

WHAT DO DATA ON MILLIONS OF U.S. WORKERS REVEAL ABOUT LIFECYCLE EARNINGS DYNAMICS?

FATIH GUVENEN

Department of Economics, University of Minnesota, FRB of Minneapolis, and NBER

FATIH KARAHAN

Microeconomic Studies Function, FRB of New York

SERDAR OZKAN

Department of Economics, University of Toronto

JAE SONG

Social Security Administration

We study individual male earnings dynamics over the life cycle using panel data on millions of U.S. workers. Using nonparametric methods, we first show that the distribution of earnings changes exhibits substantial deviations from lognormality, such as negative skewness and very high kurtosis. Further, the extent of these nonnormalities varies significantly with age and earnings level, peaking around age 50 and between the 70th and 90th percentiles of the earnings distribution. Second, we estimate nonparametric impulse response functions and find important asymmetries: Positive changes for high-income individuals are quite transitory, whereas negative ones are very persistent; the opposite is true for low-income individuals. Third, we turn to long-run outcomes and find substantial heterogeneity in the cumulative growth rates of earnings and the total number of years individuals spend nonemployed between ages 25 and 55. Finally, by targeting these rich sets of moments, we estimate stochastic processes for earnings that range from the simple to the complex. Our preferred specification features normal mixture innovations to both persistent and transitory components and includes state-dependent long-term nonemployment shocks with a realization probability that varies with age and earnings.

KEYWORDS: Earnings dynamics, higher-order earnings risk, kurtosis, skewness, non-Gaussian shocks, normal mixture.

1. INTRODUCTION

THE GOAL OF THIS PAPER is to characterize the most salient properties of individual earnings dynamics over the life cycle, focusing on *nonnormalities* and *nonlinearities*. First, by

Fatih Guvenen: guvenen@umn.edu

Fatih Karahan: fatih.karahan@ny.frb.org

Serdar Ozkan: serdar.ozkan@utoronto.ca

Jae Song: jae.song@ssa.gov

Special thanks to Gerald Ray at the Social Security Administration for his help and support. For helpful critiques and comments, we thank participants at the 2012 CASEE Conference on Quantitative Macroeconomics, 2013 and 2014 AEA Winter Meetings, 2013 RRC Conference, 2013 JDP Dilemmas on Inequality Lecture at Princeton, 2014 Sveriges Riksbank Conference on Inequality, 2014 NBER Summer Institute, 2015 NBER *EF&G* Meetings, as well as participants at research seminars at numerous universities and institutions. Special thanks to Joe Altonji, Cristian Arellano, Andy Atkeson, Mark Bilal, Richard Blundell, Chris Busch, Bob Hall, Michael Keane, Giuseppe Moscarini, Luigi Pistaferri, Fabien Postel-Vinay, Anthony Smith, and Kjetil Storesletten for comments. The views expressed herein are those of the authors and do not represent those of the Social Security Administration or the Federal Reserve Banks of Minneapolis and New York. Guvenen acknowledges financial support from the National Science Foundation (SES-1357874).

studying its higher-order moments (specifically, skewness and kurtosis), we investigate the distribution of earnings changes and whether it can be well approximated by a normal distribution. Second, we explore mean reversion patterns of earnings changes that may differ between positive and negative changes as well as by size. Finally, we study how these properties vary over the life cycle and across the earnings distribution.

The extent and nature of these nonnormalities and nonlinearities are difficult to predict beforehand, and strong parametric assumptions can mask those features, making it difficult to uncover them. With these considerations in mind, we first employ a fully nonparametric approach and take “high-resolution pictures” of individuals’ earnings histories. To this end, we use administrative panel data from the U.S. Social Security Administration (SSA) covering a long time span from 1978 to 2013, with a substantial sample size (10% random sample of males aged 25–60).¹ Next, using the facts uncovered in this descriptive analysis, we estimate nonlinear and non-Gaussian earnings processes.

Our descriptive analysis covers (i) the properties of the distributions of earnings changes, (ii) the extent of mean reversion during the 10 years following earnings changes, and (iii) workers’ long-term outcomes covering their entire working lives, such as cumulative earnings growth and the incidence of nonemployment.

Starting with the distribution of earnings changes, we find that it is left- (negatively) skewed, and this left-skewness becomes more severe as individuals get older or their earnings increase (or both). For example, workers aged 45–55 and earning about \$100,000 per year (in 2010 dollars) face a five-year log earnings change distribution with a lower tail (the gap between the 50th and 10th percentiles) 2.5 times longer than the upper tail (50th to 90th percentiles). In contrast, young low-income workers face an almost symmetric distribution. The rise in left-skewness over the life cycle is entirely due to a reduction in opportunities for large gains from ages 25 to 45 and to the increasing likelihood of a sharp fall in earnings after age 45.

In addition, earnings growth displays a very high kurtosis relative to a Gaussian density (Figure 1). There are far more people in the data with very small or with extreme earnings

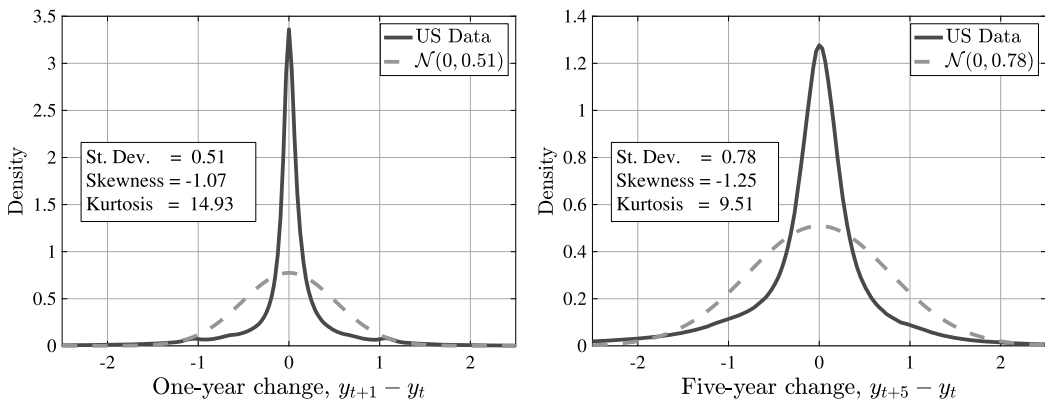


FIGURE 1.—Histograms of one- and five-year log earnings changes. Notes: This figure plots the empirical densities of one- and five-year earnings changes superimposed on Gaussian densities with the same standard deviation. The data are for all workers in the base sample defined in Section 2 and $t = 1997$.

¹In this paper, we focus on men for comparability with earlier work. Our analysis for women found qualitatively similar patterns. These results are reported in the Supplemental Material (Guvener, Karahan, Ozkan, and Song (2021)).

changes and fewer people with middling ones. For example, 31% of the annual earnings changes are less than 5%, compared to only 8% under the Gaussian distribution. Also, a typical worker sees a change larger than three standard deviations with a 2.4% chance, which is about one-ninth as likely under a normal distribution. Importantly, the *average* kurtosis masks significant heterogeneity. For example, five-year earnings growth of males aged 45–55 and earning \$100,000 has a kurtosis of 18, compared to 5 for younger workers earning \$10,000 (and 3 for a Gaussian distribution).²

To shed some light on the sources of these nonnormalities, we analyze data from the Panel Study of Income Dynamics (PSID) on work hours, hourly wages, and a rich set of additional covariates that are not available in the SSA data. We find that hourly wage changes exhibit little left-skewness but an excess kurtosis with a magnitude and lifecycle variation similar to earnings changes. Furthermore, wage changes are at least as important as changes in hours, even in the tails of the distribution. For example, on average out of an earnings decline of 165 log points, 101 log points are due to wages. Moreover, workers experiencing extreme changes are likely to have gone through nonemployment or job or occupation changes, or to have experienced health shocks, suggesting that the tails are not a statistical artifact or measurement error in the survey data.

Next, we characterize the mean reversion patterns of earnings changes by estimating nonparametric impulse response functions conditional on recent earnings and on the size and sign of the change. We find two types of asymmetry. First, fixing the size of the change, positive changes to high-earnings individuals are quite transitory, while negative ones are persistent; in contrast, the opposite is true for low-earnings individuals. Second, with a fixed level of earnings, the strength of the mean reversion differs by the size of the change: Large changes tend to be much more transitory than small ones. These asymmetries are difficult to detect in a covariance matrix, in which all sorts of earnings changes—large, small, positive, and negative—are masked by a single statistic.

Finally, we document two facts regarding long-term outcomes covering individuals' entire working lives. First, the cumulative earnings growth over the life cycle varies systematically and substantially across groups of workers with different lifetime earnings. For example, average earnings rise by 60% from age 25 to age 55 for the median lifetime earnings group, by 4.8-fold for the 95th percentile, and by 27.8-fold for the top 1% groups.³ Second, there is substantial variation in individuals' lifetime nonemployment rate—which we define as the fraction of a lifetime (ages 25 to 60) spent as (full-year) nonemployed. For example, 40% of men experience at most one year of nonemployment, while 18% spend *more than half* of their working years as nonemployed. These numbers imply an extremely high persistence in the long-term nonemployment state.

While the nonparametric approach allows us to establish key features of earnings dynamics in a transparent way, a tractable parametric process is indispensable because (i) it allows us to connect earnings *changes* to underlying *innovations* or *shocks* to earnings, and (ii) it can be used as an input to calibrate quantitative models with idiosyncratic risk.⁴

²Importantly, these non-Gaussian features of earnings growth are mostly due to *persistent* changes, which we infer from both the five-year income growth and the change between two consecutive five-year averages. The income processes we estimate in Section 6 also confirm this conclusion.

³A positive relationship between lifetime earnings and lifecycle earnings growth is to be expected. What is surprising is that the magnitudes are so large that they cannot be explained by simple processes.

⁴Although we follow the common practice in the literature of referring to innovations as income “shocks,” individuals are likely to have more information about them than what we—as econometricians—can identify from earnings data alone. Separating expected from unexpected changes requires either survey data on expectations (e.g., Pistaferri (2003)) or economic choices (e.g., Cunha, Heckman, and Navarro (2005) and Guvenen and Smith (2014)). Tackling this important question is beyond the scope of this paper.

Therefore, in Section 6, we target the empirical moments described above to estimate a range of income processes. We start with the familiar linear-Gaussian framework (i.e., the persistent plus transitory model with Gaussian shocks) and build on it incrementally until we arrive at a rich, yet tractable, benchmark specification that can capture the key features of the data. Along the way, we discuss which aspect of the data each feature helps capture so that researchers can judge the trade-offs between matching a particular moment and the additional complexity it brings.

Our benchmark process incorporates two key features to the linear-Gaussian framework: normal mixture innovations to the persistent and transitory components and, more importantly, a long-term nonemployment shock with a realization probability that depends on age and earnings. This *state-dependent* employment risk generates recurring nonemployment with scarring effects concentrated among young and low-income individuals and helps capture the lifecycle and income variation of the moments. Our empirical facts require non-Gaussian features in *persistent* innovations; these can be achieved by such income-dependent nonemployment shocks or non-Gaussian shocks to the persistent component, but not by a uniform nonemployment risk that is transitory in nature.⁵

Related Literature. The earnings dynamics literature has a long history, dating back to seminal papers by Lillard and Willis (1978), Lillard and Weiss (1979), and MaCurdy (1982). Until recently, this literature focused on linear ARMA-type time series models identified from the variance-covariance matrix, thereby abstracting away from nonlinearities and nonnormalities of the data.⁶

In an important paper, Geweke and Keane (2000) modeled earnings innovations using normal mixture distributions and found important deviations from normality.⁷ More recently, using earnings data from France, Bonhomme and Robin (2009) estimated a flexible copula model for the dependence patterns over time and a mixture of normals for the transitory component that displays excess kurtosis. Bonhomme and Robin (2010) used a nonparametric deconvolution method and found excess kurtosis in both permanent and transitory shocks. We go beyond the overall distribution and document that nonnormalities in transitory and persistent components vary substantially with earnings levels and age. Another related paper is Guvenen, Ozkan, and Song (2014), which shows that earnings growth becomes more left skewed in recessions; however, it abstracts away from lifecycle variation. We go further by analyzing kurtosis and how it varies across the income distribution, asymmetries in mean reversion, and the heterogeneity in lifecycle income growth rates and lifetime nonemployment rates, which are all absent from that paper.

In contemporaneous work, Arellano, Blundell, and Bonhomme (2017) also explored nonlinear earnings dynamics. They proposed a new quantile-based panel data framework, which allows the persistence of earnings to vary with the size and sign of the shock. They found asymmetries in mean reversion and non-Gaussian features that are consistent with

⁵A transitory nonemployment shock (or any transitory shock for that matter) cannot generate left-skewness in earnings growth, because each nonemployment spell contributes same-sized negative and positive earnings change observations (one when the worker becomes nonemployed and one when he returns to work), consequently, stretching both tails of the distribution, leaving the symmetry unaffected.

⁶Quantitative macroeconomists use income processes usually by discretizing them using Tauchen or Rouwenhorst methods. Typical implementations of the Tauchen method assume the shocks are normally distributed. The Rouwenhorst method does not rely on this assumption, but it focuses on matching the first two moments of earnings levels and ignores the mechanical implications for the distribution of changes.

⁷In an even earlier contribution, Horowitz and Markatou (1996) showed how an income dynamics model can be estimated nonparametrically and found evidence of nonnormality in the error components.

our results. They also showed that the consumption response to earnings shocks displays nonlinearities, which we do not study. Relative to that paper, we provide a more in-depth analysis of the conditional skewness and kurtosis of earnings, document how they differ between job-stayers and job-switchers, and examine systematic variation in lifecycle earnings profiles and lifetime employment rates. Overall, the two papers complement each other.

Finally, our work is also related to that of Altonji, Smith, and Vidangos (2013), who estimated a joint process for earnings, wages, hours, and job changes via indirect inference by targeting a rich set of moments. Browning, Ejrnaes, and Alvarez (2010) also employed indirect inference to estimate a process featuring “lots of heterogeneity.” However, neither paper explicitly focuses on higher-order moments, their lifecycle evolution, or asymmetries in mean reversion.

2. DATA AND VARIABLE CONSTRUCTION

2.1. *The SSA Data Set*

We draw a representative 10% panel sample of the U.S. population from the Master Earnings File (MEF) of the SSA. The MEF combines various data sets that go back as far as 1978. For our purposes, the most important variables include labor income from W-2 forms (for each job held by the employee during the year), self-employment income (from the Internal Revenue Service (IRS) tax form Schedule SE), and various demographics (date of birth, sex, and race).⁸ We focus on *total annual labor earnings*, which is the sum of total annual wage income and the labor portion (2/3) of self-employment income.

Wage income is not top coded throughout our sample, whereas self-employment income was capped at the SSA taxable limit until 1994. Although this top coding affects only a small number of individuals, we restrict our sample to the 1994–2013 period to ensure that our analysis is not affected by this issue.⁹ The only exception is our use of the entire 1978 to 2013 period in Section 5, where we analyze long-term outcomes of workers, for which a longer time series is essential. For robustness, we impute self-employment income above the cap for the years before 1994 using quantile regressions. Only a small number of individuals who make substantial income from self-employment are affected by this imputation, so the effect on our results is minimal. The details are provided in Appendix A.1 of the Supplemental Material. Finally, we convert nominal values to real values using the personal consumption expenditure (PCE) deflator, taking 2010 as the base year (see Appendix A for further details on the construction of our sample and variables).

Despite the advantages noted, the data set also has some important drawbacks, such as limited demographic information, the absence of capital income, and the lack of hours (and thus hourly wage) data. To overcome some of these limitations, we supplement our analysis with survey data whenever possible. Another important limitation is the lack of household-level data. Even though a large share of quantitative models focus on individ-

⁸The measure of earnings on the W-2 form (Box 1) includes all wages and salaries, tips, restricted stock grants, exercised stock options, severance payments, and other types of income considered remuneration for labor services by the IRS. It does not include any pre-tax payments to individual retirement accounts (IRAs), retirement annuities, child care expense accounts, or other deferred compensation.

⁹In an earlier version (available on our websites), we conducted the analysis using wage and salary income over the entire sample period and reached the same substantive conclusions.

ual earnings fluctuations (which we study), household earnings dynamics are key for some economic questions, about which we have little to say in this paper.¹⁰

2.2. Sample Selection

Our *base sample* is a revolving panel consisting of males with some labor market attachment that is designed to maximize the sample size (important for precise computation of higher-order moments in finely defined groups) and keep the age structure stable over time. First, in order for an individual-year income observation to be *admissible* to the *base sample*, the individual (i) must be between 25 and 60 years old (the working lifespan) and (ii) have earnings above the minimum income threshold $Y_{\min,t}$, equivalent to earnings from one quarter of full-time work (13 weeks at 40 hours per week) at half of the legal minimum wage in year t (e.g., approximately \$1885 in 2010). The revolving panel for year t then selects individuals that are admissible in $t - 1$ and in at least two more years between $t - 5$ and $t - 2$. This ensures that the individual was participating in the labor market and we can compute a reasonable measure of average recent earnings—a variable used extensively in the paper—which we describe next.

Recent Earnings. The average income of a worker i between years $t - 1$ and $t - 5$ is given by $\hat{Y}_{t-1}^i = \frac{1}{5} \sum_{j=1}^5 \max\{\tilde{Y}_{t-j}^i, Y_{\min,t}\}$, where \tilde{Y}_t^i denotes his earnings in year t . We then control for age and year effects by regressing \hat{Y}_{t-1}^i on age dummies separately for each year, and define the residuals as *recent earnings* (hereafter RE), \bar{Y}_{t-1}^i . In Sections 3 and 4, we will group individuals by age and by \bar{Y}_{t-1}^i to investigate how the properties of income dynamics vary over the life cycle and by income levels.

3. CROSS-SECTIONAL MOMENTS OF EARNINGS GROWTH

In this section, we study the distribution of earnings growth rates by analyzing its second to fourth moments. We start by describing our nonparametric method.

3.1. Empirical Methodology: A Graphical Construct

Our main focus is on how the moments of earnings growth vary with recent earnings and age. To this end, for each year t , we divide individuals into six groups based on their age in $t - 1$ (25–29, ..., 45–54), and then within each age group, sort individuals into 100 percentile groups by their recent earnings \bar{Y}_{t-1}^i . If these groupings are done at a sufficiently fine level, we can think of all individuals within a given age/RE group to be *ex ante* identical (or at least very similar). Then, for each such group, the cross-sectional moments of earnings growth between t and $t + k$ can be viewed as the properties of earnings uncertainty that workers within that group expect to face looking ahead (see Figure 2). In our figures, we plot the average of these moments for each age/RE group over the years between 1997 and 2013- k (for moments by age only, see Appendix C.6 of the Supplemental Material). This approach allows us to compute higher-order moments precisely as each bin contains a large number of observations (see Table A.I for sample size statistics). To

¹⁰Household earnings dynamics and hourly wage dynamics in the PSID also display non-Gaussian and non-linear features (e.g., Arellano, Blundell, and Bonhomme (2017, 2018), and De Nardi, Fella, Knoef, Paz-Pardo, and Van Ooijen (2021)).




FIGURE 2.—Timeline for rolling panel construction.

make the figures more readable, we aggregate the six age groups into three: 25–34, 35–44, 45–54.¹¹

Growth Rate Measures. In our analysis we use two measures of income change, each with its own distinct advantages and trade-offs. The first measure is *log growth rate* of income between t and $t+k$, $\Delta_{\log}^k y_t^i \equiv y_{t+k}^i - y_t^i$, where $y_t^i = \tilde{y}_t^i - d_{t,h(i,t)}$ denote the log income (\tilde{y}_t^i) of individual i in year t at age $h(i, t)$ net of age and year effects $d_{t,h(i,t)}$. $\{d_{t,h}\}_{h=25}^{60}$ are obtained by regressing \tilde{y}_t^i on a full set of age dummies separately in each year. While its familiarity makes the log change a good choice for the descriptive analysis, it has a well-known drawback that observations close to zero need to be dropped or winsorized at an arbitrary value. When we use $\Delta_{\log}^k y_t^i$, we drop individuals from the sample with earnings less than Y_{\min} in t or $t+k$, and lose information in the extensive margin.

Our second measure of income growth—*arc-percent change*—is not prone to this caveat and is commonly used in the firm-dynamics literature, where firm entry and exit are key margins (e.g., Davis, Haltiwanger, and Schuh (1996)). We define $\Delta_{\text{arc}}^k Y_t^i = \frac{Y_{t+k}^i - Y_t^i}{(Y_{t+k}^i + Y_t^i)/2}$, where earnings level $Y_t^i = \frac{\tilde{y}_t^i}{\tilde{d}_{t,h(i,t)}}$ is net of average earnings in age h and year t , $\tilde{d}_{t,h(i,t)}$. Because of its familiarity, we use the log change measure in this section and report the results for arc-percent change in Appendix C.2, which show qualitatively similar patterns.

Transitory versus Persistent Income Changes. As is well understood, longer-term earnings changes (i.e., $\Delta_{\log}^k y_t^i$ with larger k) reflect more persistent innovations. To see this intuition, consider the commonly used random-walk permanent/transitory model in which permanent (η_t^i) and transitory (ε_t^i) innovations are drawn from distributions F_η and F_ε , respectively. We denote the variance, skewness, and excess kurtosis of distribution F_x , $x \in \{\eta, \varepsilon\}$ by σ_x^2 , \mathcal{S}_x , and \mathcal{K}_x , respectively. Then the second to fourth moments of k -year log income growth $\Delta_{\log}^k y_t^i$ are given by (see Appendix B for the derivations):

$$\begin{aligned}
 \sigma^2(\Delta_{\log}^k y_t^i) &= k\sigma_\eta^2 + 2\sigma_\varepsilon^2, \\
 \mathcal{S}(\Delta_{\log}^k y_t^i) &= \frac{k \times \sigma_\eta^3}{\underbrace{(k\sigma_\eta^2 + 2\sigma_\varepsilon^2)^{3/2}}_{<1}} \mathcal{S}_\eta, \\
 \mathcal{K}(\Delta_{\log}^k y_t^i) &= \frac{k \times \sigma_\eta^4}{\underbrace{(k\sigma_\eta^2 + 2\sigma_\varepsilon^2)^2}_{<1}} \mathcal{K}_\eta + \frac{2 \times \sigma_\varepsilon^4}{\underbrace{(k\sigma_\eta^2 + 2\sigma_\varepsilon^2)^2}_{<1}} \mathcal{K}_\varepsilon.
 \end{aligned}
 \tag{1}$$

¹¹To be more precise, the first age group consists of individuals aged 28–34 in year $t-1$ due to the base sample requirement of having a minimum three years of admissible observations.

Equation (1) shows that as k increases, the variance and kurtosis of k -year log change $\Delta_{\log}^k y_t^i$ reflect more of the distribution of η_t^i than that of ε_t^i . Also, skewness is solely driven by permanent changes.¹² Finally, the distribution of $\Delta_{\log}^k y_t^i$ is closer to normal than the underlying distributions of F_η and F_ε , because as innovations η_t^i and ε_t^i accumulate, the distribution of $\Delta_{\log}^k y_t^i$ converges toward Gaussian, per the central limit theorem.

With these considerations in mind, we document the moments of one-year ($k = 1$) and five-year ($k = 5$) residual earnings growth to capture properties of transitory and persistent changes, respectively. As persistent changes have a greater effect on economic decisions compared with easier-to-insure transitory ones, we present the results for $k = 5$ in this section. The figures for $k = 1$ in Appendix C.1 show the same qualitative patterns.

3.2. Second Moment: Standard Deviation

Figure 3(a) plots the standard deviation of five-year residual earnings growth by age and recent earnings (for clarity, we use one marker for every fourth RE percentile group). In the right panel, we also report the difference between the 90th and 10th percentiles of log earnings changes, denoted by P90–P10, which is robust to outliers. Both measures show a pronounced U-shaped pattern by RE for every age group. For example, for 35- to 44-year-olds, the standard deviation falls from 1.05 for the lowest RE group to 0.6 for the 90th percentile, and then rises rapidly to 1.05 for the top 1%.

Baker and Solon (2003) and Karahan and Ozkan (2013) have estimated a U-shaped lifecycle profile for the variance of persistent shocks. Our analysis reveals a more intricate lifecycle variation, as we also condition on RE: Dispersion declines with age for the bottom one third of the RE distribution, is U-shaped until the 95th percentile, and monotonically increases for the top earners. However, notice that the variation with age is quite a bit smaller compared with the RE variation. Importantly, the highest earners (the top 5% or so) are strikingly different from other high earners—even those just below the 95th percentile. The same theme emerges again in higher-order moments.

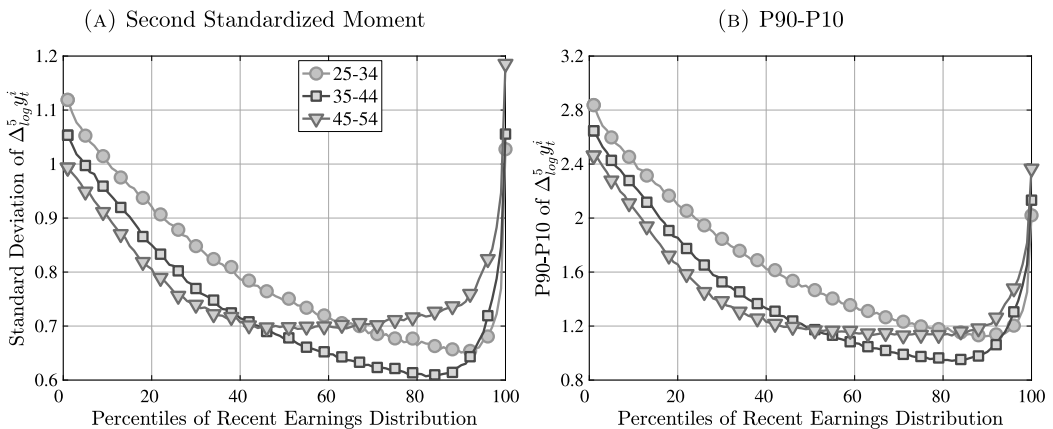


FIGURE 3.—Dispersion of five-year log earnings growth.

¹²A time-varying distribution of transitory innovations can also lead to asymmetry in earnings growth.

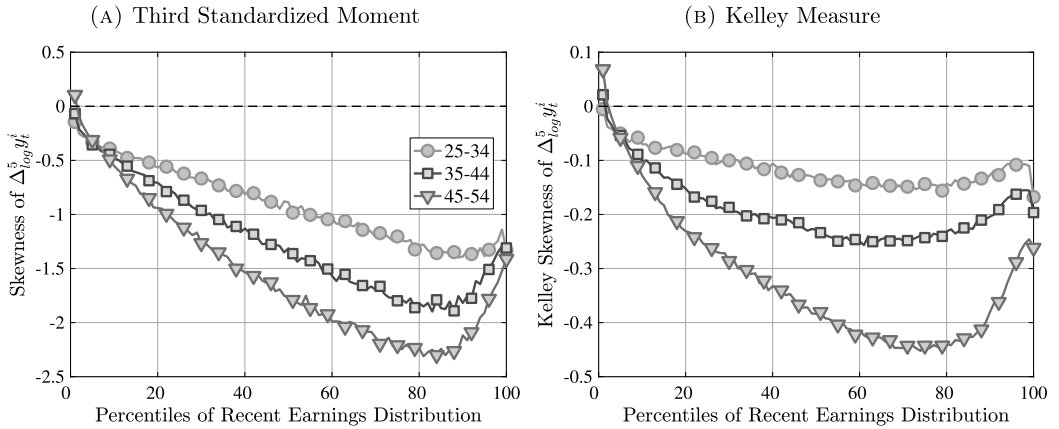


FIGURE 4.—Skewness of five-year log earnings growth.

3.3. Third Moment: Skewness (Asymmetry)

Figure 4(a) plots the skewness of five-year earnings growth, measured as the third standardized moment. First, notice that earnings changes are negatively (left) skewed at every stage of the life cycle and for (almost) all earnings groups. Second, 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. Note that the variation in skewness with age is more muted for individuals at the bottom or top of the (recent) earnings distribution (similar to the dispersion patterns above).

Is negative skewness as measured by the third central moment driven by extreme observations? While the information on tails is important (and becomes especially valuable in estimating income processes in Section 6), we also look at Kelley (1947) skewness, $S_K = \frac{(P90-P50)-(P50-P10)}{P90-P10}$, which is robust to observations above the 90th or below the 10th percentile of the distribution. Basically, S_K measures the relative fractions of the overall dispersion (P90–P10) accounted for by the upper and lower tails. Specifically, $S_K < 0$ implies that the lower tail (P50–P10) is longer than the upper tail (P90–P50).

Kelley’s skewness exhibits essentially the same pattern (Figure 4(b)). Thus, the asymmetry is prevalent across the entire distribution rather than being driven just by the tails. Furthermore, the magnitudes are substantial. For example, a Kelley measure of -0.44 (for 45- to 54-year-old workers at the 80th RE percentile) implies that P90–P50 accounts for 28% of P90–P10, far removed from the 50–50 of a normal distribution.

Another question is whether skewness becomes more negative over the life cycle because of a compression of the upper tail (fewer opportunities for large gains) or because of an expansion in the lower tail (higher risk of large declines). To answer this question, we investigate how the P90–P50 and P50–P10 change over the life cycle from their levels at ages 25–34 (Figure 5). Up until age 44, both the P90–P50 and P50–P10 decline with age across most of the RE distribution. However, the upper tail compresses more strongly than the lower tail, which leads to the increasing left-skewness. After age 45, the P90–P50 keeps shrinking, but the bottom end opens up for workers with above median RE (large declines become more likely). Top earners are again an exception to this pattern: The upper tail does not compress with age, but the bottom end opens up monotonically.

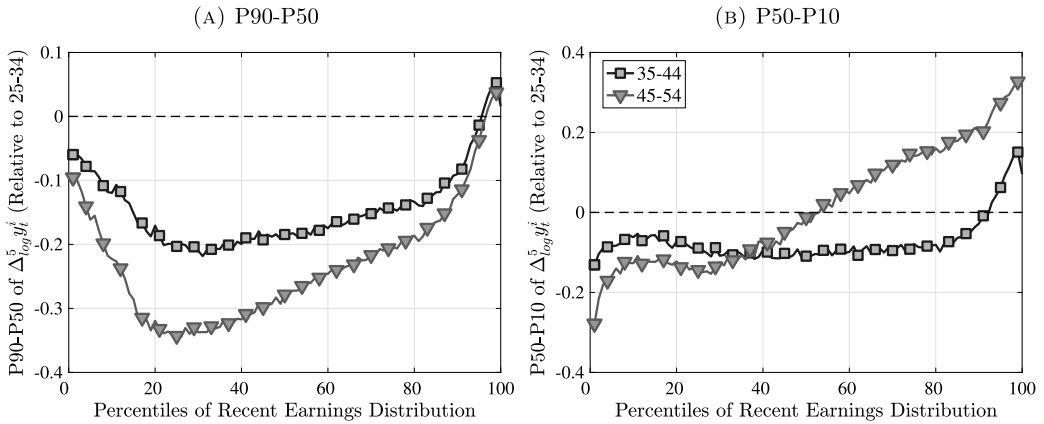


FIGURE 5.—Skewness decomposed: P90–P50 and P50–P10 relative to age 25–34. Notes: The y-axes show the change in P90–P50 and P50–P10 from the youngest age group to the two older age groups.

A natural question is whether the negative skewness is simply due to nonemployment spells. First, notice that nonemployment can generate left-skewness in earnings growth only if it has persistent effects: A transitory spell contributes one negative and one positive earnings change of similar size, leaving the symmetry unaffected (see equation (1)). Jacobson, LaLonde, and Sullivan (1993) and Von Wachter, Song, and Manchester (2009) showed that workers' earnings indeed experience large scarring effects after mass layoffs. We revisit this point in Section 6, where we link earnings changes to the underlying shocks. Second, negative skewness is stronger for upper-middle-income and older workers, for whom unemployment risk is relatively small, implying that a decline in hours is not the main driver. Finally, as noted above, the shift toward more negative skewness is mostly coming from the compression of the right tail up to age 45, which is unlikely to be related to nonemployment.

3.4. Fourth Moment: Kurtosis (Peakedness and Tailedness)

We can think of kurtosis as a measure of the tendency of a density to stay away from $\mu \pm \sigma$ (see Moors (1986)). Thus, a leptokurtic distribution typically has a sharp/pointy center, long tails, and little mass near $\mu \pm \sigma$ (relative to a Gaussian distribution). A corollary to this description is that with excess kurtosis, the usual way we interpret standard deviation—as representing the size of the typical observation—is not very useful because most realizations will be either close to the center or out in the tails.

To illustrate this point, we calculate concentration measures for earnings growth. Table I reports the fraction of individuals experiencing an absolute log earnings change less than a threshold, $|\Delta_{\log}^1 y_t| \leq 0.05, 0.10$, and so on. In the data, 31% of workers experience an earnings change of less than 5%, whereas if innovations were drawn from a Gaussian density with the same standard deviation as the data, only 8% of individuals would experience such changes. Furthermore, extreme events are more likely in the data: A typical worker experiences a change larger than three standard deviations (153 log points) once in a lifetime—with a 2.4% annual chance—whereas this probability is almost one-ninth that size under a normal distribution. These values suggest that the Gaussian assumption vastly overstates the typical earnings growth and misses the extreme changes received by a nonnegligible share of the population.

TABLE I
FRACTION OF INDIVIDUALS WITHIN SELECTED RANGES OF ANNUAL EARNINGS GROWTH^a

S:	Prob($ \Delta_{\log}^1 y_t \in S$)				
	≤ 0.05	≤ 0.10	≤ 0.20	$\geq 2\sigma (\approx 1.0)$	$\geq 3\sigma (\approx 1.5)$
Data	30.6	48.8	66.5	6.64	2.37
$\mathcal{N}(0, 0.51)$	7.7	15.4	30.2	4.55	1.46
Ratio	3.88	3.27	2.23	1.46	8.77

^aNotes: The empirical distribution used in this calculation is for 1997–1998, the same as in Figure 1.

The high likelihood of extreme events in the data motivates us to take a closer look at the tails of the earnings growth distribution by examining its empirical log density versus the Gaussian log density (which is an exact quadratic). First, in line with our previous discussion, the data have much thicker *and* longer tails compared with a normal distribution (Figure 6). Second, the tails decline almost linearly, implying a Pareto distribution at both ends. Third, they are asymmetric, with the left tail declining much more slowly than the right, which contributes to the left-skewness documented above. In fact, fitting linear lines to each tail (in the regions $\pm[1, 4]$) yields a tail index of 1.18 for the right tail and 0.40 for the left tail—the latter showing especially high thickness. We highlight that while the Pareto tail in the earnings levels distribution is well known—indeed, going as far back as Pareto (1897)—the two Pareto tails emerge here in the earnings *growth* distribution. To our knowledge, the present paper is the first to document this fact.¹³

Next, to see how kurtosis varies by age and income, we report two statistics in Figure 7 that are analogous to the ones we used for skewness: the fourth standardized moment and the quantile-based Crow and Siddiqui (1967) measure, which is defined as




FIGURE 6.—Double-Pareto tails of the U.S. annual earnings growth distribution. Notes: The empirical distribution in this figure is for 1997–1998, the same as in Figure 1 but with the y-axis now in logs.

¹³Toda and Walsh (2015) showed that consumption growth also exhibits a double Pareto tail.

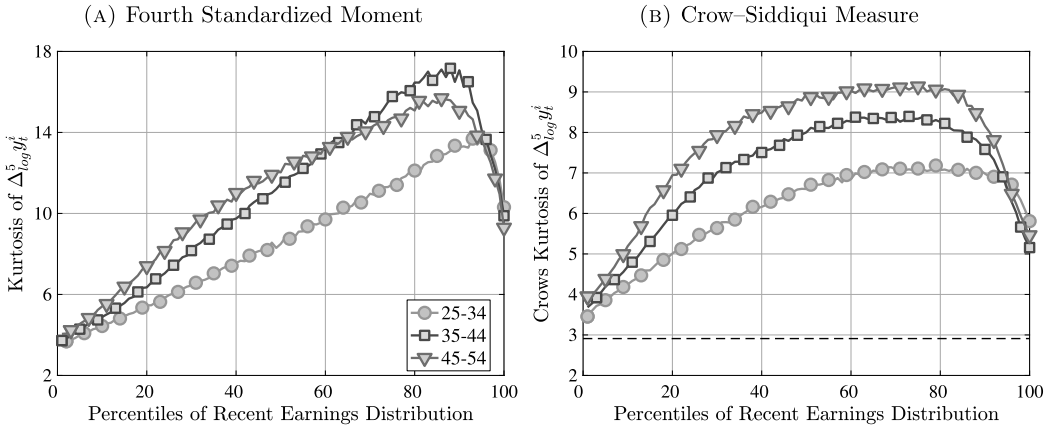


FIGURE 7.—Kurtosis of five-year log earnings growth.

$\kappa_{C-S} = \frac{P_{97.5} - P_{2.5}}{P_{75} - P_{25}}$ and is equal to 2.91 for a Gaussian distribution. As with dispersion and skewness, kurtosis varies substantially with age and recent earnings. For example, for all age groups, the fourth standardized moment increases significantly with recent earnings from 3 (Gaussian) for low-income workers to above 14 around the 90th percentile, and then declines sharply in the top 5% of the RE distribution. Kurtosis also rises with age, especially between the first two age groups. The Crow–Siddiqui measure also shows very high kurtosis levels, indicating that the excess kurtosis is not driven by outliers.

An Alternative Measure of Persistent Changes. As we noted earlier, while the five-year income growth measure reveals a good deal about persistent changes in earnings, it still contains possible transitory innovations in years t and $t + 5$, which can potentially confound the inferences we draw about persistent changes. To check the robustness of our results, we consider an alternative measure that is based on the change between two consecutive five-year averages of earnings: $\bar{\Delta}_{\log}^5(\bar{y}_t^i) \equiv \log(\bar{Y}_{t+4}^i) - \log(\bar{Y}_{t-1}^i)$, where \bar{Y}_{t+4}^i is calculated the same as \bar{Y}_{t-1}^i but over the period t to $t + 4$. Averaging earnings before differencing purges transitory changes and better isolates the persistent ones.

Figure 8 plots the standardized moments of this alternative measure, which show essentially identical patterns to their counterparts using our baseline five-year growth measure

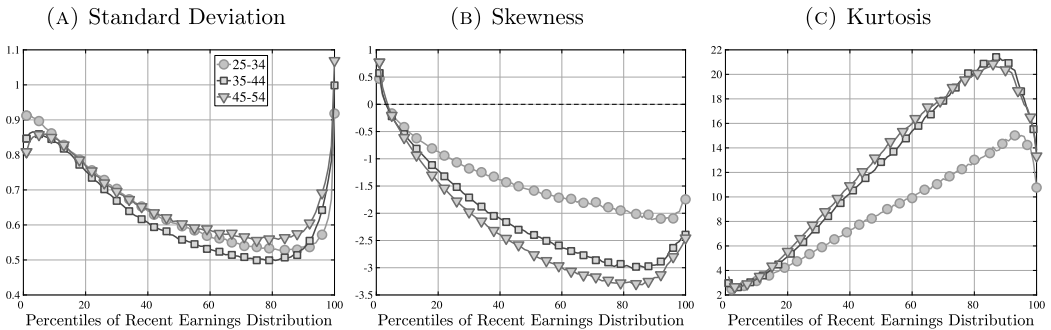


FIGURE 8.—Alternative measure of persistent changes, $\bar{\Delta}_{\log}^5(\bar{y}_t^i)$: standardized moments.

(see Appendix C.4 for quantile-based moments). In fact, if anything, this measure shows a slightly larger negative skewness and a higher excess kurtosis. These results confirm our conclusion that the nonnormalities are stronger in persistent earnings changes. A more formal analysis in Section 6 will further confirm this conclusion.

3.5. Job-Stayers and Job-Switchers

Economists have documented that the earnings changes of workers who stay with the same employer (job-stayers) are notably different from the changes of workers who switch jobs (job-switchers) (see Topel and Ward (1992), Low, Meghir, and Pistaferri (2010), and Bagger, Fontaine, Postel-Vinay, and Robin (2014)). This literature has focused on the average change, whereas we examine how the higher-order moments of earnings growth vary between job-stayers and job-switchers.

The SSA data set contains employer identification numbers (EINs) that allow us to match workers to firms. However, the annual frequency of the data, together with the fact that some workers hold multiple jobs in a given year, poses a challenge for a precise identification of job-stayers and job-switchers. We have explored several plausible definitions for stayers and switchers and found qualitatively similar results. Here, we describe one reasonable definition: A worker is said to be a job-stayer between years t and $t + 1$ if he has a W-2 form from the same firm in years $t - 1$ through $t + 2$, and that firm is the main employer by providing at least 80% 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.¹⁴

We show in Figure 9 how the quantile-based second to fourth moments of annual earnings growth for stayers and switchers vary with recent earnings. Relative to job-switchers, job-stayers experience earnings changes that have a smaller dispersion (about one-third for median-income workers), and are more leptokurtic, especially for low-RE workers. Changes are symmetric or slightly right skewed for stayers and left skewed for switchers. The age profiles are broadly similar across switchers and stayers, and figures for five-year changes and standardized moments display similar patterns (Appendix C.5).

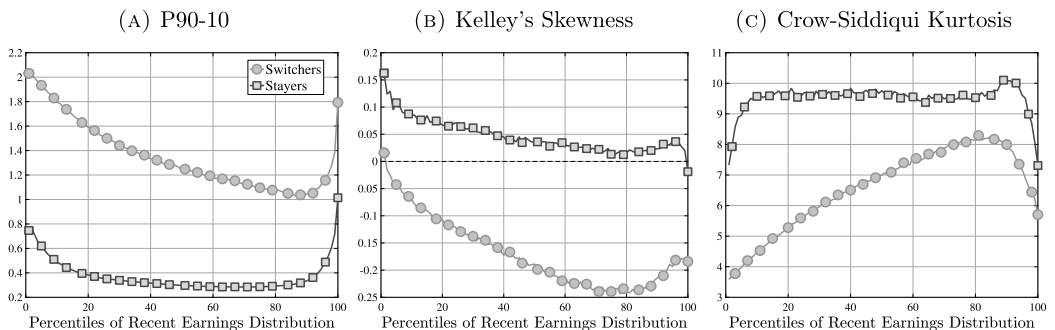


FIGURE 9.—Higher-order moments of earnings growth, $\Delta^1_{\log} y_i$: stayers vs. switchers.

¹⁴Note that this is a very conservative definition of a job-stayer. In Appendix C.5, we take the opposite approach, consider a conservative definition of a job-switcher, and reach similar conclusions.

3.6. *What Are the Sources of Nonnormalities in Earnings Growth?*

So far, our analysis has focused on the distribution of *annual earnings changes* and remained silent on what may be behind the nonnormalities. For example, are the left-skewness and excess kurtosis also present in the wage growth distribution? What are the lifecycle events associated with extreme income changes? The lack of information in the SSA data other than annual earnings does not allow us to investigate these questions, which, in turn, we study using the PSID.

3.6.1. *Separating Earnings, Wages, and Hours*

For many economic questions, it is important to know the extent to which nonnormalities in earnings dynamics are driven by wages versus hours. For example, if nonnormalities come from changes in hours and not wages, this would suggest focusing on hours to identify their underlying sources, for example, preferences for work or shocks to labor supply (health shocks, involuntary layoffs, etc.). If, instead, nonnormalities are also present in wage changes, that would point to a different set of factors on which to focus. To shed light on this question, we analyze the wage growth distribution in the PSID using a sample that closely mimics the SSA sample (see Appendix C.7 for the details).¹⁵

We start by investigating the non-Gaussian features of two-year earnings changes in the PSID (Table II). The standardized third moment and the Kelley measure point to a weakly left-skewed distribution, possibly due to added noise in the PSID to the extent that measurement error is symmetric. Excess kurtosis is a more striking feature: Both measures of kurtosis from the PSID are quite close to their SSA counterparts. The age patterns are also broadly in line with those from administrative data. In addition, De Nardi, Fella, and Pardo (2020) documented the income variation in higher-order moments of earnings growth from the PSID and found patterns similar to those in the SSA data.

Turning to hourly wage growth, negative skewness in the overall sample is even less pronounced than that of earnings. Unlike skewness, excess kurtosis of wage growth and its lifecycle variation are roughly similar to those features of earnings growth.¹⁶ This evidence

TABLE II
HIGHER-ORDER MOMENTS OF TWO-YEAR CHANGES IN THE PSID^a

	Normal	All		25–39		40–55	
		Earnings	Wages	Earnings	Wages	Earnings	Wages
Skewness	0.0	−0.26	−0.14	−0.17	−0.20	−0.34	−0.09
Kelley Skew.	0.0	−0.02	−0.02	0.03	0.016	−0.06	−0.04
Kurtosis	3.0	12.26	13.65	10.44	9.00	14.01	17.10
Crow Kurt.	2.91	6.83	5.59	6.33	5.02	7.33	6.11

^aNote: Wages are obtained by dividing annual earnings of male heads of households by their annual hours in the PSID using data over the period 1999–2013, during which data are biennial.

¹⁵Another approach is to estimate structural models of endogenous labor supply (such as those in Low, Meghir, and Pistaferri (2010) and Heathcote, Storesletten, and Violante (2010)).

¹⁶It is well known that there is significant measurement error in hours in the PSID, which leads to “division bias” when constructing hourly wages (see Bound, Brown, and Mathiowetz (2001) and Heathcote, Storesletten, and Violante (2010)). This measurement error attenuates left-skewness (if it is classical) and excess kurtosis (if it is Gaussian), implying that our estimates are lower bounds (see Halvorsen, Holter, Ozkan, and Storesletten (2020)).

TABLE III
IMPORTANT LIFECYCLE EVENTS AND EARNINGS CHANGES^a

Group $\Delta y \in$	Share %	Mean Δy	Mean Δw	Mean Δh	Δ wks not empl.	Occup. switch %	Employer switch %	Disab. Flow in %
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$(-\infty, -1)$	3.8%	-1.65	-1.01	-0.64	10.01	26.1	45.6	9.2
$[-1, -0.25)$	14.4%	-0.48	-0.34	-0.14	1.62	14.9	29.1	4.4
$[-0.25, 0)$	31.2%	-0.11	-0.08	-0.03	0.17	6.9	13.3	3.5
$[0, 0.25)$	31.1%	0.11	0.08	0.03	-0.03	5.3	9.7	2.8
$[0.25, 1)$	16.5%	0.47	0.34	0.13	-1.30	8.6	16.9	2.9
$(1, \infty)$	3.0%	1.64	1.06	0.58	-7.51	18.0	30.7	3.8

^aNotes: This table shows hours and wage growth (Δh and Δw , respectively) and the various lifecycle events for people in different biennial earnings change (Δy) groups. In column 5, “weeks not employed” is the sum of weeks unemployed and out of the labor force. Columns 6 and 7 show the fraction of workers that switch occupation and employer within each earnings change group, respectively. Column 8 shows the fraction of workers who become disabled in that period.

suggests that the leptokurtic property of earnings growth cannot be driven entirely by the hours margin. We also conducted an analogous analysis using data from the Current Population Survey (CPS), which has a larger sample size, and reached similar conclusions, specifically a weak left-skewness and strong excess kurtosis in earnings and wage growth (see Appendix C.7).

Motivated by the importance of extreme earnings changes for excess kurtosis, we investigate the roles of hours and wages in the tails of the earnings growth distribution. For this purpose, we distribute workers into six groups based on their two-year residual earnings change. As in the SSA data, most workers experience only small earnings changes (col. 1 of Table III). For each group, we compute the average change in residual earnings, hours, and wages (Table III, cols. 2–4).¹⁷ Our results show that wage changes are at least as important as hours changes. For example, the bottom group with the average earnings decline of 165 log points experiences a drop of 101 log points in wages. Clearly, extensive margin events (e.g., layoffs) can lead to large declines in hours and wages at the same time. Moreover, wage changes seem to be even more important for smaller earnings changes (e.g., more than 70% of $|\Delta y| < 0.25$ can be attributed to wages).

3.6.2. Linking Earnings Changes to Lifecycle Events

In this section, we link large earnings changes to various lifecycle events. We start with a natural suspect: nonemployment spells. The group with the largest earnings decline also reports the largest increase in the incidence of nonemployment—10 weeks (Table III, col. 5). Similarly, the group with the largest earnings increase reports the largest decline in nonemployment.¹⁸ These results underline the importance of the extensive margin for the tails of the earnings change distribution.

Next, we study occupation and job mobility, both of which are known to be associated with large changes in earnings. The likelihood of occupation and employer switches

¹⁷Log earnings, hours, and wages have been residualized by similar but separate regressions. Therefore, at the individual level, residual hours and wages do not exactly add up to residual earnings. However, as columns 2–4 of Table III show, the discrepancy is negligible.

¹⁸We reached similar conclusions analyzing unemployment and out of labor force separately. The incidence of unemployment is somewhat lower in the PSID than it is in the CPS: 6.8% of our sample reports some unemployment in the previous year, compared with around 10% in a similar CPS sample.

follows a distinct U-shaped pattern with earnings changes (Table III, cols. 6 and 7, respectively). Compared to the workers with small changes ($|\Delta y| < 0.25$), the top and bottom earnings-change groups are three to four times more likely to make these switches. The sources of mobility are possibly very different at the top and the bottom earnings-change groups. For example, the switches at the top are likely associated with promotions or outside offers, whereas moves at the bottom are probably necessitated by job losses. We also looked into involuntary geographic moves and found that they are associated with large earnings changes, too (Appendix C.7).

Finally, we investigate health shocks, which are known to have large effects on earnings (see Dobkin, Finkelstein, Kluender, and Notowidigdo (2018)). We focus on disabilities that affect individuals' work performance (see Appendix C.7 for a detailed description). We find higher transition rates into disability for workers with earnings declines, with the highest transition (9.2%) in the bottom earnings-change group (Table III, col. 8). These results suggest that the extreme earnings changes are not purely a statistical artifact or measurement error.

Disability Income. If health shocks are an important source of earnings changes, how important is disability insurance as a safety net? To answer this question, we add individuals' Social Security Disability Income (SSDI) from the SSA to their labor income and construct a "total income" measure. Our results in Appendix C.8.1 show that the cross-sectional moments of total income overlap with their labor income counterparts, mainly because the share of SSDI recipients is small, ranging from 1.3% in 1978 to 4.1% in 2013. However, SSDI makes a noticeable (albeit slight) difference only for the oldest group of workers, who constitute the majority of the recipients.

To sum up, in light of the vast micro literature that finds very small Frisch elasticities, large changes in earnings, especially declines, are much more likely to represent involuntary shocks beyond the worker's control such as health problems, reductions in hours imposed by employers, or unemployment.

Related Work. A growing literature uses administrative data from various countries to study the determinants of nonnormalities in earnings growth. Kurmann and McEntarfer (2019) showed that hourly wage growth displays high excess kurtosis for job-stayers in the United States, and wage changes constitute a substantial portion of the earnings changes, mostly for those experiencing increases. These findings are overall consistent with ours. They also argued that large declines in hours of job-stayers are involuntary and imposed by firms. Blass-Hoffmann and Malacrino (2019) used Italian data to argue that changes in weeks worked account for the procyclical left-skewness of the one-year and five-year earnings growth (first documented by Guvenen, Ozkan, and Song (2014) for the U.S.). In contrast, our analysis from the PSID also attributes an important role to wage growth. Moreover, we find that scarring effects are necessary to generate left-skewness through extensive margin fluctuations. Finally, Halvorsen et al. (2020) used Norwegian data and found that hourly wage changes exhibit left-skewness and excess kurtosis, and both the magnitudes and their lifecycle and income variation are similar to those for earnings changes. Furthermore, they showed that large earnings changes are mostly driven by wages for high-RE individuals, but the split between wages and hours is more equal for low-RE workers.

4. DYNAMICS OF EARNINGS

Having studied the distribution of earnings changes, we now turn to their persistence. Typically, earnings dynamics are modeled as an AR(1) or a low-order ARMA process,

and the persistence parameter is pinned down by the rate of decline of autocovariances with the lag order. While this linear approach might be a good first-order approximation, it imposes strong restrictions, such as the uniformity of mean reversion for positive and negative or large and small changes as well as for workers with different earnings levels.

We exploit our large sample and employ a nonparametric strategy to characterize the nonlinear mean reversion. We do so by documenting the impulse response functions of earnings changes of different sizes and signs for workers with different recent earnings. In particular, we group workers by their earnings growth between $t - 1$ and t , their recent earnings \bar{Y}_{t-1}^i , and age, and then follow their earnings over the next 10 years.

To reduce the number of graphs to a manageable level, we combine the first two age groups (ages 25 to 34) into “young workers” and the next three groups (ages 35 to 50) into “prime-age workers.” Within each age group, we rank and group individuals by \bar{Y}_{t-1}^i into the following 21 RE percentiles: 1–5, . . . , 91–95, 96–99, and 100. Next, within each age and RE group, we sort workers by the size of their log earnings change between $t - 1$ and t ($y_t^i - y_{t-1}^i$) into 20 equally sized quantiles. Hence, all individuals within a group have similar age and average earnings up to $t - 1$, and experience a similar change from $t - 1$ to t . For each such group of individuals, we then compute the *log change of their average earnings* from t to $t + k$, $\log \mathbb{E}[Y_{t+k}^i] - \log \mathbb{E}[Y_t^i]$, where Y_t^i is the income level net of age and time effects. Rather than taking the *average of log earnings change*, this approach allows us to include workers with earnings below the minimum threshold, thereby keeping the composition of workers constant for each k . The results for the alternative approach are qualitatively very similar and are available upon request.

4.1. Impulse Response Functions Conditional on Recent Earnings

In Figure 10, we show the mean reversion of different sizes of earnings changes $y_t^i - y_{t-1}^i$ for prime-age workers over a 10-year period. Specifically, we plot $\log \mathbb{E}[Y_{t+k}^i] - \log \mathbb{E}[Y_t^i]$ of each $y_t^i - y_{t-1}^i$ quantile on the y-axis against its average on the x-axis.¹⁹ This graphical construct contains the same information as a standard impulse response function but allows us to see the heterogeneous mean reversion patterns more clearly.

We start with the median-RE group ($\bar{Y}_{t-1}^i \in \text{P46–P55}$) in Figure 10(a). Even at the 10-year horizon, a nonnegligible fraction of the earnings change is still present for this group of workers, indicating a very persistent component in earnings growth. Also, negative changes tend to recover more gradually than positive ones for them. For example, workers whose earnings rise by 100 log points between $t - 1$ and t lose about 50% of this increase in the following 10 years. Almost all of this mean reversion happens after one year. Workers whose earnings fall by 100 log points recover 25% of that decline in the first year and around 50% of the total within 10 years. Finally, the degree of mean reversion varies with the magnitude of earnings changes, with stronger mean reversion for large changes: Small innovations (i.e., those less than 10 log points in absolute value) look very persistent, whereas larger earnings changes exhibit substantial mean reversion. A univariate autoregressive process with a single persistence parameter will fail to capture this behavior. In Section 6, we will show how to modify the simple income process to accommodate this variation in persistence by the size and sign of the earnings shock.

¹⁹The average mean reversion varies across the RE groups because of different earnings histories. Therefore, we normalize earnings changes on both the x- and y-axes such that their values at the median quantile of $y_t^i - y_{t-1}^i$ cross at zero.

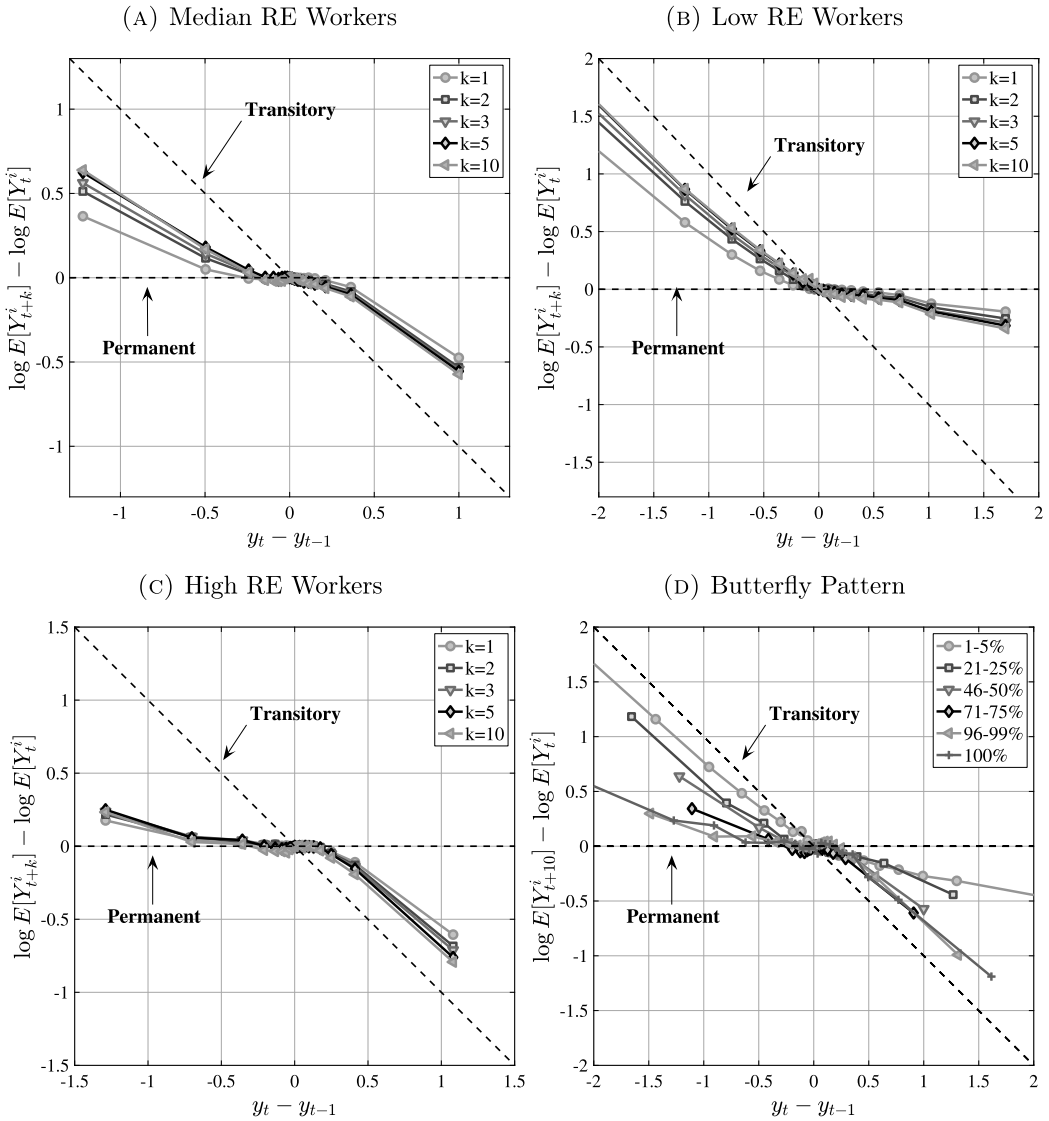


FIGURE 10.—Impulse responses, prime-age workers. Notes: Median-, low-, and high-RE in panels A, B, and C refer to workers with \bar{Y}_{t-1} in (P46–P55), (P6–P10), and (P91–P95), respectively. Prime age refers to age 35 to 50.

The analogous impulse response functions for low-income ($\bar{Y}_{t-1} \in P6-P10$) and high-income ($\bar{Y}_{t-1} \in P91-P95$) workers (Figures 10(b) and 10(c)) show that for low-income individuals, negative changes are more short-lived, whereas positive ones are more persistent, and that for high-income individuals the opposite is true.

Extending the results to the entire distribution of recent earnings, we focus on a fixed horizon and plot the cumulative mean reversion from t to $t + 10$ for the six RE groups in Figure 10(d). Starting from the lowest RE group (the bottom 5%), notice that negative changes are transitory, with an almost 75% mean reversion rate at the 10-year horizon. But positive changes are quite persistent, with only about a 25% 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 (top 1%), we have the opposite pattern: only 20 to 25% of earnings declines revert to the mean at the 10-year horizon, whereas around 80% of the increases do so at the same horizon. We refer to this shape as the “butterfly pattern.”

This butterfly pattern broadly resonates with the earnings dynamics in job ladder models. For high-RE workers—who are at the higher rungs of the ladder—a job loss leads to a more persistent earnings decline relative to low-RE workers because of search frictions. Similarly, for low-RE workers, large increases are likely due to unemployment-to-employment or job-to-job transitions, which have long-lasting effects on earnings.²⁰

5. EARNINGS GROWTH AND EMPLOYMENT: THE LONG VIEW

In this section, we turn to two questions that complete the picture of earnings dynamics over the life cycle. Both questions pertain to long-term outcomes—covering the entire working life. The first one is about average earnings growth—complementing the second to fourth moments analyzed in Section 3. In particular, how much cumulative earnings growth do individuals experience over their working life, and how does that vary across individuals with different lifetime incomes?

The second question investigates the lifetime nonemployment rate—defined as the fraction of an individual’s working life spent as full-year nonemployed. Although the incidence of *long-term* nonemployment is of great interest for many questions in economics, documenting it requires long panel data with no sample attrition, a phenomenon most common among long-term nonemployed. The administrative nature of the MEF data set and its long panel dimension provide an ideal opportunity to study this question.

5.1. Lifecycle Earnings Growth and Its Distribution

For the analysis in this section, we use the full length of the MEF panel, covering 1978 to 2013. We select individuals who (i) were born between 1951 and 1957 (hence for whom we have 33 years of data between ages 25 and 60), and (ii) had annual earnings above $Y_{\min,t}$ in at least 15 years, thereby excluding workers with very weak labor market attachment. We take a closer look at this latter group in the next subsection. We sort individuals into 100 percentiles by their lifetime earnings (LE), computed by averaging their earnings from age 25 through 60. For each LE percentile bin, denoted LE_j , $j = 1, 2, \dots, 99, 100$, we compute the growth rate between ages h_1 and h_2 by differencing the average earnings across all workers (including those with zero earnings) in those LE and age cells; that is, $\log(\bar{Y}_{h_2,j}) - \log(\bar{Y}_{h_1,j})$, where $\bar{Y}_{h,j} \equiv \mathbb{E}(\tilde{Y}_t^i | i \in LE_j, h(i, t) = h)$.

The results in Figure 11(a) show that between ages 25 and 55, the median worker (by LE) experiences a smaller earnings growth—about 60%—than a 150% mean growth estimated from a Deaton–Paxson pooled regression (see Appendix C.9). More importantly, higher-LE workers experience a much higher earnings growth over the life cycle compared with the rest of the distribution. While an upward slope per se is not surprising (as it is partly mechanical—faster growth will deliver higher LE, everything else held constant), the variation at the top end is so large and the curvature is so steep, that it turns

²⁰Lise (2012) showed that a job ladder model with a precautionary savings motive captures wealth and consumption dispersions better than the incomplete markets model with a linear-Gaussian income process.

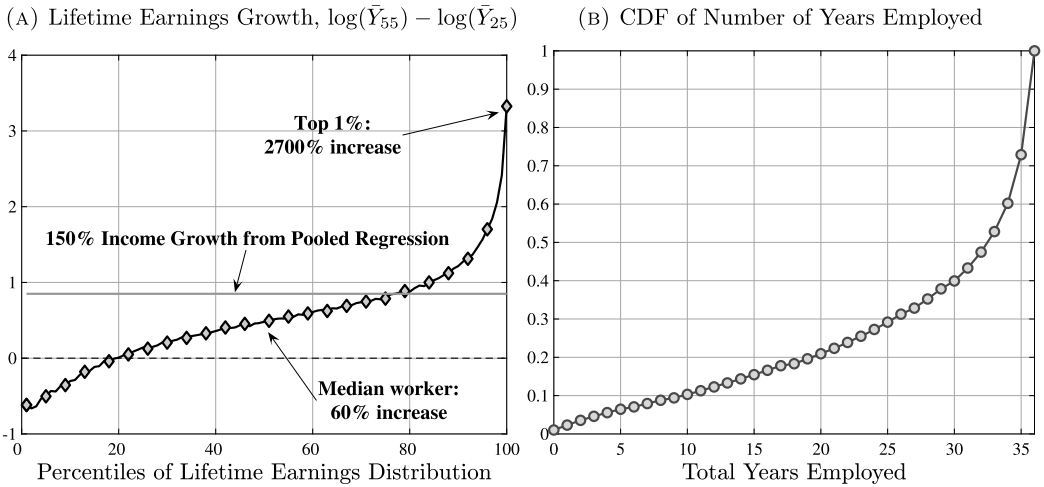


FIGURE 11.—Earnings growth and employment: the long view.

out to be difficult to capture using simple earnings processes, as we discuss in the next section. For example, average earnings grow by 1.5-fold (91 log pts) over 31 years at LE80, by 4.8-fold (157 log pts) at LE95, and by 27.9-fold (333 log pts) in the top 1%.²¹

One question is whether this extremely high growth rate at the top is driven by higher rates of school enrollment in these groups at age 25 (and thereby low earnings). While the lack of education data does not allow us to answer this question directly, several pieces of evidence are informative. First, about 21.7% of individuals in the LE100 have earnings below the Y_{\min} threshold at age 25, which is higher than the rate for half of the sample, suggesting schooling could be playing some role (see Figure C.38(a)). However, this rate drops quickly to 5.95% by age 30, which is one of the lowest in our sample. At the same time, earnings growth for this group between ages 25 and 30 is only slightly higher compared to that between ages 30 and 35 (2.9-fold vs. 2.6-fold), when schooling is unlikely to matter much. Similarly, looking at growth from 35 to 55, we still find a steep profile of earnings growth with respect to LE (see Figure C.37). These observations suggest that low labor supply at age 25 is not the major driver of these patterns.

Turning to the lower end, individuals below LE20 see their earnings *decline* from age 25 to 55. How important is disability for this decline? Adding SSDI to labor earnings has virtually no effect above the 40th LE group or so (Figure C.35). But it matters at the lower end, mitigating the decline by more than 50% for LE10 and LE5.

5.2. Lifetime Employment Rate and Its Distribution

Next, we investigate the lifetime nonemployment rates across individuals. Using the same criteria as before—working life defined as the period between ages 25 and 60, and full-year nonemployed defined as annual earnings below Y_{\min} —we examine the cumulative distribution of total lifetime years employed in Figure 11(b) (for further analysis, see Appendix C.10). The results show that, first, a large fraction of individuals are very

²¹Earnings decline from age 45 to 55 for 80% of the population, and those above LE80 experience only small increases (Figure C.37(b) in Appendix C.9). The decline in earnings later in life could be due to a decline in hours because of partial retirement (see Aaronson and French (2004)).

strongly attached to the labor market: 28% of individuals were never nonemployed during their working life, and almost half (48%) were nonemployed for less than three years. But second, the distribution has a long left tail, showing a surprisingly large fraction of men who spend half of their working life or more without employment: 18.3% of men spend 18 years—or half of their working life—as full-year nonemployed, and 12.3% spend at least 24 years as nonemployed.

To understand these magnitudes, note that the employment-to-population ratio for prime-age men in cross-sectional data (such as the one the Bureau of Labor Statistics publishes monthly) has averaged around 86% during this time period, implying a monthly nonemployment rate around 14%. Getting nonemployment at annual frequency for 24 years for 12% of the male population requires an extremely high persistence of the long-term nonemployment state. As we shall see in the next section, this statistic turns out to be very hard to match with a simple earnings process with standard parameter values.

6. ECONOMETRIC MODELS FOR EARNINGS DYNAMICS

The empirical facts documented so far can be viewed as snapshots of an earnings process taken from different angles. Each one allowed us to identify some key patterns by (partially) isolating other features. When these snapshots are combined, they provide a wealth of information that can be used to identify the underlying earnings process. In this section, we report the results of an extensive model specification search we conducted for (a class of) earnings processes that can reproduce the facts documented above.

Our specification search was guided by several considerations. First, rather than exploring entirely new frameworks for nonlinear and nonnormal dynamics, we start with a well-understood benchmark (a linear-Gaussian model) and incrementally add new, yet familiar, components to build toward richer specifications. At each incremental step, we conduct a battery of diagnostic tests to evaluate the potential of each new component for reproducing key features of the data. Second, there is the usual trade-off between the parsimony of a model and its goodness of fit. A key aspect of parsimony for the purposes of this paper is the computational burden of using an earnings process in a dynamic programming problem, including whether it requires an additional state variable or not. The benchmark process (chosen from among more than 100 models we estimated) requires only one state variable—the same as a standard persistent-plus-transitory model—while offering a good fit to the data, thereby achieving both goals.

6.1. A Flexible Stochastic Process

The models we estimate are special cases of the following general framework, which includes (i) an AR(1) process (z_t^i) with innovations drawn from a mixture of normals; (ii) a nonemployment shock whose incidence probability ($p_v^i(t, z_t)$) can vary with age or z_t or both, and whose duration (v_t^i) is exponentially distributed; (iii) a heterogeneous income profiles component (HIP); and (iv) an i.i.d. normal mixture transitory shock (ε_t^i):

$$\text{Level of earnings: } \tilde{Y}_t^i = (1 - v_t^i) e^{(g(t) + \alpha^i + \beta^i t + z_t^i + \varepsilon_t^i)}, \tag{2}$$

$$\text{Persistent component: } z_t^i = \rho z_{t-1}^i + \eta_t^i, \tag{3}$$

$$\text{Innovations to AR(1): } \eta_t^i \sim \begin{cases} \mathcal{N}(\mu_{\eta,1}, \sigma_{\eta,1}) & \text{with prob. } p_z, \\ \mathcal{N}(\mu_{\eta,2}, \sigma_{\eta,2}) & \text{with prob. } 1 - p_z, \end{cases} \tag{4}$$

Initial condition of z_t^i : $z_0^i \sim \mathcal{N}(0, \sigma_{z_0})$, (5)

Transitory shock: $\varepsilon_t^i \sim \begin{cases} \mathcal{N}(\mu_{\varepsilon,1}, \sigma_{\varepsilon,1}) & \text{with prob. } p_\varepsilon, \\ \mathcal{N}(\mu_{\varepsilon,2}, \sigma_{\varepsilon,2}) & \text{with prob. } 1 - p_\varepsilon, \end{cases}$ (6)

Nonemployment duration: $\nu_t^i \sim \begin{cases} 0 & \text{with prob. } 1 - p_\nu(t, z_t^i), \\ \min\{1, \exp(\lambda)\} & \text{with prob. } p_\nu(t, z_t^i), \end{cases}$ (7)

Prob of Nonemp. shock: $p_\nu^i(t, z_t) = \frac{e^{\xi_t^i}}{1 + e^{\xi_t^i}}$, where $\xi_t^i \equiv a + bt + cz_t^i + dz_t^i t$. (8)

In equation (2), $g(t)$ is a quadratic polynomial, where $t = (\text{age} - 24)/10$ is normalized age, that captures the lifecycle profile of earnings *common* to all individuals. The random vector (α^i, β^i) determines ex ante heterogeneity in the level and in the growth rate of earnings and is drawn from a multivariate normal distribution with zero mean and a covariance matrix to be estimated. The innovations, η_t^i , to the AR(1) component are drawn from a mixture of two normals. An individual draws a shock from $\mathcal{N}(\mu_{\eta,1}, \sigma_{\eta,1})$ with probability p_z and otherwise from $\mathcal{N}(\mu_{\eta,2}, \sigma_{\eta,2})$. Without loss of generality, we normalize η to have zero mean (i.e., $\mu_{\eta,1}p_z + \mu_{\eta,2}(1 - p_z) = 0$) and assume $\mu_{\eta,1} < 0$ for identification. Heterogeneity in the initial conditions of z_t is captured by $z_0^i \sim \mathcal{N}(0, \sigma_{z_0})$. Transitory shocks, ε_t^i , are also drawn from a mixture of two normals (eq. (6)), with analogous identifying assumptions (zero mean and $\mu_{\varepsilon,1} < 0$).

Our decision to use normal mixtures is motivated by two considerations. First, they provide a flexible way to model non-Gaussian shock distributions. In fact, by increasing the number of normals that are mixed, one can approximate almost any distribution (see McLachlan and Peel (2000)). Second, solving a dynamic programming problem with normal mixture shocks requires minimal adjustments to the computational methods commonly used with Gaussian shocks. This is appealing given our stated objectives.

The last component of the earnings process—and as it turns out, a critical one—is a nonemployment shock (eq. (7)) that is intended to primarily capture movements in the extensive margin. Specifically, a worker is hit with a nonemployment shock with probability p_ν whose duration $\nu_t > 0$ follows an exponential distribution with mean $1/\lambda$ and is truncated at 1 (corresponding to full-year nonemployment with zero annual income). This shock differs from z_t and ε_t by scaling the *level* of annual income—not its logarithm—which allows the process to capture the sizable fraction of workers who transition into and out of full-year nonemployment every year.²²

None of the components introduced so far depend explicitly on age or recent earnings, whereas variation along these dimensions is a key characteristic of the empirical patterns we saw. One promising way we found for introducing such variation was by making the nonemployment incidence p_ν depend on age t and z_t through the logistic function shown in equation (8).²³ The dependence of p_ν on z_t —which we refer to as “state dependence”—

²²It takes a -350 log point shock to z_t or ε_t for a worker earning $\$50,000$ to drop below Y_{\min} . Generating such large shocks with sufficiently high frequency to match worker exit rates makes it challenging to simultaneously match the high frequency of smaller shocks.

²³We have considered various alternative specifications such as making p_z (instead of p_ν) a function of age and recent earnings; modifying ξ_t^i by including quadratic terms in t or z_t or by introducing dependence on α^i ; and making the innovation variances or their mixture probabilities a function of t or z_{t-1} , among others. We discuss some of these specifications later.

turns out to be especially important as it *induces* persistence in nonemployment from one year to the next (despite ν_t itself being independent over time).

This completes the description of the benchmark process. We also estimated a 2-state process, which, as expected, fits the data better but increases the computational burden in a dynamic programming problem due to an extra state variable (see Appendix D.2).

Estimation Procedure

We set the model period to a year and estimate the earnings processes using the method of simulated moments (MSM), targeting seven sets of moments. The first six broadly correspond to the moments documented in Sections 3 to 5: (i) the standard deviation, skewness, and kurtosis of one-year and (ii) five-year earnings growth; (iii) impulse response moments over short-term (at one-, two-, and three-year) horizons and (iv) long-term (at five- and ten-year) horizons; (v) average earnings levels of each LE group over the life cycle (essentially a more detailed version of Figure 11(a)); and (vi) the cumulative distribution of nonemployment (Figure 11(b)). In addition, the age profile of the within-cohort variance of log earnings (Figure D.3) is a key moment that has been extensively studied in previous research. For both completeness and consistency with earlier work, we include these variances as a seventh set of moments. See Appendix D.1 for the full list of moments and their details.²⁴

The MSM objective function is the (weighted) sum of squared arc-percent deviations between the data and simulated moments. Using arc-percent (as opposed to level) deviations is a natural way to deal with the (large) differences between the scales of different moments—it is essentially reweighting the moments to prevent those with large absolute levels to mechanically receive larger weights. It is also preferable to using percentage deviations because it is more well-behaved when data moments are close to zero. Finally, in terms of weighting, we assign equal weight (1/7) to each of the seven sets of moments and weigh moments within each set equally.²⁵

The resulting high-dimensional objective function has a challenging geometry, presenting a difficult global optimization problem. We use an efficient and parallelizable global search algorithm that makes the estimation feasible on large parallel clusters, although it still requires days or weeks.²⁶ Appendix D.3 presents the details of the estimation.

²⁴The targeted moments differ from those in the descriptive analysis in two ways. First, some additional statistics that are discussed in the descriptive analysis but not reported above to save space are targeted in estimation (e.g., the moments of one-year changes). Second, while the log growth measure is convenient for exposition, it requires dropping very low earnings observations and results in a loss of valuable information along the extensive margin. Thus, in the estimation, we target the arc-percent analogs of these moments (shown in Appendix C.2). Using the arc-percent measure also allows us to target the dynamic effects of nonemployment as part of the impulse response moments.

²⁵Although we found similar estimates with the equal weighting (identity) matrix (for the benchmark process), we prefer our approach because the somewhat arbitrary number of moments within each set does not artificially affect its importance. For example, instead of targeting the moments of one- and five-year growth rates, we could also add two-, three-, and four-year growth rates to the list of targets. These additional moments are highly correlated with the original ones, so they do not provide substantially more information, yet adding them would increase the number of moments in the set (and the weight received) by 2.5 times. Our approach avoids this outcome. We have also experimented with other weighting matrices; for example, in earlier versions of this paper, we assigned weights based on our subjective views about the importance of each moment in typical incomplete markets applications and found similar estimates. We have chosen not to use the optimal weighting matrix because our moments are very precisely estimated thanks to the sheer size of our data set; therefore, efficiency is not a concern.

²⁶The algorithm is a version of the TikTak algorithm described in Arnaud, Guvenen, and Kleineberg (2019). The latest version of this algorithm incorporates improvements we made while using it for this paper.

6.2. Results: Estimates of Stochastic Processes

We now present the estimation results for six different specifications (Table IV). We start from the canonical linear-Gaussian model and add new features step by step until we reach our preferred benchmark process. We discuss along the way which aspect(s) of the data each feature helps capture. Figure 12 plots the fit of each model to the six sets of moments targeted in the estimation. We also show the fit to selected impulse response functions separately in Figure 13.²⁷ We exclude two specifications from these figures for readability and show them in Figures D.12 and D.13 in Appendix D.5 of the Supplemental Material.

In *Model (1)*, we start with a simple but widely used linear-Gaussian specification: the sum of an individual fixed effect, an AR(1) process, and an i.i.d. transitory shock, all drawn from Gaussian distributions (i.e., $\sigma_\beta = 0$, $p_z = 1$, $p_\nu = 0$, and $p_\varepsilon = 1$ in equations (2) to (8)). The estimates of key parameters are unusually large—the standard deviations of the fixed effect and the transitory shock ($\sigma_\alpha = 1.18$ and $\sigma_\varepsilon = 0.70$) are several times larger than the typical estimates in the literature (cf., Storesletten, Telmer, and Yaron (2004) and Heathcote, Storesletten, and Violante (2010)). The persistence parameter is slightly above 1 ($\rho = 1.005$), implying a nonstationary process, with the effects of shocks being amplified over time.

The fact that these parameter values are quite different from previous estimates should not come as a surprise, since most of the moments targeted here have not been used in previous analyses. That said, it turns out that one set of moments is responsible for most of these differences—the CDF of lifetime employment rates, which shows that nonemployment is an extremely persistent state for a significant fraction of men.²⁸ To match this large fraction of persistently nonemployed individuals, the estimation chooses a wide dispersion of fixed effects, placing many individuals closer to the minimum threshold. Combined with the large transitory and nonstationary persistent shocks, the model manages to match the lifetime (non)employment distribution very well (Figure 12(e)).

However, the model fails in most of the other dimensions. First, it generates essentially zero skewness and no excess kurtosis (Figures 12(b) and 12(c)), which is not surprising given the Gaussian structure. Second, it vastly overstates lifecycle income growth—for example, implying a 3.1-fold rise at the median compared with only a 60% rise in the data (Figure 12(d)). Finally, it overshoots both the level of inequality and its rise over the life cycle (Figure 12(f)).²⁹ Overall, this process does not offer a good fit to the data.

Introducing Nonemployment Shocks

We first introduce nonemployment shocks (ν_t) to Model (1) to improve the fit by allowing the model to match the employment CDF more easily and leaving more flexibility for matching other moments. Another motivation is to investigate a common conjecture that the negative skewness could be entirely due to unemployment (disaster) shocks. Finally,

²⁷Given the very large number of moments used in the estimation, it is not feasible to show and discuss all of them here. We show the fit of models to the remaining moments in Appendix D.5.

²⁸Reestimating the same process without targeting the employment CDF yields $\sigma_\alpha = 0.42$, $\rho = 1.0$, $\sigma_\eta = 0.155$, and $\sigma_\varepsilon = 0.59$. These estimates are similar to those in previous studies cited above, with the exception of σ_ε being on the high side.

²⁹Notice that the variance of log earnings at age 25 (85 log points) is less than half of what would be predicted from the estimated parameters ($\sigma_\alpha^2 + \sigma_\varepsilon^2 + \sigma_{z,0}^2 + \sigma_{\eta,1}^2 = 202$ log points). This difference arises because the variance is computed excluding workers below the minimum threshold.

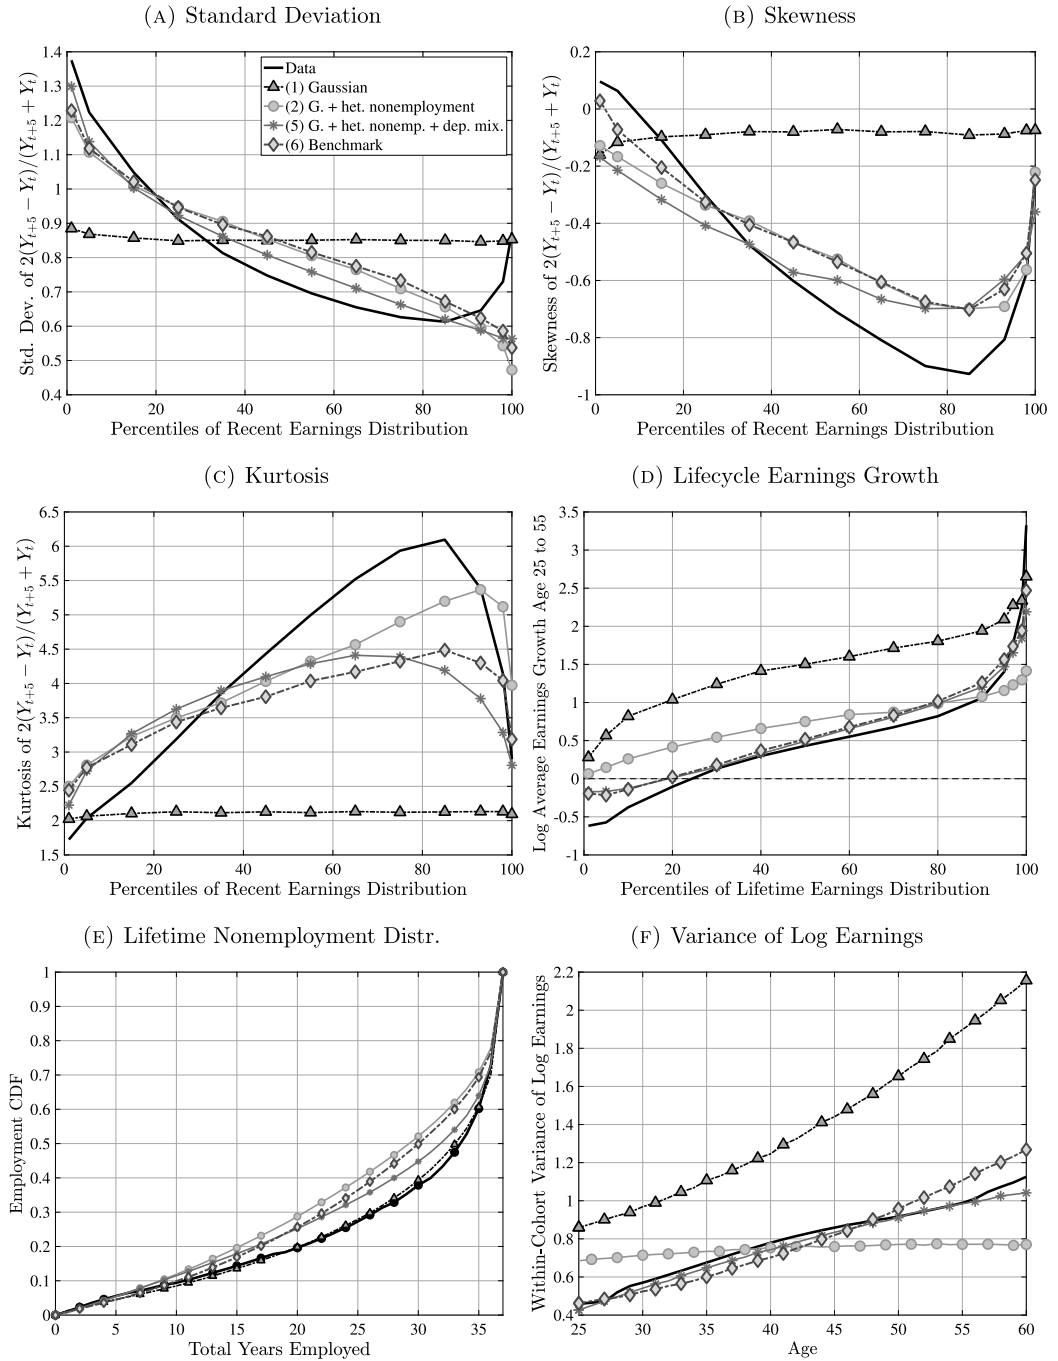


FIGURE 12.—Estimated model versus data: key moments.

unemployment shocks are a common feature in quantitative models, so it is instructive to understand their contribution to the fit to the data.

Heterogeneity and State-Dependence in Nonemployment Risk

In *Model (2)*, we allow p_ν to depend on age and the persistent component z_t (equation (8)).³⁰ The estimated value of $\lambda = 0.03$ implies that 97% of nonemployment spells last for the entire year, so ν is best thought of as a full-year nonemployment shock. Importantly, there is striking heterogeneity in nonemployment probability: It varies from 18.2% for the bottom 10% RE workers, to 5.8% for the median, and further down to 0.8% for the top 10% (when averaged across age groups).³¹ Put differently, those in the bottom decile experience a full-year nonemployment spell about every 5 years, whereas the majority of top decile workers do not experience this at all (i.e., every 125 years). Hence, *Model (2)* captures the extreme concentration of nonemployment at the bottom of the income distribution. Further, nonemployment risk falls with age, from 7.9% for the first age group to 6.1% for the last age group. Finally, with these state-dependent nonemployment shocks, the model fits the data with a smaller dispersion of fixed effects ($\sigma_\alpha = 0.655$) and transitory shocks ($\sigma_\epsilon = 0.163$), and with a lower persistence ($\rho = 0.97$).

Overall, *Model (2)* provides a much better fit to the data than *Model (1)*. The objective value falls by half, from 74.9 in *Model (1)* to 35.7, with improved fit to all sets of moments, with the exception of the nonemployment CDF. The most obvious improvement is in the variation in cross-sectional moments by RE levels, which does a fairly good job of capturing the general patterns in the data (first three panels of Figure 12).³² Although this result may seem a bit obvious—given the explicit age and income variation through p_ν —it is notable for several reasons. First, during the specification search, we have experimented with various other formulations to introduce such heterogeneity (e.g., in the means or variances of innovations in normal mixtures) and did not find them to perform nearly as well. (*Model (4)* is a specific example of this.)

Second, *Model (2)* not only generates more realistic heterogeneity in skewness and kurtosis; it also captures their *levels* better even relative to a specification with uniform nonemployment risk. While not reported in Table IV, we estimated a simpler model with *uniform nonemployment risk* across workers, that is, $b, c, d \equiv 0$ in equation (8) (see Table D.IV). This model manages to generate some excess kurtosis but very little negative skewness. To understand why this is the case, note that nonemployment is basically a large but *fully* transitory shock when the risk is uniform: Every worker whose income goes down due to nonemployment in the current period bounces back to his previous income level (on average) in the next period. Consequently, uniform nonemployment stretches both tails of the income change distribution, leaving its symmetry largely unaffected but generating some excess kurtosis in annual earnings growth (less so over longer horizons). As a result, the fit improves marginally (73.4 versus 74.9 before), reflecting improvements

³⁰We have also considered an alternative specification where p_ν depends on α , z_t , age, and their interaction terms. We found modest improvement in the objective value. This result is not surprising because more than 80% of ex ante heterogeneity is captured by differences in the initial conditions of the (highly) persistent component, z_0 , which can capture ex ante heterogeneity in nonemployment risk.

³¹Two recent papers attribute an important role to heterogeneity in unemployment risk. Jarosch (2021) showed heterogeneity in job security across firms to be important in accounting for the scarring effects of unemployment. Karahan, Ozkan, and Song (2019) found ex ante differences in unemployment risk to be important drivers of lifetime earnings differences among the bottom half of the distribution.

³²Because the age variation is smaller than the RE variation in the data, we omit the age patterns here and report them in Figures D.7 and D.8. The fit improves along this dimension as well.

TABLE IV
ESTIMATES OF STOCHASTIC PROCESS PARAMETERS^a

<i>Model:</i>	(1)	(2)	(3)	(4)	(5)	(6)	
	<i>Gaussian process</i>					<i>Benchmark process</i>	
						Parameters	Std. Err.
<i>AR(1) Component</i>	G	G	mix	mix	mix	mix	mix
\hookrightarrow <i>Prob. age/income</i>	—	—	no/no	yes/yes	no/no	no/no	no/no
<i>Nonemployment shocks</i>	no	yes	no	no	yes	yes	yes
\hookrightarrow <i>Prob. age/inc.</i>	—	yes/yes	—	—	yes/yes	yes/yes	yes/yes
<i>Transitory shocks</i>	G	G	mix	mix	mix	mix	mix
<i>HIP</i>	no	no	no	no	no	yes	yes
<i>Parameters</i>							
ρ	1.005	0.967	1.010	0.992	0.991	0.959	0.0001
p_z			5.0%	— [†]	17.6%	40.7%	0.0005
$\mu_{\eta,1}$			−1.0*	−1.0*	−0.524	−0.085	0.0006
$\sigma_{\eta,1}$	0.134	0.197	1.421	1.070	0.113	0.364	0.0004
$\sigma_{\eta,2}$			0.010	0.032	0.046	0.069	0.0002
σ_{z_0}	0.343	0.563	0.213	0.446	0.450	0.714	0.0005
λ		0.030			0.016	0.0001	0.0003
p_e			11.8%	8.8%	4.4%	13.0%	0.0004
$\mu_{e,1}$			−0.826	0.311	0.134	0.271	0.0009
$\sigma_{e,1}$	0.696	0.163	1.549	0.795	0.762	0.285	0.0006
$\sigma_{e,2}$			0.020	0.020	0.055	0.037	0.0003
σ_α	1.182	0.655	0.273	0.473	0.472	0.300	0.0009
σ_β						0.196	0.0002
$\text{corr}_{\alpha\beta}$						0.768	0.0015
<i>Objective value</i>	74.87	35.69	56.66	40.97	27.22	22.64	
<i>Decomposition:</i>							
(i) Standard deviation	9.48	6.28	7.66	6.11	5.65	5.77	
(ii) Skewness	43.03	14.12	23.75	15.02	14.12	9.96	
(iii) Kurtosis	26.53	5.90	12.95	9.22	5.83	5.81	
(iv) Impulse resp. short	18.35	13.51	20.04	16.70	9.85	8.65	
(v) Impulse resp. long	28.34	22.87	32.86	27.00	16.27	12.12	
(vi) Lifetime inc. growth	37.12	10.33	24.21	11.82	7.96	6.89	
(vii) Age-ineq. profile	20.70	8.24	16.42	12.78	1.52	3.32	
(viii) Nonempl. CDF	3.63	10.97	9.05	4.42	6.95	8.13	
<i>Model Selection p-val.</i>							
Test 1	0.000	0.000	0.000	0.000	0.000	0.000	
Test 2	0.000	0.000	0.000	0.000	0.000	—	

^aNotes: The top panel provides a summary of the features of each specification, the middle panel shows the estimated values of key parameters (the rest are reported in Table D.III), and the bottom panel reports the weighted percentage deviation between the data and simulated moments for each set of moments (the total objective value is the square root of the sum of the squares of objective values of each component) as well as the p-values for model selection. The *'s indicate that in columns 3 and 4, the value of $\mu_{\eta,1}$ is constrained by the lower bound we impose in the estimation. [†]: p_z is not a number but a function in this specification and reported in Table D.III. The standard errors of parameter estimates (using a parametric bootstrap with 100 repetitions) are extremely small, thanks to the very large sample size. Hence, we do not report them except for the benchmark process.

mainly for kurtosis. Thus, uniform nonemployment risk—the way nonemployment is commonly modeled—cannot generate the significant left-skewness and very high excess kurtosis in earnings growth, especially in the persistent component of earnings.

Two additional features are worth noting. First, the fit to the lifecycle earnings growth moments improves significantly (the objective falls from 37.12 in Model (1) to 10.33 here). Second, in contrast to Model (1), which substantially overstated the rise in inequality with

age, Model (2) generates a very flat age-variance profile (Figure 12(f)). Although, strictly speaking, this is a failure of the model, it actually represents an important step in the right direction: The model manages to match the large dispersion of earnings changes without overstating the rise in the variance of log earnings with age, which has proved challenging for the workhorse linear-Gaussian model, as we have seen for Model (1) (see also Heathcote, Perri, and Violante (2010) and Daly, Hryshko, and Manovskii (2016) for further discussion).

Model (2) yields these improvements by attributing a sizable fraction of earnings volatility to state-dependent nonemployment shocks, which generate systematic (rather than uniform) nonemployment risk. Furthermore, these shocks (through their dependence on z_t) have persistent but nonpermanent effects, thereby facilitating a better fit to impulse response moments as well.

Introducing Normal Mixtures

We next investigate the potential of modeling η_t and ε_t as normal mixtures. To isolate their effects, we shut down nonemployment risk for the time being ($p_v \equiv 0$). Model (3) allows a mixture in persistent and transitory shocks. As before, we begin by restricting the mixture probabilities to be the same for all individuals.³³

The fit improves appreciably relative to Model (1), with the objective value falling from 74.9 to 56.7 in Model (3).³⁴ The improvements are largest for skewness, kurtosis, and lifetime income growth moments.³⁵ The normal mixtures have very similar features in both models and for both η_t and ε_t . One of the two normals looks like a rare disaster shock: It has a low probability (5%–6% for η_t and 12% for ε_t), large negative mean (as low as -1 , the lower bound), and very large standard deviation (in excess of 140 log points). The other normal represents the “typical” earnings change, with a very high probability, slightly positive mean, and very small standard deviation. Put together, the estimated normal mixtures correspond to a world where, in most years, workers experience only small shocks to their income but, every once in a while, they are hit with a potentially very large negative shock. Notice that this description is not too far from the nonemployment shocks.

Not surprisingly, with constant mixture probabilities, Model (3) does not capture the age and income variation in higher-order moments and also overstates the slope of the age-variance profile more than any other model (Figure 12). Thus, in Model (4), we allow p_z to vary with age and z_t with the same functional form used before for p_v (eq. (8)) but replace ξ_t^i with ξ_{t-1}^i (to avoid circularity). Model (4) generates the second best objective value (41.0) among the models so far, only behind Model (2). A large part of this improvement comes from a better match to the age and income variation in cross-sectional moments, with the exception of kurtosis for top earners (Figure D.12).

³³Nonemployment shocks can be mimicked by the mixture of normals if one of them has a very large negative mean and a small variance. To distinguish the two modeling tools from each other, we impose a lower bound of -1 (or -63% mean shock) on the smaller mean of both mixture distributions ($\mu_{\eta,1}$ and $\mu_{\varepsilon,1}$). We also estimated with a bound of -5 (or -99.3%). The objective value for the latter is significantly lower but the fit is still worse than that in Model (2) (Table D.V).

³⁴We have also estimated a specification which has a mixture only in the persistent component that yields an objective value of 62.1 (see Table D.IV for parameter estimates). Thus, the data demand non-Gaussian features in persistent innovations more than in transitory ones.

³⁵We have also estimated a specification that adds uniform nonemployment risk to Model (3) and found its probability to be almost zero (0.01%) (Table D.V).

The most natural specification to which to compare Model (4) is Model (2), as they both feature age and income variation in the shock probability. That said, Model (2) is more parsimonious without the mixture in ε and fewer parameters for the exponential shock compared with a normal mixture. Yet, it provides a better fit for six of the seven sets of moments, with the employment CDF being the only exception.

In *Model (5)*, we combine the two most promising features we found so far: the normal mixture specifications for η_t and ε_t from Model (3) with the state-dependent nonemployment risk from Model (2).³⁶ The resulting model improves the fit quite a bit, with an objective value of 27.2 compared with 41.0 for Model (4) and 35.7 for Model (2). Moreover, with nonemployment shocks capturing the negative tail shocks, the normal mixture for η_t no longer features a rare disaster shock: The negative mean is no longer at the lower bound, its probability is higher (17.6%), and its standard deviation is much smaller (0.11 versus between 1.07 and 1.58 in previous models).

Compared with Model (2), the improvements are mostly in the impulse response moments (Figure 13), the age-variance profile (showing by far the best fit of any model so far), and the nonemployment CDF. Hence, while the cross-sectional moments can be generated to a large extent with either heterogeneous nonemployment risk or a mixture of normals, the nonlinear earnings persistence is better captured by combining the two features. In particular, the moderate mean reversion following extreme earnings changes (Figure 10) cannot be explained by fully transitory or permanent shocks. Nonemployment shocks are better able to generate this pattern with their long-lasting but nonpermanent effects.

Introducing Heterogeneous Income Profiles

Last, the most general specification, *Model (6)*, adds a HIP component to Model (5), further improving the objective value to 22.6 from 27.2. The fit improves for the skewness and long-term impulse response moments; is comparable for the standard deviation, kurtosis, and lifetime income growth moments, and the age-variance profile; and is slightly worse for the age-inequality profile and nonemployment CDF.³⁷ The better fit to impulse response moments mainly arises from a less persistent z_t ($\rho = 0.959$ vs. 0.991), with the half-life of η_t shocks declining to 17 years from 77 years in Model (5).³⁸

Overall, we believe the benchmark process offers a reasonable trade-off between a good fit to the data and the need for a parsimonious process that can be implemented in models without increasing the computational complexity.³⁹ Nonetheless, one could make an argument in favor of Model (5) as well. An even simpler process with a reasonable fit is

³⁶We found that the alternative—in which p_v is constant and p_z is age and income dependent—was outperformed by Model (5). Further, adding heterogeneous p_z to Model (5) improves the fit very little.

³⁷The better fit is not visible in Figure 13 as it shows a small number of impulse response moments.

³⁸The fact that adding HIP to a process lowers the estimated ρ is well understood (Guenen (2009)). Basically, the age profile of the variance of earnings is close to linear, which can be generated by an AR(1) process only if $\rho = 1$. HIP provides flexibility to generate the linear age profile, allowing ρ to fall to values more consistent with the mean reversion patterns.

³⁹We have also considered introducing individual-specific heterogeneity in innovation variances, a feature that Browning, Ejrnaes, and Alvarez (2010) found important in their estimation. While allowing for this heterogeneity improved the fit in simpler specifications we estimated, adding it to the richer specifications made a very small difference to the fit of the models. This suggests that a significant part of the heterogeneity in earnings volatility across individuals we see in the data might be attributable to variation by characteristics such as age and income history.

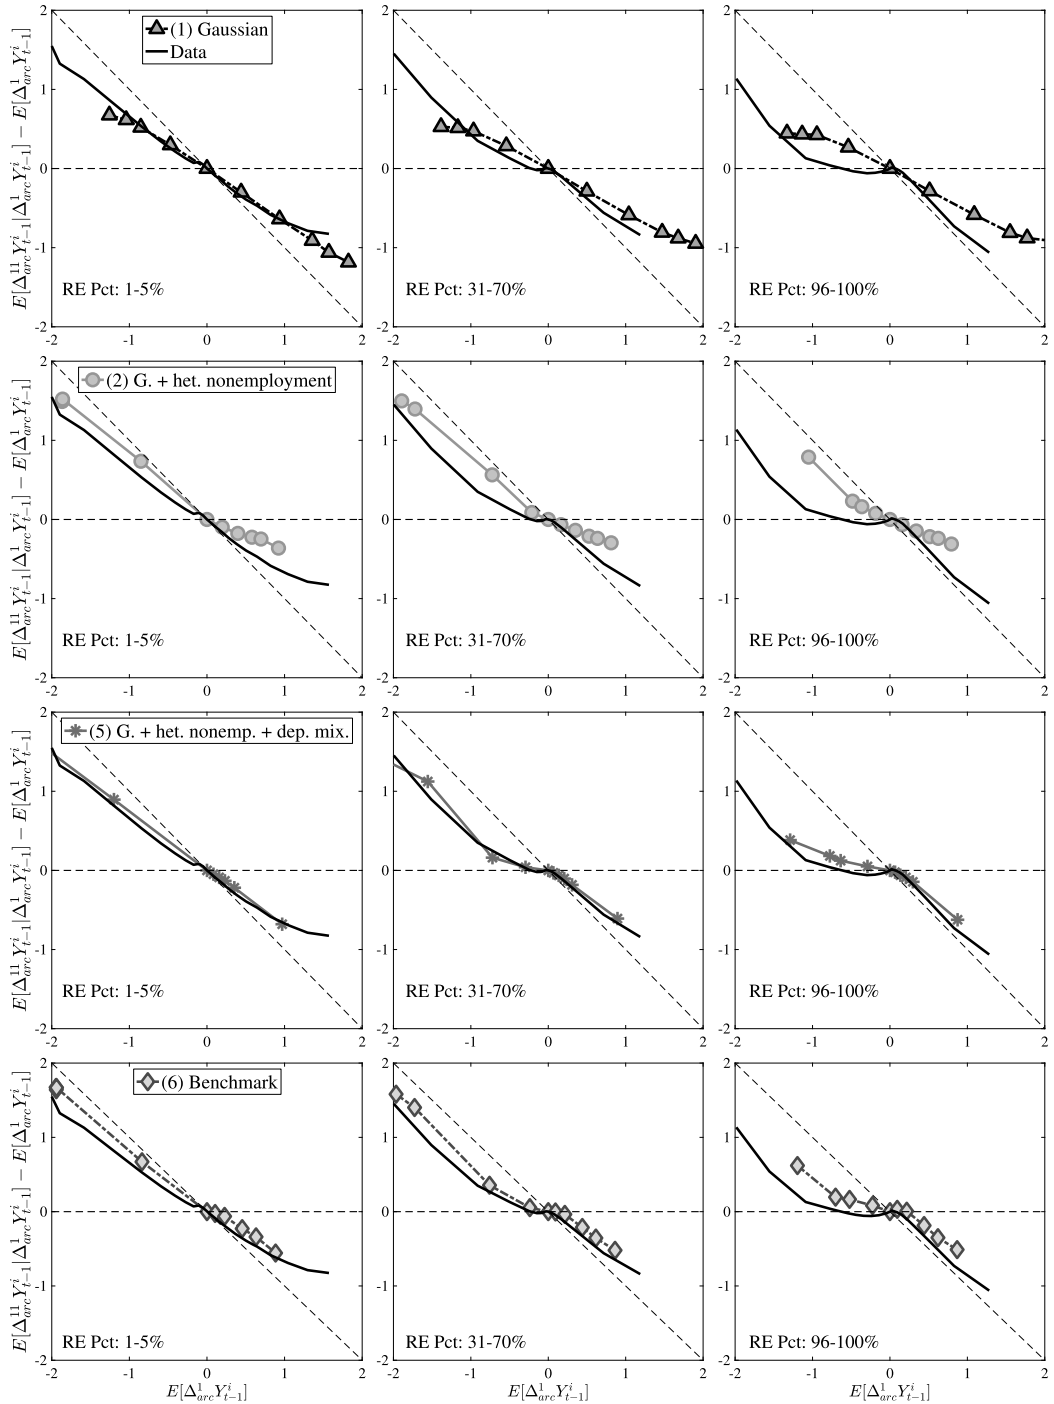


FIGURE 13.—Estimated model versus data: selected impulse response moments. Notes: In the estimation, we target arc-percent earnings growth between $t + k$ and $t - 1$, which allows us to keep the composition of workers constant for each k (for details, see Appendix D.1). To keep concepts analogous to what is shown in Section 4 (income change between t and $t + k$), we plot $\mathbb{E}[\Delta^{k+1}_{arc} Y^i_{t-1} | \Delta^1_{arc} Y^i_{t-1}] - \mathbb{E}[\Delta^1_{arc} Y^i_{t-1}]$ for $k = 10$. The solid line in each panel shows the data counterparty.

Model (2), which does not have the normal mixture in η_t . But the computational burden introduced by a normal mixture is often very small, so in our subjective opinion, the trade-off present in Model (2) is not likely to be worth it in many applications.

For completeness, we have also carried out a procedure for model selection. Testing different specifications and selecting one of them requires knowing the distribution of the objective value. These tests are in general carried out using the J-statistic, which has a chi-squared distribution if the efficient weighting matrix is employed. The J-test cannot be used in our setting as our weighting matrix is not the efficient one. Nevertheless, we propose two tests (see Appendix D.4 for further details). The first one tests the null hypothesis that a given specification is the true data-generating process. We apply this test to the six specifications and reject all of them (bottom panel of Table IV).

The second one tests the null hypothesis that a given specification provides as good a fit statistically as the 1-state benchmark. This test utilizes the distribution of the objective value for Model (6), which we construct through Monte Carlo simulations. We reject this hypothesis for the first five specifications in Table IV and conclude that the improved fit of the benchmark specification is also statistically significant.

6.3. Parameter Estimates of the Benchmark Process

We now turn to the parameter estimates from the benchmark process and its fit to some untargeted moments. Starting with the AR(1) process, the persistent shock is drawn about every 2.5 years ($p_z = 40.7\%$) from an “unfavorable” distribution—with a negative mean and large standard deviation ($\mu_{\eta,1} = -0.085$ and $\sigma_{\eta,1} = 0.364$)—and in the other years from a “favorable” one—with a positive mean and small standard deviation ($\mu_{\eta,2} = 0.058$, $\sigma_{\eta,2} = 0.069$). This mixture of normals implies that innovations to the persistent component are both strongly left skewed (skewness of -0.87) and leptokurtic (kurtosis around 6.3). In contrast, transitory shocks (ε_t) are typically smaller: In most years (with 87% probability), they are drawn from a tight distribution, $\mathcal{N}(-0.041, 0.037^2)$, and every eight years or so ($p_\varepsilon = 13.0\%$) from a distribution with large positive mean and dispersion, $\mathcal{N}(0.271, 0.285^2)$. Consequently, transitory shocks feature a skewness of 3.2 and a kurtosis of 15.4. However, η_t and ε_t are not the only sources of higher-order moments; workers also face nonemployment risk, which allows the model to generate a leptokurtic density for arc-percent changes with spikes at both ends (Figure 14(a)). The state dependence in nonemployment risk also leads to age- and income-varying skewness and kurtosis in *persistent* earnings changes. Thus, as expected, Models (3) and (4) (without nonemployment risk) find persistent shocks to be more strongly leptokurtic and left skewed than Models (5) and (6).⁴⁰ We conclude that persistent innovations are key drivers of non-Gaussian features in the data.

The initial heterogeneity in earnings is captured by (i) the permanent fixed effect σ_α (since $\beta^i \times 0 = 0$), and (ii) σ_{z_0} , the initial dispersion in z_t , whose effect declines at rate ρ . Since the latter is twice as large as the former ($\sigma_{z_0} = 0.71$ versus $\sigma_\alpha = 0.30$), a large part of earnings inequality at age 25 reflects persistent but not permanent differences. Finally, our estimate $\sigma_\beta \simeq 2\%$ is close to earlier estimates from the PSID (e.g., Haider (2001) and Guvenen and Smith (2014)), despite using different data sets and moments.

⁴⁰In fact, in Model (3), in which the only sources of nonnormalities are the persistent and transitory innovations, we find persistent innovations to be more left skewed and leptokurtic than transitory ones.

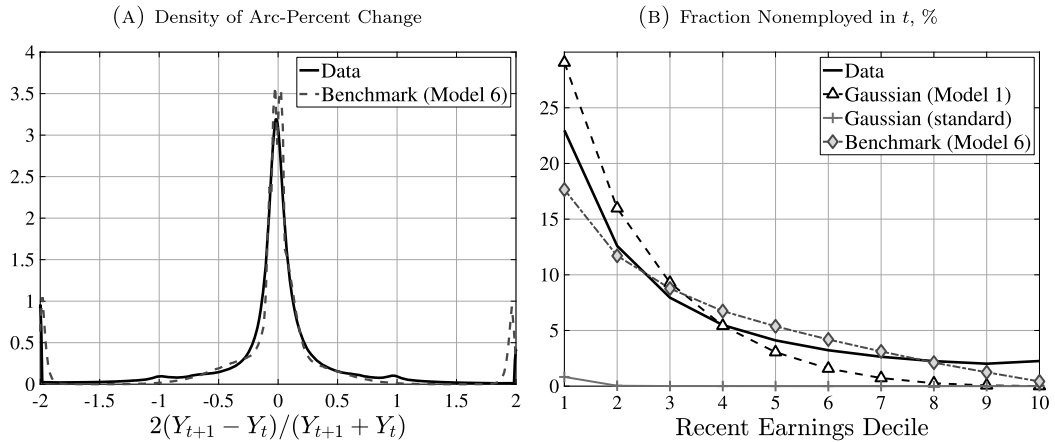


FIGURE 14.—Model fit: nontargeted statistics. Notes: The parameters of the Gaussian process come from Model 1 in Table IV, whereas Gaussian (standard) is the same process estimated without targeting the employment CDF (see footnote 28 for its parameter estimates). The data series on Panel (B) is conditional on past 2 years’ income (in $t - 1$ and $t - 2$).

State-Dependent Nonemployment Shocks. Almost all workers hit by nonemployment shocks experience *full-year* nonemployment ($\lambda = 0.0001$).⁴¹ How often are these shocks realized? Similarly to Model (2), our benchmark process also displays striking heterogeneity in nonemployment risk. The probability function p_v (equation (8)) is hard to interpret on its own (see Figure D.14 for a 3D graph), so instead, we investigate this probability for various age and RE quantiles for workers who satisfy the conditions of the RE sample: Nonemployment risk declines modestly over a working life, from 6.9% over 25–34 to 6.1% in the next 10 years, and to 5.5% over 45–54. There is less variation in the data, with the corresponding figures being 6.8%, 6.2%, and 6.6%, respectively.

Differences in nonemployment risk between income groups are much more pronounced (Figure 14(b)). In the bottom RE decile, workers experience nonemployment almost once every six years (with 17.6% probability). This probability declines sharply to 5.4% for the median quintile and to 0.4% for the top decile. The fit to the data is fairly good, though the RE variation is slightly less pronounced compared with the data.⁴² Perhaps surprisingly, the linear-Gaussian income process also matches this feature of the data reasonably well. As noted earlier, however, this specification’s ability to capture the nonemployment CDF comes at the cost of an implausibly steep inequality profile (Figure 12(f)). In fact, under a more plausible parameterization, only less than 4% of individuals ever experience a full year of nonemployment over their working life.

How does the model capture the nonlinear earnings dynamics with a single AR(1) component? The autocorrelation of persistent shocks is not precisely captured by ρ , because modeling p_v as a function of z_t implies that nonemployment is autocorrelated even though v_t is drawn in an i.i.d. manner.⁴³ Moreover, since this function is highly nonlinear, how income responds on impact to a given shock η_t , and how persistent this response is de-

⁴¹Recall that shorter-term nonemployment spells are captured by transitory and persistent shocks.

⁴²These patterns are also consistent with the evidence from the Survey of Income and Program Participation on income and age variation in unemployment risk (see Karahan, Ozkan, and Song (2019)).

⁴³Krusell, Mukoyama, Rogerson, and Şahin (2011) also showed that persistent idiosyncratic productivity shocks play a key role in matching the persistence of the nonemployment found in the CPS data.

TABLE V
PERSISTENCE OF NONEMPLOYMENT RISK, BENCHMARK PROCESS (MODEL (6))

Nonemp. at $t \rightarrow$	25–35		36–45		46–55	
	$t+1$	$t+5$	$t+1$	$t+5$	$t+1$	$t+5$
<i>RE Groups</i>						
1–10	0.543	0.481	0.583	0.500	0.623	0.519
41–60	0.455	0.426	0.544	0.475	0.604	0.503
91–100	0.334	0.352	0.491	0.440	0.582	0.498

pend very much on the persistent component z_t and the sign and magnitude of the shock. This feature, along with the normal mixture shocks, generates the asymmetric mean reversion in impulse responses in Figure 13.

To illustrate its persistence, we examine the future nonemployment risk of workers who are nonemployed in t . As usual, we further condition workers on their RE in $t-1$. Nonemployment is fairly persistent overall and more so for low-income workers: Between ages 25 and 35, 48% of the workers in the bottom RE decile experience another nonemployment spell five years after the initial nonemployment (Table V). This number declines monotonically over the RE distribution to 35% for the top decile. Furthermore, nonemployment risk becomes more persistent over the working life, particularly for high earners. For example, in the top decile, the conditional probability of nonemployment in $t+5$ increases from 35% in the first 10 years to 50% in the last 10 (see Guvenen, Karahan, Ozkan, and Song (2017) for an empirical investigation of nonemployment persistence).

Autocovariance Structure of Earnings. Another set of moments that has been widely used when estimating income processes is the autocovariance matrix of *log* earnings in levels or changes (e.g., Abowd and Card (1989)). An important drawback of this approach, especially through the lens of our specification, is that it ignores small earnings observations below the minimum income threshold. Our estimation targets the dynamics of earnings in a more complete manner through impulse response functions, which capture the heterogeneity in the dynamics of income changes by their size and sign as well as for workers with different earnings levels (Figure 13).

For completeness and consistency with earlier work, we report these moments and show the benchmark specification's fit to them in Appendix D.6. In our data, consistent with earlier work using survey data (Meghir and Pistaferri (2004)), the autocovariance of earnings growth approaches zero quickly after a couple of lags. Our model generates a similar pattern, although the autocovariances for smaller lags tend to be closer to zero than in the data. Modeling the transitory component as a moving average of order q may fix this shortcoming (see Meghir and Pistaferri (2004)). Furthermore, MaCurdy (1982) noted that if a HIP component is present ($\sigma_\beta^2 > 0$), the autocovariance of one-year log earnings growth should turn positive at longer lags. In the data, this test would not reject $\sigma_\beta^2 = 0$. Interestingly, MaCurdy's test reaches a similar conclusion in data simulated from the benchmark process, which features a sizable HIP component (see Guvenen (2009) for a discussion on the reliability of this test).

Discretizing the Benchmark Income Process. Another advantage of building on the linear-Gaussian framework is that existing methods can be adopted to discretize our income process. After all, the benchmark specification is still within the realm of first-order

Markov processes, with the two important additions being the normal mixture shocks in transitory and persistent components and the state-dependent nonemployment shocks. The standard Tauchen (1986) methodology does not take a stance on the distribution of shocks. Nevertheless, there are complications that arise due to the nonnormalities. The key issue is the placement of grid points in the discretized space. The negative skewness and high kurtosis require points to be irregularly placed (see Civale, Diez-Catalan, and Fazilet (2016) for the optimal grid placement).

Moreover, another crucial task is to capture the nonlinearities in the nonemployment probability. How well these are captured by a discretization depends on how fine the grid points of the persistent component are in the region where nonemployment risk is highly nonlinear. Thus, a given level of precision requires many more grid points relative to the standard process. Therefore, in an earlier version of the paper (Guvener, Karahan, Ozkan, and Song (2016)), we chose not to discretize the shock distribution and use quadrature methods instead.

7. CONCLUDING THOUGHTS

Using earnings histories of millions of U.S. workers, we have studied non-Gaussian and nonlinear earnings dynamics and reached the following conclusions: The distribution of earnings growth is left skewed and leptokurtic. Critically, these features vary substantially over the life cycle and across the earnings distribution: Higher-income and older workers on average face a more left-skewed and leptokurtic earnings growth. Finally, earnings changes display asymmetric persistence: Increases for high earners are quite transitory, whereas declines are very persistent; the opposite is true for low earners.

By targeting these nonparametric facts, we estimated an earnings process that allows for normal mixture innovations to the persistent and transitory components and a long-term nonemployment shock with a realization probability that varies with age and earnings.⁴⁴ We found that this state-dependent employment risk generates systematic recurring nonemployment with scarring effects, which is important to match the data.

Our empirical findings are broadly consistent with job ladder models, in which workers see little earnings change in most years but once in a while experience a large change due to unemployment or an outside offer. However, whether existing models can be quantitatively consistent with the large lifecycle and income variation in earnings dynamics is an open question (see Hubmer (2018) and Karahan, Ozkan, and Song (2019)).

Are non-Gaussian features important for household consumption-savings behavior? Consider a household with constant relative risk aversion preferences that pays a risk premium π to avoid consumption risk by a random proportion $(1 + \tilde{\delta})$.⁴⁵ When $\tilde{\delta}$ is Gaussian with zero mean and a variance of 0.01, π is 4.9%, compared to 22.2% when $\tilde{\delta}$ has a skewness of -2 and a kurtosis of 30. While this back-of-the-envelope calculation is meant to be illustrative (e.g., what is relevant for consumption is the dynamics of disposable household income, whereas our facts pertain to male earnings), Guvener, Karahan, Ozkan, and Song (2016), De Nardi, Fella, and Pardo (2020), and Arellano, Blundell, and Bonhomme

⁴⁴The descriptive facts documented here (including those in the Appendix) add up to more than 10,000 empirical moments. Adding analogous moments for women, as mentioned in footnote 1, doubles this number. The richness of this information is far beyond what we are able to fully utilize in the estimation in this paper. Furthermore, different questions may require different subsets of these moments to focus on. With these considerations in mind, we make all the moments available online on our websites.

⁴⁵In particular, π solves $U(c \times (1 - \pi)) = \mathbb{E}[U(c \times (1 + \tilde{\delta}))]$, where $U(c) = \frac{c^{1-\gamma}}{1-\gamma}$ with $\gamma = 10$.

(2017) found that the substantial welfare costs of non-Gaussian income risk carry over to richer and more realistic models.

That said, incorporating nonlinear, non-Gaussian earnings dynamics into quantitative models is still in its infancy. Building on the evidence in Guvenen, Ozkan, and Song (2014), Constantinides and Ghosh (2016) and Schmidt (2016) showed that an incomplete markets model with procyclical skewness generates plausible asset pricing implications. McKay (2017) and Busch and Ludwig (2020) studied cyclical consumption dynamics of these skewness fluctuations and Campbell, Clara, and Cocco (2021) investigated mortgage design features for macroeconomic stabilization in their presence. Targeting the moments documented here, Golosov, Troshkin, and Tsyvinski (2016) showed that a non-Gaussian process implies a substantially higher optimal top marginal income tax compared with a traditional Gaussian calibration, and Kaplan, Moll, and Violante (2018) introduced leptokurtic idiosyncratic risk in a New Keynesian model to generate a realistic portfolio distribution. Finally, using our benchmark process, Catherine, Miller, and Sarin (2020) simulated Social Security retirement wealth and revisited trends in wealth inequality. We hope our findings will further feed back into economic research and policy analyses.

REFERENCES

- AARONSON, D., AND E. FRENCH (2004): "The Effect of Part-Time Work on Wages: Evidence From the Social Security Rules," *Journal of Labor Economics*, 22 (2), 329–352. [2322]
- ABOWD, J. M., AND D. CARD (1989): "On the Covariance Structure of Earnings and Hours Changes," *Econometrica*, 57 (2), 411–445. [2335]
- ALTONJI, J., A. A. SMITH, AND I. VIDANGOS (2013): "Modeling Earnings Dynamics," *Econometrica*, 81 (4), 1395–1454. [2307]
- ARELLANO, M., R. BLUNDELL, AND S. BONHOMME (2017): "Earnings and Consumption Dynamics: A Non-linear Panel Data Framework," *Econometrica*, 85 (3), 693–734. [2306,2308,2336,2337]
- (2018): "Nonlinear Persistence and Partial Insurance: Income and Consumption Dynamics in the PSID," in *AEA Papers and Proceedings*, Vol. 108, 281–286. [2308]
- ARNAUD, A., F. GUVENEN, AND T. KLEINEBERG (2019): "Benchmarking Global Optimizers," Working Paper 26340, National Bureau of Economic Research. [2325]
- BAGGER, J., F. FONTAINE, F. POSTEL-VINAY, AND J.-M. ROBIN (2014): "Tenure, Experience, Human Capital and Wages: A Tractable Equilibrium Search Model of Wage Dynamics," *American Economic Review*, 104 (6), 1551–1596. [2315]
- BAKER, M., AND G. SOLON (2003): "Earnings Dynamics and Inequality Among Canadian Men, 1976–1992: Evidence From Longitudinal Income Tax Records," *Journal of Labor Economics*, 21 (3), 289–321. [2310]
- BLASS-HOFFMANN, E., AND D. MALACRINO (2019): "Employment Time and the Cyclicity of Earnings Growth," *Journal of Public Economics*, 169 (2), 160–171. [2318]
- BONHOMME, S., AND J. ROBIN (2010): "Generalized Non-Parametric Deconvolution With an Application to Earnings Dynamics," *Review of Economic Studies*, 77 (2), 491–533. [2306]
- BONHOMME, S., AND J.-M. ROBIN (2009): "Assessing the Equalizing Force of Mobility Using Short Panels: France, 1990–2000," *Review of Economic Studies*, 76 (1), 63–92. [2306]
- BOUND, J., C. BROWN, AND N. MATHIOWETZ (2001): "Measurement Error in Survey Data," in *Handbook of Econometrics*, Vol. 59, ed. by J. Heckman and E. Leamer. Elsevier, 3705–3843. [2316]
- BROWNING, M., M. EJRNAES, AND J. ALVAREZ (2010): "Modelling Income Processes With Lots of Heterogeneity," *Review of Economic Studies*, 77, 1353–1381. [2307,2331]
- BUSCH, C., AND A. LUDWIG (2020): "Higher-Order Income Risk Over the Business Cycle," No. 274, SAFE Working Paper. [2337]
- CAMPBELL, J. Y., N. CLARA, AND J. F. COCCO (2021): "Structuring Mortgages for Macroeconomic Stability," *Journal of Finance*. [2337]
- CATHERINE, S., M. MILLER, AND N. SARIN (2020): "Social Security and Trends in Inequality," Jacobs Levy Equity Management Center for Quantitative Financial Research Paper, Available at SSRN: <https://ssrn.com/abstract=3546668>. [2337]
- CIVALE, S., L. DIEZ-CATALAN, AND F. FAZILET (2016): "Discretizing a Process With Non-Zero Skewness and High Kurtosis," Available at SSRN: <https://ssrn.com/abstract=2636485>. [2336]

- CONSTANTINIDES, G. M., AND A. GHOSH (2016): "Asset Pricing With Countercyclical Household Consumption Risk," *Journal of Finance*, 72, 415–460. [2337]
- CROW, E. L., AND M. SIDDIQUI (1967): "Robust Estimation of Location," *Journal of the American Statistical Association*, 62 (318), 353–389. [2313]
- CUNHA, F., J. HECKMAN, AND S. NAVARRO (2005): "Separating Uncertainty From Heterogeneity in Life Cycle Earnings," *Oxford Economic Papers*, 57 (2), 191–261. [2305]
- DALY, M., D. HRYSHKO, AND I. MANOVSKII (2018): "Improving the Measurement of Earnings Dynamics," Working Paper 22938, National Bureau of Economic Research. [2330]
- DAVIS, S. J., J. C. HALTIWANGER, AND S. SCHUH (1996): *Job Creation and Destruction*. Cambridge, MA: MIT University Press. [2309]
- DE NARDI, M., G. FELLA, M. G. KNOEF, G. PAZ-PARDO, AND R. VAN OOIJEN (2021): "Family and Government Insurance: Wage, Earnings, and Income Risks in the Netherlands and the US," *Journal of Public Economics*, 193, 104327. [2308]
- DE NARDI, M., G. FELLA, AND G. P. PARDO (2020): "Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare," *Journal of the European Economic Association*, 18 (2), 890–926. [2316,2336]
- DOBKIN, C., A. FINKELSTEIN, R. KLUENDER, AND M. J. NOTOWIDIGDO (2018): "The Economic Consequences of Hospital Admissions," *American Economic Review*, 108 (2), 308–352. [2318]
- GEWEKE, J., AND M. KEANE (2000): "An Empirical Analysis of Earnings Dynamics Among Men in the PSID: 1968–1989," *Journal of Econometrics*, 96, 293–356. [2306]
- GOLOSOV, M., M. TROSHKIN, AND A. TSYVINSKI (2016): "Redistribution and Social Insurance," *American Economic Review*, 106 (2), 359–386. [2337]
- GUVENEN, F. (2009): "An Empirical Investigation of Labor Income Processes," *Review of Economic Dynamics*, 12 (1), 58–79. [2331,2335]
- GUVENEN, F., AND A. A. SMITH (2014): "Inferring Labor Income Risk and Partial Insurance From Economic Choices," *Econometrica*, 82 (6), 2085–2129. [2305,2333]
- GUVENEN, F., F. KARAHAN, AND S. OZKAN, AND J. SONG (2016): "What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?," Working paper, University of Toronto. [2336]
- (2017): "Heterogeneous Scarring Effects of Full-Year Non-Employment," *American Economic Review P&P*, 107 (5), 369–373. [2335]
- (2021): "Supplement to 'What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?'," *Econometrica Supplemental Material*, 89, <https://doi.org/10.3982/ECTA14603>. [2304]
- GUVENEN, F., S. OZKAN, AND J. SONG (2014): "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 122 (3), 621–660. [2306,2318,2337]
- 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. [2333]
- HALVORSEN, E., H. A. HOLTER, S. OZKAN, AND K. STORESLETTEN (2020): "Dissecting Idiosyncratic Income Risk," Available at SSRN: <https://ssrn.com/abstract=3669146>. [2316,2318]
- HEATHCOTE, J., F. PERRI, AND G. L. VIOLANTE (2010): "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States: 1967–2006," *Review of Economic Dynamics*, 13 (1), 15–51. [2330]
- HEATHCOTE, J., K. STORESLETTEN, AND G. L. VIOLANTE (2010): "The Macroeconomic Implications of Rising Wage Inequality in the United States," *Journal of Political Economy*, 118 (4), 681–722. [2316,2326]
- HOROWITZ, J. J., AND M. MARKATOU (1996): "Semiparametric Estimation of Regression Models for Panel Data," *Review of Economic Studies*, 63 (1), 145–168. [2306]
- HUBMER, J. (2018): "The Job Ladder and Its Implications for Earnings Risk," *Review of Economic Dynamics*, 29, 172–194. [2336]
- JACOBSON, L. S., R. J. LALONDE, AND D. G. SULLIVAN (1993): "Earnings Losses of Displaced Workers," *American Economic Review*, 83 (4), 685–709. [2312]
- JAROSCH, G. (2021): "Searching for Job Security and the Consequences of Job Loss," Working Paper 28481, National Bureau of Economic Research. [2328]
- KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): "Monetary Policy According to HANK," *American Economic Review*, 108 (3), 697–743. [2337]
- KARAHAN, F., AND S. OZKAN (2013): "On the Persistence of Income Shocks Over the Life Cycle: Evidence, Theory, and Implications," *Review of Economic Dynamics*, 16 (3), 452–476. [2310]
- KARAHAN, F., S. OZKAN, AND J. SONG (2019): "Anatomy of Lifetime Earnings Inequality: Heterogeneity in Job Ladder Risk vs. Human Capital," FRB of New York Staff Report (908). [2328,2334,2336]
- KELLEY, T. L. (1947): *Fundamentals of Statistics*. Harvard University Press. [2311]
- KRUSELL, P., T. MUKOYAMA, R. ROGERSON, AND A. ŞAHİN (2011): "A Three State Model of Worker Flows in General Equilibrium," *Journal of Economic Theory*, 146 (3), 1107–1133. [2334]

- KURMANN, A., AND E. MCENTARFER (2019): "Downward Wage Rigidity in the United States: New Evidence From Administrative Data," Census Working Paper No. CES-19-07. [2318]
- LILLARD, L. A., AND Y. WEISS (1979): "Components of Variation in Panel Earnings Data: American Scientists 1960–70," *Econometrica*, 47 (2), 437–454. [2306]
- LILLARD, L. A., AND R. J. WILLIS (1978): "Dynamic Aspects of Earnings Mobility," *Econometrica*, 46 (5), 985–1012. [2306]
- LISE, J. (2012): "On-the-Job Search and Precautionary Savings," *Review of Economic Studies*, 80 (3), 1086–1113. [2321]
- LOW, H., C. MEGHIR, AND L. PISTAFERRI (2010): "Wage Risk and Employment Risk Over the Life Cycle," *American Economic Review*, 100 (4), 1432–1467. [2315,2316]
- 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. [2306,2335]
- MCKAY, A. (2017): "Time-Varying Idiosyncratic Risk and Aggregate Consumption Dynamics," *Journal of Monetary Economics*, 88, 1–14. [2337]
- MCLACHLAN, G., AND D. PEEL (2000): *Finite Mixture Models*. John Wiley and Sons. [2324]
- MEGHIR, C., AND L. PISTAFERRI (2004): "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72 (1), 1–32. [2335]
- MOORS, J. J. A. (1986): "The Meaning of Kurtosis: Darlington Reexamined," *The American Statistician*, 40, 283–284. [2312]
- PARETO, V. (1897): *Cours d'Economie Politique*. Geneva: Droz. [2313]
- PISTAFERRI, L. (2003): "Anticipated and Unanticipated Wage Changes, Wage Risk, and Intertemporal Labor Supply," *Journal of Labor Economics*, 21 (3). [2305]
- SCHMIDT, L. D. W. (2016): "Climbing and Falling off the Ladder: Asset Pricing Implications of Labor Market Event Risk," Available at SSRN: <https://ssrn.com/abstract=2471342>. [2337]
- STORESLETTEN, K., C. I. TELMER, AND A. YARON (2004): "Consumption and Risk Sharing Over the Life Cycle," *Journal of Monetary Economics*, 51 (3), 609–633. [2326]
- TAUCHEN, G. (1986): "Finite-State Markov Chain Approximations to Univariate and Vector Autoregressions," *Economics Letters*, 20, 177–181. [2336]
- TODA, A. A., AND K. WALSH (2015): "The Double Power Law in Consumption and Implications for Testing Euler Equations," *Journal of Political Economy*, 123 (5), 1177–1200. [2313]
- TOPEL, R. H., AND M. P. WARD (1992): "Job Mobility and the Careers of Young Men," *The Quarterly Journal of Economics*, 107 (2), 439–479. [2315]
- VON WACHTER, T., J. SONG, AND J. MANCHESTER (2009): "Long-Term Earnings Losses due to Mass Layoffs During the 1982 Recession: An Analysis Using U.S. Administrative Data From 1974 to 2004," Tech. rep, Columbia University. [2312]

Co-editor Giovanni L. Violante handled this manuscript.

Manuscript received 4 August, 2016; final version accepted 26 February, 2021; available online 19 April, 2021.