

Anatomy of Lifetime Earnings Inequality: Heterogeneity in Job-Ladder Risk versus Human Capital

Serdar Ozkan

Federal Reserve Bank of St. Louis and University of Toronto

Jae Song

Social Security Administration

Fatih Karahan

Amazon

We study the determinants of lifetime earnings (LE) inequality in the United States by focusing on latent heterogeneity in job-ladder dynamics and on-the-job learning. We use administrative data to document a novel set of moments on job mobility and earnings growth across the LE distribution. We then estimate a structural model featuring a rich set of worker types and firm heterogeneity. We find vast ex ante differences in job-loss, job-finding, and contact rates across worker types. These differences account for 75% of the lifetime wage growth differential among the bottom half of the LE distribution. Above the median, almost all lifetime wage growth differences are a result of Pareto-distributed learning ability.

Special thanks to Gerald Ray and Pat Jonas for their help and support. This project started jointly with Samuel Kapon, to whom we are grateful. We also thank Brendan Moore and Rachel Schuh for excellent research assistance. For helpful critiques and comments, we thank participants at the 2016 and 2018 Society for Economic Dynamics annual meetings, the 2018 Barcelona School of Economics Summer Forum, the 2020 American Economic Association winter meetings, and the 2020 National Bureau of Economic Research Economic

Electronically published August 8, 2023

Journal of Political Economy Macroeconomics, volume 1, number 3, September 2023.

© 2023 The University of Chicago. All rights reserved. Published by The University of Chicago Press.

<https://doi.org/10.1086/725790>

I. Introduction

Large differences in lifetime earnings (LE) are evident among workers in the United States (Guvenen et al. 2017). Even though inequality starts early in life, the striking differences in earnings growth over the life cycle are key for understanding the LE distribution. In this paper, we study these differences, using administrative balanced-panel data, by focusing on heterogeneities in two important forces that the previous literature has deemed important for earnings growth: (i) the ability to accumulate human capital (Huggett, Ventura, and Yaron 2011) and (ii) the ability to climb the job ladder (Topel and Ward 1992). We aim to quantify the importance of each of these mechanisms throughout the LE distribution by identifying and investigating different types of career paths.¹

We use a confidential employer-employee matched panel of the earnings histories of male workers between 1978 and 2013 from the US Social Security Administration (SSA). Using a 10% sample of workers born between 1953 and 1960, we first compute workers' total wage and salary income over the ages of 25–55 and rank them into 50 LE quantiles. The top 2% group earns about 7.5 times as much as median-LE workers, who earn 3.5 times as much as bottom-quantile earners. The vast majority of these differences are a result of earnings growth heterogeneity: top-LE individuals see their incomes rise by more than 17-fold between the ages of 25 and 55, median-LE workers experience more than a two-fold increase, and those at the bottom see a mere 16% earnings growth.

We employ a job-ladder model with two-sided heterogeneity in the spirit of Cahuc, Postel-Vinay, and Robin (2006) and Bagger et al. (2014) as a measurement device to quantify the relative roles of latent heterogeneity in job-ladder dynamics and human capital accumulation throughout the LE distribution. The model features learning on the

Fluctuations and Growth meetings, as well as participants at other conferences and research seminars at numerous universities and institutions. Special thanks to Fatih Guvenen, Chris Huckfeldt, Henry Hyatt, Gregor Jarosch, Rasmus Lentz, Jeremy Lise, Giuseppe Moscarini, Richard Rogerson, Ayşegül Şahin, Ludo Visschers, and David Wiczer for comments and discussions. An earlier version of this paper was circulated under the titles “Sources of Earnings Growth Heterogeneity” and “Sources of Inequality in Earnings Growth over the Life Cycle.” The views expressed herein are those of the authors and do not represent those of the Social Security Administration, the Federal Reserve Bank of St. Louis, or the Federal Reserve System. This paper was edited by Joseph Vavra.

¹ A long line of literature, dating back to seminal papers by Mincer (1974), Heckman (1976), and Deaton and Paxson (1994), studies the fanning out of inequality over the life cycle. Some explanations of wage growth heterogeneity include human capital accumulation (e.g., Caucutt and Lochner 2020), learning about workers' ability (e.g., Jovanovic 1979; Pastorino 2019), and workers selecting into positions via “tournaments” (Lazear and Rosen 1981) or according to their comparative advantage (Lise, Meghir, and Robin 2016). See Neal and Rosen (2000) for a comprehensive review of this literature.

job, on-the-job search, employer competition, and idiosyncratic shocks to worker productivity. To bring the model closer to the data, we add to this framework a life-cycle structure in the form of perpetual youth. Importantly, we allow for rich ex ante worker heterogeneity in unemployment risk, the job-finding rate, and the contact rate for employed workers, as well as the ability to learn on the job. Finally, the model also features recalls for unemployed workers by their most recent employers (Fujita and Moscarini 2017).

The key insight for identifying the importance of human capital and job-ladder risk throughout the LE distribution relies on differences in job-switching patterns and earnings changes of job stayers and switchers. In the data, about 30% of the bottom-LE workers stay with the same employer for two full consecutive years, compared to around 60% of those above the median. Relatedly, bottom earners work for about 12 employers between the ages of 25 and 55, more than twice as many as those above the median.

Our key novel empirical finding is that average annual earnings growth for job stayers is surprisingly similar, around 2% in the bottom two-thirds of the LE distribution, whereas for job switchers it rises almost linearly, from 0% for the bottom earners to around 3% for those in the 65th percentile.² This large heterogeneity indicates that the nature of job switches is very different across the LE distribution. By exploiting the distribution of earnings changes for job switchers, we argue that more than 35% of job switches are a result of a significant unemployment spell for bottom earners, compared to only around 15% in the top tercile. Finally, earnings growth of job stayers and switchers increases steeply in the top tercile, reaching around 10% for the highest earners.

These facts imply that differences in earnings growth in the bottom half of the LE distribution are coming from growth differences of job switchers, suggesting strong heterogeneity in job-ladder risk among them. Job stayers' growth differences, however, should be the main culprit in the upper half, because high-LE workers rarely switch employers, hinting at differences in returns to experience. Inference is more complicated, because, for example, wage growth of job stayers is also affected by outside offers, and a higher incidence of unemployment can stem from higher ex ante risk or bad ex post luck. We estimate our model with this rich set of facts to obtain an exact quantitative assessment of the importance of different economic forces. Specifically, we target the fraction and average earnings growth of job stayers and switchers as well as the

² Heterogeneity in life-cycle earnings growth has been well documented in the literature (see, e.g., most recently Guvenen et al. 2021). However, to the best of knowledge, we are the first to decompose LE growth differences into job-stayer and job-switcher components.

higher-order moments of their earnings changes by LE groups and over the life cycle.

One of our major contributions is to quantify the vast *ex ante* heterogeneity in job-ladder risk. We estimate a quarterly job-loss risk of 9% for bottom-LE workers, compared to 2% above the median. Quarterly job-finding rates also display large differences, ranging from 30% at the bottom to 50% above the median LE. Given the annual nature of the SSA data, we cannot directly test these estimates. Instead, we use the Survey of Income and Program Participation (SIPP) to document large differences in job-loss and job-finding rates across workers with different past earnings and over the life cycle that are quantitatively consistent with the estimated model. Turning to the contact rate for employed workers, we find that bottom-LE workers have a 30% probability of being contacted in a quarter, relative to a 55% probability at the top. However, the SIPP data show that high earners are less likely to make job-to-job transitions. Our model matches this feature of the data as well: despite getting more outside offers, high-LE workers tend to work for more productive firms and can rarely be poached. We directly test this mechanism, using data from the Survey of Consumer Expectations (SCE), and find more contacts for people with higher past earnings, consistent with our estimates.

The estimated model provides a good account of the career trajectories of workers by LE groups, so we use it to decompose the differences in LE. First, we find that wage—rather than employment—differences explain the vast majority of LE inequality. The only exception is inequality in the bottom half, where employment differences also play some role: bottom-LE workers work about 25% less than those at the median. A higher *ex ante* job-loss rate and—to a lower extent—a lower job-finding rate for bottom-LE workers explain almost all of these employment differences. Employment differences among those above the median are negligible in comparison.

Turning to differences in lifetime wages, we find them to be mainly driven by wage growth over the life cycle, resonating with our empirical findings on earnings inequality and earnings growth. In a series of experiments, we isolate the relative roles of *ex ante* differences in job-ladder risk and the returns to experience. Heterogeneity in unemployment risk accounts for more than 50% of the wage growth differences between workers at the bottom and those at the median. High unemployment rates among low-LE workers reduce wage growth by preventing these workers from accumulating human capital and from climbing the job ladder; the former channel accounts for about 70% of the total effect. Differences in contact rates also have a nonnegligible effect on wage growth heterogeneity. Eliminating them closes an additional 20% of the wage growth gap between workers at the bottom and those at the median by allowing low-LE workers to move to better firms.

While *ex ante* heterogeneity in job-ladder risk is important in the bottom half of the LE distribution, it explains very little of the differences above the median. In contrast, heterogeneity in learning ability drives almost all earnings growth differences above the median but only about 20% of these differences among the lower half. This is because learning ability is Pareto distributed, which implies that average returns to experience are relatively similar in the bottom half and increase steeply toward the top of the LE distribution, essentially mirroring average stayer earnings growth. For workers who enjoy high wage growth regardless of job switching—the top LE group in the data—the model assigns a high level of returns to experience.

A key conclusion of our study is that different economic forces are driving the inequality in different parts of the LE distribution. While bottom-LE workers experience low wage growth relative to the median throughout their working life, primarily because of poor labor market experience (i.e., high unemployment risk and fewer outside offers), workers at the top see high wage growth, mainly because they enjoy a very high level of returns to experience. These quantitative findings resonate with the empirical patterns of job mobility and income growth of job stayers and switchers across the LE distribution.

Our findings also shed new light on the relative roles of initial conditions and *ex post* shocks in determining lifetime inequality (Keane and Wolpin 1997; Huggett, Ventura, and Yaron 2011). In our model, 81% of the variation in LE is a result of *ex ante* heterogeneity in initial conditions, which is substantially higher than the corresponding 61% figure Huggett, Ventura, and Yaron (2011) find from a similar exercise. They use a calibrated Ben-Porath (1967) human capital model that features *ex ante* heterogeneity in initial human capital and learning ability but not job-ladder dynamics. Thus, the greater role for initial conditions stems from the rich worker heterogeneity in job-ladder risk in our estimation, which can more precisely capture the source of inequality in the bottom half.

Related literature and our contributions.—The broad contribution of our paper is to quantify the heterogeneity in economic forces across the income distribution. Recently, Guvenen, Ozkan, and Song (2014) and Guvenen et al. (2021) use the SSA data to document the nature of idiosyncratic risk, with a focus on higher-order moments, across the income distribution and over the business cycle and the life cycle, respectively. Both papers are mostly descriptive and estimate reduced-form income processes. We use moments similar to those in Guvenen et al. (2021), along with new ones we document, to structurally estimate a job-ladder model and identify the economic forces behind earnings growth.

Hubmer (2018) shows that a reasonably calibrated job-search model, as in Burdett and Mortensen (1998), can capture the higher-order moments

of earnings growth documented in Guvenen et al. (2021). We make use of these insights for estimation, but our focus is on quantifying the latent heterogeneity in job-ladder risk and the ability to learn and their roles in life-cycle earnings growth differences, all of which are absent in Hubmer (2018). Another closely related paper is by Bagger et al. (2014), who estimate for Denmark a job-ladder model similar to ours. They investigate educational differences, while we study the entire LE distribution. Finally, to estimate their model they use information on firm productivity, which we lack in our dataset. Instead, we develop an identification scheme that exploits the earnings growth distributions of job stayers and job switchers.

Finally, there is a growing body of research that estimates the latent heterogeneity in job-ladder risk (e.g., Ahn and Hamilton 2020; Morchio 2020). Most recently, Gregory, Menzio, and Wiczer 2021 employ k -means clustering to group workers according to their job-ladder risk in the Longitudinal Employer-Household Dynamics data; Ahn, Hobbijn, and Şahin (2023) and Hall and Kudlyak (2022) use Current Population Survey data to estimate a hidden-state Markov model of labor force transitions. All three papers reach conclusions similar to ours.

The rest of the paper is organized as follows. Section II presents the data and the stylized facts. Section III describes the model, section IV discusses its estimation, and section V presents the estimation results. Section VI provides the decomposition of LE, and section VII discusses policy implications of our findings and concludes.

II. Empirical Analysis

In this section, we document several stylized facts that motivate and guide our analysis of LE inequality. Our aim is to identify the heterogeneities in the ability to accumulate human capital and climb the job ladder by investigating the career paths of different LE groups over the working life. Our analysis is based on administrative annual data from the SSA. We support these findings by using monthly panel data from the SIPP as well.

A. *The SSA Data and Sample Selection*

Our data are drawn from the Master Earnings File (MEF) of the SSA records, which includes every individual with a Social Security number (SSN). Basic demographic variables available are date of birth, place of birth, sex, and race. The earnings data are derived from the employee's W-2 forms. The measure of labor earnings is annual and includes all wages and salaries, bonuses, and exercised stock options as reported on the W-2 form (box 1). The MEF has a small number of extremely (uncapped) high earnings observations; therefore, we winsorize observations above

the 99.999th percentile in each year. We convert nominal earnings into real values by using the personal consumption expenditure deflator, taking 2005 as the base year. See Panis et al. (2000) and Olsen and Hudson (2009) for detailed documentation of the MEF.

W-2 forms contain another crucial piece of information for our purpose: an employer identification number (EIN), which identifies firms at the level at which they file their tax returns with the Internal Revenue Service. We use this variable to follow each worker's career path at an annual frequency. Note that an EIN is a different concept from an "establishment," which typically represents a single geographic facility of the firm. Two caveats are worth mentioning regarding the use of EINs to identify firms. First, an EIN does not always identify the parent firm, because some large firms choose to file taxes at a level lower than the parent firm (see Song et al. 2019). Second, firms may change their EINs, for example, as a result of ownership changes (see Haltiwanger et al. 2014). As a result, we may be overcounting the number of job switches.

Sample selection.—We construct a 10% sample based on the randomly assigned last four digits of (a confidential transformation of) the SSN. We select individuals born between 1953 and 1960, for whom we therefore have data between the ages of 25 and 55 (referred to as a worker's lifetime). Furthermore, we work with a sample of wage and salary workers with strong labor market attachment, because the mechanisms we investigate speak to labor market participants. One drawback is that the MEF does not have a direct measure of labor force participation. We address this problem by excluding individuals with earnings below a time-varying minimum earnings threshold $Y_{\min,t}$ —25% of a full-year, full-time salary at half the minimum wage, for example, $\approx \$1,885$ in 2010—for (i) at least one-fourth of their working life or (ii) two or more consecutive years. These two criteria help us exclude early retirees, the disabled, and those who are out of the labor force for other reasons.³ We also drop workers who are self-employed—those with self-employment income above the minimum earnings threshold $Y_{\min,t}$ and more than 10% of annual total earnings—(iii) for more than one-eighth of their working life or (iv) for two or more consecutive years.⁴ These restrictions exclude workers who choose self-employment as their career path and yet keep those who rely on self-employment income during unemployment spells, as well as payroll workers with a small self-employment income on the side. This

³ Note that a nonemployment spell of at least two full calendar years implies a significantly longer actual spell. Given the duration dependence of job-finding rates (Jarosch and Pilossoph 2019), a worker with such a long nonemployment spell is unlikely to have been looking for jobs the entire time.

⁴ Of course, self-employment is an important option in workers' careers. However, investigating this additional endogenous decision is outside the scope of this paper.

procedure reduces our sample from 1,845,640 individuals to 840,194 for whom we have at least 31 years of earnings data.⁵

B. Stylized Facts on LE Inequality and Growth

We compute LE as the sum of individuals' W-2 earnings from age 25 to age 55.⁶ This measure is then used to rank workers into 50 equally sized quantiles, LE_j for $j = 1, \dots, 50$. Individuals around the 90th percentile (LE_{45}) earn 3.7 times as much as those around the 10th percentile (LE_5 ; see fig. 1A and table A.2). This inequality is roughly half the annual earnings inequality, for which the ratio of the 90th percentile to the 10th percentile hovered around 8 throughout our sample period (Güvenen, Ozkan, and Song 2014). Inequality is more pronounced at the top, with LE_{50} earning almost 4 times as much as LE_{45} , whereas LE_5 earned almost twice as much as LE_1 .

It is well known that in the United States differences in earnings growth over the life cycle are key for understanding the inequality in LE (see, e.g., Huggett, Ventura, and Yaron 2011 and Kaplan 2012). To illustrate this point, figure 1 shows the log growth of average earnings between different ages over the LE distribution; that is, $\log \bar{Y}_{h_2,j} - \log \bar{Y}_{h_1,j}$, where $\bar{Y}_{h,j}$ is the average earnings of workers in LE_j at age h (see Güvenen et al. 2021 for a similar figure from a broader sample). This growth measure allows us to include workers with zero earnings.⁷ Earnings growth is positively related to the level of LE, which is not surprising, since, all else the same, one should expect the higher-growth individuals to rank at the top of the distribution. However, the quantitative magnitudes are striking: the top LE earners (LE_{50}) see their earnings rise by more than 17-fold between the ages of 25 and 55, and median workers experience a twofold increase, whereas those at the bottom see little to no earnings growth (around 16%).⁸ As we quantify in section VI, large differences in earnings growth make an unmistakable contribution to the LE inequality.

⁵ Clearly, our final sample is highly selective (for more detail, see table A.1). Appendix A.4 documents the key empirical findings for a broader sample, which are qualitatively similar to our baseline results.

⁶ Career choices are likely based on discounted future earnings. However, there are no market prices for human capital from which we can infer the appropriate discount rate (see Huggett and Kaplan 2016). We prefer to treat earnings at all ages equally and measure lifetime income in this way (see also Güvenen et al. 2017). However, we discount future periods in the utility function.

⁷ The results are qualitatively similar for log earnings growth, which excludes earnings below $Y_{\min,t}$.

⁸ Further striking differences exist among top earners. Earnings grow around 700% for those around the 98th percentile, compared to more than 5,000% in the top 0.2% group (fig. A.5).

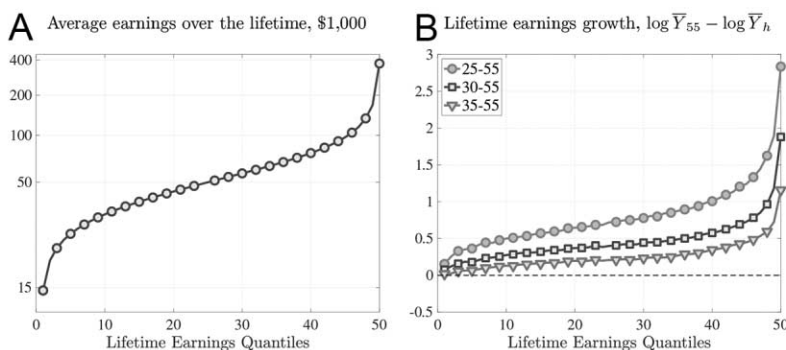


FIG. 1.—Heterogeneity in LE growth. *A*, Average annual earnings over the life cycle for each LE group. *B*, Log difference of average earnings \bar{Y} between age 55 and various ages over the LE distribution. For clarity, we use one symbol for every other LE quantile.

Some of this steep rise in earnings growth at the top could simply be due to transition from school to employment in the labor market. For example, top-LE individuals might be pursuing graduate degrees around earlier ages. While the lack of education data does not allow us to answer this question directly, figure 1*B* plots earnings growth between the ages of 30 and 55 and 35 and 55, when schooling is unlikely to matter much. While the magnitudes change, we still find a steep profile of earnings growth with respect to LE, suggesting that low labor supply at age 25 is not the major driver of these patterns.

Top-earnings inequality: a brief digression.—As we discussed above, inequality is more pronounced at the top of the LE distribution. For example, the average annual earnings in the top 0.2% is over \$1,000,000, compared to \$200,000 for those around the 98th percentile. In fact, the right-hand tail of the LE distribution follows a power law (table A.2) with a Pareto tail slope of -2.13 (fig. A.1). It is also already established that the population earnings distribution has Pareto tails (Piketty and Saez 2003; Atkinson, Piketty, and Saez 2011). Importantly and interestingly, we find that this power law also holds in the cross-sectional distribution of earnings at each age (see app. A.2.1). Log density is linear in the tails at all ages, and the slope gets closer to 1 in absolute value, which points to rising income concentration over the life cycle (figs. A.3, A.4).

C. Career Paths by LE

A natural immediate question is, What accounts for the large differences in earnings growth? To this end, we investigate the differences in labor market experiences between LE groups. Earlier work has shown that job mobility is important for earnings growth over the life cycle (Topel

and Ward 1992). Therefore, we start by investigating how the number of (distinct) employers over the working life differs between LE groups.

Individuals at the bottom of the LE distribution work for almost five different employers on average between the ages of 25 and 34, whereas the number of unique employers drops sharply to around three in the upper half of the LE distribution (fig. 2A). As workers age, job switching declines throughout the LE distribution but much more so in the upper half. While top-LE workers work for around 1.5 different employers per decade after age 35, bottom-LE workers still end up working for 3.5 employers on average, not much lower than the number of employers before age 35. At first glance, one might think that low-LE individuals switch jobs very often and experience large earnings growth as a result. As we will see next, the nature of switches is very different across the LE groups.

We now document the average earnings growth across LE groups for workers who stay with the same employer and for those who change jobs. Given the annual frequency of the data, it is possible for job switchers to have more than one W-2 in a given year. Moreover, some workers may hold multiple jobs concurrently. These issues pose a challenge for a precise classification of job stayers and switchers. There is more than one plausible definition for a job stayer, and we opt for a conservative one. Specifically, we call a worker a *job stayer* between years t and $t + 1$ if (i) he has income from the same employer in years $t - 1$, t , $t + 1$, and $t + 2$; (ii) his income in years t and $t + 1$ is above the minimum income threshold for that year; and (iii) this employer accounts for at least 90% of his total labor income in years t and $t + 1$.⁹ This definition ensures that the main employer was the same firm in years t and $t + 1$. We label all other workers as *job switchers*. Note that according to this definition, switchers are a very heterogeneous group and consist of people who make direct job-to-job transitions, those who experience nonemployment, and those who come out of nonemployment. We return to this heterogeneity later.

Figure 2B shows the fraction of job stayers within each LE group, averaged over the working life. Resonating with the large differences in the number of different jobs held over the life cycle, bottom-LE individuals stay with the same firm, on average, for 30% of their working life, compared to around 60% for individuals above the median.¹⁰

How much earnings growth does a worker experience when he stays with the same employer versus when he switches jobs? The answer differs widely across the LE distribution (fig. 2C). For job stayers, log average

⁹ We find similar results when we impose the condition that the main employer accounts for at least 50% of the total income. Results are available upon request.

¹⁰ The decline in the fraction of job stayers at the top of the LE distribution can be due to some of the workers leaving the labor force temporarily for schooling (e.g., MBA or law school).



FIG. 2.—Job stayers and switchers. *A*, Number of distinct employers over the working life by LE. *B*, Fraction of workers in each LE group who are job stayers according to our definition, calculated for each age and averaged over the working life. *C*, Log growth of average earnings \bar{Y} between t and $t + 1$ for job stayers and switchers separately, again averaged over t over the working life.

earnings growth (between t and $t + 1$) is surprisingly similar, at around 2% in the bottom two-thirds of the LE distribution, whereas average earnings growth for stayers increases sharply in the top LE tercile, reaching around 10% for LE_{50} . Turning to job switchers, we find that average annual earnings growth is essentially zero at the bottom of the LE distribution and rises almost linearly to around 4% for LE_{45} , after which it accelerates to 10% for the top-LE individuals. This large heterogeneity indicates that the nature of job switches is very different throughout the LE distribution, which we will investigate shortly.¹¹

These differences in earnings growth between job stayers and switchers are key for understanding the different forces behind the LE growth over the LE distribution. For example, given the little heterogeneity among job stayers below the median LE, it is clear that the differences in LE growth are due to the differences in the frequency and nature of job switches among these workers, suggesting large heterogeneity in job-ladder risk for them. Inasmuch as workers above the median LE are much more likely to stay with the same employer, large differences in earnings growth of job stayers should therefore be the main culprit behind the LE growth differences among them, suggesting large heterogeneity in returns to experience. That being said, wage growth of job stayers is also determined by outside offers they do not accept, and returns to experience is still a factor for earnings changes of job switchers. Thus, the exact decomposition of the importance of different economic

¹¹ Guvenen et al. (2021) also investigate how job stayers and switchers across the income distribution differ in earnings growth. However, they document patterns on only higher-order moments of earnings growth (but not on average earnings growth). Thus, we are the first to decompose average earnings growth differentials into job-stayer and job-switcher components.

forces requires a structural model estimated to match these salient features of the data (sec. III).

As we discussed above, job switchers are a very heterogeneous group, because they include workers who switch jobs directly or due to a job loss (or a quit). The annual nature of the data does not allow us to separate these directly. Yet we argue that the earnings growth distribution of switchers is informative about the nature of switches. For example, switchers who see their earnings decline by more than 25% have most likely experienced some nonemployment spell in $t + 1$. Thus, we classify such workers as “U-switchers” and the remaining job switchers as “E-switchers.” The latter category contains workers who make direct job switches as well as those coming out of nonemployment in $t + 1$.¹²

More than 35% of job switchers are U-switchers for bottom-LE workers (fig. 3A). This share declines sharply over the LE distribution and reaches a low of 15% for LE_{40} , before increasing to 20% for top-LE workers. Thus, on average, higher-LE individuals are more likely to make job switches involving earnings increases. Investigating the average earnings growth associated with each type of switch, we find large differences between E- and U-switches but little variation across the LE distribution (except for the bottom and top ends). On average, an E-switch is associated with an earnings increase of larger than 15%, whereas a U-switch is associated with a decline of more than 60%.

The annual nature of our data limits the analysis of the earnings changes of job switchers.¹³ To investigate the role of annual aggregation in our results, we construct (normalized/average) earnings growth between the years when a worker is full-year employed in the same firm before and after the switch. Our substantive conclusions hold when we analyze this measure of earnings growth (fig. A.10). In addition, we investigate the monthly SIPP data, which allow us to construct direct measures of job-loss (EU), job-finding (UE), and job-to-job (EE) transition rates by income and age (see app. B for details). We find that low-income workers lose their jobs more often and that the job-finding rate is significantly lower for them (fig. B.1); therefore, they are more likely to make U-switches. In section V, we use these direct measures of quarterly job-flow rates to quantitatively test our estimation results in an external validation exercise.

¹² Jolivet, Postel-Vinay, and Robin (2006) show that a sizable portion of direct job-to-job transitions indeed involve wage cuts. Sorkin (2018) argues that some of these differences can be traced to amenity differences across firms. Tanaka, Warren, and Wiczer (2023) link the earnings declines from direct job switchers to labor force dynamics at both the origin and destination firms.

¹³ For example, if a worker becomes unemployed some time in year t or $t + 1$, then his earnings in $t + 1$ may reflect earnings from a short-term job in that year. Our approach throughout the paper to dealing with such issues is using the estimated model where we aggregate simulated quarterly earnings to annual and construct moments in a similar fashion.

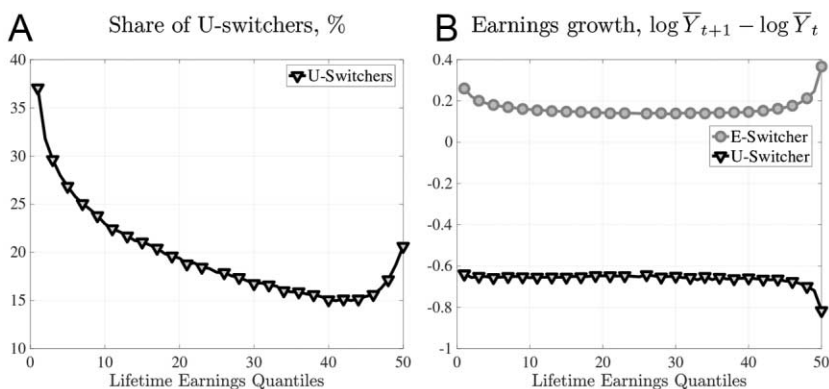


FIG. 3.—E-switchers and U-switchers. A, Share of U-switchers among job switchers averaged over the life cycle. B, Log growth of average earnings \bar{Y} between t and $t + 1$ for U- and E-switchers.

Life-cycle variation.—Significant age variation in job switching and earnings growth rates has been extensively documented (Topel and Ward 1992). We contribute to this literature by investigating differences in these life-cycle profiles between income groups. Figure 4 plots the fraction of job stayers and the earnings growth of stayers and switchers for three stages of the working life over the LE distribution. The fraction of workers who stay with the same firm increases in a concave fashion over the life cycle for all LE groups. This increase is consistent with the declining unemployment risk and job mobility previously documented (Jung and Kuhn 2018). Interestingly, the concavity is more pronounced above the median LE, resonating with the larger decline in the number of distinct employers over the life cycle (fig. 2A).

Turning to the average earnings growth of job stayers, we find a flat profile below LE_{30} at all ages. Moreover, consistent with the existing literature, the rate of earnings growth declines with age. Similarly, the average

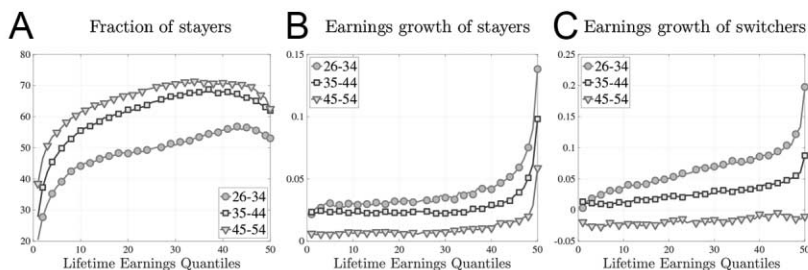


FIG. 4.—Life-cycle variation.

earnings growth for job switchers also declines sharply over the life cycle, especially for higher-LE workers, and becomes negative for the oldest age group throughout the LE distribution (fig. 4C). Interestingly, being a job switcher or a stayer has limited effects on earnings growth of top-LE workers before age 44, but it matters quite a bit for the oldest group. These individuals keep experiencing large earnings gains even after 44 when they stay with the same employer, but their earnings decline if they switch jobs. In other words, in the top-LE group earnings growth falls sharply for both job stayers and switchers over the life cycle but much more so for job switchers. This is because, first, U-switches become more likely among job switchers as the top-LE workers get older. Second, and more importantly, average earnings growth for U-switches falls sharply for the oldest top-LE workers (fig. A.8), which implies that unemployment spells become more costly for them. Our model is able to capture this feature of the data, and we investigate it further in section V.

We have documented several facts regarding the careers of individuals who end up in different parts of the LE distribution. While these facts are useful for describing the various components of earnings growth heterogeneity, they do not suffice to provide an exact decomposition of the importance of the underlying economic forces or to separate *ex ante* heterogeneity from *ex post* luck. In what follows, we develop and estimate a structural model of wages and job turnover with heterogeneity in returns to experience and job-ladder risk as well as *ex post* productivity and job-ladder shocks. In the end, this quantitative model will allow us to disentangle the various economic forces that shape the distribution of wage changes of job stayers and switchers.

III. Model

We build on Bagger et al. (2014), because it features a tractable framework to study the role of job search and learning on the job in generating wage growth.¹⁴ Despite endogenously generating some age variation in job mobility and earnings dynamics, this model falls short of capturing the magnitudes in the data. Thus, we incorporate stochastic aging into this framework, à la Blanchard (1985). Furthermore, motivated by large negative earnings changes for job stayers, we also allow for recalls of unemployed workers by their most recent employers, à la Fujita and Moscarini (2017). Next, we present our theoretical model and motivate each ingredient by linking them to specific empirical facts.

¹⁴ Bowlus and Liu (2013) incorporate endogenous job search in a Ben-Porath human capital model to quantify the relative contributions of each mechanism to life-cycle earnings growth of male high school graduates. They allow for two worker types in job-search efficiency and learning ability but assume uniform unemployment risk.

A. Environment

The economy is populated by heterogeneous workers and firms that produce a single consumption good sold in a competitive market. Workers can be employed or unemployed and search for jobs in a frictional labor market, both on and off the job. They start life as young (y) and become old (o) with probability γ . They have preferences with log per-period utility over consumption and discount future periods at rate ρ :

$$U(\{c_t\}_{t=0}^{\infty}) = \sum_{t=0}^{\infty} \left(\frac{1}{1 + \rho} \right)^t \log c_t.$$

There is no intertemporal savings technology that allows workers to smooth their consumption. This assumption, along with the log preferences and log-linear production function—which we introduce shortly—greatly simplifies computation.

1. Worker Productivity

Each worker enters the labor market with no experience and accumulates human capital as he gains actual experience from employment. The human capital of worker i in period t is given by

$$h_{it} = \tilde{h}_{it} + \epsilon_{it}, \quad \tilde{h}_{i0} = \alpha_i,$$

$$\tilde{h}_{it} = \begin{cases} \tilde{h}_{it-1} - \varsigma & \text{if unemployed,} \\ \tilde{h}_{it-1} + \beta_i + \zeta [\tau_{it}^2 - (\tau_{it} - 1)^2] & \text{if employed.} \end{cases} \quad (1)$$

Here, \tilde{h}_{it} denotes the deterministic component of human capital. Its level at $t = 0$, \tilde{h}_{i0} , is determined by the worker's type α_i , which reflects permanent heterogeneity in productivity due to differences in initial conditions such as innate ability, education, and labor market experience before $t = 0$. Human capital accumulates as the worker gains actual experience τ_{it} through employment. Note that human capital is not specific to the firm, consistent with Gathmann and Schönberg (2010), who show labor market skills to be quite portable. Motivated by the large differences in average earnings growth for job stayers by LE, the rate of human capital accumulation has a worker-specific linear component β_i , potentially correlated with α_i , and a common quadratic component ζ .¹⁵ In Huggett, Ventura, and Yaron (2011), individual-specific growth rates of human capital arise as a result of different investment choices due to the heterogeneity in productivity in the production of human capital. Our model captures

¹⁵ We have also considered a version with individual-specific ζ but have not found significant heterogeneity.

this heterogeneity through exogenous differences in returns to experience. When a worker is unemployed, his human capital \tilde{h}_{it} depreciates at a constant rate ζ .¹⁶ Finally, ϵ_{it} is an idiosyncratic shock to worker productivity that captures the residual sources of variation in earnings not modeled in our framework, such as bonuses. Motivated by the large variation in the distribution of earnings changes for job stayers (fig. D.2), we assume that its distribution depends on worker type α_i and age. We specify the process for ϵ_{it} in section IV in detail.

2. Firm Distribution and Production Technology

Productivity of a firm is constant over time and is drawn from a distribution $F(p)$ with a support of $[\underline{p}, \infty]$ common to all workers. A worker with human capital h_{it} , who works for a firm with productivity $p_{j(i,t)}$, produces a homogeneous good according to a log-linear production function, $e^{y_{it}} = p_{j(i,t)} + h_{it}$.

3. Heterogeneity in Search and Matching

Unemployment risk.—A job dissolves exogenously with probability $\delta^a(\alpha_i)$, in which case the worker searches for a job. We model separation rates to be heterogeneous across workers of different types and ages. This heterogeneity is needed to capture the declining unemployment risk by the wage and age of workers discussed in section II.

Job-finding rate.—An unemployed worker of age $a \in \{y, o\}$ with permanent ability α_i meets a firm with probability $\lambda_0^a(\alpha_i)$, which captures ex ante heterogeneity in job-finding rates. This heterogeneity is motivated by our findings from the SSA and SIPP data and is potentially important for wage growth over the life cycle, because workers with a high job-finding rate will work for more years, end up accumulating more human capital, and, on average, work for more productive firms. To account for the sources of earnings growth, we explicitly model the differences in job-finding rates. Furthermore, workers who are hit by separation shocks find a job immediately with probability $\xi \lambda_0^a(\alpha_i)$. As we discuss below, our model period is a quarter, and a nonnegligible fraction of laid-off workers find a job within 3 months (Ábrahám, Álvarez-Parra, and Forstner 2016). Moreover, there is evidence of transitions that look like direct job-to-job switches but are actually involuntary (Jolivet, Postel-Vinay, and Robin 2006). Thus, we allow for the possibility of finding a job within the same period.

¹⁶ Unemployed workers also lose search capital, negotiation rents, and the forgone opportunity of accumulating experience, but these fall short of capturing the magnitude of the scarring effects of unemployment.

Recalls.—In the data, there are many job stayers who experience large declines in annual earnings. Strongly left-skewed idiosyncratic productivity shocks could, in principle, account for these large losses; in turn, we would have erroneously assigned a bigger role for ex post productivity shocks, especially for low-income workers (as opposed to their higher ex ante unemployment risk). Instead, we allow for a recall option of unemployed workers by their most recent employers, which is fairly prevalent in the data (Fujita and Moscarini 2017). Specifically, we assume that with probability λ_r the offers for unemployed workers come from their most recent employers. The recall option can alter the wages, because it affects the value of a job to a worker. However, we assume that the option value of recall is not considered in the wage-bargaining process.¹⁷ This assumption keeps the estimation computationally feasible, because it allows us to derive the wage equation analytically.

Search on the job.—While employed, workers search for better jobs and with probability $\lambda_1^a(\alpha_i)$ receive an outside offer from another employer, whose productivity is drawn from the distribution $F(p)$, triggering a renegotiation between two firms that we explain below. As figures 3A and B.1C show, workers differ in the types and rates of job switches. Our framework can generate qualitatively similar patterns without explicit differences in the contact rates: high-wage workers—employed, on average, by more productive firms—are less likely to get an offer that beats that of their current employer. This reduces their EE transition rate even if they receive counteroffers at the same rate as low-wage workers. Similarly, as workers get older, they settle into higher-paying jobs and are less likely to move. However, our estimation shows that this endogenous mechanism is insufficient to explain the quantitative differences in the data.

Timing of events.—At the beginning of each period, the productivity shocks are drawn and workers' human capital is updated according to equation (1). Next, output is produced and wages are paid. There is no intertemporal savings device, so workers consume their wages. At the end of the period, search and matching shocks are realized: unemployed workers who find jobs negotiate their wage, workers who receive an outside offer renegotiate their wages or switch employers, and employed workers who draw separation shocks become unemployed. They may find

¹⁷ The option value of recall is higher for workers with higher unemployment risk. Therefore, low-LE workers would have accepted lower starting wages when switching to a more productive firm, which would, in turn, imply stronger wage growth on the same job. So, allowing for the value of the recall option in the wage-bargaining process would be an additional factor why earnings growth of low-LE workers is higher for job stayers than for job switchers and further strengthen our main conclusion that ex ante heterogeneity in job-ladder risk is key for understanding differences in wage dynamics between low- and median-LE workers.

a job immediately or have to wait for the next period to search. Aging occurs stochastically at the end of the period with probability γ and is mutually exclusive from the labor market shocks.

B. Wage Determination

We now briefly explain the bargaining protocol by focusing on the key equations and how the life-cycle structure affects them. See appendix C for derivations.

Wages are specified as piece-rate contracts. In particular, if a worker with human capital h works for a firm of productivity p at a piece rate of $R = e^r \leq 1$, he receives a log wage w of $w = r + p + h$. Here R , the contractual piece rate, is determined endogenously. Upon meeting with a firm, the worker bargains over this piece rate R , which is not updated until the worker meets with another firm.

We now describe how this piece rate is determined for workers with different labor market states. First, let us define $\mathbb{I}_i \equiv \{\alpha_i, \beta_i\}$ as the vector of individual-specific state variables capturing ex ante (fixed) heterogeneity. Note that as we discussed above, \mathbb{I}_i pins down the individual-specific worker flow rates as well as the firm distribution, that is, $\{\delta^y(\alpha_i), \delta^o(\alpha_i), \lambda_0^y(\alpha_i), \lambda_0^o(\alpha_i), \lambda_1^y(\alpha_i), \lambda_1^o(\alpha_i)\}$. The value functions introduced below are individual specific and thus a function of \mathbb{I}_i in addition to other state variables.

1. Hires from Unemployment

Let $V_0^a(h; \mathbb{I}_i)$ and $V^a(r, h, p; \mathbb{I}_i)$ denote the expected lifetime utility of an unemployed worker i with human capital h at age a and that when he is employed at a firm with productivity p at piece rate e^r , $r < 0$, respectively. We define $V^a(r, h, p; \mathbb{I}_i)$ below and assume that the value of unemployment is equivalent to employment in the least productive firm of type p_{\min} extracting the entire match surplus; that is, $V_0^a(h; \mathbb{I}_i) = V^a(0, h, p_{\min}; \mathbb{I}_i)$. This assumption—typical for this class of models and justified by the high empirical job acceptance rate of the unemployed (van den Berg 1990)—implies that unemployed workers accept any job offer and simplifies the problem.

The wage-bargaining protocol dictates that unemployed workers receive share θ of the expected match surplus, where θ captures the worker's bargaining power.¹⁸ More specifically, the piece rate of a hire from unemployment, r_0 , solves

¹⁸ Cahuc, Postel-Vinay, and Robin (2006) argue that this protocol can be micro founded as the equilibrium of a strategic bargaining game adapted from Rubinstein (1982).

$$\mathbb{E}V^a(r_0, h', p; \mathbb{I}_i) = V_0^a(h; \mathbb{I}_i) + \theta\mathbb{E}[V^a(0, h', p; \mathbb{I}_i) - V_0^a(h; \mathbb{I}_i)]. \quad (2)$$

The worker's surplus from the match is the increase in expected lifetime utility from unemployment to a state where he is paid his entire output ($r = 0$). Thus, when an unemployed worker is hired, the firm offers a piece rate that increases his expected lifetime utility by a share θ of this surplus. In equation (2), the expectation is with respect to ϵ_{t+1} .

2. Poaching

When a worker is contacted by a firm with productivity p' , the incumbent firm and the poacher compete. The more productive firm outbids the less productive one and hires the worker. We now discuss the wage that arises as a result of this competition.

There are several cases to consider. First, suppose that the poacher has higher productivity, $p' > p$. Then, the poacher hires the worker by paying a piece rate r' that increases the worker's value by share θ of the surplus generated by the match:

$$\mathbb{E}V^a(r', h', p'; \mathbb{I}_i) = \mathbb{E}[V^a(0, h', p; \mathbb{I}_i) + \theta(V^a(0, h', p'; \mathbb{I}_i) - V^a(0, h', p; \mathbb{I}_i))]. \quad (3)$$

Note that, as Postel-Vinay and Robin (2002) have shown, job switches may result in workers accepting wage losses, because they anticipate faster wage growth in higher-productivity firms. Wage losses upon job switches are a prominent feature of the data.

Second, let us consider a case in which the poacher has lower productivity than the current employer. Bertrand competition implies that the incumbent firm retains the worker, possibly by adjusting the worker's piece rate. This new piece rate offers the worker the maximum value he could attain working at firm p' , that is, the value associated with $r = 0$ ($R = 1$), and a share θ of the additional surplus generated by the offer. In this case, the new piece rate r' solves the following equation:

$$\mathbb{E}V^a(r', h', p; \mathbb{I}_i) = \mathbb{E}[V^a(0, h', p'; \mathbb{I}_i) + \theta(V^a(0, h', p; \mathbb{I}_i) - V^a(0, h', p'; \mathbb{I}_i))]. \quad (4)$$

Note that in contrast to other models of on-the-job search, such as Burdett and Mortensen (1998) and Hubmer (2018), this model generates potentially large and leptokurtic increases in wages for job stayers, which is prevalent in the data (Güvenen et al. 2021).

In some cases, the productivity of the poacher may be so low that the new offer does not generate any additional surplus and therefore does not trigger a change in the piece rate. Then, the worker discards the offer. Let $q^a(r, h, p; \mathbb{I}_i)$ denote this threshold firm productivity, such that offers from firms with $p' \leq q^a(r, h, p; \mathbb{I}_i)$ are discarded; q^a solves

$$\mathbb{E}V^a(r, h', p; \mathbb{I}_i) = \mathbb{E}[V^a(0, h', q^a; \mathbb{I}_i) + \theta(V^a(0, h', p; \mathbb{I}_i) - V^a(0, h', q^a; \mathbb{I}_i))]. \quad (5)$$

IV. Estimation

We now use this model to estimate the contributions of the heterogeneity in the worker flow rates and the ability to accumulate human capital to the differences in earnings growth over the life cycle. To this end, we first exogenously set four parameters: The quarterly discount rate ρ is set to 0.005 to match the annual rate of 2%; workers' bargaining power θ is set to 0.4, following Bagger et al. (2014); the quarterly aging probability γ is set to 1/60, so that a worker becomes old, on average, in 15 years; and the reallocation probability ξ is set to 0.4, following Ábrahám, Álvarez-Parra, and Forstner (2016).

We estimate the remaining parameters by using the simulated method of moments (SMM). We simulate quarterly data for 100,000 individuals and aggregate them to annual observations to create a matched employer-employee panel mimicking the SSA sample. Importantly, we subject the model to the same sample selection criteria in section II.A and compute the model counterparts of targeted moments. Recall that our sample consists of workers between the ages of 25 and 55 (sec. II.A). However, estimation does not assume that workers start their careers at age 25. Instead, in both the data and the model we are agnostic about workers' labor market experiences before age 25. Thus, previous labor market experience or time spent in school would show up as ex ante heterogeneity in our estimation. In our simulations, each individual starts unemployed at the age of 23, and we discard the first two years to allow workers to find a job before age 25.

A. Targeted Moments

We target five sets of moments. The first two are about the cross-sectional distribution of earnings changes for job stayers and switchers. The third and fourth have to do with the fractions of job stayers, E-switchers, and U-switchers and their average annual earnings growth, respectively. Finally, we target average earnings at age 25 by LE group. We choose not to target the heterogeneity in lifetime income growth. As we argue in the next section, the model is already identified using these five sets of moments.

Grouping workers.—We condition each targeted moment on LE and age groups. Specifically, we calculate workers' LE as in section II and rank them into 12 percentile groups: 1–4, 5–10, 11–20, . . . , 81–90, 91–96, 97–100. Furthermore, we group workers into three age groups: 25–34, 35–44, and 45–54.

Cross-sectional moments of earnings growth.—As documented in Guvenen et al. (2021), earnings changes are highly leptokurtic and left skewed. This shape of the earnings change distribution is broadly consistent with job-ladder models. Most workers see little change, but a small share experience a large swing due to unemployment, an EE transition, or an outside offer, which in turn may lead to a left-skewed and leptokurtic distribution. On the basis of these insights, we target the mean, standard deviation, skewness, and kurtosis of annual earnings changes for job stayers and switchers separately.¹⁹ We also condition workers on the basis of their LE because of the large variation in these moments by income (Guvenen et al. 2021).

An issue when computing growth rates is dealing with zero earnings. Recall that in our sample, we drop workers with two or more consecutive years of zero earnings. However, there are still observations with no income in a given year. We would like to keep them, because they contain information about the importance of search frictions. For this purpose, we use the arc percent growth measure, defined as $2(Y_{t+1} - Y_t)/(Y_{t+1} + Y_t)$, where Y_t is annual earnings. Targeted cross-sectional moments are shown in figure D.2.

Average income growth moments.—Next, we target the fraction and average income growth of job stayers, E-switchers, and U-switchers by three age and 12 LE groups. The details of how these moments are constructed are discussed in section II.C. Figures D.1 and D.3 show these moments by age and the targeted LE groups.

Average earnings at age 25.—Finally, we target the average earnings by LE group at age 25. This moment of the data is shown in figure D.4.

B. Identification

Below, we provide an informal discussion of identification of our model. We acknowledge that all parameters are determined jointly within the SMM estimation, because most parameters affect more than one aspect of the data. In this section, our goal is to show that each feature of the model has a pronounced effect on at least one unique moment targeted in the estimation. Namely, there is at least one unique feature of the data that informs each ingredient of the model. This identification discussion also justifies the selected targeted moments presented in the previous section.²⁰

¹⁹ An alternative is to target percentile-based moments (90-10 differential, Kelley's skewness, Moors's kurtosis, etc.), which we have experimented with and with which we have found similar results. We target centralized moments, because they are less costly to compute and do not overlook valuable information in the tails.

²⁰ Following Andrews, Gentzkow, and Shapiro (2017), we compute the sensitivity, $\Lambda = -(G'WG)^{-1}G'W$, of parameter estimates $\hat{\theta}$ to moment conditions, $F(\hat{\theta})$, where G and W

Ex ante worker productivity (α, β).—The concave average life-cycle profile of earnings growth is informative about the average experience profile of worker productivity, driven in the model by the mean of the joint (α, β) distribution and the common quadratic term ζ . The differences in the initial earnings levels of LE groups and their stayer earnings growth (fig. 2C) help us pin down the variance-covariance matrix of the joint distribution of α and β . Note that the distribution of firm productivities also has a first-order effect on the initial earnings dispersion as well as on the earnings growth of job stayers through outside offers. As we discuss next, we use other features of the data to identify the distribution of firm productivities.

Firm productivity distribution.—In the estimation of job-ladder models, identifying the distribution of firm productivities is a key challenge. There are several approaches to estimate this distribution using matched employer-employee data.²¹ For example, Postel-Vinay and Robin (2002), Cahuc, Postel-Vinay, and Robin (2006), and Bagger et al. (2014) use data on firms' value added or profitability to back out the firm distribution. We cannot implement this method, because our dataset does not contain any direct information on value added or profitability. Barlevy (2008) shows that under appropriate conditions the wage gains of job switchers could identify the offer distribution nonparametrically, even in the presence of unobserved worker heterogeneity. Bagger and Lentz (2019) use poaching patterns between firms to rank firms with respect to their productivity. More recently, Bonhomme, Lamadon, and Manresa (2019) use k -means clustering to classify firms into discrete groups.

The key insight for our approach of identifying the firm productivity distribution relies on differences in earnings growth between job stayers and switchers, with stayer growth exhibiting relatively little heterogeneity in the bottom two-thirds of the LE distribution and switchers showing much larger differences throughout the LE distribution (fig. 2C). If there was no job ladder to be climbed (i.e., a degenerate firm distribution), then the average earnings growth of switchers and stayers would look very similar (especially in the upper half of the LE distribution), because they would both be mainly driven by the differences in β . Job-ladder dynamics through the shape of the firm distribution, on the other hand, help the model generate different profiles of earnings growth for stayers and switchers. We confirm this insight by investigating the sensitivity of the firm distribution to switcher earnings growth moments, à la Andrews, Gentzkow, and Shapiro (2017; see app. D.2).

are the Jacobians of $F(\hat{\theta})$ and the weighting matrix, respectively. This formal sensitivity analysis confirms our informal identification strategy below.

²¹ Some papers have used only worker-side data to identify the firm distribution by relying on the distribution of wages coming out of unemployment (e.g., Lise 2013). This approach is not reliable in an environment with worker heterogeneity, as shown by Barlevy (2008).

Heterogeneity in worker flow rates ($\lambda^a(\alpha)$, $\lambda_0^a(\alpha)$, $\lambda_1^a(\alpha)$, λ_r).—Our strategy relies on identifying these flow rates separately for each LE and age group and then linking the LE groups to ex ante worker type α . U-switches, those that involve an earnings loss larger than 25% (fig. 3B), are intimately linked to the job-loss rate δ . Moreover, their frequency is not affected by the rate of EE transitions, because such transitions result in either wage increases or wage losses smaller than 25% and are therefore counted among E-switches.²² Turning to the job-finding rate λ_0 , this rate determines how long a given unemployment spell lasts. Therefore, it has a pronounced effect on the average earnings loss of U-switchers, along with the possible wage decline associated with falling off the job ladder. The latter is determined by the shape of the firm distribution, whose empirical underpinning is discussed above. Finally, the stayer probability is given by a combination of the job-loss rate δ and the offer arrival rate for the employed λ_1 as well as the recall rate λ_r . The key feature that identifies the recall rate is the left skewness of earnings growth for job stayers. In the model, stayer growth distribution is dramatically right skewed in the absence of recalls. Having already identified δ and λ_r , stayer probability can now be used to pin down λ_1 .

Idiosyncratic shocks (ϵ).—They are residuals of earnings growth not explained by the structural features of the model. Our simulations show that the endogenous mechanisms can explain well the earnings dynamics of job switchers. Thus, we use the higher-order moments of earnings changes for job stayers to identify the idiosyncratic risk.

Age dependence in parameters.—Targeted moments identifying the flow rates and the distribution of idiosyncratic shocks have strong age variation in the data (see sec. II), which we exploit to identify the age dependence in these parameters.

C. Estimation Methodology

In this section, we first explain the functional-form assumptions concerning the worker and firm distributions as well as the flow rates. While our identification strategy does not require specific functional forms, these assumptions allow us to have more statistical power and keep the estimation computationally feasible. Next, we describe the SMM objective function, along with the computational method used for estimation.

Functional forms.—The worker fixed effect α is normally distributed, with mean μ_α and standard deviation σ_α ; β is Pareto distributed with shape and scale parameters χ_w and ψ_w , respectively, and is correlated with α by the coefficient $\rho_{\alpha\beta}$. We also estimated a version of our model

²² In our simulations, less than 0.2% of direct job-to-job switches lead to a wage cut larger than 25%.

with Gaussian β and have found that a fat-tailed distribution such as Pareto helps the model better match the very large earnings growth of top-LE groups relative to the median and the relatively smaller differences between the median- and bottom-LE groups.²³ We revisit this choice in the context of estimation results in section V. Firm productivity is also assumed to be Pareto distributed, with shape and scale parameters χ_f and ψ_f , respectively.²⁴ We normalize the scale parameter ψ_f to 1, because one cannot separately identify ψ_f and the mean of the α distribution.

We model the heterogeneity in worker flow rates as a function of worker type α and age. In particular, we use a cubic spline to model unemployment risk, the job-finding rate, and the contact rate as a function of $\alpha_i - \mu_\alpha$ for each age group. We experimented with the number of points for each flow rate and concluded that three points for each age group was flexible enough for job-finding and contact rates, whereas unemployment risk required five points for each age to fit the heterogeneity in the data.

Finally, we assume that the independently and identically distributed idiosyncratic shocks hit only job stayers once a year with some probability $\pi(\alpha)$ (because endogenous mechanisms can explain well the earnings dynamics of job switchers). Innovations are normally distributed, with standard deviation σ_ϵ , and $\pi(\alpha)$ is modeled as a cubic spline separately for each age group.²⁵

SMM objective function.—Let d_n for $n = 1, \dots, N$ denote a generic empirical moment, and let $m_n(\theta)$ be the corresponding model moment that is simulated for a given vector of model parameters, θ . The scales of the moments vary largely; thus, we measure the distance between the data and the simulated moments by arc percentage deviation, $F_n(\theta) = 2[(m_n(\theta) - d_n)/(m_n(\theta) + d_n)]$. Our SMM estimator is then defined by $\hat{\theta} = \arg \min_\theta \mathbf{F}(\theta)' \mathbf{W} \mathbf{F}(\theta)$, where $\mathbf{F}(\theta) = [F_1(\theta), \dots, F_N(\theta)]^T$. The weighting matrix \mathbf{W} reflects our beliefs about the importance of each set of moments in identifying the economic forces behind earnings growth.²⁶ We target a total of 380 moments to estimate 41 parameters.

²³ Polachek, Das, and Thamma-Apiroam (2015) also estimate a fat-tailed distribution of learning ability in a human capital production function. Gabaix et al. (2016) argue that the “high-growth” worker types, as opposed to a random growth mechanism, are key for explaining the rising top income inequality.

²⁴ We have experimented with lognormally distributed firm productivity and found that a Pareto fits the data better. Hubmer (2018) uses a different search model and reaches a similar conclusion.

²⁵ We experimented with AR(1) productivity shocks and found the persistence to be low, around 0.5. We have also considered alternative distributions for innovations, such as those exhibiting skewness and did not find significant improvement in the objective value.

²⁶ The weighting matrix, \mathbf{W} , assigns a 15% weight to the first two sets of moments (i.e., cross-sectional moments of job stayers and switchers), a 30% weight to the third and fourth sets (i.e., the fraction of job stayers, EE-switchers, and EU-switchers and their average wage growth), and a 10% weight to the moments on average earnings at age 25 for each LE group. We chose not to use the optimal weighting matrix because efficiency is not a concern, because

Numerical method for estimation.—We employ a multistart global optimization algorithm by Guvenen and Ozkan (2021). In particular, we generate 15,000 uniform Sobol (quasi-random) points, compute the objective value for each of these, and select the best 1,000 (ranked by the objective value), each of which is used as an initial guess for the local minimization stage. This stage is performed with a mixture of Nelder-Mead’s downhill simplex algorithm and the derivative-free nonlinear least squares algorithm of Zhang, Conn, and Scheinberg (2010). In the end, we pick the best parameter estimates out of 1,000 local minima.

V. Estimation Results

A. Parameter Estimates

We first discuss the key parameter estimates by relating them to the moments that inform them the most. The full set of estimates are presented in appendix D.

1. Distribution of α and β

We start by investigating the heterogeneity in permanent ability α_i and the returns to experience β_i (fig. 5A). Permanent ability α increases almost linearly throughout the LE distribution. Top-LE individuals have an α that is more than 60 log points higher than that of those at the bottom. Moreover, there is a sizable variation within each LE group. The interquartile range (dashed lines in fig. 5A) is around 10 log points. Together with this, the standard deviation of α in the entire population is 0.25. Return to experience, β , also increases with LE—not surprising, given its positive correlation with α of $\rho_{\alpha\beta} = 0.44$ —but with a different shape: β is relatively flatter in the bottom two-thirds and increases steeply toward the top.²⁷ Clearly, this variation of β by LE is dictated largely by the shape of its distribution, which is assumed to be Pareto to match the average earnings growth differences of job stayers (fig. 2C).

The Pareto-distributed β , along with a Pareto firm productivity distribution, implies that earnings distribution exhibits a power law throughout the life cycle, as in the data (fig. A.2). While not targeted in the estimation,

our moments are precisely estimated thanks to the sheer sample size. This also applies to Altonji and Segal’s (1996) identity matrix, which is about minimizing small-sample bias. We have yet experimented with an equal-weighting matrix and come to roughly similar conclusions (available upon request).

²⁷ Note that the interquartile range of β also increases from less than 0.005 at the bottom to more than 0.03 at the top, which is a direct feature of the fat tail of the Pareto distribution. The standard deviation of β in the population is estimated to be 0.017, in line with the estimates in the literature using different methodologies and datasets (Huggett, Ventura, and Yaron 2011; Guvenen et al. 2021).

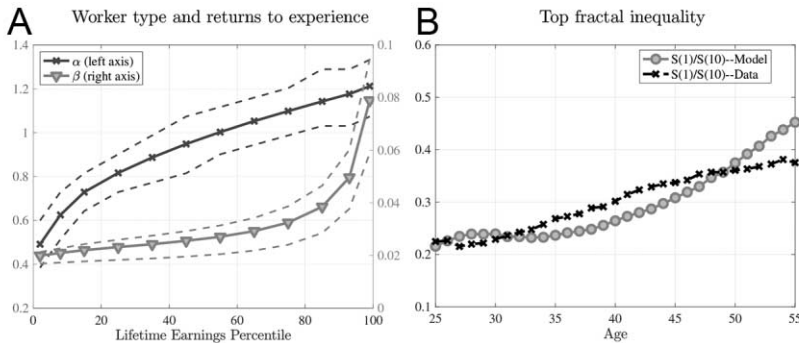


FIG. 5.—Worker types and fractal inequality. *A*, Mean of the distributions of α and β by LE groups. Both distributions have been de-meaned to have mean zero in the overall distribution. *B*, Ratio of income earned by the top 1% earners ($S(1)$) relative to that earned by the top 10% earners ($S(10)$).

the model tracks the relative earnings share of the top 1% in the top 10% fairly well from age 25 to 50, after which the relative share in the model increases faster, driven by the growing importance of the return heterogeneity (fig. 5*B*). Note that typical models of top-income inequality deliver a Pareto distribution through the accumulation of random returns over long periods of time; therefore, log income is exponentially distributed in the entire population (e.g., see Gabaix et al. 2016; Jones and Kim 2018). However, the distribution of log income within each age is Gaussian in the random-growth setting or in a process with normally distributed “growth types” (see Guvenen, Kaplan, and Song 2021, who also argue that several other features of the MEF data are not consistent with this mechanism of top inequality).

2. Human Capital Depreciation

We estimate human capital depreciation to be around 1.5% on a quarterly basis, larger in magnitude than that estimated by Jarosch (2023) using German data. This is not the only channel in our model that contributes to scars from unemployment, which are large and persistent (von Wachter, Song, and Manchester 2009; Krolikowski 2017). An unemployed worker also loses search capital and negotiation rents as well as the forgone opportunity of accumulating experience.

3. Heterogeneity in Flow Rates

Figure 6 plots in three panels how the quarterly unemployment risk, the job-finding rate, and the contact rate vary with calendar age and LE

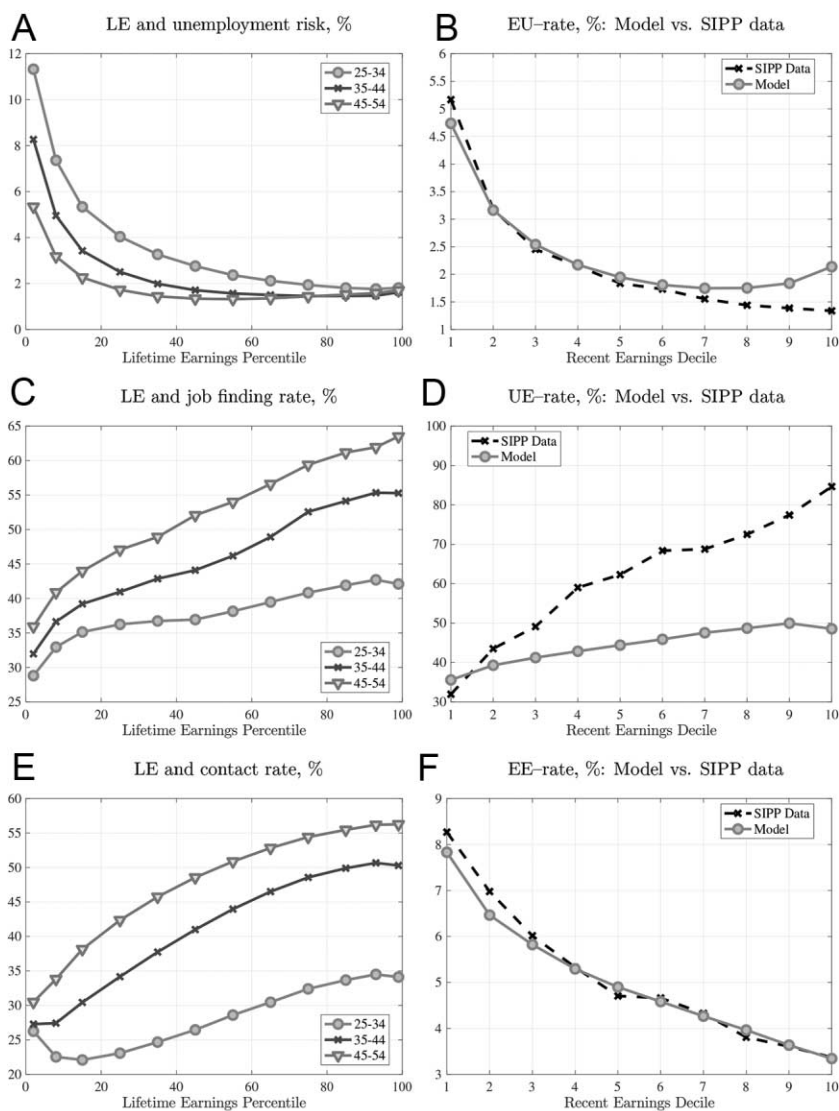


FIG. 6.—Labor market flows.

groups.²⁸ Unemployment risk, $\delta^a(\alpha)$, declines sharply with LE up to the median LE and is essentially flat for individuals above the median (fig. 6A). The job-loss rate for bottom-LE workers is around four times as high as

²⁸ We remind the reader that the ages “old” (o) and “young” (y) in the model do not correspond to the calendar age in our simulations. As a result of the stochastic aging process, there are old (o) workers in the model, even at earlier ages in the simulation.

that for above-median workers. For example, for the youngest age group it declines from around 12% for the bottom earners to less than 3% for median workers. Consistent with previous work, we find the unemployment risk to be significantly higher for younger workers (see Shimer 1998 and Jung and Kuhn 2018). However, the life-cycle variation in unemployment risk is dwarfed by the differences between income groups, which we also observe in the SIPP data (fig. B.1).²⁹ For example, for median workers the job-loss rate declines from around 3% to less than 2% over the life cycle. Furthermore, even though they see a significant decline in their job-loss rate, bottom-LE workers never achieve the job stability above-median workers enjoy.

Given the annual nature of the SSA data, we cannot directly test these estimates. Instead, we investigate how the model fits the evidence on heterogeneity in job-flow rates from the high-frequency SIPP data. SIPP contains monthly observations in overlapping panels with lengths between 2.5 and 4 years. We select a sample of males (aged 25–55) with strong labor force attachment (see details in app. B). We rank them into 10 equally sized deciles within each age group (25–34, 35–44, and 45–55) on the basis of their recent earnings (RE) over the past 3 years. Next, we compute the EU, UE, and EE transition rates for each group over the next 4 months. We also follow the exact same sample construction in the model-generated data. Figure 6B shows how the unemployment risk varies with RE in the SIPP data averaged over the life cycle, along with its model counterpart (for separate age groups, see fig. D.7). While the extent of variation in the data is not explicitly targeted in the estimation, the model captures it remarkably well, except for the top decile, where there is a slight uptick in the model-based EU rate but not in its empirical counterpart.

We estimate the job-finding rate to be increasing with LE and age (fig. 6C). For example, the quarterly job-finding rate increases from around 30% at the bottom for workers aged 25–34 to above 60% at the top for workers aged 45–54. These estimates imply that the youngest bottom-LE workers stay unemployed for around 3 quarters, compared to less than 2 quarters for the oldest top-LE individuals.³⁰ Coupled with an especially high unemployment risk for low-LE workers, these estimates imply large differences in actual experience over the life cycle (fig. 9B). In particular, quarters worked over the working life range from 90 for low-LE individuals to 120 at the top, which then has implications for earnings growth differences that we discuss below. The increasing

²⁹ Jarosch (2023) develops a two-dimensional job-ladder model with jobs also differing in unemployment risk and uses the life-cycle variation in unemployment risk to estimate this model.

³⁰ Although these differences are significant, they are not as large as job-loss heterogeneity. Cairó and Cajner (2017) also show that more educated workers have similar job-finding rates but much lower and less volatile separation rates relative to their less educated peers.

job-finding rate across the income distribution is qualitatively consistent with the evidence from the SIPP (fig. 6D), however, it does not increase as much as in the data.³¹ There is also almost no age variation in the data on job-finding rates, whereas the model estimates are systematically higher for older workers.

We estimate that 12.5% of unemployed workers are recalled by their most recent employer ($\lambda_r = 0.125$). This recall probability is lower than the 40% measured by Fujita and Moscarini (2017) for the United States. They measure recalls directly, using survey data from the SIPP, whereas we infer them indirectly to match the left-hand tail of the earnings growth distribution of job stayers.

Turning to the contact rate for employed workers, we find this to be increasing with LE and age, with a range between 25% and 55% (fig. 6E).³² While the increasing contact rate with LE and age seems contradictory to a declining EE transition rate by RE and age in the SIPP data, the model actually captures both of these patterns well endogenously (figs. 6F, D.7).³³ This is because high-LE and older workers get more offers but work for high-productivity firms that are hard to poach from. Therefore, they reject most of the contacts, whereas low-LE or younger workers make EE-switches more often with fewer offers.

To validate this finding, we analyze data from the SCE, which is a monthly, nationally representative survey of roughly 1,300 individuals that asks respondents about their expectations about various aspects of the economy as well as their employment status, prior work history, and job-search behavior (see Faberman et al. 2022 for more details). Importantly for our purposes, it asks about the number of employer contacts and job offers received. To keep the analysis similar, we take a sample of employed respondents between the ages of 25 and 55 and group them into five bins based on their wages over the past year. We find that contacts received from other potential employers increase in previous wages and are quite high at the top (table 1). People in the highest group (workers above the 95th percentile) are contacted around 0.43 times per month, versus 0.18 contacts per

³¹ Recall that the job-finding rate has a pronounced effect on the average earnings loss of U-switchers. The model can successfully capture its variation along the LE distribution (fig. 8D), with slightly smaller losses for high-LE workers than in the data. Thus, if their job-finding rate were higher, we would have seen even smaller earnings losses for high-LE U-switchers. The failure to capture their higher job-finding rates, however, is not very consequential for our main results, because they are much less likely to be unemployed to begin with.

³² Interestingly, in contrast to the literature that estimates a much smaller offer arrival rate for employed workers than for the unemployed (e.g., Schaal 2017; Jarosch 2023), our estimates for λ_0 and λ_1 are comparable in level, which has implications for unemployment insurance. For example, a high level of λ_1 increases the value of work and reduces moral hazard, thereby increasing the optimal replacement rate (Chetty 2008).

³³ Using Danish data, Bagger et al. (2014) and Lentz, Piyapromdee, and Robin (2018) estimate contact rates to be increasing in worker type too.

TABLE 1
SUBJECTIVE CONTACT RATE

Recent-Earnings Groups	1%–25%	26%–50%	51%–75%	76%–94%	95+%
Total number of contacts	.18	.18	.13	.26	.43
Unsolicited contacts	.09	.02	.04	.11	.43

NOTE.—Respondents are between the ages of 25 and 55. Individuals who report 25 or more contacts in the past 4 weeks are dropped from the sample. We assign zero contacts for those reporting a positive number of contacts but none corresponding with either (i) an employer directly online or through email, (ii) an employer directly through other means, including in-person, or (iii) an employment agency or career center.

month for the lowest quartile, consistent with the underlying mechanism in the model. Moreover, inspecting unsolicited contacts, those that were not initiated by the employee, we find much larger differences. For top earners, contacts are almost five times as likely as contacts for those at the bottom (0.43 vs. 0.09).

4. Idiosyncratic Shocks

We estimate that the probability of experiencing productivity shocks, $\pi(\alpha)$, increases with LE from around 5% to 20%. Recall that their distribution is identified from the second-to-fourth-order moments of the earnings growth for job stayers. Figure 7 shows that the variance increases above the 20th percentile and kurtosis decreases above the 40th percentile of the LE distribution. These patterns require shocks to be more likely for higher-LE workers, whereas for low-LE individuals the earnings dynamics for job stayers are mainly driven by the endogenous job-ladder mechanisms, including recalls. This finding is consistent with evidence from Norway that for high earners large earnings changes are mostly driven by wage changes rather than movements in extensive margin of labor hours (Halvorsen et al. 2019).

B. Model's Fit to the Data

We now show the model's performance in fitting the targeted moments. In doing so, we also discuss the economic forces behind the higher-order moments of earnings changes as well as earnings growth patterns for job stayers and switchers.

1. Cross-Sectional Moments

Figure 7 shows the fit of the model to cross-sectional moments. For clarity of exposition, we suppress the life-cycle variation and plot averages over three age groups. The fit along the life cycle is shown in appendix D.3.

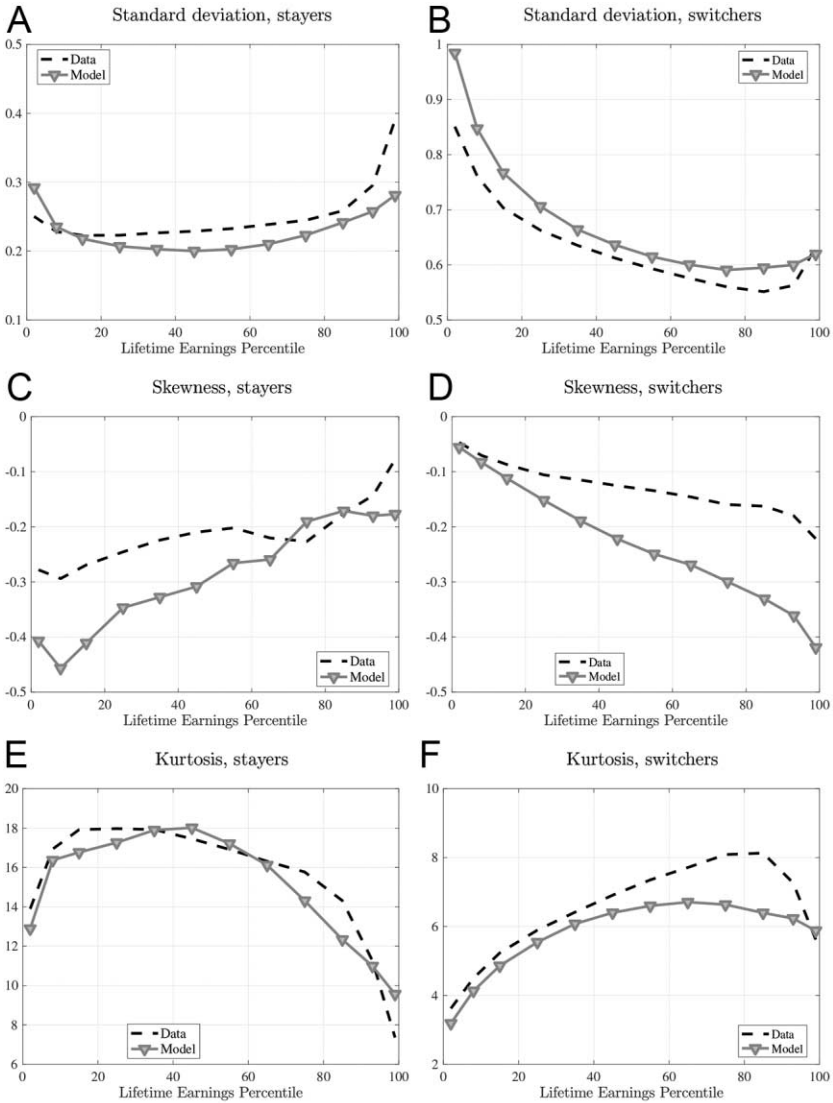


FIG. 7.—Model's fit to cross-sectional moments of $(Y_{t+1} - Y_t)/[(Y_{t+1} + Y_t)/2]$.

The model captures well the standard deviation of earnings changes for job stayers and switchers (fig. 7A). Both in the data and in the model, job switchers have a higher standard deviation throughout the LE distribution. In the model, big changes to earnings happen when people switch jobs because of a job loss. The declining unemployment risk (fig. 6A), combined with an increasing poaching rate (fig. 6E), implies

that a higher share of job switchers at the bottom go through unemployment, as opposed to direct job switches, and explains why the standard deviation is higher at the bottom than in the rest of the distribution. The profile flattens out because there is much less variation in the unemployment risk above the median.

For job stayers, earnings changes are driven by job loss followed by a recall, an outside offer that leads to renegotiation, and idiosyncratic productivity shocks. Because of the high job-loss rates at the bottom, the share of recalls is highest there, which tends to push up the standard deviation at the bottom. As we move to the right along the LE distribution, unemployment risk fades, the prevalence of outside offers increases, and a larger share of such offers result in the worker staying with the same employer and getting a large raise (figs. 6E, 6F). Moreover, idiosyncratic shocks become more frequent and contribute to the increasing standard deviation for job stayers above the median.

Turning to skewness, we find that the model captures well the essential features of the data (figs. 7C, 7D). First, earnings changes are negatively skewed for both job switchers and stayers. For switchers, the negative skewness is mostly a result of flows into unemployment, which result in the worker losing the position on the job ladder and human capital depreciation throughout the spell of unemployment. The decreasing profile of skewness (increasing negative skewness) is a result of two offsetting forces. On the one hand, human capital depreciation is stronger for low-LE individuals because of longer unemployment durations, pushing skewness down at the bottom. On the other hand, job loss is less frequent but more costly for high-LE individuals, because they have more search capital and negotiation rents to lose. The latter force dominates and causes the skewness of earnings changes to be more negative for job switchers among high-LE individuals.

As for job stayers, recalls generate large earnings declines within the same firm. In the absence of recalls, the model cannot generate a negative skewness for job stayers. As we move to the right along the LE distribution, the left-hand tail shrinks as temporary layoffs become less frequent. The right-hand tail expands, because outside offers arrive more often and are more likely to result in wage renegotiation. The two forces combined result in a milder negative skewness for job stayers at higher LE percentiles.

The model is quite successful in matching the extent of kurtosis and its variation over the LE distribution. Infrequent events that lead to large changes, such as outside offers and unemployment spells followed by recalls, are the leading sources of excess kurtosis for job stayers. In fact, they are so strong that without idiosyncratic shocks, earnings changes would be a lot more leptokurtic. The idiosyncratic shocks, despite being leptokurtic themselves, help the model bring down the kurtosis of job

stayers closer to values in the data. Earnings changes of job switchers are also leptokurtic in the model and the data, but to a lesser degree, compared to those for job stayers.

Finally, we investigate the model's fit on cross-sectional moments along the life-cycle dimension. Figure D.2 shows how the higher-order moments of earnings changes for stayers and switchers vary between three age groups. As in the data, life-cycle variation in the model is less pronounced than the variation between LE groups. Overall, we conclude that the model does fairly well in capturing the essential moments of earnings changes for job stayers and switchers across the LE distribution and over the life cycle.

2. Income Growth Moments

Next, we study job stayers and switchers. The model reproduces remarkably well the increasing share of job stayers by LE quantile in the data (fig. 8). There are few job stayers at the bottom because of high flow rates into unemployment. The share of job stayers essentially follows the

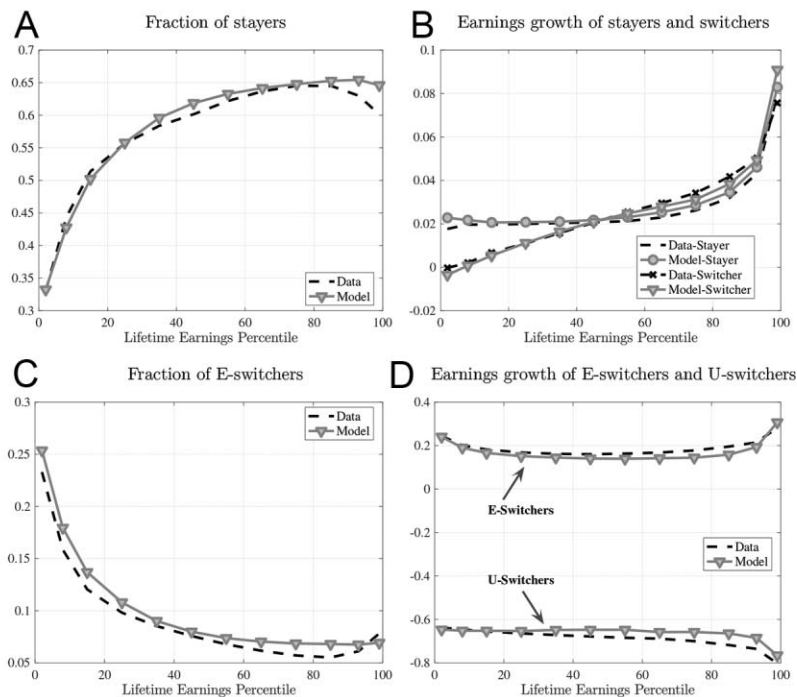


FIG. 8.—Model's fit to income growth moments.

unemployment risk along the LE distribution, increasing up to around the 70th percentile and stabilizing thereafter.

The model also generates overall a realistic average earnings growth for job stayers and switchers throughout the LE distribution (fig. 8B). In particular, there is little heterogeneity among job stayers in the bottom two-thirds of the LE distribution, which, as discussed above, is in part due to the relatively flat average profile of returns to experience (β) in each LE group. Earnings growth of job stayers has a component due to human capital accumulation, governed by β , and a component due to the job ladder, through outside offers that lead to wage increases on the job. As figure 5A shows, the former component is basically flat for two-thirds of the distribution, with a very small positive slope. Yet the earnings growth of stayers in the model is higher at the low end of the distribution, compared to that at the 20th percentile. This feature has to do with the second component, which is stronger at the low end. This result may seem surprising, because bottom-LE individuals have the lowest contact rates when employed. However, given their high unemployment risk, employed workers at the bottom tend to also have a lower piece rate, because they frequently lose their job before they receive many outside offers and can negotiate a better piece rate. A lower piece rate implies that, conditional on the worker staying with the same firm (which is the group we consider in fig. 8B), an outside offer is more likely to lead to wage renegotiation. Thus, there are two competing forces determining the effect of the job-ladder risk at the bottom: a lower contact rate and a higher share of those contacts that lead to wage growth. It turns out that the latter is stronger at the bottom than at the 20th percentile of the LE distribution.

Turning to job switchers, the model captures well their average earnings growth (fig. 8B). In particular, there is a large variation throughout the LE distribution, ranging from zero at the bottom to 9% at the top. Moreover, consistent with the data, most of this heterogeneity is due to compositional differences among job switchers. The share of E-switchers among all workers declines from 25% to around 5% over the LE distribution (fig. 8C). However, their share among only switchers increases sharply, from around 65% at the bottom of the LE distribution to above 80%. These shares are slightly below those in the data but capture remarkably well the variation along the LE dimension. Finally, consistent with the data, there is much less between-group heterogeneity in the earnings growth of E-switchers and U-switchers (fig. 8D).

Recall that being a job switcher or stayer has limited effects on earnings growth of top-LE workers younger than 44, but it matters quite a bit in the oldest age group (fig. 4). This is both because U-switches become more likely as they get older and, more importantly, because average earnings growth for U-switches falls sharply for the oldest top-LE workers (fig. A.8). Our model can capture this feature of the data (fig. D.3), and

thus we are able to investigate it further: top-LE workers have already climbed to the top of the job ladder in the oldest age group, so they are less likely to make voluntary E-switches. Therefore, the relative likelihood of U-switches increases dramatically for them. Furthermore, since they are at the higher end of the job ladder, if they make a U-switch (e.g., lose a job), their earnings decline sharply because they lose search capital (a job in a high-productivity firm), negotiation rents, and human capital.

Thus, we conclude that the estimated job-ladder model captures quite well the key features of the careers of individuals in different parts of the LE distribution.

VI. Decomposing LE Inequality

The model matches well the distribution of LE (fig. 9A). At the top of the distribution, LE_{50} earns around 4.19 times as much as LE_{45} in the model, slightly overstating the data (3.83). The fit is much better below LE_{45} : LE_{45} earns 1.97 times as much as LE_{25} in the model, compared to 1.94 in the data. Moreover, the ratio of LE_5 to LE_1 is 1.80 in the model, slightly below its empirical counterpart of 1.92 (table 2).

A. Earnings Differences: Wages versus Employment

To what extent are these large differences in LE driven by differences in wages, as opposed to differences in employment rates over the life cycle? Figure 9A plots the earnings and wage differences in the model by normalizing the median group to 1. This figure shows that wage—rather

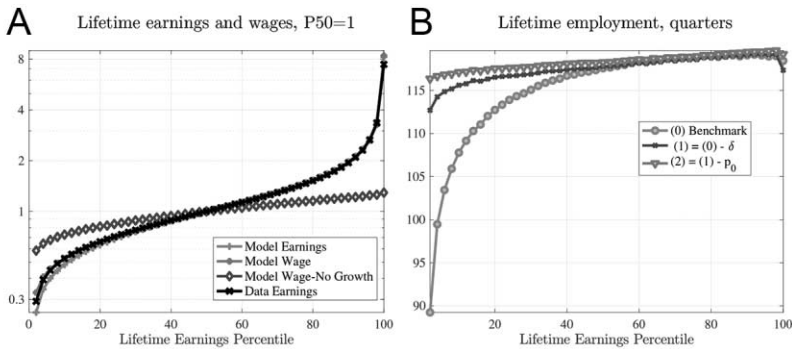


FIG. 9.—LE, wages, and employment. “Model Wage–No Growth” corresponds to an experiment that shuts down the heterogeneity in β , eliminates search frictions ($\delta = 0$, $\lambda_0 = 1$, $\lambda_1 = 0$), removes idiosyncratic shocks, and makes the firm distribution degenerate. In this specification, the only source of wage (and earnings) differences is permanent ability, α . Each series is normalized so that it takes a value of 1 for the median group.

TABLE 2
LE DIFFERENCES ACROSS LE GROUPS

	LE_{50}/LE_{45}	LE_{45}/LE_{25}	LE_{25}/LE_5	LE_5/LE_1
Data	3.83	1.94	1.90	1.80
Model	4.19	1.97	2.07	1.92

NOTE.— LE_i/LE_j is the ratio of average LE of the individuals in LE_i to that of individuals in LE_j .

than employment—differences explain the vast majority of LE inequality. Differences in average wages over the life cycle are remarkably similar to the LE differences, except below the 25th percentile, where differences in employment (measured as the number of quarters worked over the working life) play an important role. For example, employment of workers at the bottom of the LE distribution is about 25% lower than that of the median workers. Employment differences above the median are negligible in comparison (fig. 9B).

Before investigating the sources of lifetime wage differences, we briefly discuss the sources of employment differences below the median. These differences arise as a result of ex ante heterogeneity in unemployment risk and job-finding rates as well as the ex post job-ladder risk, that is, ex ante similar workers experiencing different job-loss and job-finding shocks. To measure their relative roles, we first shut down ex ante heterogeneity in job-loss risk by endowing all individuals with $\delta^a(0)$, the job-loss risk of workers with $\alpha_i - \mu_\alpha = 0$, which is roughly the average value for median-LE workers (fig. 5A), and then compute the resulting distribution of total lifetime employment. In doing so (and in all experiments that follow), we keep the rankings of workers, and thus the composition of LE groups, unchanged from the baseline. Therefore, the differences between this experiment and the baseline are due only to the differences in ex ante job-loss risk, δ .

We find that employment differences between the bottom- and top-LE workers decline sharply, from around 25% to 7%, when all workers have the same job-loss rate (fig. 9B). When we further eliminate differences in job-finding rates by setting $\lambda_0^a(\alpha)$ to $\lambda_0^a(0)$ for all workers, employment differences decline further, albeit to a smaller extent, with bottom-LE individuals working only 3% less than those at the top. The remaining differences are entirely due to the ex post realizations, that is, luck. Our estimation thus attributes little role to luck in generating sizable lifetime employment differences. One caveat is that in our model unemployment does not beget future unemployment, as in Jarosch (2023).

Life-cycle wage differentials: the roles of initial conditions versus wage growth.—Figure 9A shows that when all sources of wage growth have been turned off and only the differences in permanent ability α are allowed for, the model generates a wage inequality that is an order of magnitude

smaller. Wage differentials are largely shaped by wage growth heterogeneity rather than by the initial differences in levels. Therefore, it is essential to understand why some workers have a much steeper wage profile than others.

B. *Decomposing Life-Cycle Wage Growth Differences*

Recall that in the model wage growth can differ across individuals as a result of differences in the ability to accumulate human capital, ex ante and ex post differences in unemployment risk, and the quality and quantity of offers on and off the job. To assess the relative roles of these factors in lifetime wage growth differences, we shut down each component one after another until we eliminate all differences, again keeping the composition of the LE groups the same as in our benchmark (fig. 10A).³⁴

We start by eliminating the differences in unemployment risk, which we accomplish by shutting off differences in job-loss and job-finding rates together ($\delta^a(\alpha_i) = \delta^a(0)$ and $\lambda_0^a(\alpha_i) = \lambda_0^a(0)$). Heterogeneity in unemployment risk has a marked effect on wage growth differences between bottom- and median-LE workers, and to a lesser degree those above the median (series 1 in fig. 10A). Specifically, slightly more than 50% of wage growth differences between the bottom- and median-LE workers would diminish if the workers at the bottom had job-loss and job-finding rates similar to those of the median-LE workers.³⁵

High unemployment rates of low-income individuals (fig. 9B) not only prevent them from accumulating human capital but also lead to depreciation in human capital during unemployment. Furthermore, a higher incidence of unemployment prevents bottom-LE workers from climbing the job ladder. Figure 10B shows the contributions of human capital, search capital, and negotiation rents to the wage growth differences between bottom- and median-LE workers.³⁶ Differences in human capital accumulation account for almost 70% of the wage growth differences between these two groups. And lower human capital accumulation at the bottom of the LE distribution is not because these workers have much lower returns to experience but because they accumulate less experience

³⁴ We decompose LE growth into wage and hours growth components (fig. D.8A). Earnings growth is less dispersed than wage growth between LE groups, especially in the bottom half of the LE distribution, because of the higher lifetime employment growth of bottom-LE workers.

³⁵ Recall that the model cannot capture the higher job-finding rates for higher-LE workers seen in the data (fig. 6D). Even though differences in $\lambda_0(\alpha)$ are less important than those in $\delta(\alpha)$ for our results, larger variation in $\lambda_0(\alpha)$ would have surely generated stronger role for heterogeneity in job-ladder risk in the bottom half of the LE distribution.

³⁶ The growth in search capital is measured as the log change in firm productivity $\mathbb{E}[p_{j(i,55)} - p_{j(i,25)}]$, and the growth in negotiation rents is measured as the change in the log piece rate $\mathbb{E}[r_{i,55} - r_{i,25}]$.

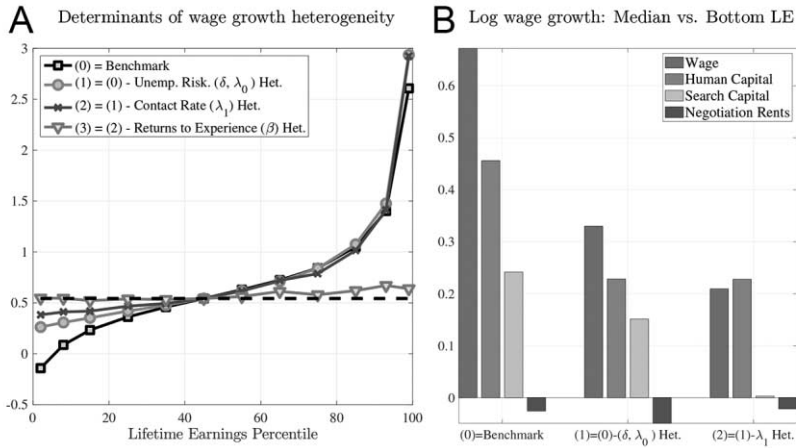


FIG. 10.—Decomposing wage growth between the ages of 25 and 55. *A*, Decomposition of the differences in lifetime wage growth. The solid black line with squares shows the benchmark wage growth; the lines with circles, crosses, and triangles cumulatively eliminate ex ante heterogeneity (Het.) in job-loss and job-finding rates, contact rate, and returns to experience, respectively. The dashed line shows the lifetime wage growth after all ex ante heterogeneity is eliminated and shocks are shut down. *B*, Decomposition of the average lifetime log wage growth differences between bottom- and median-LE workers into human capital, search capital, and negotiation rents.

as a result of their higher prevalence of unemployment spells. The remaining difference in wage growth between bottom- and median-LE workers is essentially due to the accumulation of search capital (i.e., working for more productive firms). The contribution of the negotiation capital is very small and negative, meaning that workers at the bottom experience larger growth in their piece rate compared to those at the median. This is because higher-LE workers are employed at more productive firms, which are hard to poach from, and, in turn, increases in wages due to outside offers are smaller for them.³⁷ Eliminating unemployment risk brings down the differences in human and search capital accumulation, and thus differences in wage growth, by around 65%.

Job-loss and job-finding differences matter much less in the upper half of the distribution, because these workers have fairly low unemployment risk to begin with. The only exception is the top earners, who have a slightly higher job-loss risk; thus, eliminating this difference would actually raise their income growth further, by around 20 log points.

³⁷ As we discussed above, to maintain tractability for the model we assume that the value of the recall option is not considered in the wage-bargaining process. If we had allowed for it, it would have contributed to wage growth differences through the piece rate. However, as we show here, the contribution of the negotiation capital over the life cycle is small relative to the human capital and search capital components. So, the role of the recall option in the wage-bargaining process would be quantitatively small.

Recall that LE groups also display sizable differences in their contact rates λ_1 . In contrast to the job-loss and job-finding rates, these differences have a smaller effect on lifetime wage growth differences (fig. 10A). Specifically, eliminating differences in contact rates would close an additional 20% of the wage growth gap between the bottom and the median, with essentially no effect at the top. All of this effect is due to the closing of search capital differences. Namely, endowing bottom-LE individuals with the contact rate of median-LE individuals allows bottom-LE workers to climb to better jobs. These two experiments show that eliminating differences in job-ladder risk can go a long way in ameliorating the labor market experiences of bottom-LE workers and eliminate more than 70% of the differences in wage growth with median-LE workers.

Next, we turn to the role of heterogeneity in returns to experience. To this end, we assign all workers' β to the average, after which we are left with only idiosyncratic productivity shocks and the random realizations of labor market shocks (fig. 10A). Recall that β is relatively flatter in the bottom two-thirds and increases steeply toward the top of the LE distribution (fig. 5A). Thus, eliminating differences in returns to experience has an effect across the entire LE distribution, but the largest effect by far is on the top LE earners. Together with the fact that job-ladder risk plays a smaller role for top earners, the reason why top earners experience a much larger wage growth than the median is primarily because they have much higher returns to experience, thereby accumulating human capital at a higher pace. In the bottom half of the LE distribution, β heterogeneity can explain only around 25% of wage growth differences, and the primary source of heterogeneity is in job-ladder risk.

Intuition behind the quantitative results.—Which feature of the data tells the model that human capital accumulation is more important in the upper half of the LE distribution and vice versa in the bottom half? While all targeted moments are informative, we argue that the differences between the income growth of job stayers and that of switchers are key. To see this, note that human capital is capitalized into wages in all firms. Therefore, wage growth always reflects a worker's human capital accumulation, regardless of whether he stays with the current employer or switches to a new one. If the data show high wage growth for a group of workers relative to median workers regardless of job switching, as is the case in the data for higher-LE individuals (fig. 2C), the model infers a high return to experience.

The difference in earnings growth between stayers and switchers is informative about the role of job-ladder risk. If a group of workers experience lower growth when switching than they do when they stay with the same employer, the model rationalizes this by inferring a poor job ladder due to a high job-loss or a low job-finding rate. At the bottom of the LE distribution, job switchers experience much smaller earnings growth

compared to stayers, consistent with our finding that the job-ladder component is more important for explaining their lackluster lifetime growth relative to the median.

Initial conditions versus shocks.—Our findings also shed new light on the relative roles of ex ante heterogeneity in initial conditions and ex post shocks (Keane and Wolpin 1997; Huggett, Ventura, and Yaron 2011). Figure 10A shows a negligible role for luck—idiosyncratic productivity and job-ladder shocks—in lifetime wage growth differences after all ex ante heterogeneity is removed. On average, above-median-LE workers are somewhat more lucky, but this has a very small quantitative effect. However, figure 10A shows differences between LE groups and averages out possible within-group differences from ex post shocks.

For a more precise measure, we decompose the variance of log LE into “ex ante heterogeneity” and “ex post shocks” components. Of the variation in LE, 81% is due to ex ante heterogeneity in initial conditions, substantially higher than the 61% that Huggett, Ventura, and Yaron (2011) find from a similar exercise. Their Ben-Porath model does not feature job-ladder dynamics and attributes income changes due to job-ladder shocks (e.g., job loss) to productivity shocks. Through the lens of their model, bottom-LE workers are just unlucky and often draw negative productivity shocks. In our model, they have high ex ante unemployment risk. Thus, the higher role for initial conditions in our model stems from richer worker heterogeneity, which more precisely captures the source of inequality in the bottom half.

VII. Conclusion

We investigated the determinants of LE inequality by focusing on heterogeneity in job-ladder dynamics and on-the-job learning. Empirically, we showed that (i) lower-LE workers switch jobs more often, mainly driven by higher nonemployment, (ii) earnings growth for job stayers is similar in the bottom two-thirds of the LE distribution, while (iii) rising strongly with LE for job switchers. Estimating a model featuring rich worker and firm heterogeneity, we found large differences in ex ante job-ladder risk across workers. These differences account for 75% of the lifetime wage growth differential below the median of the LE distribution. Above the median, almost all lifetime wage growth differences are a result of Pareto-distributed learning ability.

We conclude that different economic forces are driving the inequality in different parts of the LE distribution. These differences have important implications for the design of insurance policies. For example, Golosov, Maziero, and Menzio (2013) show that optimal redistribution looks very different when differences in labor income emanate from search frictions, as opposed to differences in workers’ productivity. Similarly, the effects of

monetary policy can be heterogeneous, because it may affect workers' job-ladder risk differently across the income distribution, thereby leading to differences in wage growth.

An emerging literature studies the effects of firms' power in setting wages. Firms can hire and retain workers at wages lower than the competitive fringe if they are large in a market (Jarosch, Nimczik, and Sorkin 2019; Berger, Herkenhoff, and Mongey 2022) or if they do not face much competition from other employers, because of either contractual restrictions on job mobility (Johnson, Lavetti, and Lipsitz 2019) or other frictions. One interpretation of the estimated differences in outside contacts is about employers' ability to restrict poaching. Through this lens, our results suggest that firms are better able to restrict poaching for low-skill workers and have more power over them. This interpretation is consistent with Caldwell and Danieli (2018), who find much less competitive pressure for low-skill workers.

Finally, our analysis has focused on worker differences. However, some of the differences could be a characteristic of jobs rather than of workers. Jarosch (2023) focuses on firm heterogeneity in job stability. More broadly, firms might contribute to wage growth heterogeneity by providing different learning environments (e.g., Herkenhoff et al. 2018; Gregory 2019). To fully understand the role of firms and workers in wage dynamics, a unified approach is necessary, which we leave for future work (e.g., Koffi et al. 2022).

Data Availability

We provide a replication kit that generates all the results in this article in Ozkan, Song, and Karahan (2023) in the Harvard Dataverse, <https://doi.org/10.7910/DVN/FLAFLQ>. Some of the results are based on data from the SSA, which does not allow us to share the micro data or the intermediate files obtained from it. To facilitate replication for researchers with access to the MEF at the SSA, we provide all the Stata programs that process the micro data and the Matlab programs that produce the results in the paper. For the results from the SIPP, we provide the compiled data as well as the Stata code that processes the SIPP panels from scratch. Finally, we provide all files that are necessary to estimate and simulate our structural model, which includes the Fortran source code. For further details, please see Ozkan, Song, and Karahan (2023).

References

- Ábrahám, Árpád, Fernando Álvarez-Parra, and Susanne Forstner. 2016. "The Effects of Moral Hazard on Wage Inequality in a Frictional Labor Market." Working paper.

- Ahn, Hie Joo, and James D. Hamilton. 2020. "Heterogeneity and Unemployment Dynamics." *J. Bus. and Econ. Statist.* 38 (3): 554–69.
- Ahn, Hie Joo, Bart Hobijn, and Ayşegül Şahin. 2023. "The Dual U.S. Labor Market Uncovered." Working Paper no. 31241 (May), NBER, Cambridge, MA.
- Altonji, Joseph G., and Lewis M. Segal. 1996. "Small-Sample Bias in GMM Estimation of Covariance Structures." *J. Bus. and Econ. Statist.* 14 (3): 353–66.
- Andrews, Isaiah, Matthew Gentzkow, and Jesse M. Shapiro. 2017. "Measuring the Sensitivity of Parameter Estimates to Estimation Moments." *Q.J.E.* 132 (4): 1553–92.
- Atkinson, Anthony B., Thomas Piketty, and Emmanuel Saez. 2011. "Top Incomes in the Long Run of History." *J. Econ. Literature* 49 (1): 3–71.
- Bagger, Jesper, François Fontaine, Fabien Postel-Vinay, and Jean-Marc Robin. 2014. "Tenure, Experience, Human Capital, and Wages: A Tractable Equilibrium Search Model of Wage Dynamics." *A.E.R.* 104 (6): 1551–96.
- Bagger, Jesper, and Rasmus Lentz. 2019. "An Empirical Model of Wage Dispersion with Sorting." *Rev. Econ. Studies* 86 (1): 153–90.
- Barlevy, Gadi. 2008. "Identification of Search Models Using Record Statistics." *Rev. Econ. Studies* 75 (1): 29–64.
- Ben-Porath, Yoram. 1967. "The Production of Human Capital and the Life Cycle of Earnings." *J.P.E.* 75 (4): 352–65.
- Berger, David, Kyle Herkenhoff, and Simon Mongey. 2022. "Labor Market Power." *A.E.R.* 112 (4): 1147–93.
- Blanchard, Olivier J. 1985. "Debt, Deficits, and Finite Horizons." *J.P.E.* 93 (2): 223–47.
- Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa. 2019. "A Distributional Framework for Matched Employer Employee Data." *Econometrica* 87 (3): 699–739.
- Bowlus, Audra J., and Huju Liu. 2013. "The Contributions of Search and Human Capital to Earnings Growth over the Life Cycle." *European Econ. Rev.* 64:305–31.
- Burdett, Kenneth, and Dale T. Mortensen. 1998. "Wage Differentials, Employer Size, and Unemployment." *Internat. Econ. Rev.* 39 (2): 257–73.
- Cahuc, Pierre, Fabien Postel-Vinay, and Jean-Marc Robin. 2006. "Wage Bargaining with On-the-Job Search: Theory and Evidence." *Econometrica* 74 (2): 323–64.
- Cairó, Isabel, and Tomaz Cajner. 2017. "Human Capital and Unemployment Dynamics: Why More Educated Workers Enjoy Greater Employment Stability." *Econ. J.* 128 (609): 652–82.
- Caldwell, Sydnee, and Oren Danieli. 2018. "Outside Options in the Labor Market." Working paper.
- Caucutt, Elizabeth M., and Lance Lochner. 2020. "Early and Late Human Capital Investments, Borrowing Constraints, and the Family." *J.P.E.* 128 (3): 1065–147.
- Chetty, Raj. 2008. "Moral Hazard versus Liquidity and Optimal Unemployment Insurance." *J.P.E.* 116 (2): 173–234.
- Deaton, Angus, and Christina Paxson. 1994. "Intertemporal Choice and Inequality." *J.P.E.* 102 (3): 437–67.
- Faberman, R. Jason, Andreas I. Mueller, Ayşegül Şahin, and Giorgio Topa. 2022. "Job Search Behavior among the Employed and Non-employed." *Econometrica* 90 (4): 1743–79.
- Fujita, Shigeru, and Giuseppe Moscarini. 2017. "Recall and Unemployment." *A.E.R.* 107 (12): 3875–916.
- Gabaix, Xavier, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2016. "The Dynamics of Inequality." *Econometrica* 84 (6): 2071–111.
- Gathmann, Christina, and Uta Schönberg. 2010. "How General Is Human Capital? A Task-Based Approach." *J. Labor Econ.* 28 (1): 1–49.

- Golosov, Mikhail, Pricila Maziero, and Guido Menzio. 2013. "Taxation and Redistribution of Residual Income Inequality." *J.P.E.* 121 (6): 1160–204.
- Gregory, Victoria. 2019. "Firms as Learning Environments: Implications for Earnings Dynamics and Job Search." Working paper, Fed. Reserve Bank, St. Louis.
- Gregory, Victoria, Guido Menzio, and David G. Wiczer. 2021. "The Alpha Beta Gamma of the Labor Market." Working Paper no. 28663 (April), NBER, Cambridge, MA.
- Güvenen, Fatih, Greg Kaplan, and Jae Song. 2021. "The Glass Ceiling and the Paper Floor: Changing Gender Composition of Top Earners since the 1980s." *NBER Macroeconomics Ann.* 35:309–73.
- Güvenen, Fatih, Greg Kaplan, Jae Song, and J. Weidner. 2017. "Lifetime Incomes in the United States over Six Decades." Working Paper no. 23371 (April), NBER, Cambridge, MA.
- Güvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song. 2021. "What Do Data on Millions of U.S. Workers Reveal about Lifecycle Earnings Dynamics?" *Econometrica* 89 (5): 2303–39.
- Güvenen, Fatih, and Serdar Ozkan. 2021. "TikTak: A Multistart Global Optimization Algorithm." <https://github.com/serdarozkan/TikTak>.
- Güvenen, Fatih, Serdar Ozkan, and Jae Song. 2014. "The Nature of Countercyclical Income Risk." *J.P.E.* 122 (3): 621–60.
- Hall, Robert E., and Marianna Kudlyak. 2022. "Churn and Stability: The Remarkable Heterogeneity of Flows among Employment, Job Search, and Non-Market Activities in the US Population." Working paper.
- Haltiwanger, John, Henry R. Hyatt, Erika McEntarfer, Liliana Sousa, and Stephen Tibbets. 2014. "Firm Age and Size in the Longitudinal Employer-Household Dynamics Data." Working Paper no. CES-WP-14-16, Center for Econ. Studies, US Census Bur., Washington, DC.
- Halvorsen, Elin, Hans Holtzer, Serdar Ozkan, and Kjetil Storesletten. 2019. "Dissecting Idiosyncratic Income Risk: Evidence from Norwegian Registry Data." Working paper.
- Heckman, James J. 1976. "A Life-Cycle Model of Earnings, Learning, and Consumption." *J.P.E.* 84 (4, pt. 2): S9–S44.
- Herkenhoff, Kyle, Jeremy Lise, Guido Menzio, and Gordon M. Phillips. 2018. "Production and Learning in Teams." Working Paper no. 25179 (October), NBER, Cambridge, MA.
- Hubmer, Joachim. 2018. "The Job Ladder and Its Implications for Earnings Risk." *Rev. Econ. Dynamics* 29:172–94.
- Huggett, Mark, and Greg Kaplan. 2016. "How Large Is the Stock Component of Human Capital?" *Rev. Econ. Dynamics* 22:21–51.
- Huggett, Mark, Gustavo Ventura, and Amir Yaron. 2011. "Sources of Lifetime Inequality." *A.E.R.* 101 (7): 2923–54.
- Jarosch, Gregor. 2023. "Searching for Job Security and the Consequences of Job Loss." *Econometrica* 91 (3): 903–42.
- Jarosch, Gregor, Jan Sebastian Nimczik, and Isaac Sorkin. 2019. "Granular Search, Market Structure, and Wages." Working Paper no. 26239 (September), NBER, Cambridge, MA.
- Jarosch, Gregor, and Laura Pilossoph. 2019. "Statistical Discrimination and Duration Dependence in the Job Finding Rate." *Rev. Econ. Studies* 86 (4): 1631–65.
- Johnson, Matthew S., Kurt Lavetti, and Michael Lipsitz. 2019. "The Labor Market Effects of Legal Restrictions on Worker Mobility." Working paper. <https://doi.org/10.2139/ssrn.3455381>.

- Jolivet, Gregory, Fabien Postel-Vinay, and Jean-Marc Robin. 2006. "The Empirical Content of the Job Search Model: Labor Mobility and Wage Distributions in Europe and the U. S." In *Structural Models of Wage and Employment Dynamics*, edited by Henning Bunzel, Bent Christensen, George R. Neumann, and Jean-Marc Robin, 269–308. Bingley, UK: Emerald.
- Jones, Charles I., and Jihee Kim. 2018. "A Schumpeterian Model of Top Income Inequality." *J.P.E.* 126 (5): 1785–826.
- Jovanovic, Boyan. 1979. "Job Matching and the Theory of Turnover." *J.P.E.* 87 (5): 972–90.
- Jung, Philip, and Moritz Kuhn. 2018. "Earnings Losses and Labor Mobility over the Life Cycle." *J. European Econ. Assoc.* 17 (3): 678–724.
- Kaplan, Greg. 2012. "Inequality and the Life Cycle." *Quantitative Econ.* 3 (3): 471–525.
- Keane, Michael P., and Kenneth I. Wolpin. 1997. "The Career Decisions of Young Men." *J.P.E.* 105 (3): 473–522.
- Koffi, Marlène, Serdar Ozkan, Sergio Salgado, and Marco Weißler. 2022. "High Risk Workers and High Risk Firms." Working paper.
- Krolukowski, Pawel. 2017. "Job Ladders and Earnings of Displaced Workers." *American Econ. J. Macroeconomics* 9 (2): 1–31.
- Lazear, Edward P., and Sherwin Rosen. 1981. "Rank-Order Tournaments as Optimum Labor Contracts." *J.P.E.* 89 (5): 841–64.
- Lentz, Rasmus, Suphanit Piyapromdee, and Jean-Marc Robin. 2018. "On Worker and Firm Heterogeneity in Wages and Employment Mobility: Evidence from Danish Register Data." PIER Discussion Paper no. 91, Puey Ungphakorn Inst. Econ. Res., Bangkok.
- Lise, Jeremy. 2013. "On the Job Search and Precautionary Savings." *Rev. Econ. Studies* 80 (3): 1086–113.
- Lise, Jeremy, Costas Meghir, and Jean-Marc Robin. 2016. "Matching, Sorting and Wages." *Rev. Econ. Dynamics* 19:63–87.
- Mincer, Jacob. 1974. *Schooling, Experience, and Earnings*. New York: Columbia Univ. Press (for NBER).
- Morchio, Iacopo. 2020. "Work Histories and Lifetime Unemployment." *Internat. Econ. Rev.* 61 (1): 321–50.
- Neal, Derek, and Sherwin Rosen. 2000. "Theories of the Distribution of Earnings." In *Handbook of Income Distribution*, vol. 1, edited by Anthony B. Atkinson and François Bourguignon, 379–427. Amsterdam: North-Holland.
- Olsen, Anya, and Russell E. Hudson. 2009. "Social Security Administration's Master Earnings File: Background Information." *Social Security Bull.* 69 (3): 29–46.
- Ozkan, Serdar, Jae Song, and Fatih Karahan. 2023. "Replication Kit for: 'Anatomy of Lifetime Earnings Inequality: Heterogeneity in Job Ladder Risk vs. Human Capital.'" Harvard Dataverse: <https://doi.org/10.7910/DVN/FLAFLQ>.
- Panis, Conastantijn, Roald Euler, Cynthia Grant, Melissa Bradley, Christine E. Peterson, Randall Hirscher, and Paul Steinberg 2000. *SSA Program Data User's Manual*. Project Memorandum PM-973-SSA. Baltimore: Social Security Admin.
- Pastorino, Elena. 2019. "Careers in Firms: The Role of Learning and Human Capital." Working paper.
- Piketty, Thomas, and Emmanuel Saez. 2003. "Income Inequality in the United States: 1913–1998." *Q.J.E.* 118 (1): 1–39.
- Polachek, Soloman W., Tirthatanmoy Das, and Rewat Thamma-Apiroam. 2015. "Micro- and Macroeconomic Implications of Heterogeneity in the Production of Human Capital." *J.P.E.* 123 (6): 1410–55.
- Postel-Vinay, Fabien, and Jean-Marc Robin. 2002. "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity." *Econometrica* 70 (6): 2295–350.

- Rubinstein, Ariel. 1982. "Perfect Equilibrium in a Bargaining Model." *Econometrica* 50 (1): 97–109.
- Schaal, Edouard. 2017. "Uncertainty and Unemployment." *Econometrica* 85 (6): 1675–721.
- Shimer, Robert. 1998. "Why Is the U.S. Unemployment Rate So Much Lower?" *NBER Macroeconomics Ann.* 13:11–61.
- Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter. 2019. "Firming Up Inequality." *Q.J.E.* 134 (1): 1–50.
- Sorkin, Isaac. 2018. "Ranking Firms Using Revealed Preference." *Q.J.E.* 133 (3): 1331–93.
- Tanaka, Satoshi, Lawrence Warren, and David Wiczer. 2023. "Earnings Growth, Job Flows and Churn." *J. Monetary Econ.* 135:86–98.
- Topel, Robert H., and Michael P. Ward. 1992. "Job Mobility and the Careers of Young Men." *Q.J.E.* 107 (2): 439–79.
- van den Berg, Gerard J. 1990. "Search Behaviour, Transitions to Non-participation and the Duration of Unemployment." *Econ. J.* 100 (402), 842–65.
- von Wachter, Till, Jae Song, and Joyce 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." Manuscript, Columbia Univ.
- Zhang, Hongchao, Andrew R. Conn, and Katya Scheinberg. 2010. "A Derivative-Free Algorithm for Least-Squares Minimization." *SIAM J. Optimization* 20 (6): 3555–76.