^i = \log(Y_{t+5}^i) - \log(\bar{Y}_{t-1}^i).$$

Case V considers two methods of trimming the tails. The first method excludes the top and bottom 1% of earnings growth observations (referred to as Case V.A in figures). The second method changes the lower threshold for sample exclusion from $Y_{\min,t}$ to be individual-specific and equal to 5% of each worker's own recent earnings (Case V.B). For all case numbers indicated below, see Section 3.5 in the main text contains more detailed descriptions.



(A) Case III: Bi-Annual



(B) Case IV: Relative to Usual

FIGURE A.7 – Standard Deviation of Annual Earnings Growth




FIGURE A.8 – Standard Deviation of Bi-Annual Five-Year Earnings Growth, Case III




FIGURE A.9 – Standard Deviation of Five-Year Earnings Growth, Usual, Case IV




FIGURE A.10 – Skewness of (Bi-Ann) Annual Earnings Growth, Case III




FIGURE A.11 – Skewness of Annual Earnings Growth, Usual, Case IV




FIGURE A.12 – Skewness of (Bi-Ann) Five-Year Earnings Growth, Case III

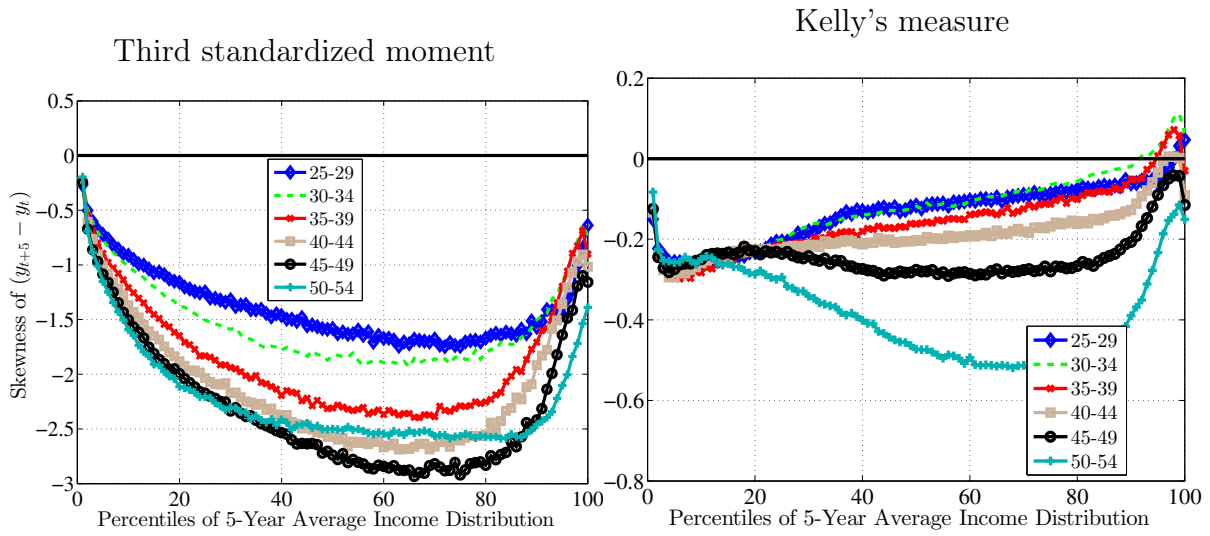


FIGURE A.13 – Skewness of Five-Year Earnings Growth, Usual, Case IV

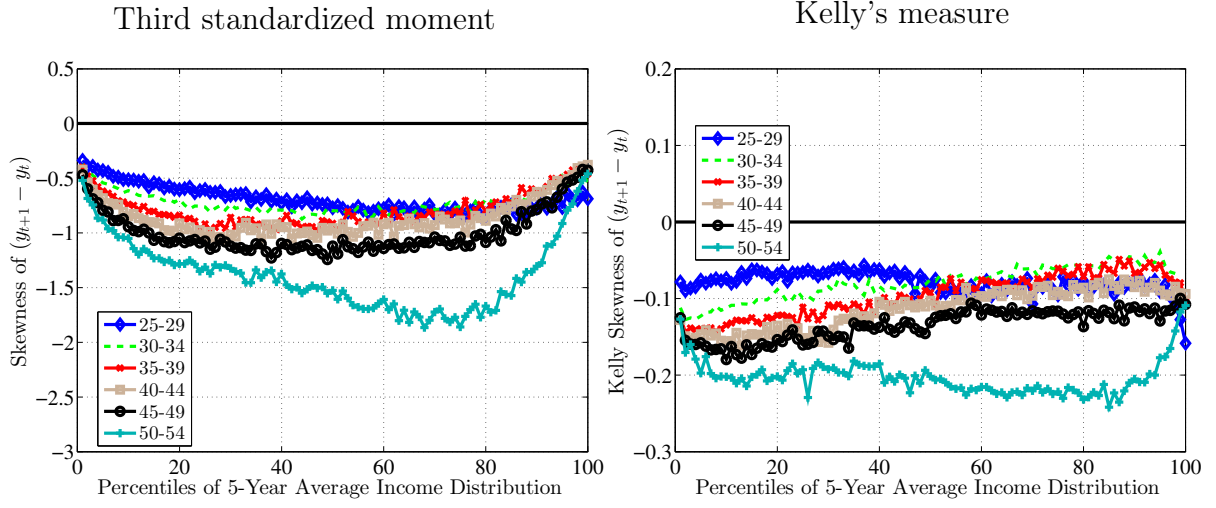


FIGURE A.14 – Skewness of One-Year Earnings Growth, Trimming Tails, Case V.A




FIGURE A.15 – Skewness of Five-Year Earnings Growth, Trimming Tails, Case V.A




FIGURE A.16 – Skewness of One-Year Earnings Growth, Trimming Tails, Case V.B




FIGURE A.17 – Skewness of Five-Year Earnings Growth, Trimming Tails, Case V.B

FIGURE A.19 – Kurtosis of Annual Earnings Change, Usual, Case IV




FIGURE A.18 – Kurtosis of Bi-Annual Earnings Change, Case III




FIGURE A.20 – Kurtosis of Five-Year Earnings Change, Bi-Annual, Case III




FIGURE A.21 – Kurtosis of Five-Year Earnings Change, Usual, Case IV




FIGURE A.22 – Kurtosis of One-Year Earnings Change, Trimming Tails, Case V.A




FIGURE A.23 – Kurtosis of Five-Year Earnings Change, Trimming Tails, Case V.A






FIGURE A.24 – Kurtosis, Robustness

(A) Kurtosis of One-Year Earnings Change, Excluding those with Earnings lower than 5% of RE, Case V.B



(B) Kurtosis of Five-Year Earnings Change, Excluding those with Earnings lower than 5% of RE, Case V.B

B Details of Estimation Method

B.1 Moment Selection and Aggregation

1. 1. Cross-sectional moments of earnings changes. In Section 3, we have dropped individuals with low earnings ($\exp(y_{it}) < \bar{Y}_{\min}$) in t or $t + k$ (or both) when constructing the moments of the cross-sectional distribution. We now modify this procedure slightly so as to allow the incorporation of these previously dropped observations. The basic idea is to transform these observations by mapping them into earnings values that are slightly above \bar{Y}_{\min} . Specifically, the transformed earnings values are obtained as

$$\tilde{y}_{i,t} = \ln \left(\bar{Y}_{\min} + \frac{\bar{Y}_{\min}}{10(\bar{Y}_{\min} + 10)} (\exp y_{i,t} + 10x) \right),$$

where $x \sim U(0, 1)$ follows a uniform distribution.

In order to have a smooth earnings change distribution, we replace earnings observations lower than our minimum earnings threshold by mapping them to levels above. This imputation has the following three goals: First, there is a mass of individuals with an earnings observation of zero, which would otherwise be dropped when constructing log earnings changes. We would like to use these observations as they contain information about the nature of earnings risk. Second, we would like to avoid having any mass in the distribution to avoid possible jaggedness in the objective function. Third, we want to preserve the ranking among individuals with less than threshold earnings.

In order to capture the variation in the cross-sectional moments of earnings changes along the age and recent earnings dimensions, we condition the distribution of earnings changes on these variables. For this purpose, we first group workers into 6 age bins (five-year age bins between 25 and 54) and within each age bin into 13 selected groups of RE percentiles in age $t - 1$. The RE percentiles are grouped as follows: 1, 2–10, 11–0, 21–30, ..., 81–90, 91–95, 96–99, 100. Thus, we compute the three moments of the distribution of one- and five-year earnings changes for $6 \times 13 = 78$ different groups of workers. We aggregate these 6 age bins into 2 age groups, A_{t-1}^i . The first age group is defined as young workers between ages 25 through 34, whereas the second age group is defined as prime-age workers between the ages of 35 and 54. Consequently, we end up with $2 \times 2 \times 13 \times 3 = 156$ cross-sectional moments.

2. Mean of log earnings growth. The second set of moments captures the heterogeneity in log earnings growth over the working life across workers that are in different percentiles of the LE distribution. We target the average dollar earnings at 8 points over the life cycle: ages 25, 30, ..., and 60 for different LE groups. We combine LE percentiles into larger groups to keep the number of moments at a manageable number, yielding 15 groups consisting of percentiles of the LE distribution: 1, 2–5, 6–10, 11–20, 21–30, ..., 81–90, 91–95, 96–97, 98–99, and 100. The total number of moments we target in this set is $8 \times 15 = 120$.

3. Impulse response functions. We target average log earnings growth over the next k years for $k = 1, 2, 3, 5, 10$, that is, $\mathbb{E}[y_{t+k} - y_t]$, conditional on groups formed by crossing age, \bar{Y}_{t-1} , and $y_t^i - y_{t-1}^i$. We impute income observations lower than our minimum income threshold following the same procedure as in the smoothing of the cross-sectional moments.

We aggregate age groups into two: young workers (25-34) and prime-age workers (35-55). In each year, individuals are assigned to groups based on their ranking in the age-specific RE distribution. The following list defines these groups in terms of the RE distribution: 1–5, 6–10, 11–30, 31–50, 51–70, 71–90, 91–95, 96–100. We then group workers based on the percentiles of the age- and RE-specific income change distribution, $y_t^i - y_{t-1}^i$: 1–2, 3–5, 6–10, 11–20, 21–30, 31–40, 41–50, 51–60, 61–70, 71–80, 81–90, 91–95, 96–98, 99–100. As a result, we have a total of $2 \times 8 \times 14 \times 5 = 1120$ moments based on impulse response.

4. Variance of log earnings. We target the life cycle profile of the variance of log earnings for income observations larger than our minimum income threshold. As a result, we have a total of 36 moments based on the variance of log earnings.

B.2 Numerical Method for Estimation

The estimation objective is maximized as described now. The global stage is a multi-start algorithm where candidate parameter vectors are uniform Sobol' (quasi-random) points. We typically take about 10,000 initial Sobol' points for pre-testing and select the best 2000 points (i.e., ranked by objective value) for the multiple restart procedure. The local minimization stage is performed with a mixture of Nelder-Mead's downhill simplex algorithm (which is slow but performs well on difficult objectives) and the DFNLS algorithm of [Zhang et al. \(2010\)](#), which is much faster but has a higher tendency to be stuck at local minima. We have found that the combination balances speed with reliability and provides good results.

B.3 Derivation of the Risk Premium

$$U(c(1 - \pi)) = \mathbb{E} \left(U(c(1 + \tilde{\delta})) \right)$$

Taking a first-order Taylor-series approximation to the left hand side and fourth order approximation to the right hand side we get:

$$U(c) - U'(c)c\pi = \mathbb{E} \left(U(c) + U'(c)c\tilde{\delta} + \frac{1}{2}U''(c)c^2\tilde{\delta}^2 + \frac{1}{6}U'''(c)c^3\tilde{\delta}^3 + \frac{1}{24}U''''(c)c^4\tilde{\delta}^4 \right).$$

Observing that the second term on the right hand side is zero when $\mathbb{E}(\tilde{\delta}) = 0$, and rearranging yields:

$$\pi = -\frac{1}{2} \frac{U''(c)c}{U'(c)} \times m_2 - \frac{1}{3} \frac{U'''(c)c}{U''(c)} \times m_3 + \frac{1}{4} \frac{U''''(c)c}{U'''(c)} \times m_4, \quad (11)$$

where m_n denotes the n^{th} central moment of $\tilde{\delta}$. To convert these into statistics that are reported in the paper, we write $m_2 = \sigma_\delta^2$, $m_3 = s_\delta \times \sigma_\delta^3$, where s_δ is the skewness coefficient, and $m_4 = k_\delta \times \sigma_\delta^4$, where k_δ is kurtosis. With this notation, and assuming a CRRA utility function with curvature θ , we get:

$$\pi^* = \frac{\theta}{2} \times \sigma_\delta^2 - \frac{(\theta + 1)\theta}{6} \times s_\delta \times \sigma_\delta^3 + \frac{(\theta + 2)(\theta + 1)\theta}{24} \times k_\delta \times \sigma_\delta^4,$$


which can also be written as:

$$\pi^* = \frac{\theta}{2} \times \sigma_\delta^2 \times \left[1 + \frac{1}{3}(\theta + 1) \left(-s_\delta \times \sigma_\delta + \frac{1}{4}(\theta + 2)k_\delta \times \sigma_\delta^2 \right) \right].$$


C Estimation Results

FIGURE A.25 – 3-D Plot of Mixing Probabilities


(A) Mixing Probability p_z for z_t



(B) Mixing Probability p_x for x_t



(C) Mixing Probability p_ν for ν_t



C.1 Model Fit: Additional Figures

FIGURE A.26 – Fit of Estimated Model to Key Impulse Response Data Moments for the 8 RE Groups

