



The job ladder and its implications for earnings risk [☆]

Joachim Hubmer

Department of Economics, Yale University, 28 Hillhouse Avenue, New Haven, CT 06511, USA



ARTICLE INFO

Article history:

Received 13 September 2016

Received in revised form 10 January 2018

Available online xxxx

JEL classification:

E24

J24

J31

J64

Keywords:

Earnings dynamics

Life-cycle earnings risk

Non-Gaussian shocks

On-the-job search

ABSTRACT

This paper analyzes the ability of a job ladder framework to explain recent evidence on life-cycle earnings dynamics. Using administrative data, [Guvenen et al. \(2015\)](#) document several new facts about the distribution of earnings growth, most notably large negative skewness and high excess kurtosis, rejecting the frequently used log-normal framework. I show that these new facts can be well explained by a standard structural representation of a frictional labor market, a life-cycle version of the job ladder model, in combination with a simple human capital process. Furthermore, I identify endogenous search effort, risk aversion and wealth accumulation, and skill loss in unemployment as key model features that interact with the labor market friction to jointly reconcile the evidence.

© 2018 Elsevier Inc. All rights reserved.

1. Introduction

Incorporating heterogeneous agents has been one of the key advances in macroeconomic equilibrium modeling in the past decades. Idiosyncratic earnings risk, which is quantitatively significant and to a large extent uninsurable, is the major ingredient in these models. Capturing the salient features of earnings risk is therefore essential. While previously earnings dynamics have been mainly described by parsimonious log-normal earnings processes estimated from survey data, recent evidence based on administrative data reveals several new facts that contradict the assumption of log-normality.¹ In particular, [Guvenen et al. \(2015\)](#) utilize a rich dataset from the U.S. Social Security Administration to document that earnings growth exhibits large negative skewness and very high kurtosis.² Furthermore, they find sizable variation in these higher order moments across different population groups defined by age and the level of earnings. Negative skewness means that negative changes are typically larger than positive changes. Excess kurtosis implies that most individuals experience rather small (or no) earnings changes, while a few are hit by very large shocks. Both features are in stark contrast to the log-normal framework, which implies symmetric changes and rather mild risk.

[☆] I am grateful to Giuseppe Moscarini and Tony Smith for invaluable guidance and support. I thank Fatih Guvenen, Christopher Huckfeldt, Per Krusell, Aleh Tsyvinski, and seminar participants at the 2016 SED Meeting in Toulouse as well as at Yale for helpful comments and suggestions. I thank in particular the editor, Gianluca Violante, and an anonymous referee for their generous input, which greatly improved the paper.

E-mail address: joachim.hubmer@yale.edu.

¹ Administrative data has several advantages over survey data, including the absence of measurement error and often much larger sample sizes, thus allowing for estimating higher order moments and in general applying nonparametric methods.

² Earlier work on earnings dynamics going beyond the log-normal framework includes [Geweke and Keane \(2000\)](#) and [Bonhomme and Robin \(2010\)](#).

This paper aims to show that a standard structural representation of a frictional labor market, the job ladder model, can explain these new facts. Moreover, it sheds light on which model features have to be added to the basic job ladder model outlined by [Burdett \(1978\)](#) for this purpose. From a statistical point of view, these new facts can also be modeled by error component models, in the tradition of a rich literature going back to [MaCurdy \(1982\)](#) and [Abowd and Card \(1989\)](#), among many others. It is, however, from an economist's point of view more satisfactory to explore to what extent a behavioral model can account for the new evidence. The job ladder framework is a natural starting point as its defining element is expected to generate negative skewness and excess kurtosis: In the absence of a separation shock, workers only switch to more productive firms, moving up the ladder relatively slowly. A layoff implies falling down the ladder all the way, hence on-the-job search implies negative skewness. Moreover, both job-to-job transitions as well as movements between employment and unemployment are potentially large and infrequent changes, thus generating excess kurtosis.

In the proposed life-cycle model, heterogeneous and risk averse workers search for job opportunities at heterogeneous firms while unemployed as well as on-the-job. To smooth consumption and partially insure against labor market shocks, they have access to a risk-free asset. Workers differ by age, wealth, and human capital, modeled as a Gaussian random walk. Firms are characterized by their productivity level. Crucially, I do not calibrate the model to match the evidence on earnings dynamics. Instead, I target aggregate labor market transition rates and the earnings distribution in levels, plus the average long-term earnings loss associated with displacement from a job. In turn, I evaluate the performance of the model against the earnings dynamics moments documented in [Guvenen et al. \(2015\)](#). Overall, the model is successful in explaining the magnitude of negative skewness and excess kurtosis. Moreover, it captures that negative skewness and kurtosis are increasing in the level of earnings and over the life-cycle.

I identify three key model features that are essential to fit the data by way of shutting each one of them down separately, re-estimating the respective alternative model, and comparing the resulting earnings dynamics to the benchmark model. The first one is endogenous search. In the benchmark model, search is costly and agents optimally choose search intensity. It turns out that if contact rates are exogenous (and uniform across agents), then the model generates much less negative skewness as long-term unemployment (and thus very large negative shocks) becomes extremely rare. Furthermore, the model is no longer able to replicate that negative skewness is increasing in age. Primarily, this pattern holds in the benchmark model because older agents optimally search less. Thus, they re-climb the ladder more slowly after a separation shock.

The second key feature is risk aversion and wealth accumulation. If workers are risk neutral, I find that both negative skewness as well as kurtosis decline substantially. While the risk neutral model is explaining the data well for the bottom 40% of the earnings distribution, it fails to match the increase in negative skewness and kurtosis for higher earners. In the benchmark version, there is a wealth effect: higher earners tend to be richer, thus they optimally search at lower intensity, which prolongs non-employment spells and thus generates larger negative earnings changes.

The third important ingredient concerns the dynamics of human capital. The benchmark model assumes that the growth rate of human capital depends on the employment status of a worker. In particular, human capital depreciates in unemployment in the spirit of [Ljungqvist and Sargent \(1998\)](#). The size of this decline is disciplined by targeting the average earnings loss, in present value terms, associated with displacement from a job as estimated by [Davis and von Wachter \(2011\)](#). If the growth rate of human capital is required to be the same for employed and unemployed workers, the model can no longer generate large enough earnings losses, while simultaneously matching aggregate labor market transition rates. Consequently, negative earnings changes are smaller and skewness is less negative. Similarly, as the largest changes are negative, kurtosis decreases.

There are some aspects of the [Guvenen et al. \(2015\)](#) data that the benchmark model is missing. First, the magnitude of earnings fluctuations, as measured by the standard deviation of earnings changes, is too low. At least to some extent, this is to be expected as the benchmark model excludes some sources of transitory fluctuations such as multiple jobholding, varying hours worked, and bonus payments. Structurally modeling all those would be daunting.³ Second, the model is creating too much negative skewness and excess kurtosis for the highest earners, those in the top 10% and especially even farther out in the right tail. This failure of the model is related to two features of the benchmark model that are otherwise crucial in explaining earnings dynamics for the bottom 90% of earners: risk aversion and wealth accumulation, as well as endogenous search effort. Very high earners tend to be very rich, and thus in the model they have very little incentive to search for a new job once hit by a separation shock, creating very large negative earnings changes. There is, perhaps, an analogy to the literature on explaining wealth inequality that struggles with explaining wealth concentration at the very top. In that literature, the standard behavioral model has trouble explaining why the very rich are not consuming more of their wealth. Here, the model has trouble explaining why the very rich are not searching less, and working less. One could imagine that a taste for work is at play for top athletes and artists who enjoy performing, and

³ One can, of course, enrich the dynamics of human capital by adding a transitory component that statistically accounts for these missing elements. This avenue is pursued in an extended model presented in [Appendix A](#). Relative to the benchmark, this model generates more earnings fluctuations. As the added exogenous transitory shock is symmetric, skewness is increasing (i.e., less negative), thus providing a better fit to the data (the baseline model is overshooting). The kurtosis of earnings growth is not affected materially.

likewise for top managers who like exerting power.⁴ Third, depending on the classification of workers into job stayers and switchers, the benchmark model is not creating any, respectively not enough, excess kurtosis for stayers. This is because in the model, stayers' earnings are proportional to human capital, which is evolving smoothly over time. Introducing lumpy adjustment of earnings to human capital, job promotions and demotions, or recall unemployment are potential solutions.

While I argue that the proposed structural model provides a reasonably good fit to new evidence on earnings dynamics, one can surely provide a better statistical representation by fitting an exogenous earnings process (an avenue that is successfully pursued by [Guvenen et al., 2015](#)). The latter approach, however, runs the risk of conflating genuine earnings risk with earnings changes that are the outcome of choices, thus overstating the welfare cost of labor market risk. I illustrate this in the last part of this paper, where I use the calibrated model to simulate counterfactual earnings paths, shutting down some of the agents' choices. In particular, I find that while the standard deviation of log earnings changes is not changing dramatically, the higher order moments are shrinking drastically (by 30%–90%) in two counterfactuals: In the first, all agents search at full intensity and accept every job; thus, earnings dynamics are fully exogenous. In the second, agents are still risk averse but do not have access to the risk-free asset to smooth consumption and partially insure against idiosyncratic shocks. In both cases, longer non-employment spells are extremely rare, and so are very large (negative) shocks.

This paper is related to [Postel-Vinay and Turon \(2010\)](#), who also explore the ability of a structural job search model to replicate salient features of earnings dynamics. Relative to that paper, the focus on recent evidence on especially the higher order moments of earnings growth is novel. [Karahan et al. \(2016\)](#) offer a complimentary contribution that focuses especially on explaining heterogeneity in lifetime earnings growth.

The proposed model follows [Lise \(2013\)](#), who develops a job ladder model with risk averse agents that optimally choose search effort and savings. While Lise analyzes optimal behavior analytically in more detail, the present paper adds several elements of heterogeneity to the model. In particular, human capital dynamics as well as a non-trivial reservation wage decision are added, and the life-cycle is explicitly modeled. These features help the model in reconciling the data, at the cost of losing analytical tractability. Joint modeling of human capital accumulation and labor market frictions is shared with [Bagger et al. \(2014\)](#), see also the references therein for earlier contributions.

The rest of the paper is organized as follows. The baseline model is outlined in Section 2, its calibration explained in Section 3 and the main results presented in Section 4. Section 5 discusses the findings and especially key model features that help explain the data, mainly by comparing the baseline to alternative versions of the model. Section 6 contains counterfactuals that shut down agents' choices, thus using the model to distinguish genuine earnings risk from earnings changes caused by choices. Finally, Section 7 concludes.

2. Model

In this section, I describe the model economy. The framework builds on [Lise \(2013\)](#), featuring risk averse workers that search for job opportunities at heterogeneous firms both in non-employment as well as on-the-job and can smooth consumption by means of a risk-free asset. While structural search models typically assume risk neutrality (or risk aversion without any access to financial markets), this paper aims to provide a framework for studying the causes and consequences of earnings risk, connecting to the macroeconomic literature on incomplete markets in the tradition of [Aiyagari \(1994\)](#).⁵ Hence, in the benchmark version workers are risk averse and have limited access to financial markets. On top of that, worker heterogeneity is added as the model is calibrated to match certain features of the cross-sectional earnings distribution (in levels). Moreover, the distribution of earnings growth will be evaluated against [Guvenen et al. \(2015\)](#)'s estimates, who do not condition on fixed worker characteristics.

The economy is populated by a continuum of finitely-lived risk averse agents. An agent can work in period $t = 1, \dots, T$, followed by retirement. Once retired, agents face a constant probability of death. A newborn worker starts his working life non-employed in $t = 1$ with an initial endowment of assets a and productive human capital h . There is no aggregate uncertainty and the model is not set in an equilibrium framework. Specifically, the asset market is not required to clear and firm ownership is not specified. These simplifications do not appear to be crucial for explaining individual earnings dynamics.

2.1. Non-employment

A non-employed worker receives unemployment benefits $b(h)$ and chooses (i) savings $a' \geq \underline{a}$ as well as (ii) search effort $s \in [0, 1]$. Note that unemployment benefits $b(h)$ depend only on current human capital h , but not explicitly on past wages, in order to simplify the problem. Consumption $c = (1 + r)a + b(h) - a'$ yields CRRA utility $u(c)$ and search

⁴ See e.g. [Carroll \(2000\)](#), who proposes a "Capitalist Spirit" model to explain high saving rates at the top. Another candidate explanation for this failure to match earnings dynamics at the very top end of the earnings distribution concerns the assumption of uniform, exogenous, separation rates. [Jarosch \(2015\)](#), among others, finds that higher paying jobs are more stable. However, I have experimented with adding heterogeneous separation rates to the model and found that the fit of the model does not improve.

⁵ A notable exception is [Krusell et al. \(2010\)](#), who introduce labor market frictions into an incomplete markets model.

entails a convex search cost $\chi_s \frac{s^{1+\phi}}{1+\phi}$ for some $\phi > 0$, entering additively. At the end of the period, a non-employed worker's human capital is updated to $h' \sim F_u(\cdot|h)$, following a Markovian process. Subsequently, the agent samples a job offer with probability $\lambda_w s$, where $\lambda_w \in (0, 1)$ is the probability of finding a job when searching full-time. Job offers are characterized by the one-dimensional productivity of a firm p , drawn from some distribution F_p . Matching with firms is random. The worker can choose between accepting the job and starting period $t + 1$ working at firm p , or staying in non-employment.⁶

2.2. Employment

An agent that starts period $t \in \{1, \dots, T\}$ employed at a p -firm first collects a wage $w(h, p) = \gamma hp$ for some fixed constant $\gamma \in (0, 1]$. In the background, the worker-firm match produces output hp and a fraction γ of this rent accrues to the worker. Working (full-time, as there is no intensive margin of labor supply) entails a leisure penalty $\chi_e \geq 0$ as well as higher search costs $(\chi_s + \chi_{se}) \frac{s^{1+\phi}}{1+\phi}$, where χ_{se} parametrizes the additional dis-utility of searching when working. Note that [Lise \(2013\)](#) does not feature a dis-utility of working ($\chi_e = 0$) and also $\chi_{se} = 0$. Consequently, workers accept any job (of course, if and only if it pays more than unemployment benefits), which has the desirable implication of ensuring concavity of the value function. Allowing for $\chi_{se} > 0$ enables the model to match both aggregate unemployment-to-employment as well as job-to-job transition rates. Moreover, if search were costly but there were no dis-utility of working (i.e., $\chi_e = 0$), then reservation wages would be decreasing in wealth above a certain threshold.⁷ At the end of t , h is updated to $h' \sim F_e(\cdot|h)$, i.e., human capital can potentially evolve differentially in employment and non-employment. Subsequently, she samples a new job offer with probability $\lambda_w s$ and independently her job is destroyed with probability $\delta \in (0, 1)$. Consequently, in the best case she can choose between staying at the current employer, accepting the new job offer, or going to non-employment. At worst, she is forced into non-employment.

2.3. Value functions

In recursive structure, denote the value of being employed in period t as $W^t(a, h, p)$, the value of being non-employed as $U^t(a, h)$ and the maximum of the two as $V^t(a, h, p) = \max\{W^t(a, h, p), U^t(a, h)\}$. Then

$$U^t(a, h) = \max_{a' \geq a, s \in [0, 1]} \left\{ u((1+r)a + b(h) - a') - \chi_s \frac{s^{1+\phi}}{1+\phi} + \beta \int_{h'} \left[(1 - s\lambda_w)U^{t+1}(a', h') + s\lambda_w \int_{p'} V^{t+1}(a', h', p') dF_p(p') \right] dF_u(h'|h) \right\}, \tag{1}$$

where current period utility is derived from consumption and search effort. The two terms in the continuation value reflect having to stay in non-employment, and drawing a job offer, respectively.

The value of employment is given by

$$W^t(a, h, p) = \max_{a' \geq a, s \in [0, 1]} \left\{ u((1+r)a + w(h, p) - a') - \chi_e - (\chi_s + \chi_{se}) \frac{s^{1+\phi}}{1+\phi} + \beta \int_{h'} \left[(1 - \delta)(1 - s\lambda_w)V^{t+1}(a', h', p) + (1 - \delta)s\lambda_w \int_{p'} V^{t+1}(a', h', \max\{p, p'\}) dF_p(p') + \delta(1 - s\lambda_w)U^{t+1}(a', h') + \delta s\lambda_w \int_{p'} V^{t+1}(a', h', p') dF_p(p') \right] dF_e(h'|h) \right\}. \tag{2}$$

The continuation value in (2) sums over four terms, reflecting all possible combinations of whether or not a separation shock hits, and whether or not a new job offer arrives (these events are assumed to be independent).⁸

Reservation wages differ from unemployment benefits for a number of reasons: On the one hand, the dis-utility of working, and search being more costly on-the-job, both increase reservation wages.⁹ The former impetus is stronger for rich agents, as consumption utility is concave. Thus, reservation wages increase in wealth. The latter creates an option value of

⁶ Since job search is a continuous variable in this setup, the notation does not differentiate between unemployment and out-of-labor-force.

⁷ To see this, note that a very rich agent would choose a very low search intensity. Consequently, if offered any job that pays more than unemployment benefits, she would accept it, anticipating that she would search very little in the future, and thus not draw better job offers. Note that this is not true with $\chi_e > 0$.

⁸ In the second of these terms, $V^{t+1}(a', h', \max\{p, p'\})$ already incorporates the trivial choice between two firms.

⁹ As a result, the value functions are generally not globally concave. [Appendix B](#) discusses the computational approach, in particular how the non-concavities are dealt with.

non-employment that vanishes as agents approach retirement. On the other hand, to the extent that human capital evolves more favorably for employed workers, this creates a force that depresses reservation wages. Again, this effect is stronger for young agents. Theoretically, the relation between reservation wages and age is thus ambiguous.

Optimal search intensity as governed by similar forces. Richer agents search less due to diminishing marginal utility of consumption. Furthermore, search effort decreases with age. Search is a costly investment into future earnings; its payoff diminishes as retirement approaches.

3. Calibration

In this section, the calibration of the benchmark model is described. I target the cross-sectional distribution of earnings (in levels) over the life-cycle as well as aggregate labor market stocks and flows. Moreover, I explicitly target estimates of residual wage dispersion (interpreted as frictional in the model) and the present value of earnings losses following displacement. To generate an empirically plausible amount of the latter, I use differential human capital evolution on and off the job in the spirit of [Ljungqvist and Sargent \(1998\)](#). Crucially, the distribution of earnings growth is not targeted, and thus used to evaluate the performance of the model.

3.1. Preliminaries

The model is solved at monthly frequency. The discount factor $\beta = 0.95^{\frac{1}{12}}$ and the interest rate $r = 0.75(\frac{1}{\beta} - 1)$ are set to standard values. Borrowing is not allowed ($\underline{a} = 0$) and $u(c) = \log(c)$. Unemployment benefits are given by $b(h) = \min\{\chi_b h, \bar{b}\}$, where the cap \bar{b} is set to half of mean earnings in the economy and χ_b is chosen to yield a mean replacement rate of 0.4.¹⁰ Thus, unemployment benefits do not directly depend on past wages in order to simplify the problem. The convexity of the search cost function is set to $\phi = 0.2$, following an estimate by [Lise \(2013\)](#) for a comparable job ladder model with risk averse workers. Retired workers receive a lumpsum retirement benefit, equaling a fraction 0.35 of mean earnings.¹¹

3.2. Labor market stocks and flows

As the model-generated earnings dynamics will be evaluated against the measurements in [Guvenen et al. \(2015\)](#), it is important to mirror their sample, consisting of all US males aged 25 to 60 who have ever been issued a Social Security number. In particular, this includes individuals that are (temporarily) out of the labor force. In the model, there is no discrete distinction between unemployment and out of the labor force as search effort is a continuous choice. To begin with, the job offer probability for full search effort λ_w is set to 0.43 and jobs are exogenously separated with probability $\delta = 0.03$.¹² In turn, the internal calibration will target the non-employment (NE) rate of prime-age males, averaged over the sample period 1978–2010 and equaling 12.4%, as well as the aggregate job to job (J2J) transition rate of employed workers. The transition rates from unemployment, respectively out of the labor force, to employment are not explicitly targeted. Given that over the sample period 38% of all non-employed males were unemployed (4.7% of the full sample), we can, however, order the non-employed in the model according to their search intensity and think of the 38% with highest search intensity as representing the unemployed (and the remaining 62% as being out of the labor force). Then, the monthly model generated transition rates from unemployment, respectively out-of-labor-force, to employment are 0.38, respectively 0.10.¹³

3.3. Human capital

Initial human capital h_0 is log-normally distributed, with standard deviation σ_{h_0} . The dynamics of human capital over the life-cycle follow a simple Gaussian unit root process. The drift term is allowed to depend on an agent's employment status. Furthermore, a reflecting lower boundary \underline{h} is imposed to prevent human capital from depreciating without bounds:

$$\log(h_{t+1}) = \begin{cases} \max\{\log(\underline{h}), \log(h_t) + \mu_e + \sigma_h \epsilon_{t+1}\} & \text{in employment,} \\ \max\{\log(\underline{h}), \log(h_t) + \mu_u + \sigma_h \epsilon_{t+1}\} & \text{in non-employment,} \end{cases}$$

¹⁰ For the replacement rate, see [Shimer \(2005\)](#) and [Hornstein et al. \(2011\)](#); for the cap see [Department of Labor \(2008\)](#).

¹¹ The average benefit amount is based on [Social Security Administration \(2015\)](#).

¹² These are standard values used in the literature ([Shimer, 2012](#); [Hornstein et al., 2011](#)). Note that the model features in addition endogenous quits into non-employment, which would suggest that it is necessary to calibrate δ internally. On the other hand, if a workers that is hit by an exogenous job destruction shock draws a new job offer simultaneously, he can accept it, generating a job-to-job transition. These two effects are quantitatively small and offsetting; the employment exit probability equals 0.0297 in the model, very close to the target of 0.03. Hence, I abstain from targeting this moment in the internal calibration.

¹³ The low convexity of the search cost function implies that non-employed workers typically search either very little or at full intensity $s = 1$. Thus, the unemployment-to-employment transition rate is only slightly below λ_w , justifying the choice of $\lambda_w = 0.43$, which corresponds to the observed job finding probability in the data.

Table 1
Jointly calibrated parameters and targeted moments.

Moment	Target	Model	Estimate
NE rate	0.124 ^a	0.123	$\chi_s = 1.03$ ($= \chi_e(1 + \phi)$)
J2J transition rate	0.027 ^b	0.027	$\chi_{se} = 0.57$
Variance of log earnings at 25	0.59 ^c	0.59	$\sigma_{h_0} = 0.89$
Variance of log earnings at 60	1.05 ^c	1.05	$\sigma_h = 0.037$
Mean life-cycle earnings growth	0.82 ^c	0.82	$\mu_e = 0.0035$
20Y earnings loss from U-shock	0.12 ^d	0.12	$\mu_u = -0.035$
Mean–min ratio residual wages	1.75 ^e	1.75	$\eta = 3.87$

^a Non-employment rate, US males aged 25–54, averaged 1978–2010 (Organisation for Economic Co-operation and Development, 2016).

^b Average of estimates ranging from 0.022–0.032 discussed in Hornstein et al. (2011).

^c Guvenen et al. (2015).

^d Davis and von Wachter (2011).

^e Hornstein et al. (2007).

where ϵ_{t+1} is an i.i.d. normal shock. In particular, if $\mu_u < \mu_e$, then human capital depreciates on average in non-employment. The floor is chosen such that a worker with human capital \underline{h} , employed at the lowest productivity firm ($p = 1$), earns half the minimum wage.¹⁴

3.4. Internal calibration

Table 1 displays the seven targeted moments and associated seven jointly calibrated parameters. Intuitively, the level of search costs χ_s and the additional search cost for employed workers χ_{se} control the fraction of non-employed workers and the aggregate J2J transition rate (of course, given the separation rate δ and all other parameters). To economize on the number of estimated parameters, I impose that $\chi_e = \frac{\chi_s}{1+\phi}$, i.e., the dis-utility of searching full-time, for a non-employed agent, equals the one for working full-time. Note that searching on the job has to be roughly 1.5 times as costly as searching in non-employment to account for the fact that the J2J transition rate is an order of magnitude smaller than the overall NE2E transition rate. While not targeted explicitly, the latter equals 0.21 in the model.

The calibration also matches the variance of log annual earnings at age 25 and age 60, as well as the average life-cycle earnings growth. These values are adopted from Guvenen et al. (2015) in order to render the subsequent analysis of earnings risk comparable. In the model simulation, the agents start their working life at age 24 in non-employment; the first twelve months are used as a burn-in period. The standard deviation of initial human capital σ_{h_0} is tightly linked to the cross-sectional earnings variability at age 25. In turn, the standard deviation of the innovation term to human capital σ_h affects the variance of earnings at age 60.

Life-cycle earnings growth results mostly from accumulating human capital on the job, which happens at rate μ_e . While I target only average earnings growth, Guvenen et al. (2015) document that earnings growth increases dramatically in the level of life-time earnings. As Fig. 1 reveals, the model captures this relation well for the bottom 90% earners, even though there is no ex-ante heterogeneity in human capital growth. Instead, the positive relation between earnings levels and earnings growth in the model is driven by the pure random walk component as well as differential human capital evolution on the job, relative to non-employment.¹⁵ For the top 10% earners, the model fails to match the magnitude of earnings growth over the life-cycle, indicating that a richer form of heterogeneity in human capital growth, or a more involved job ladder (e.g. the top rungs being open to high-skilled workers only), would be needed to account for this particular feature of the data.

Davis and von Wachter (2011) find that average earnings losses associated with an exogenous job destruction event amount to 12% of counterfactual earnings over a 20 year horizon, in present value terms. The model is able to match this average loss primarily by allowing for human capital depreciation in unemployment, i.e., $\mu_u < \mu_e$.¹⁶ Fig. 2 illustrates this by decomposing the losses. The dash-dotted line at the bottom reports earnings losses resulting from the employment channel only. That is, it assumes that when exogenously separated workers regain employment, they do so at the firm they worked at previously, and moreover their human capital does not depreciate. Consequently, earnings losses are virtually zero, except for the year of the layoff. The dashed line in the middle adds the job ladder effect, which generates some persistence. When

¹⁴ The minimum wage is computed as a fraction 0.31 of the average wage, this number is the average over 1978–2010 for the U.S., see Organisation for Economic Co-operation and Development (2017). Guvenen et al. (2015) are also using a threshold for annual earnings that corresponds to three month of employment at half the minimum wage.

¹⁵ This positive relation is primarily mechanical, in the sense that it remains almost as steep when shutting down all choices, so that the resulting dynamics of earnings are completely exogenous.

¹⁶ Several other mechanisms have recently been proposed to account for the cost of job loss: Jarosch (2015) proposes heterogeneous separation rates, based on the observation that hours recover slowly in Germany. In the U.S. labor market, however, hours have been shown to recover rapidly, see Altonji et al. (2013). Relatedly, Jung and Kuhn (2017) emphasize job stability at the top. Huckfeldt (2016) points out that losses are concentrated among workers who (have to) switch to lower-paying occupations.

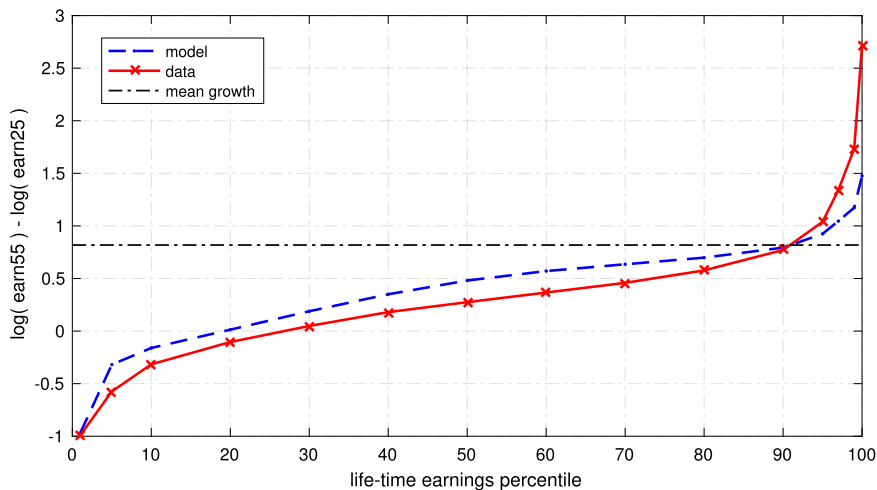


Fig. 1. Life-cycle earnings growth by life-time earnings percentile.

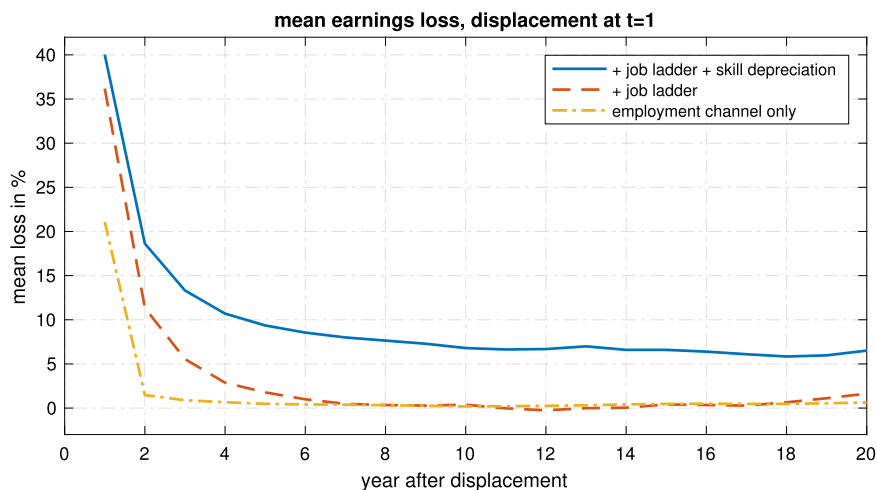


Fig. 2. Decomposition of earnings losses from job displacement.

an unemployed worker regains employment, the productivity of his new employer will on average be lower than the one of the firm he has been employed at before being hit by the displacement shock. Finally, the solid line on top shows the total earnings loss, including skill depreciation. Five years after the shock, persistence of earnings losses is solely resulting from permanent loss of human capital while non-employed.

Finally, I assume that the job offer distribution (over firm productivity p) is a Pareto distribution, with shape coefficient η . Wages are proportional to hp , the product of human capital and firm productivity. I define residual wage dispersion in the model as the dispersion in earnings conditional on age and human capital, i.e., the dispersion in p , and calibrate η to match the mean–min ratio of residual wages, following [Hornstein et al. \(2011\)](#).

4. Quantitative results: earnings dynamics

The model was calibrated to match aggregate labor market statistics as well as the earnings distribution in levels. Earnings risk was not explicitly targeted. Hence, we can now meaningfully assess the performance of the model in matching the evidence on, especially higher order, earnings dynamics as documented in [Guvenen et al. \(2015\)](#). After visually inspecting the distribution of earnings growth, the model is evaluated by comparing the standard deviation, skewness and kurtosis generated by it to the respective data moments.

Fig. 3 displays the histogram of log annual earnings changes in the model and compares it to the data as well as to a normal distribution with the same mean and standard deviation. The left panel displays changes over a one year horizon, the right shows five year changes. The model goes a long way in replicating the peakedness of the data. Peakedness and thick tails are the characteristic visual expressions of high kurtosis: while most workers' earnings change very little (or not




Fig. 3. Histogram of log earnings changes.




Fig. 4. Log density of log earnings changes.

at all), few experience dramatic changes. In the model, the workers that stay at the same firm are the ones with relatively stable earnings. As their human capital evolves too, their earnings are not constant, but the risk they are exposed to is rather mild. In contrast, those agents that become unemployed, regain employment, or switch jobs are exposed to heavy risk.

Fig. 4 shows the same evidence as Fig. 3, the only difference being that the natural logarithm of the densities is taken. As such, it highlights the tail behavior of earnings growth. While it is well known that the earnings distribution – in levels – has a Pareto tail, Guvenen et al. (2015) document that earnings growth exhibits linear tails too in logarithmic scale. I.e., the distribution of earnings growth has a double Pareto tail.¹⁷

The model successfully replicates the left tail, both year-to-year and over a five year horizon. We observe that the left is fatter than the right tail, indicating negative skewness. The model fails to fully match the thickness of the right tail. This failure is related to not explicitly targeting the Pareto tail of the earnings distribution in levels (only the cross-sectional variance of earnings over the life-cycle is matched by construction). Table 2 shows top earnings shares. While the model does reasonably well in replicating them, the fact that ratios of top shares (e.g. the ratio of the top 0.01% share to the top 0.1% share) are larger in the data indicates that the asymptotic Pareto tail of the earnings distribution is not quite as thick

¹⁷ To see this, consider a random variable X that is Pareto distributed with shape coefficient η . The pdf of X is then given by $f_X(x) = \frac{\eta}{x^{\eta+1}}$. In this setting, X is earnings growth, i.e., $X = \frac{\$ \text{ earnings in } t+1}{\$ \text{ earnings in } t}$. Consider the right tail of X , which we claim is Pareto distributed. Fig. 3 displays the histogram of log earnings changes, i.e., the pdf of $Y = \log(X)$. Assuming that X is Pareto distributed, we obtain the pdf of Y by a simple change of variables: $f_Y(y) = \exp(y) f_X(\exp(y)) = \exp(y) \frac{\eta}{\exp(y)^{\eta+1}} = \frac{\eta}{\exp(\eta y)}$. Then, Fig. 4 is plotting the same data, but with the vertical axis now in logs. Taking the logarithm of the density of log earnings changes, we derive $\log(f_Y(y)) = \log(\eta) - \eta y$, i.e., a linear relationship between log earnings changes and the log density of log earnings changes. Hence, the (absolute value) of the slope in Fig. 4 is the tail index of earnings growth.

Table 2
Top earnings shares.

	Top 10%	Top 1%	Top 0.1%	Top 0.01%
Data	32.10	9.29	3.16	1.14
Model	43.95	12.78	3.06	0.64

Shares displayed in %. Data: updated top wage income series by [Piketty and Saez \(2003\)](#), average over the sample period 1978–2010.

Table 3
Main results: earnings dynamics in the baseline model.

	1 year change			5 year change		
	Std. dev.	Skewness	Kurtosis	Std. dev.	Skewness	Kurtosis
Data	0.722	−0.977	16.120	1.112	−1.033	7.700
Baseline model	0.440	−1.431	13.568	0.789	−1.278	7.814
Relative difference	0.391	0.465	0.158	0.290	0.237	0.015
Mean error	0.391	0.469	0.263	0.290	0.264	0.176

Data moments are from [Guvenen et al. \(2015\)](#). Each moment is computed separately for each age by recent earnings percentile cell, and population weighted averages are displayed in the first two rows. There are two age groups, 25–35 and 35–55, and thirteen recent earnings groups, with cut-offs at percentiles 1, 10, 20, 30, 40, 50, 60, 70, 80, 90, 95 and 99. Let m_i^k for $k \in \{data, model\}$ be moment m for cell i , with \bar{m}^k being the respective average (over all cells). Then the relative difference is defined as $\frac{|\bar{m}^{data} - \bar{m}^{model}|}{|\bar{m}^{data}|}$, while the mean error takes into account heterogeneity across different groups: it is defined as $\frac{\sum_i n_i |m_i^{data} - m_i^{model}|}{(\sum_i n_i) |\bar{m}^{data}|}$, where n_i is the size of group i . By construction, the mean error is weakly greater than the relative difference, with equality if and only if all deviations are of the same sign.

in the model. In a job ladder model, as long as the firm productivity distribution has a thicker tail than the distribution of worker productivity, the tails of the earnings distribution in levels and growth rates are asymptotically of equal size.¹⁸

4.1. Parametric results

Having established that the model appears to replicate the overall log earnings change distribution in a satisfactory manner, we now look at the standard deviation, skewness and kurtosis in more detail. For a random variable X with mean μ and standard deviation σ , skewness (the third standardized moment) is defined as $\mathbb{E} \left[\left(\frac{X - \mu}{\sigma} \right)^3 \right]$ and kurtosis (the fourth standardized moment) as $\mathbb{E} \left[\left(\frac{X - \mu}{\sigma} \right)^4 \right]$. Here X is the log change in annual earnings, which are computed as deviations from a deterministic age profile. Besides variation over the life-cycle, [Guvenen et al. \(2015\)](#) analyze the distribution of log earnings changes between age t and $t + k$, for $k = 1, 5$, as a function of recent earnings. Recent earnings are defined as the average over earnings in $t - 5, t - 4, \dots, t - 1$. I include very low (or zero) earnings observations and replace them by a positive lower bound, consistent with the construction of the data moments used.¹⁹

[Table 3](#) summarizes the results. Each moment is first computed separately for each age group by recent earnings percentile cell, and subsequently the average is displayed, both for the data and the model. Two different measures of fit are used: the third line reports the relative difference between the data average and the model average. The last line reports the mean error across age by recent earnings cells. For example, the five year kurtosis is almost perfectly replicated by the model on average, but this masks that the model is overstating the kurtosis for certain population groups and understating it for others. The benchmark model is understating the size of deviations, as measured by the standard deviation, while actually generating too much negative skewness and a comparable amount of excess kurtosis. The fit is better over a longer horizon, hinting that the model is missing out on some short-term dynamics. The remainder of this section discusses each of these moments in detail.

4.2. Standard deviation

[Fig. 5](#) plots the standard deviation of log earnings changes as a function of the recent earnings percentile, separately for two age groups (25–35 and 35–55). Both at short and long horizons, a U-shaped pattern emerges. Naturally, top earners face very large downside risk (falling off the ladder). At the lower end, comprised of workers with low human capital and working at low-productivity firms, the chance of experiencing high earnings growth is greater. While the model captures

¹⁸ A formal argument is offered in [Appendix C](#). A previous version of this paper contained a more complicated model, which explicitly matched top earnings shares by adding an additional parameter for the firm productivity distribution. In that version, the right tail of earnings growth was also matched. However, beyond that, the added complexity did not improve the performance of the model.

¹⁹ See [Appendix B.1.1](#) in [Guvenen et al. \(2015\)](#). The threshold amounts to the equivalent of 1.5 months of minimum wage earnings, working full-time, and a small amount of random noise is added. Note that unemployment benefits are not included in the definition of earnings.




Fig. 5. Standard deviation of log earnings changes.

the shape of the relation between earnings levels and earnings variability, it does not generate enough earnings risk. What could the model be missing? First, the calibration does successfully target transitions into and out of employment. This leaves variation in hours worked within employment spells (including changes in multiple jobholding) and changes in hourly wages. The former is not modeled, and thus one would expect the model to understate earnings volatility to some extent. As for changes in hourly wages, in the model they can result from job to job transitions and human capital dynamics. The firm productivity distribution is tied to measures of residual wage dispersion in the model calibration. Human capital is modeled as a simple unit root process, where the standard deviation of the innovation term is disciplined by the increase in the cross-sectional earnings variance over the life-cycle. Furthermore, human capital depreciates in non-employment, tied to estimates on earnings losses following job displacement.

In sum, the model appears to be missing transitory fluctuations in earnings, including varying hours worked, multiple jobholding, and bonus payments. Explicitly incorporating all these features in a structural model is daunting. Moreover, the idea of the benchmark model is to see how far a relatively simple job ladder model can go in explaining new evidence on earnings dynamics – in particular, in light of the calibration, which does not directly target earnings dynamics.²⁰

4.3. Skewness

If log earnings changes were normally distributed, skewness would be zero across the board. However, as Fig. 6 illustrates, for all but those workers at the very bottom of the recent earnings distribution, there is a significant amount of negative skewness. To a certain extent, this is naturally implied by a job ladder model: the presence of on the job search implies that movements up the ladder are on average smaller than movements down the ladder (which, in the absence of compensating differentials, consist of transitions from employment to non-employment only). The dynamics of human capital formation, which is continuously built up during employment spells and depreciates in non-employment, reinforce the asymmetry.

Skewness is decreasing both in age and recent earnings. The model reproduces both trends: both high earners as well as more senior workers have on average climbed up the job ladder farther, thus there is more room to fall for them. Moreover, optimal search effort is higher for younger workers, thus non-employment spells are longer for older workers and correspondingly negative shocks are larger.

While the amount of negative skewness generated by the model approximately fits the data for most workers, for those at the very top of the earnings distribution, the model exaggerates it. Those very high earners also tend to be very rich. Thus, in the model, they have very little incentive to search for a new job if hit by a separation shock, or if they choose to quit employment. This magnifies the size of negative earnings changes, which are also mechanically larger than for an average worker (at the top, there is more room to fall). Abandoning endogenous job search or risk aversion would help to explain the dynamics at the very top, but at the cost of missing out on explaining the dynamics for the bulk of the population.²¹ This suggests that top earners derive positive utility from their work, relative to other workers, either because of ex-ante differences in worker types or because their jobs are more enjoyable. A possible alternative solution might be to allow for a combination of lower separation and higher job finding rates for highly skilled workers.

²⁰ A different approach is pursued in Appendix A: In an extended model, I add a transitory component, which does not have a clear structural interpretation, and directly target earnings change moments.

²¹ The implications of assuming exogenous contact rates, respectively risk neutrality, are discussed in greater detail in Section 5.1, respectively Section 5.2.




Fig. 6. Skewness of log earnings changes.




Fig. 7. Kurtosis of log earnings changes.

4.4. Kurtosis

The kurtosis of a random variable can be seen both as a measure of peakedness and a measure of tail thickness. Normally distributed log earnings changes would imply a kurtosis of three. Fig. 7 illustrates that male U.S. workers are exposed to earnings risk with considerable excess kurtosis. Overall, the model is successful in generating a comparable amount of excess kurtosis. It also captures the increase as a function of recent earnings for the bottom 90%, though this relation is steeper in the data. For top earners, the model overshoots again.

Furthermore, kurtosis is increasing over the life-cycle, at least for annual changes. Partly, this holds because older workers tend to be employed at better firms as they had more time to climb the ladder. Hence, the probability of finding an even better job decreases and falls down the job ladder (into unemployment and subsequent re-employment at an on average worse firm) are more severe, implying less frequent but larger changes. Again, endogenous search effort amplifies the age differential, as older workers optimally search at lower intensity.

5. Discussion of model features

In this section, I interpret the main results by comparing the benchmark model to alternative versions. The model in section 5.1 features exogenous contact rates instead of endogenous search. Section 5.2 discusses a risk neutral model. Finally, section 5.3 removes human capital depreciation in non-employment. Each of these three alternative models deviates from the baseline in only one aspect, and otherwise agrees with it, in order to isolate the effects. The section concludes with a separate assessment of earnings dynamics generated by the baseline model for job stayers and job switchers.

Table 4 displays re-estimated model parameters for all three alternative model versions. The number of internally calibrated parameters (seven) is unchanged from the baseline, except for the model without human capital depreciation, which

Table 4
Alternative models: jointly calibrated parameters.

	Exogenous search	Risk neutrality	No skill depreciation
χ_s	–	0.684	0.000 ^a
χ_{se}	–	2.376	0.629
$\lambda_u \times 10^2$	19.540	–	–
$\lambda_e \times 10^2$	8.502	–	–
σ_{h_0}	0.685	0.698	1.209
$\sigma_h \times 10^2$	4.541	3.483	4.133
$\mu_e \times 10^3$	1.442	2.477	0.759
$\mu_u \times 10^3$	–6.955	–28.848	0.759 ^b
η	3.683	3.992	3.992

^a At boundary of parameter space, $\chi_s \geq 0$.

^b Constrained to equal μ_e .

Table 5
Alternative models: targeted moments.

Moment	Data	Exogenous search	Risk neutrality	No skill depreciation
NE rate	0.124	0.124	0.124	0.147
J2J transition rate	0.027	0.027	0.027	0.028
Variance of log earnings at 25	0.59	0.59	0.59	0.56
Variance of log earnings at 60	1.05	1.04	1.05	1.01
Mean life-cycle earnings growth	0.82	0.82	0.82	0.81
20Y earnings loss from U-shock	0.12	0.12	0.12	0.06
Mean–min ratio residual wages	1.75	1.75	1.75	2.03

Table 6
Alternative models: results.

	1 year change			5 year change		
	Std. dev.	Skewness	Kurtosis	Std. dev.	Skewness	Kurtosis
Data	0.722	–0.977	16.120	1.112	–1.033	7.700
Benchmark model	0.440 (0.391)	–1.431 (0.469)	13.568 (0.263)	0.789 (0.290)	–1.278 (0.264)	7.814 (0.176)
Exogenous search	0.650 (0.153)	–0.395 (0.601)	9.202 (0.448)	0.876 (0.214)	–0.084 (0.923)	5.471 (0.327)
Risk neutrality	0.438 (0.393)	–0.702 (0.299)	9.048 (0.449)	0.710 (0.362)	–0.405 (0.611)	4.667 (0.413)
No skill depreciation	0.717 (0.257)	–0.480 (0.607)	8.348 (0.553)	0.939 (0.204)	–0.366 (0.754)	5.843 (0.382)

Log earnings change moments as documented by [Guvonen et al. \(2015\)](#), the benchmark model, and alternative model versions. Mean errors, defined as the average difference between model and data moments across age by recent earnings percentile cells, are displayed in parentheses.

restricts human capital to evolve independently of the employment status. [Table 5](#) reports the set of targeted moments, which is also unchanged from the baseline. Earnings change moments, which are not targeted, are displayed in [Table 6](#), along with the relative distance between data and model moments, averaged over age and recent earnings groups. Some additional results and figures are relegated to [Appendix D](#).

5.1. Endogenous search effort

How important is endogenous search effort? To answer this question, I re-calibrate a version of the model with exogenous contact rates instead of endogenous search effort. Non-employed agents are offered a job with (monthly) probability λ_u , employed workers sample a new job offer with probability λ_e . While aggregate transition rates agree with the baseline model by virtue of the calibration, it is the heterogeneity in transition rates that is lost. As a result, long-term unemployment is extremely rare, and so are very large negative shocks. Thus, negative skewness and kurtosis decrease significantly. [Table 6](#) documents that skewness shrinks from –1.43 to –0.40 for one year changes and from –1.28 to –0.08 for five year changes. Kurtosis decreases from 13.57 to 9.20, respectively 7.81 to 5.47. [Fig. 8](#) also illustrates that with exogenous contact

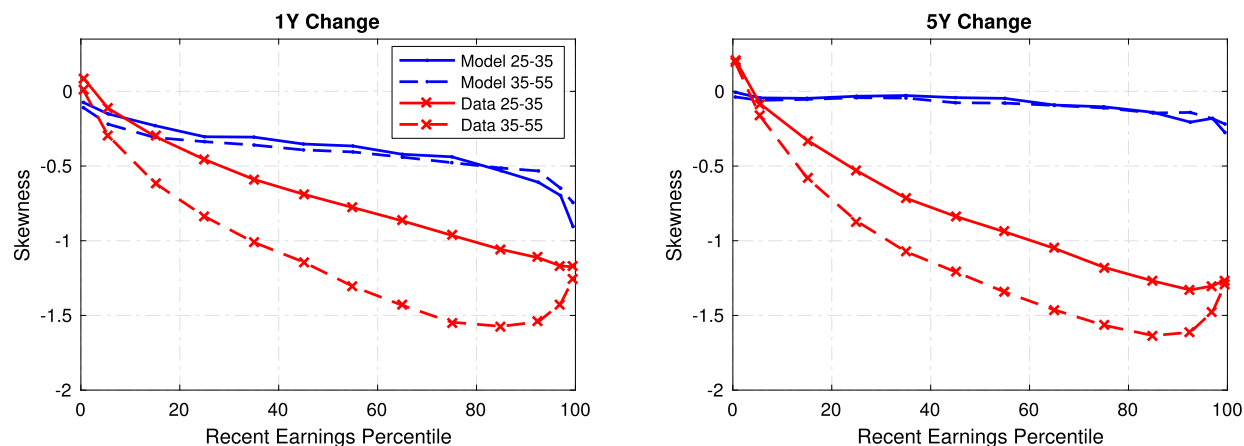


Fig. 8. Exogenous search effort: skewness of log earnings changes.

rates, skewness hardly varies with age. As optimal search effort is lower for more senior workers, being hit by a separation shock entails more dire consequences for them when search is endogenous.²²

One (indirect) way of assessing whether or not the heterogeneity in transition rates is realistic in the benchmark model is to look at the unemployment exit hazard. Shimer (2008), using US micro data from the CPS, finds that it declines from 39% in the first month of an unemployment spell to 22% after six months and 18% after a year. In the benchmark model, the non-employment exit rate similarly decreases from 39% initially to 22% after six months and 8% after a year. These numbers are not exactly comparable as the notion of non-employment is broader in the model. Still, they indicate that the heterogeneity in transition rates generated by the benchmark model is reasonable.²³ This duration dependence of the job finding probability stems mostly from a dynamic selection mechanism, i.e., the pool of non-employed workers is increasingly composed of workers who search at low intensity. With exogenous contact rates, on the other hand, the non-employment exit probability is constant at 20%, clearly at odds with the evidence.

Relatedly, the fact that the standard deviation of log earnings changes is higher, thus closer to the data moment in this specification, has to be interpreted with caution. While in both model versions aggregate transition rates (and thus the non-employment rate) agree, under the benchmark of endogenous search there is a small group of workers who are very long-term unemployed or even permanently separated from the labor force. After the initial very large negative shock following a displacement, their earnings are already at zero and thus do not change anymore (until they possibly regain employment). At the same time, the majority of workers regains employment relatively quickly once hit by an unemployment shock. Consequently, earnings fluctuate less overall compared to the specification with exogenous contact rates, where the (expected) duration of non-employment spells is moderately large for everyone. In other words, the model with exogenous search is doing better on that dimension for the wrong reason.

To summarize, endogenizing search effort is promising in the sense that it endogenously explains earnings dynamics better without expanding the parameter space.

5.2. Risk aversion

The baseline model features risk aversion and a risk-free asset to partially insure against idiosyncratic shocks. This incomplete markets setting is the standard macroeconomic framework for studying the implications of earnings risk. While analyzing the consequences of earnings risk under risk neutrality is not fully satisfactory, we can learn about the effects of risk aversion on choices and consequently earnings dynamics by comparing the benchmark to a risk neutral version of the model.

As documented in Table 6, negative skewness is two to three times larger in the benchmark model compared to the specification with risk neutrality. Moreover, the kurtosis of log earnings changes is also considerably smaller in the latter. Fig. 9 illustrates that under risk neutrality skewness is no longer decreasing beyond the 4th decile of recent earnings, in contrast to the evidence, even though higher earners have certainly more room to fall. But they also optimally search at higher intensity, as their human capital is on average higher, which shortens non-employment spells and thus decreases downside risk. These two counteracting forces cancel each other out. On the other hand, in the benchmark model with risk aversion, higher earners have been able to accumulate more wealth. With concave consumption utility, richer agents search

²² Another point that follows from this discussion concerns explicit modeling of the life-cycle vs. a perpetual youth model: as in the latter search effort trivially cannot depend on age, such a more parsimonious model would have an equally hard time of explaining the age differential.

²³ If anything, one would expect that there is more heterogeneity if the sample consists both of unemployed workers and agents that are out of the labor force.

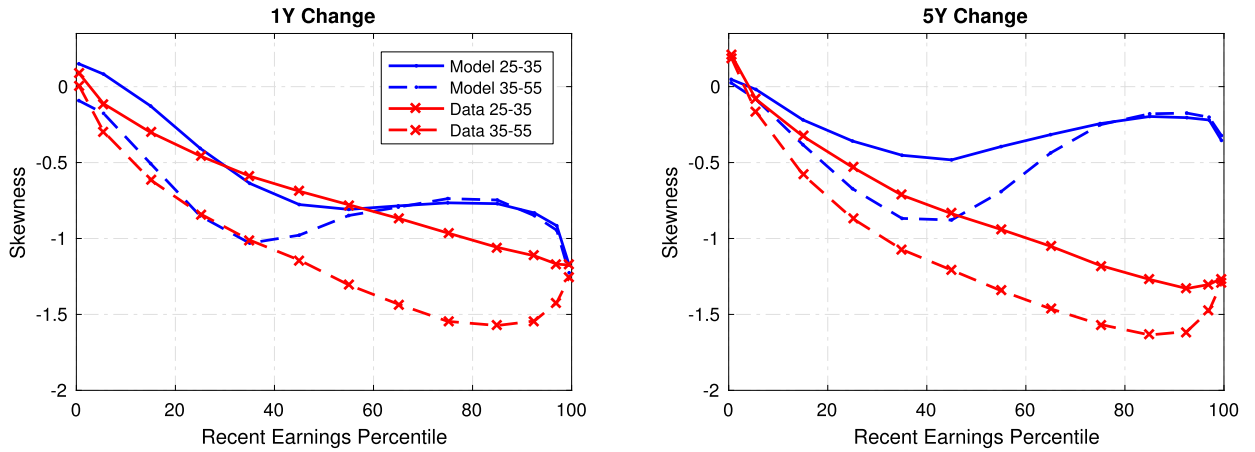


Fig. 9. Risk neutrality: skewness of log earnings changes.




Fig. 10. Risk neutrality: kurtosis of log earnings changes.

Table 7
Top wealth shares.

	Top 10%	Top 1%	Top 0.1%	Top 0.01%
Data	68.1	30.0	12.9	5.4
Benchmark model	65.3	24.2	6.5	1.4

Shares in %. Data: Saez and Zucman (2016), average over 1978–2010.

at lower intensity and, furthermore, the increase in search effort as a function of human capital is flatter. In sum, negative shocks are larger for high earners in the benchmark model with risk aversion. The same reasoning applies to the relation between kurtosis and recent earnings, depicted in Fig. 10: without the wealth effect, this relation is mostly flat.

Since risk aversion is an important ingredient to the model and agents have access to a risk-free asset to smooth consumption, it is interesting to see how much wealth inequality the model generates and how that compares to the data. As wealth is highly concentrated, top wealth shares, displayed in Table 7, are particularly informative. The model fit is satisfactory for the top 10% and the top 1%. The fit is getting worse farther out in the tail, which is to be expected from a model without any heterogeneity in discount rates or capital returns. Overall, the performance of the model in this aspect is reassuring.

5.3. Human capital depreciation in non-employment

In the benchmark version, human capital drifts at a different rate $\mu_u < \mu_e$ in non-employment, enabling the model to match empirical estimates on the consequences of job loss. To illustrate the necessity of this feature, I re-estimate the model subject to the constraint $\mu_u = \mu_e$. As the set of targeted moments is unchanged but the number of calibrated parameters decreases by one, the model is now overidentified. As shown in Table 5, there is a tension between generating

Table 8

Model moments for job stayers vs. switchers.

	1 year changes			5 year changes		
	Std. dev.	Skewness	Kurtosis	Std. dev.	Skewness	Kurtosis
full output	0.44	−1.43	13.57	0.79	−1.28	7.81
stayers (12 month)	0.10	0.03	3.04	0.27	0.07	3.01
switchers (12 month)	0.56	−1.29	10.41	0.83	−1.59	8.98
stayers (6 month)	0.18	−0.26	6.25	0.31	0.06	3.71
switchers (6 month)	0.65	−1.09	8.11	0.84	−1.56	8.76
stayers (recall adj.)	0.24	−0.18	15.89	0.35	−0.01	6.46
switchers (recall adj.)	0.67	−0.99	7.71	0.86	−1.47	8.36

Log earnings change moments computed from the baseline model. The first line reports moments computed from the full panel of simulated workers; agents are split in job stayers and job switchers in the rest of the table. The label '12 month' refers to classifying a worker as a stayer if she is continuously employed at the same firm for the full sample period (i.e., for changes between t and $t+k$ from the first month of t to the last month in $t+k$). The label '6 month' refers to classifying her as stayer if she is employed at the same firm for at least six month in t as well as $t+k$. The last two lines report results when adjusting for recall, see the main text for details.

large earnings losses following displacement on the one hand, and matching the non-employment rate as well as frictional wage dispersion on the other hand. Without skill depreciation in unemployment, earnings losses can only be large if the job finding probability is too low (i.e., the non-employment rate too high) or if frictional wage dispersion is too high.

As the average earnings loss associated with a displacement event is lower without skill depreciation, one would expect skewness to be less negative. This is indeed the case. Table 6 shows that negative skewness is roughly three times smaller. Relatedly, as very large earnings changes tend to be downward movements, kurtosis decreases by a third on average. This model version is actually generating a higher standard deviation of earnings changes, closer to the data moment on average. Again, this has to be interpreted with caution. The average duration of non-employment spells is too long and, similarly to the case with exogenous search, the unemployment exit hazard is relatively flat, declining from 17% to 15% over the course of a year.²⁴

5.4. Job stayers vs. switchers

Are the deviations from log-normality entirely caused by transitions between non-employment and employment as well as job to job transition? Guvenen et al. (2015) find that this is not the case. In particular, log earnings changes for job stayers exhibit a large amount of excess kurtosis.

The presence of multiple job holders complicates the definition of a job stayer in the data. Guvenen et al. (2015) use a definition that classifies a worker as a stayer if more than 90% of his earnings over the relevant time period were stemming from the same firm. As the data on stayers still contains e.g. workers that are temporarily unemployed or working reduced hours, variation in hours worked might be important in explaining deviations from log-normality for stayers. The model entirely abstracts from these considerations and therefore the statistics are of limited comparability. In particular, I naturally define a job stayer in the model as a worker that is continually employed at the same firm throughout the relevant time period (i.e., for changes between year t and year $t+k$ from the first month of t to the last month of $t+k$). For comparison, I also report results when defining a stayer as somebody that works at the same firm for at least six months in every year of the relevant time period. The model output is summarized in Table 8 for both definitions. The last two rows report results from a simulation exercise that adjusts for recall spells in a simply way: in this scenario, I assume that 40% of all exogenously separated workers enter a recall state, associated with an exogenous monthly transition rate of 0.43 into re-employment at the same firm.²⁵ For this exercise, I define a stayer as a worker that is continually employed at the same firm, possibly with recall spells, throughout the relevant time periods.

Keeping these caveats in mind, we can evaluate the performance of the model with respect to three main facts: First, as expected the standard deviation of log earnings changes for stayers is only about half as large as for those workers that switch jobs, denoted as (job) switchers. As illustrated in Table 8, this trend is even more extreme in the model output, where for annual changes the standard deviation is more than five times higher for switchers. Defining a job stayer as a worker that stays at the same firm for at least six month in every year, this ratio decreases a bit. At least part of this discrepancy is due to the absence of any variation in hours worked. Indeed, when adjusting for recall the standard deviation of log earnings changes for stayers increases further relative to the one for switchers.

Second, skewness is close to zero or even slightly positive for stayers, in contrast to switchers. The model captures this pattern well, as log human capital follows a unit root process with normal (and thus symmetric) innovations. Adjusting

²⁴ Table 5 shows that the non-employment rate is too high. While contact rates are not assumed to be exogenous here, the internal calibration sets the search cost in unemployment to the lower boundary at zero (Table 4).

²⁵ These two numbers are based on evidence from the Survey of Income and Program Participation (SIPP), as documented by Fujita and Moscarini (2017). Note that I do not re-compute the decision rules for this experiment.

Table 9
Dependence of model moments on choices vs. economic risk.

	1 year changes			5 year changes			Aggregate flows			
	Std. dev.	Skewness	Kurtosis	Std. dev.	Skewness	Kurtosis	NE rate	NE2E	E2NE	J2J
benchmark	0.44	−1.43	13.56	0.79	−1.28	7.82	0.123	0.207	0.030	0.027
AAJ	0.39	−1.32	13.30	0.70	−1.18	7.98	0.096	0.260	0.028	0.029
FSI	0.52	−1.56	10.97	1.02	−1.01	5.14	0.177	0.085	0.020	0.064
AAJ + FSI	0.39	−0.65	7.09	0.63	−0.07	3.64	0.038	0.430	0.017	0.071
no saving	0.38	−0.61	9.78	0.57	−0.20	4.93	0.071	0.381	0.029	0.026

Log earnings change moments for the baseline model and counterfactuals. In scenario AAJ (accept all jobs), agents accept all jobs. In counterfactual FSI (full search intensity), agents always search at full intensity, both in employment and non-employment. Line (AAJ + FSI) combines the previous two restrictions. Counterfactual 'no saving' shuts down (only) the saving choice: agents are still risk averse, but they do not have access to the risk-free asset.

for recall, or using a less strict definition of stayers, implies that some workers with unemployment spells are classified as stayers, and consequently skewness decreases.

Finally, a striking fact is that even job stayers face massive kurtosis risk. In particular, [Guvenen et al. \(2015\)](#) find that kurtosis is higher for stayers than for switchers. This indicates that within a job spell, earnings are changing very little for most workers, while some experience large changes (e.g. a promotion or a temporary layoff). The simple process for human capital used in the model does not generate any excess kurtosis for stayers, at least when using the strict stayer definition. However, adjusting the results for recall substantially increases the kurtosis measure for stayers.²⁶

6. Implications: risk vs. choice

So far, the discussion did not distinguish between actual risk and predictable changes. Empirically, this distinction requires data on expectations or on consumption (along with a model of consumption). Through the lens of the model, we can to some extent distinguish genuine earnings risk from earnings changes caused by choices.

[Table 9](#) compares the benchmark model output to counterfactual scenarios in which some choices are shut down (i.e., agents are making suboptimal choices). The line labeled AAJ (accept all jobs), reports the results from a counterfactual simulation in which workers accept every job, i.e., reservation wages are below the wage offered by the least productive firm. The line labeled FSI (full search intensity) reports results for a counterfactual scenario in which all employed as well as non-employed agents search at full intensity.²⁷ Neither of these scenarios decreases measured earnings risk significantly. However, combining both modifications (AAJ + FSI) does indeed significantly reduce the higher order moments of log earnings changes. Strikingly, negative skewness decreases by more than half for annual changes and almost completely vanishes for five year changes. In the model, negative skewness is driven by job losses, followed by non-employment spells, human capital depreciation, and subsequent re-employment at on average lower-productivity firms. As human capital depreciates, agents optimally decrease their search effort and increase their reservation wages (per unit of human capital) as the cost of searching and working is independent of human capital. These endogenous responses magnify the size of negative earnings changes. Similarly, kurtosis is roughly cut in half, partly for the same reason, and partly because job flows increase, shifting away mass from the center of the distribution.

The last counterfactual shuts down the saving channel. Agents are still risk averse, but they are forced to live hand-to-mouth. Without the option of self-insurance, non-employment is a dire state for all agents. Consequently, the job finding probability almost doubles, manifesting itself in shorter non-employment spells and much lower negative skewness. The magnitude of the decline in the higher order moments of earnings growth is striking. Overall, these simulation results illustrate that a major fraction of higher order earnings risk could be attributed to choices.

7. Concluding remarks

The present paper aims to show that the new non-parametric evidence on life-cycle earnings dynamics can be well explained by a relatively standard structural representation of a frictional labor market, the job ladder model. Its defining elements, firm heterogeneity and the corresponding dynamics of the job ladder, generate large and negatively skewed earnings changes. These are magnified by the joint effects of endogenous job search and labor supply, risk aversion and

²⁶ A simple way of increasing the kurtosis of stayers, which allows for keeping the parsimonious unit root process for human capital, would be to specify that innovations to human capital arrive only with a certain fixed probability every period (analogously to a Poisson process in continuous time). Alternatively, introducing promotions/demotions could also increase stayer's kurtosis (see e.g. [Shi, 2016](#) for a framework that features both a job ladder across firms as well as a job upgrading within firms).

²⁷ It might appear counterintuitive that in this counterfactual both negative skewness (for annual changes) and the non-employment rate are higher. Indeed, increasing the search effort of all non-employed workers directly decreases the NE rate (and diminishes negative skewness). However, what is dominating in this scenario is the increase in the search effort of employed workers, which is a more drastic change as they search on average at lower intensity. It implies that they work on average at much better firms. Hence, separation shocks can have even more dire consequences, and the wealth effect increases reservation wages.

an associated saving choice, and human capital depreciation in unemployment. As a quantitative exercise, I calibrate the model to match several salient features of the U.S. labor market and the associated earnings levels distribution. The exercise is successful in the sense that it replicates key features of the distribution of log earnings changes (even if not targeted in the calibration). These are, as documented by Guvenen et al. (2015), in particular large negative skewness and very high excess kurtosis, both strikingly at odds with the log-normal framework that is commonly used to model earnings processes and thus earnings risk. However, in the proposed model that replicates these facts, the dynamics of human capital itself are described by a simple univariate log-normal process, augmented with skill loss in unemployment as a way of capturing empirical estimates on the consequences of job loss. This fairly conventional, and parsimonious, representation of human capital, in combination with a frictional view of the labor market, suffices to generate deviations from normality that fit the evidence qualitatively and by and large also quantitatively.

A natural question that follows from the positive analysis in this paper concerns the welfare aspects of the new evidence on earnings dynamics. Guvenen et al. (2015) show that when fitting a rich stochastic earnings process to the data, the welfare costs of idiosyncratic earnings fluctuations are much higher than in a conventional Gaussian framework. The analysis in this paper suggests that it is the interaction between the job ladder and optimal search and job acceptance behavior of workers that drives deviations from the Gaussian world. In particular, the counterfactual analysis in the previous section indicates that a major fraction of especially the higher order moments of earnings growth could be attributed to choices. Hence, an estimated exogenous earnings process could be suboptimal from the point of view of a searching worker, thus overstating the welfare costs of idiosyncratic earnings fluctuations. On the other hand, for welfare analysis, the utility cost of search and labor supply has to be taken into account. A thorough normative appraisal of the welfare costs of uninsurable earnings fluctuations in this setting is beyond the scope of the present paper and left for future work.

Finally, while for the purpose of this paper it suffices to analyze the decision problem of a worker, one can easily extend the proposed model to general equilibrium, if one considers the safe asset to be the portfolio of all firms in the economy. In such a setting, avenues for future research include normative questions such as the optimal design of social insurance and institutions such as the minimum wage. For these important questions, modeling earnings dynamics correctly, including higher-order moments, is fundamental.²⁸ At the same time, the distinction between genuine shocks and changes induced by choices is crucial. An important omission in the present paper that is relevant for normative applications concerns wage bargaining. This is another great avenue for future research.

Appendix A. Extended model

This section contains an extended model that differs from the baseline version in two aspects: First, a transitory shock to earnings is introduced. Second, higher order moments of earnings growth are explicitly targeted in the calibration. The aim of these changes is to explore how much better the model can explain earnings dynamics once these moments are directly targeted. The transitory shock should be viewed as a reduced form, statistical way of capturing some sources of earnings fluctuations that are excluded from the baseline model (varying hours, bonus payments, etc.).²⁹ As such, it breaks with the structural paradigm of this paper.

Specifically, in the extended model the monthly earnings of a worker with human capital h , employed at a firm with productivity p , are given by γhpv , where

$$\log v = \begin{cases} \overset{i.i.d.}{\sim} N(-\sigma_v^2/2, \sigma_v) & \text{with probability } 1/12, \\ 0 & \text{else.} \end{cases}$$

Thus, the transitory shock materializes on average once per year, an assumption that is particularly realistic when thinking about an annual bonus payment. The one additional parameter σ_v is jointly calibrated with the other seven parameters (Table A.10) to match the second to fourth moment of log earnings changes (both for one year and for five year changes), in addition to the seven moments targeted in the baseline model (Table A.12). In sum, eight parameters are calibrated to minimize the sum of squared relative differences between data and model moments, where all thirteen moments are weighted equally.

The results are displayed in Table A.11. Naturally, the standard deviation of log earnings changes increases, relative to the baseline model, due to the introduction of the transitory component. However, it is still too small (see also Fig. A.11). As the transitory component by itself is not skewed, one would expect that skewness is increasing. This is indeed the case, and Fig. A.12 reveals that the extended model is providing a nearly perfect fit to the data for the bottom 90% of earners. At the very top of the earnings distribution, however, the model is still overshooting. Turning to kurtosis, the fit has already been quite good in the baseline. Consequently, the extended model can improve over it only to a small extent (Fig. A.13).³⁰

²⁸ See e.g. Golosov et al. (2016) on how higher-order moments of income shocks affect optimal taxation.

²⁹ Unlike when working with survey data, measurement error is not a concern here.

³⁰ The assumption of the transitory shock hitting only once per year on average is relevant for the kurtosis generated by the extended model. If the transitory shock were to hit every month, kurtosis would decrease relative to the baseline (in contrast, standard deviation and skewness are not affected by this assumption).

Table A.10
Extended model: jointly calibrated parameters.

	Benchmark	Extended
χ_s	1.031	1.040
χ_{se}	0.567	0.741
σ_{h_0}	0.888	0.788
$\sigma_h \times 10^2$	3.698	3.635
$\mu_e \times 10^3$	3.475	3.253
$\mu_u \times 10^2$	-3.473	-3.007
η	3.873	3.717
σ_v	-	3.152

Table A.11
Extended model: earnings dynamics moments.

	1 year change			5 year change		
	Std. dev.	Skewness	Kurtosis	Std. dev.	Skewness	Kurtosis
Data	0.722	-0.977	16.120	1.112	-1.033	7.700
Benchmark model	0.440 (0.391)	-1.431 (0.469)	13.568 (0.263)	0.789 (0.290)	-1.278 (0.264)	7.814 (0.176)
Extended model	0.526 (0.272)	-0.967 (0.153)	14.717 (0.222)	0.868 (0.219)	-1.133 (0.147)	7.619 (0.162)

Log earnings change moments in the data (Guvenen et al., 2015), the benchmark model, and the extended model. Mean errors, defined as the average difference between model and data moments across age by recent earnings percentile cells, are displayed in parentheses.

Table A.12
Extended model: other targeted moments.

Moment	Data	Benchmark model	Extended model
NE rate	0.124	0.123	0.140
J2J transition rate	0.027	0.027	0.025
Variance of log earnings at 25	0.59	0.59	0.61
Variance of log earnings at 60	1.05	1.05	1.08
Mean life-cycle earnings growth	0.82	0.82	0.80
20Y earnings loss from U-shock	0.12	0.12	0.12
Mean-min ratio residual wages	1.75	1.75	1.77




Fig. A.11. Extended model: standard deviation of log earnings changes.

In sum, the results from the extended model indicate that some sources of earnings fluctuations are missing from the structural baseline model, and incorporating those can further improve the performance of the model. However, structurally modeling those is beyond the scope of this paper.




Fig. A.12. Extended model: skewness of log earnings changes.




Fig. A.13. Extended model: kurtosis of log earnings changes.

Appendix B. Numerical solution method

Solving the decision problem of the worker is complicated by the non-convexity and non-smoothness of the value function (as a function of assets) at points where the endogenous labor supply decision changes. I use an insight by Fella (2014), who generalizes the endogenous grid point method (EGPM), originally proposed by Carroll (2006), to non-convex and not globally smooth problems. My setup fulfills the assumptions used in Fella (2014). I will only provide brief intuition here: The key is to observe that the per-period objective function is smooth and strictly concave. Since the value function “points inwards” at kink points, a kink point can never be an optimal savings choice. Consequently, the first order condition (FOC) is still necessary. The FOC is not sufficient, however, in the non-concave region of the value function. Hence, the algorithm proceeds in two steps: (i) use the standard EGPM to identify candidate pairs of cash-on-hand today x_i and savings a'_i for a grid on future assets $(a'_i)_{i=1}^I$; (ii) for those pairs (x_i, a'_i) that lie in the non-concave region, verify that they are indeed global optima using discretized value function iteration. Using this method as opposed to brute-force global discretized value function iteration yields substantial speed and accuracy gains.

Finding the optimal search choice is simplified by the additive nature of search costs. This property implies that one can solve the multi-dimensional optimization step sequentially. First, compute optimal search effort s^* as a function of (h, p, t) (i.e., the current state except for current assets a) and next period assets a' ; (ii) given $s^*(a', h, p, t)$, use the EGPM-modification as described above to find optimal pairs (x, a') for every (h, p, t) grid point.

Otherwise, the dynamic programming problem is fairly standard. I use non-linear interpolation in the asset dimension and linear interpolation in all other dimensions (i.e., h and p). There are 300 grid points for asset holdings a , 50 for firm productivity p and 100 for human capital h . Since the n -th moment of a Pareto distributed random variable with shape

coefficient $\eta < n$ does not exist, I truncate the firm productivity distribution at the 99.999% quantile. To generate model moments, I simulate a population of 500,000 agents from age 24–65 (i.e., they retire at 66) at monthly frequency, starting out all agents at zero initial wealth and unemployed. Simulated earnings data from age 25–60 is used to compute model moments, in accordance with [Guvenen et al. \(2015\)](#).

Appendix C. Pareto tails

In this section, I show that if job offers are drawn from a Pareto distribution, then the right tail of the distribution of earnings growth inherits its tail index. To this end, I will make a few simplifying assumptions: ignore human capital h , and, to avoid the time aggregation problem, consider period-to-period earnings growth (the full model is solved at monthly frequency whereas the object of interest is annual earnings). Furthermore, assume exogenous contact rates $\lambda_u \in (0, 1)$, respectively $\lambda_e \in (0, 1)$, for unemployed and employed workers. Abstract from endogenous labor supply, i.e., workers accept every job offer (above the current wage) and jobs are destroyed at exogenous probability $\delta \in (0, 1)$. Consider a steady state characterized by a fraction $u = \frac{\delta}{\delta + \lambda_u}$ of unemployed workers and a wage distribution $G_p(p)$ per unit of human capital (which, as usual, differs from the job offer distribution $F_p(p) = 1 - p^{-\eta}$ as long as $\lambda_e > 0$). Earnings are given by $e = \gamma p$ for employed workers (implicitly, h is normalized to one). Unemployed workers have zero earnings, replaced by a lower bound \underline{e} when computing earnings growth $x = \frac{e'}{e}$. I show that if job offers are drawn from a Pareto distribution with shape coefficient η , then the right tail of the distribution of earnings growth inherits the tail index. More precisely, there exists an $\bar{x} \geq 1$ such that $\mathbb{P}(X > x) \propto x^{-\eta}$ for all $x \geq \bar{x}$: Let $x \geq \bar{x}$ and denote by I_e (I_j) an indicator variable for being employed (drawing a new job offer), then

$$\begin{aligned}
 \mathbb{P}(X > x) &= \mathbb{E} \left[\mathbb{1} \left(\frac{e'}{e} > x \right) \right] \\
 &= (1 - u) \mathbb{E} \left[\mathbb{1} \left(\frac{e'}{\gamma p} > x \right) | I_e = 1 \right] + u \mathbb{E} \left[\mathbb{1} \left(\frac{e'}{\underline{e}} > x \right) | I_e = 0 \right] \quad \text{by L.I.E.} \\
 &= (1 - u) \lambda_e \mathbb{E} \left[\mathbb{1} \left(\frac{\gamma \max\{p', p\}}{\gamma p} > x \right) | I_e = 1 \wedge I_j = 1 \right] + \\
 &\quad u \lambda_u \mathbb{E} \left[\mathbb{1} \left(\frac{\gamma p'}{\underline{e}} > x \right) | I_e = 0 \wedge I_j = 1 \right] \quad \text{by L.I.E.} \\
 &= (1 - u) \lambda_e \mathbb{E}_p \left[\mathbb{P}(p' > px | p) \right] + u \lambda_u \mathbb{P} \left(p' > \frac{e x}{\gamma} \right) \quad \text{by L.I.E. and since } x \geq 1 \\
 &= (1 - u) \lambda_e \mathbb{E}_p \left[(px)^{-\eta} \right] + u \lambda_u \min \left\{ 1, \left(\frac{e x}{\gamma} \right)^{-\eta} \right\} \tag{C.1} \\
 &= \left((1 - u) \lambda_e \mathbb{E}_p \left[p^{-\eta} \right] + u \lambda_u \left(\frac{e}{\gamma} \right)^{-\eta} \right) x^{-\eta} \propto x^{-\eta} \quad \text{for } x \geq \frac{\gamma}{e}.
 \end{aligned}$$

Hence $\bar{x} = \frac{\gamma}{e}$. The derivation used the law of iterated expectations (L.I.E.) multiple times; \mathbb{E} refers to the cross-sectional expectation taken over the measure of workers, while \mathbb{E}_p denotes the expectation over the wage distribution $G_p(p)$. Considering the two terms in (C.1) separately, note that the derivation implies that both for currently employed and unemployed workers, earnings growth has a Pareto tail (that is decaying at uniform rate η). How one deals with zero earnings (or in general whether one includes unemployment benefits) affects only how far out in the distribution the Pareto tail starts, but not its thickness.

Intuitively, even when human capital h is evolving over time, as in the full model, the additional term $\frac{h'}{h}$ that is not canceling out does not affect the *tail behavior* of earnings growth (as long as innovations to human capital are not fat-tailed itself). In a similar vein, time aggregation, such as a worker being employed only part of the year, or changing jobs multiple times, does affect the distribution of earnings growth, but not its asymptotic tails.

Pareto tails, also referred to as power laws, are found in remarkably many economic variables (see [Gabaix, 2009](#) for a review). The job ladder framework is attractive because a single distributional assumption can explain multiple empirical observations: if firm productivity is Pareto distributed, then both the earnings levels as well as the earnings growth distribution inherits this Pareto tail.

Appendix D. Additional figures and results (Figs. D.14–D.19)




Fig. D.14. Exogenous search effort: standard deviation of log earnings changes.




Fig. D.15. Exogenous search effort: kurtosis of log earnings changes.




Fig. D.16. Risk neutrality: standard deviation of log earnings changes.




Fig. D.17. No human capital depreciation: standard deviation of log earnings changes.

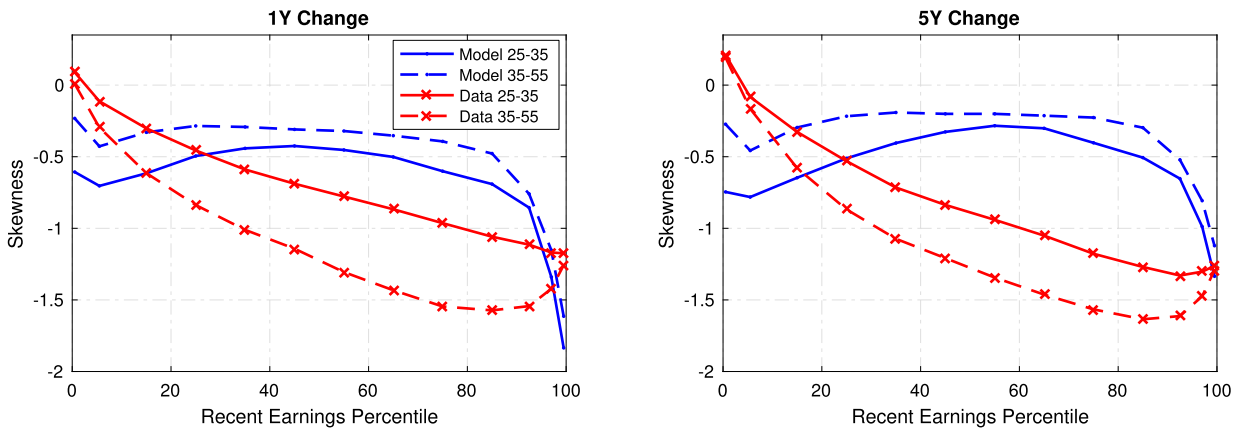


Fig. D.18. No human capital depreciation: skewness of log earnings changes.




Fig. D.19. No human capital depreciation: kurtosis of log earnings changes.

References

Abowd, J.M., Card, D., 1989. On the covariance structure of earnings and hours changes. *Econometrica* 57 (2), 411–445. <http://www.jstor.org/stable/1912561>.
 Aiyagari, S.R., 1994. Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics* 109 (3), 659–684. <http://www.jstor.org/stable/2118417>.
 Altonji, J.G., Smith, A.A.J., Vidangos, I., 2013. Modeling earnings dynamics. *Econometrica* 81 (4), 1395–1454. <https://ideas.repec.org/a/ectm/emetrp/v81y2013i4p1395-1454.html>.

- Bagger, J., Fontaine, F., Postel-Vinay, F., Robin, J.-M., 2014. Tenure, experience, human capital, and wages: a tractable equilibrium search model of wage dynamics. *The American Economic Review* 104 (6), 1551–1596. <http://www.aeaweb.org/articles?id=10.1257/aer.104.6.1551>.
- Bonhomme, S., Robin, J.-M., 2010. Generalized non-parametric deconvolution with an application to earnings dynamics. *The Review of Economic Studies* 77 (2), 491–533. <https://ideas.repec.org/a/oup/restud/v77y2010i2p491-533.html>.
- Burdett, K., 1978. A theory of employee job search and quit rates. *The American Economic Review* 68 (1), 212–220. <http://www.jstor.org/stable/1809701>.
- Carroll, C.D., 2000. Why do the rich save so much? In: Slemrod, J.B. (Ed.), *The Economic Consequences of Taxing the Rich*. Harvard University Press.
- Carroll, C.D., 2006. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters* 91 (3), 312–320. <http://ideas.repec.org/a/eee/econlet/v91y2006i3p312-320.html>.
- Davis, S.J., von Wachter, T., 2011. Recessions and the costs of job loss. *Brookings Papers on Economic Activity* 43 (2 (Fall)), 1–72. <https://ideas.repec.org/a/bin/bpeajo/v43y2011i2011-02p1-72.html>.
- Department of Labor, 2008. *Comparison of State Unemployment Insurance Laws*.
- Fella, G., 2014. A generalized endogenous grid method for non-smooth and non-concave problems. *Review of Economic Dynamics* 17 (2), 329–344. <https://ideas.repec.org/a/red/issued/11-275.html>.
- Fujita, S., Moscarini, G., 2017. Recall and unemployment. *The American Economic Review* 107 (12), 3875–3916. <http://www.aeaweb.org/articles?id=10.1257/aer.20131496>.
- Gabaix, X., 2009. Power laws in economics and finance. *Annual Review of Economics* 1 (1), 255–294. <http://ideas.repec.org/a/anr/reveco/v1y2009p255-294.html>.
- Geweke, J., Keane, M., 2000. An empirical analysis of earnings dynamics among men in the PSID: 1968–1989. *Journal of Econometrics* 96 (2), 293–356. <https://ideas.repec.org/a/eee/econom/v96y2000i2p293-356.html>.
- Golosov, M., Troshkin, M., Tsyvinski, A., 2016. Redistribution and social insurance. *The American Economic Review* 106 (2), 359–386. <http://www.aeaweb.org/articles?id=10.1257/aer.20111550>.
- Guvenen, F., Karahan, F., Ozkan, S., Song, J., 2015. What Do Data on Millions of U.S. Workers Reveal About Life-Cycle Earnings Risk? Working Paper 20913. National Bureau of Economic Research. <http://www.nber.org/papers/w20913>.
- Hornstein, A., Krusell, P., Violante, G.L., 2007. Frictional Wage Dispersion in Search Models: A Quantitative Assessment. Working Paper 13674. National Bureau of Economic Research. <http://www.nber.org/papers/w13674>.
- Hornstein, A., Krusell, P., Violante, G.L., 2011. Frictional wage dispersion in search models: a quantitative assessment. *The American Economic Review* 101 (7), 2873–2898. <http://www.aeaweb.org/articles.php?doi=10.1257/aer.101.7.2873>.
- Huckfeldt, C., 2016. Understanding the Scarring Effects of Recessions. Working paper.
- Jarosch, G., 2015. Searching for Job Security and the Consequences of Job Loss. Working paper.
- Jung, P., Kuhn, M., 2017. Earnings Losses and Labor Mobility Over the Life Cycle. CESifo Working Paper 6552. CESifo Group Munich. https://ideas.repec.org/p/ces/ceswps/_6552.html.
- Karahan, F., Ozkan, S., Song, J., 2016. Sources of Inequality in Earnings Growth Over the Life Cycle. Working paper.
- Krusell, P., Mukoyama, T., Şahin, A., 2010. Labour-market matching with precautionary savings and aggregate fluctuations. *The Review of Economic Studies* 77 (4), 1477–1507. <http://restud.oxfordjournals.org/content/77/4/1477.abstract>.
- Lise, J., 2013. On-the-job search and precautionary savings. *The Review of Economic Studies* 80 (3), 1086–1113. <http://ideas.repec.org/a/oup/restud/v80y2013i3p1086-1113.html>.
- Ljungqvist, L., Sargent, T.J., 1998. The European unemployment dilemma. *Journal of Political Economy* 106 (3), 514–550. <https://ideas.repec.org/a/ucp/jpolec/v106y1998i3p514-550.html>.
- MaCurdy, T.E., 1982. The use of time series processes to model the error structure of earnings in a longitudinal data analysis. *Journal of Econometrics* 18 (1), 83–114. <http://EconPapers.repec.org/RePEc:eee:econom:v:18:y:1982:i:1:p:83-114>.
- Organisation for Economic Co-operation and Development, 2016. Employment rate: aged 25–54: males for the United States. LREM25MAUSM156S. <https://fred.stlouisfed.org/series/LREM25MAUSM156S>. Retrieved from FRED, Federal Reserve Bank of St. Louis, July 27, 2016.
- Organisation for Economic Co-operation and Development, 2017. Minimum relative to average wages for full-time workers. <https://stats.oecd.org/Index.aspx?DataSetCode=MIN2AVE#>. Retrieved from OECD.stat, February 23, 2017.
- Piketty, T., Saez, E., 2003. Income inequality in the United States, 1913–1998. *The Quarterly Journal of Economics* 118 (1), 1–41. <https://doi.org/10.1162/00335530360535135>.
- Postel-Vinay, F., Turon, H., 2010. On-the-job search, productivity shocks, and the individual earnings process. *International Economic Review* 51 (3), 599–629. <https://doi.org/10.1111/j.1468-2354.2010.00594.x>.
- Saez, E., Zucman, G., 2016. Wealth inequality in the United States since 1913: evidence from capitalized income tax data. *The Quarterly Journal of Economics* 2, 519–578.
- Shi, S., 2016. Efficient job upgrading, search on the job and output dispersion. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2729019.
- Shimer, R., 2005. The cyclical behavior of equilibrium unemployment and vacancies. *The American Economic Review* 95 (1), 25–49. <http://www.aeaweb.org/articles.php?doi=10.1257/0002828053828572>.
- Shimer, R., 2008. The probability of finding a job. *The American Economic Review* 98 (2), 268–273. <http://www.aeaweb.org/articles.php?doi=10.1257/aer.98.2.268>.
- Shimer, R., 2012. Reassessing the ins and outs of unemployment. *Review of Economic Dynamics* 15 (2), 127–148. <https://ideas.repec.org/a/red/issued/11-293.html>.
- Social Security Administration, 2015. *Fast Facts & Figures About Social Security*.