

Cyclical Earnings, Career and Employment Transitions ^{*}

Carlos Carrillo-Tudela [†]

University of Essex,
CEPR, CESifo and IZA

Ludo Visschers [‡]

University of Edinburgh,
U. Carlos III Madrid
CESifo and IZA

David Wiczer [§]

Stony Brook University, USA.

September 2022

Abstract

This paper studies the cyclical behaviour of earnings risk and career changes. We document that the procyclical skewness of the earnings growth distribution arises mostly from the earnings changes of employer and occupation switchers. To uncover their relative importance in driving cyclical earnings changes and whether this arises from changes in the returns to mobility or mobility shocks, we propose a multi-sector business cycle model with on-the-job search and endogenous occupational mobility. Idiosyncratic occupational mobility is the main driver of cyclical earnings risk, mainly due to cyclical shifts in the returns to this mobility. This is the main reason why the sulling effects of recessions are long-lasting. These effects manifest themselves through a collapse of the job ladder and forgone lifetime earnings gains, especially for low-paid workers, and through large lifetime earnings losses among high-paid workers who experience forced occupational mobility and poor re-employment outcomes.

Keywords: Earnings, Unemployment, Business Cycle, Search, Occupational Mobility.

JEL: E24, E30, J62, J63, J64

^{*}We would like to thank Joe Altonji, Serdar Ozkan, Sedar Birinci, Richard Rogerson, Matthew Shapiro, Ija Trapeznikova, Gianluca Violante as well as seminar participants at the University of Bristol, Copenhagen, Florida, Edinburgh, Essex, Kansas State, Konstanz, Carlos III, and the St. Louis Fed, Kansas City Fed, Bank of Portugal, the Chilean Central Bank, as well as conference participants at the EEA-ESEM 2018, AEA 2019, SED 2019, SaM annual conference, 2021, Midwest Macro 2022, Labor Market with Frictions Conference 2022, and the Vienna Macro workshop, 2022. The usual disclaimer applies.

[†]Correspondence: Department of Economics, University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, UK. Email: cocarr@essex.ac.uk. He acknowledges financial support from the UK Economic and Social Research Council (ESRC), award reference ES/V016970/1.

[‡]Correspondence: School of Economics, U. Edinburgh, UK. Email: ludo.visschers@ed.ac.uk. He acknowledges financial support from the UK Economic and Social Research Council (ESRC), award references ES/V016970/1, and ES/L009633/1, and Fundacion BBVA (Ayudas a Equipos de Investigación Científica Sars-Cov-2 y Covid-19)

[§]Correspondence: Department of Economics, Stony Brook University. Email: david.wiczer@stonybrook.edu.

1 Introduction

Understanding the nature of earnings risk is important because it impacts individuals' career path decisions, their consumption, savings and investment decisions, and ultimately inequality. In this paper we study the joint behaviour of earnings risk and career changes over the business cycle. It is well documented that during recessions those individuals who separate directly to another job face lower earnings gains, while those who lose their jobs suffer larger earnings losses relative to expansions (see Huckfeldt, 2022). These patterns are consistent with increasing evidence showing that the distribution of earnings growth exhibits procyclical skewness (see Guvenen, Ozkan and Song, 2014 and Busch, Domeij, Guvenen and Madera, 2021).¹ Less is known, however, about the type of job mobility that drives cyclical earnings risk and whether the latter is mainly caused by cyclical changes in the returns to mobility (i.e. the earnings change conditional on a transition), cyclical changes in the frequency of job loss and job finding and how workers' mobility decisions interacts with them both. Studying these features is important as they help determine the sources of idiosyncratic earnings risk and why this risk changes over the cycle. Without this understanding it remains difficult to evaluate, for example, whether governments should emphasise labour market policies that aim to bring individuals back to work quickly over longer duration re-training schemes that help individuals improve the type of their re-employment jobs.²

This paper shows that career changes, which we observe as occupational mobility, are the main driver behind the cyclical patterns of the earnings growth distribution. Moreover, it is occupational mobility due to workers' evolving idiosyncratic career prospects rather than occupation-wide productivity differences that makes occupational mobility the more important component. We also show that employer mobility on its own does not contribute as much as occupational mobility in shaping cyclical earnings risk. Further, worsening returns to (idiosyncratic) occupational mobility during recessions explain to a large extent the observe procyclical skewness of earnings changes. Our analysis demonstrate that the sullyng effects of recessions are long-lived. They manifest themselves through a collapse of the job ladder, particularly for occupation switchers, and forgone lifetime earnings gains especially for low-paid workers, and through the large lifetime earnings losses faced by high-paid workers who experience forced occupational mobility and poor re-employment outcomes.

To motivate our approach we use the Survey of Income and Program Participation (SIPP) to construct the annual earnings growth distribution both in the cross-section and over the business cycle. After showing that the earnings growth distribution in our data is also characterised by procyclical skewness as in Guvenen et al. (2014), and exhibits the same cross-sectional properties as documented

¹Procyclical skewness implies that the interquartile range of the distribution of earnings growth remains stable over the cycle, while the right tail compresses and the left tail expands in recessions and vice versa in expansions (Busch et al., 2021).

²For example, during the recent Covid-19 pandemic the UK Government ended up implemented unemployment benefits cuts to individuals who did not actively search for jobs outside their occupations after three months into their unemployment spell. This is similar to the Hartz reforms aimed at tackling high unemployment in Germany. These reforms imposed severe penalties on the level of unemployment benefits individuals could claim if they rejected a suitable job offer irrespectively of the industry/occupation.

in Guvenen, Karahan, Ozkan and Song (2021), we present novel evidence showing the importance of occupational mobility driving earnings growth. First we show that, among those individuals who changed employer, there is an increasing relationship between the size of the earnings change (positive or negative) and the probability of an occupational switch. This feature is observed among individuals who made a direct employer transition and among those who changed employers through a spell of unemployment. We then document our main empirical finding. The procyclical skewness of the earnings growth distribution arises mostly from the earnings changes of those individuals who switched employers and occupations at the same time. We find considerably stronger evidence of procyclical skewness among those who made an occupation switch and changed employers directly through a job-to-job transition or through unemployment relative to those who changed employer but did not switch occupation.

To interpret these data patterns, we propose a model of on-the-job search in which a job consists of two dimensions: the employer dimension (where the job is done) and the occupation dimension (the tasks involved in the job). We use this framework to investigate (i) whether cyclical changes in the earnings risk arise from the occupation or from the employer dimension, and (ii) whether the importance of each dimension arise from cyclical changes in the returns to mobility or from cyclical changes on worker flows. A simple decomposition of cyclical changes in observed worker flows and observed earnings changes conditional on flows does not allow us to answer these questions. Observed worker flows reflect worker choices in the face of changing returns, where workers could opt not to reallocate. Observed earnings changes reflect workers' acceptance decisions in the face of alternative employment possibilities, shaped also by the flows, and returns elsewhere. To untangle these forces we extend the canonical on-the-job search model originally proposed by Burdett (1978) to a multi-sector business cycle economy with endogenous employer and occupational mobility and cyclical fluctuations in job loss and job finding probabilities and offer distributions.

In our model, worker and earnings heterogeneity arises principally from two idiosyncratic productivity shocks. Essentially, workers search to improve on the two dimensions of job match quality, occupation- and firm-specific. Occupations are distinguished from one another by their workers' specific match quality and specific human capital as well as occupation-wide productivity differences. A key feature of our framework is that the decision to change employer or occupations is fundamentally different in terms the workers' information. As is standard in job ladder models, meetings with employers are treated akin to "inspection" goods. Once a worker encounters a firm, he knows enough to make a job acceptance decision. This assumption captures that typically many job interviews reveal enough information to the worker (and firm) to accept or reject the employment offer. As in many multi-sector models following Lucas and Prescott (1974), the decision to change occupations is instead treated akin to an "experience" good (see Carrillo-Tudela and Visschers, 2021, Pilossoph, 2021, Wiczer, 2015, among others). The worker does not know his labour market opportunities before starting the search for a job in the new occupation. This captures that when changing careers workers typically are able to re-build their employment contacts, learn about new employment opportunities and their job finding prospects only when they start searching for jobs in the new occupation. We

show that this structure is able to replicate the data very well.

We use simulated methods of moments to structurally estimate our model. The estimation reveals that workers' endogenous decisions to change employers and occupations with or without intervening spells of unemployment can reproduce the observed cyclical behaviour of the earnings growth distribution, the cyclical behaviour of worker mobility as well as a wide range of cross-sectional patterns that characterise earnings risk and worker mobility. In particular, our model reproduces the (i) job loss, re-employment and direct employer transition probabilities in the cross-section and over the business cycle; (ii) the gross occupational mobility patterns among employer stayers and movers in the cross-section and over the cycle, including the increasing relationship between the size of the earnings change and the probability of an occupation switch; (iii) the bilateral net mobility flows across occupations; (iv) the earnings growth distributions conditional on workers' employer and occupation transitions; (v) the cyclical change in the earnings growth distribution, characterised by its procyclical skewness; and (vi) the cyclical changes of the earnings growth distribution conditional on employer and occupation transitions, showing that the procyclical skewness arises from simultaneously changes in employers and occupations.

Our approach allow us to disentangle the separate effects of firm-worker and occupation-worker shocks in explaining the cyclicity of earnings growth. We find that returns to idiosyncratic occupation mobility can explain nearly the entirety of the difference between the expansion and recessions earnings growth distributions, particularly between the 10th and 90th percentiles. Outside this range two major conclusions emerge. At the bottom tail, about 53% of the largest earnings losses we observe in recessions arise from those individuals who lost their jobs *and* had to change occupations. This type of events are similar to the “disaster” shocks in Guvenen et al. (2021) or “obsolescence” shocks in Huckfeldt (2022). At the top tail, the more abundant number of opportunities to improve the occupational dimension of a job becomes the more important, and can explain 26% of the largest earnings gains observed in expansions. The remainder is explained by cyclical changes in returns to occupational mobility.

Given the estimates in which both flows and returns respond to the cycle, we investigate the cost of business cycles. In particular, our model implies that cyclical changes in the returns to mobility and probabilities of job loss and job finding affect workers differently based on their positions in the employer and occupation dimensions of the job ladder. For example, by affecting their ability to climb the ladder and their outcomes after falling from it, workers might suffer differently from the sullyng effects of recessions (see Barlevy, 2002). We are interested in whether this sullyng effects persists over time. We find that low-paid workers suffer disproportionately more in terms of lifetime earnings during recessions from the reduced opportunities to climb the job ladder and the lower returns to mobility than do high-paid workers. However, we also find that it is high-paid workers who suffer disproportionately more in recessions from job transition through spells of unemployment than low-paid workers. A key result is that these costs arise primarily from the occupation dimension rather than the employer dimension of a job. This suggests that labour market policies that enable workers to achieve higher returns to occupation mobility through access to higher quality jobs could have a

significant impact in reducing cyclical earnings risk and the cost of recessions.

Related Literature This paper contributes to the literature investigating the properties of idiosyncratic earnings risk. Measures of the latter have been typically obtained using log-normal earnings processes and used in a wide variety of heterogeneous agent macroeconomic models. A set of recent papers, however, show that idiosyncratic earnings risk present non-Gaussian features. For example, Guvenen et al., (2021) using administrative data for the US find that the earnings growth distribution is left-skewed and has excess kurtosis (see also Arellano, Blundell and Bonhomme, 2017). Di Nardi et al. (2019), Halvorsen, Holter, Ozkan and Storesletten (2020) and Busch (2020) find similar features using data from the Netherlands, Norway and Germany. The former two also analyse the roles of changes in hours relative to hourly wages in shaping the earnings growth distribution. The main difference between these papers and ours is that we focus on understanding cyclical earnings risk.

Over the business cycle, Guvenen et al. (2014), Busch et al. (2021), Harmenberg (2018) and Kramer (2022) among others, show that earnings risk is mainly characterised by its procyclical skewness.³ Building on these results, our main empirical contribution is to show that the procyclical skewness characterising the earnings growth distribution arises primarily from those who changed employers and occupations at the same time. To the best of our knowledge this feature of the data is novel. As in Di Nardi et al. (2019) and Halvorsen et al. (2020) we find that changes in hours are important for cyclical earnings risk. However, our data show that the importance of simultaneous occupation and employer mobility in shaping cyclical earnings risk is strongly visible when considering hourly wages.

Our model focuses on ex-post worker heterogeneity through firm and occupation match-specific productivities, occupation-wide productivities and human capital accumulation. This complements the work of Gregory, Menzio and Wiczer (2021), Karahan, Ozkan, and Song (2019) and Kudlyak and Hall (2019) who emphasise ex-ante worker heterogeneity, particularly in terms of the job loss and job finding probabilities. An abiding difficulty in the literature highlighting the role of ex-ante heterogeneity is distinguishing between fixed types and long-lasting shocks. In occupations, and the careers they represent, our paper captures one such long-lasting shock and shows how it can create higher-order effects in cyclical earnings risk. The reality is certainly neither polar assumption but our evidence suggests occupational disruptions are an important feature.

We also contribute to the growing literature that uses job search models to explain earnings risk. In particular, the canonical job ladder models can naturally generate some of the left-skewness and excess kurtosis observed in the cross-sectional earning growth distribution. This is because in this model the majority of workers experience small wage changes, while some workers experience large losses due to job loss. Postel-Vinay and Turon (2010), for example, shows that a version of such a model in which firms match workers' outside offers is able to replicate many salient properties of earnings dynamics as well as the annual earnings growth distribution in the cross-section. Lise (2012)

³This is in contrast to Storesletten, Telmer and Yaron (2004) who find that earnings risk is instead characterised by its countercyclical variance.

shows that a version of the job ladder model where risk averse workers chose their search intensities and accumulate wealth is consistent with the greater dispersion of the wealth distribution relative to the earnings distribution. More recently, Hubmer (2018) further consider human capital accumulation in the latter environment and shows that his model is consistent with the main properties of the cross-sectional earnings growth distribution (see also Karahan et al., 2020). Our model not only is able to reproduce the properties of the cross-sectional earnings growth distribution at the aggregate and by type of employment and occupation transition, but it also replicates their cyclical properties. We show that a version of our model without occupational mobility akin to the canonical job ladder model reveals an important trade-off. Either it is able to replicate the cyclical behaviour of the earnings growth distribution or the cyclicalities of worker transition flows, but not both at the same time.⁴

Kamborouv and Manovskii (2009a,b) and Busch (2020) have already highlighted the role of occupational mobility in explaining earnings dynamics through the importance of returns to occupational tenure and long run changes in earnings inequality. Carrillo-Tudela and Visschers (2021) show the importance of occupational mobility in determining the joint behaviour of cyclical unemployment and its duration distribution. Huckfeldt (2022) emphasises the role of occupational mobility during recessions in shaping the long term costs of job loss. The novelty here is to show that occupational mobility is the main driver behind cyclical earnings risk and to highlight the importance of the returns to occupational mobility relative to the probabilities of job loss and job finding. In addition to these papers, our analysis evaluates the sullyng effects of recessions on lifetime earnings focusing on the different outcomes between high- and low-paid workers.

The rest of this paper is organized as follows. In Section 2 we present empirical work highlighting the role of occupational mobility in the dispersion of earnings changes and on the procyclical skewness of the earnings growth distribution. Sections 3 and 4 present our job ladder model and provide a careful discussion of its estimation and fit to the data. In Section 5 we use our estimated model to decompose the change in the earnings growth distribution over the cycle. Section 6 evaluates a version of our model without occupational mobility and explores its implications. Section 7 considers the sullyng effect of recessions. Section 8 concludes. Details of our data construction, model derivations, its estimation procedure, identification and robustness exercises are relegated to an Online Appendix.

2 The Earnings Growth Distribution

2.1 Data

We use data from the SIPP from the 1990 to 2008 panels, covering the 1990-2013 period. The advantage of using this dataset is that each of its panels follows a large number of workers for up to four years. Within each panel individuals are divided into four rotation groups, where each group

⁴Hamenberg (2021) investigate whether an on-the-job search model can reproduce the time series relationship between the skewness and mean of the earnings growth distribution. Like his, our model generates a high skewness-mean correlation over time. Kramer (2022) also investigates the cyclical properties of the earnings growth distribution, but in an environment without on-the-job search.

is interviewed in waves of four months. At the end of each wave individuals report information on their current and previous employment status, occupations, industries and earnings (hourly wages and hours worked), covering the last four months. Using this information, we define employer, occupation and earnings changes based on a worker’s main job for each period.⁵

Labor market flows Within a panel we identify for each individual whether he/she experienced an employer and/or an occupational transition. Employer changes that occurred without an intervening full month of unemployment are labeled *EE* transitions and those that occurred through unemployment are labeled *EUE* transitions. In the latter case, we include all transitions in which the worker returned to employment within the sample, even if the worker did not report actively searching. Since we only consider unemployment spells completed within the survey period, to minimize the potential under-representation of *EUE* relative to *EE* transitions we consider only transitions with at least 4 waves remaining in the panel. *EUE* transitions, however, would remain affected by our choice to exclude “temporary recalls”, as these workers do not seem to face the same reallocation and search frictions as those workers who do not expect to return to their previous employers (see Fujita and Moscarini, 2017). We further detail our procedure in Online Appendix A.

To measure occupation changes we homogenise the occupation classification across SIPP panels using the crosswalk translation scheme created by IPUMS based on the 1990 Standard Occupational Classification (1990 SOC). We then aggregate the resulting three-digit occupational codes. For our benchmark analysis we use the aggregation based on four task-based occupational categories: Non-routine Cognitive, Routine Cognitive, Non-routine Manual and Routine Manual. We chose this coarse aggregation to focus on those changes in an individual’s line of work that also involve a change in the main tasks performed. We then compare the task-based occupation for a given individual across waves. In Online Appendix A we discuss measurement error in occupational mobility, while in Online Appendix B we present robustness using two-digit occupational codes.

For an individual to be labeled an “employer/occupational stayer” in a given wave he/she should not have changed either of these dimensions in the previous or in the posterior year relative to this wave.⁶ Under this categorisation we find that over 75% of the transitions are made up of employer/occupation stayers. The reminder observations contain at least one form of transition. From the latter we label an individual to be an “employer/occupation mover” in a given wave when he/she reported a simultaneous employer and occupational change. As mentioned above we distinguish on whether the employer change takes the form of a *EE* or *EUE* transition. An individual is labeled as an “employer stayer / occupation mover” (“employer mover / occupation stayer”) when we only observe a change in the occupation (employer) dimension of a job.

Earnings To study earnings we deflate nominal monthly earnings in the SIPP by the Personal Consumption Expenditure price index. Consistent with the recent literature on earnings risk, we use as our measure of earnings the residuals obtained from regressing log real earnings on a quadratic on

⁵Since the SIPP records up to two jobs at a time for any individual, we define the main job as the one in which the worker spent the most hours, and break ties using earnings.

⁶See Card et al. (2013), Halvorsen et al. (2020) and Guvenen et al. (2021), for a similar definition of employer stayers.

potential experience, education and month dummies. After this step we clean reporting errors in the residual earnings data by dropping the bottom and top 2% of the wave-frequency earnings sample and drop imputed earnings (see Online Appendix A for further discussion about measurement error in earnings changes). Following the literature we focus on year-to-year earnings growth. We construct annual earnings as the sum of all (residual) monthly earnings observed during the past 12 months. This measure therefore includes any zero earnings associated with the months in which an individual was unemployed.

For employed workers who do not change employers or occupation and are continually at work, earnings growth are computed by comparing the one year earnings to the next. In the case of employer and/or occupation movers, we compare earnings differences in the year prior to a transition to earnings in the year following the wave in which the transition occurred.⁷ As mentioned above, some of these transitions include either an *EU* or *UE* observation. Because some *EU* transitions entail re-employment after an unemployment spell that lasted for more than a year and hence these workers are associated with very low or zero earnings, we use the inverse hyperbolic sine differences $\Delta_{i,t+1} = \log(w_{i,t+1} + (w_{i,t+1}^2 + 1)^{1/2}) - \log(w_{i,t} + (w_{i,t}^2 + 1)^{1/2})$ rather than log differences to compute annual earnings changes, where $w_{i,t}$ denotes the earnings of individual i at time t . The inverse hyperbolic sine differences are approximately the same as log differences except in the case of very low and zero earnings.

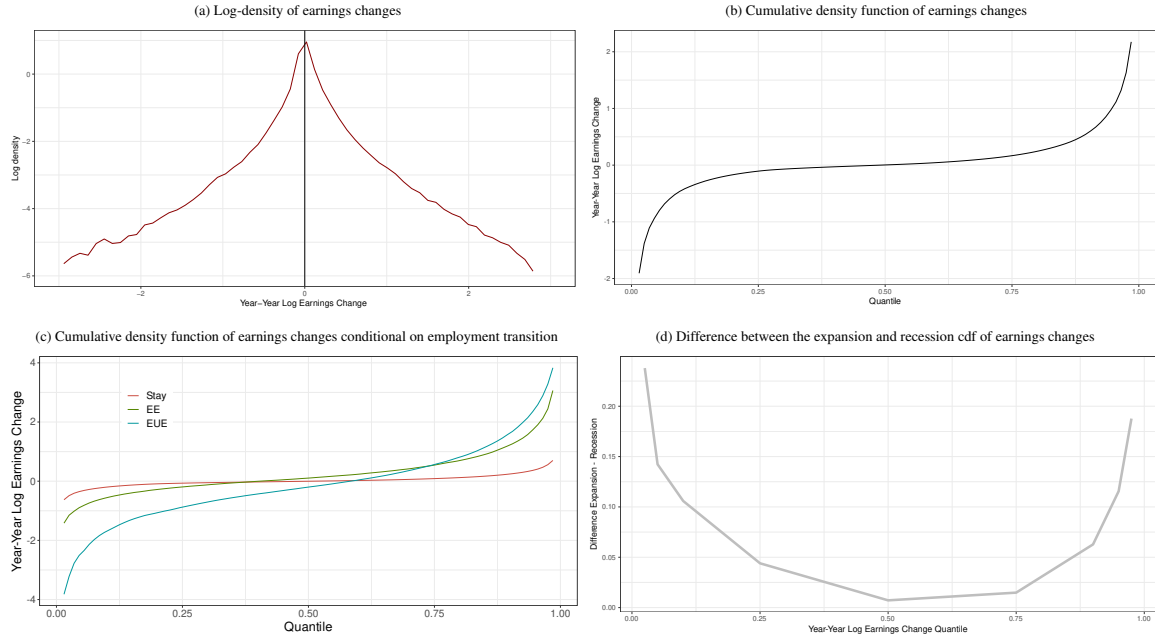
2.2 The earnings distribution in the cross-section and over the cycle

Figure 1a depicts the distribution of annual earnings growth pooling all years in our sample. Guvenen et al. (2021) shows that key features of this distribution is that it is left skewed and very leptokurtic, with approximately Pareto-distributed tails, hence we plot its log density to better visualize this latter property and show that both properties also are present in our sample. In Online Appendix B we present the density in levels to highlight its leptokurtosis. Figure 1b instead presents the earnings growth cdf, where we graph the earnings changes in the y-axis and the quantiles of the distribution in the x-axis. This figure further emphasises that large earnings changes lie at the bottom and top tails of the distribution. Figure 1c conditions this cdf by whether the earnings change was associated with the worker staying with the same employer or it was obtained through an *EE* or *EUE* transition. We observe that employer stayers (who represent the vast majority of workers in our sample) exhibit positive and negative earnings changes that are concentrated around zero. In contrast, employer movers have much more dispersed earnings changes and are primarily the ones behind the distribution's fat tails, as shown in Figure 1a. The large negative earnings changes are mainly due to workers who experienced *EUE* transitions, while the large positive earnings changes are due to workers who experienced *EE* transitions or came back into employment to complete an *EUE* transition.⁸

⁷In this case earnings in the reference period are measured without the aforementioned stability restriction that the respondent was at work every week of the reference period.

⁸In Online Appendix B we show that we arrive to a similar conclusion by analysing earnings growth conditional on previous earnings. The negative earnings/wage changes among workers who made an *EE* transitions have been

Figure 1: Earnings growth distribution in the cross-section and over the cycle



Note: The annual earnings growth distribution is constructed for the sample period 1990-2013. It is based on residual earnings after controlling for potential experience, education and month dummies. Section 2.1 presents the details of the definition of earnings and worker transitions. Recessions are defined as periods in which the HP-filtered unemployment rate is in the top 20% of realizations.

Figure 1d depicts the cyclical changes of the cdf of the earnings growth distribution.⁹ This graph is the main data pattern of our analysis. It shows how the earnings growth distribution depicted in Figure 1b changes over the cycle by subtracting for a given quantile (x-axis) the earnings changes in periods of recessions from those in expansions (y-axis). As the median of the expansion and recession CDFs are essentially zero (0.006 and -0.003, respectively), the bottom half of these distributions represent earnings losses and the top half represent earnings gains. The pronounce U-shape depicted in Figure 1d implies that earnings losses (the left of the distribution) are larger in recession and earnings gains (the right of the distribution) are larger in expansions. For example, at the 25th percentile of the earnings growth distribution recessions have an earnings loss that is about four percentage points higher than in expansions, while at the 75th percentile expansions have an earnings gain that is about two percentage points higher than in recessions. This property has already been documented by Guvenen et al. (2014) and shows that the annual earnings growth distribution is mainly characterised by its procyclical skewness. If recessions were to bring a level-change in the observed earnings growth distribution then we would observe a horizontal line at the average loss in earnings. If recessions instead were to bring countercyclical variance, we would observe a downward sloping curve crossing zero, such that losses in recessions would be worse meaning positive values and gains in recessions

documented previously in Jolivet et al. (2006), among many others. Using the SIPP information on the reasons for leaving a job, we find that these changes are associated with both voluntary and involuntary transitions. Below we discuss this feature in the context of our model.

⁹Here recessions are defined as periods in which the HP-filtered unemployment rate is in the top 20% of realizations. We also analyse the cyclicity of the earnings change distribution by defining recessions as periods defined by the NBER, without a meaningful change to our conclusions.

would be higher meaning negative values. In Online Appendix B we show that the same U-shape that characterises procyclical skewness in earnings growth also arises when considering changes in hourly wages. We show that although changes in hours work do matter in determining cyclical earnings changes, particularly for the larger earnings losses during recessions, changes in hourly wages can explain a large part of the U-shape.

The key message from Figure 1 is therefore that the earnings risk as measured by the earnings growth distribution is characterized by long, thick tails that exhibit very large fluctuations over the business cycle and that this arise from changes in both hourly wages and hours worked. Hence, to understand cyclical changes in the earnings growth distribution, we need to understand its tails and the transitions that comprise it. We now present novel evidence showing the importance of occupational mobility in accounting for the behaviour of these tails.

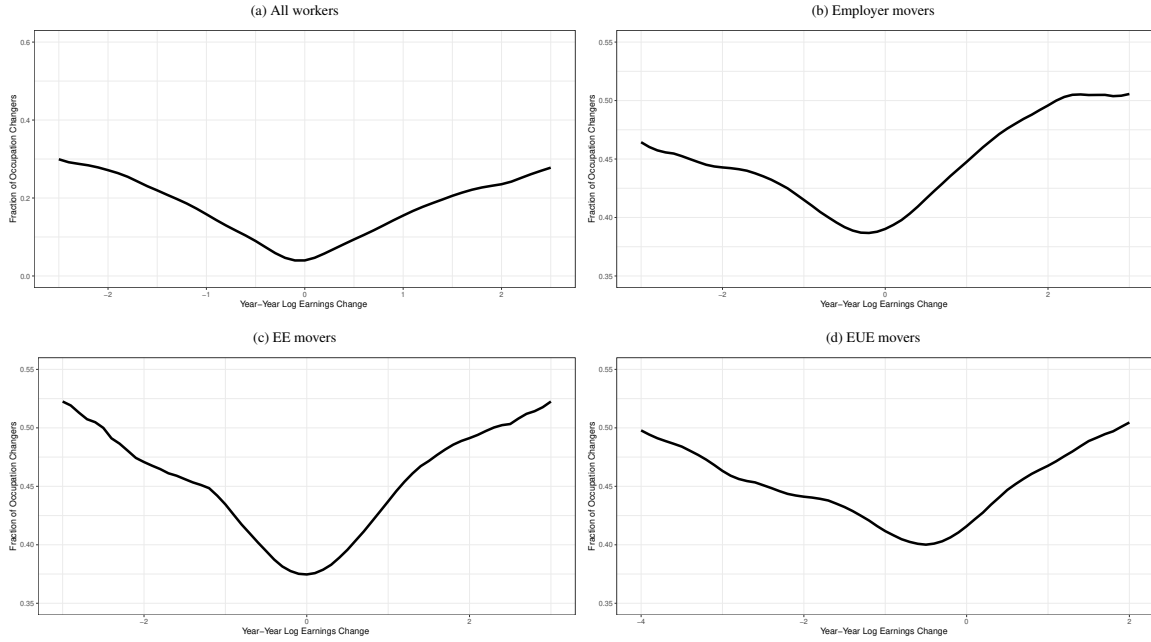
2.3 The importance of occupational mobility

Occupation switching and cross-sectional earnings growth Figure 2 depicts the probability of an occupational change associated with a given value of earnings growth. Figure 2a shows that when pooling together employer movers and stayers the fat tails of the annual earnings growth distribution depicted in Figure 1a are associated with a high probability of an occupational change (see also Guvenen et al., 2021). Small earnings changes centred around zero, however, are associated with a much smaller probability of an occupational change. Although the large difference in the probability of an occupational move can be accounted for by the larger propensity to change occupations among employer movers, we observe that the same pattern remains when only considering employer movers, as shown in Figure 2b. Figures 2c and 2d further show that this pattern holds even when analysing separately those individuals who changed employers through an *EE* or *EUE* transition. Among *EE* transitions we observe a near symmetric rise in the probability of occupational mobility and the size of the earnings losses and earnings gains. Among *EUE* transitions, however, we observe a faster rise in the probability of an occupational change among those workers who had positive earnings growth relative to those with negative earnings growth.¹⁰ This evidence thus shows that large negative or positive earnings changes are associated with a higher propensity to change occupations than smaller earnings changes, even after an employer change has been taken into account.

To investigate further the role of occupational change on the the tails of the earnings growth distribution, we calculate the variance of this distribution as the proportion of the sum of squared deviations $\sum_K \sum_{o \in K} (\Delta w_o - E_{pop}[\Delta w])^2$ that originates from a group K of workers who share an occupation and employer transition (for example, the set of workers with an *EE* transition and an occupation switch), and divide it by the overall sum of squared deviations $\sum_{pop} (\Delta w_o - E_{pop}[\Delta w])^2$. We find that occupation movers contribute about 50% of the overall variance of earnings growth,

¹⁰Note that for the *EUE* transitions plot we have top censored the earnings gains at +2 as the probability of occupational mobility for workers with larger earnings gains is estimated from a small number of observation in this region. The resulting pattern presents an essentially constant relationship between very large earnings gains (above +2) and the probability of an occupational change. This can be somewhat visualised in Figure 2b.

Figure 2: Earnings growth distribution and occupation mobility



Note: The probability of occupational change is computed as the proportion of workers with a given earnings change who took a job in a different task-based occupation: Non-routine cognitive, Routine cognitive, Non-routine manual and Routine manual. Earnings changes are based on residual earnings after controlling for potential experience, education and month dummies.

even though the share of occupation movers in our sample is about 17%, where the biggest share of this contribution arises from *EUE* transitions. As we move away from the tails and consider progressively the variance among those between the 0.95-0.05, 0.9-0.1 and 0.75-0.25 percentiles, the contribution of occupation movers and employer movers diminishes, reaching 15% when considering the interquartile range.¹¹

Occupation ladder Earnings changes associated with occupational mobility can arise from workers moving from occupations with higher average earnings to occupations with lower average earnings, and vice versa, and from workers changing occupations due to idiosyncratic factors. To investigate the extent of these two sources of earnings growth, we derive conditional occupational earnings averages by estimating an earnings regression, which includes a quadratic for potential experience, dummies for education, gender and race and a set of dummies for occupational categories. We treat the coefficient on the occupational dummies as the occupation-wide earnings effect. For each occupation switcher we then calculate the difference between these coefficients at their source and destination occupation. A positive (negative) difference can be considered as climbing up (falling down) the occupational ladder. We perform this exercise using the 4 task-based categories and the 22 occupation categories of the two-digit 1990 SOC. Here we present the results based on the task-based categorisation to keep the benchmark analysis based on the same level of aggregation. Online Appendix

¹¹Note that workers on average do seem to gain by switching occupations. In particular, employer stayers and *EE* movers who switched occupations gain on average about 4% and 26%, while employer stayers and *EE* movers gain on average about 1% and 16% when not switching occupations. For those workers who experienced *EUE* transitions, we find that occupation movers and stayers lose a very similar amount (around -34%).

B shows the results are very similar when using two-digit occupations, implying that the relative importance of an occupational ladder is not due to aggregation.

Figure 3: Occupational ladder



Note: Occupational switchers are ranked by their earnings growth in the horizontal axis. For each rank the vertical axis depicts the mean, median, 90th and 10th percentiles of the distribution of the differences in occupational earnings effects. A similar pattern also holds at two-digit level.

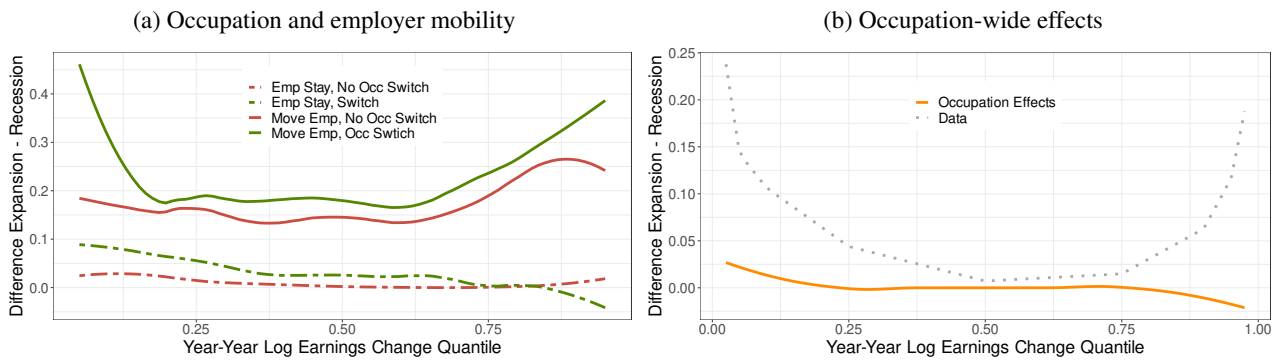
Figure 3 ranks occupational switchers by their earnings growth (x-axis) and relates each of these workers' rank to the associated distribution of the differences in occupational earnings effects (y-axis). For each rank we show the mean, median, 90th and 10th percentiles of the latter distribution. The bottom and top curves show the sequence of 10th and 90th percentiles obtained from each of these distributions. The middle two curves show the sequence of median and means. This exercise is done for each type of labour market transition. Figure 3b shows an occupation ladder for *EUE* movers, where the sequence of mean, median, 90th and 10th percentiles are all upward sloping. This implies that workers with higher earnings growth are more likely to move to higher paying occupations. Figure 3c also shows a similar ladder for employer stayers. For *EE* movers, however, the estimates depicted in Figure 3a only show a noticeable occupational ladder at the mean and the 10th percentile. Although occupational ladders are visible, the key feature of these figures is that these ladders have a very subdued effect on the distribution of changes in the occupational earnings effect. That is, large earnings changes are associated with movements both up or down the occupational ladder across all types of labour market transitions. Further, Figure 3d shows that the magnitude of the difference in the occupation effects is small compared to the magnitude of earnings growth. It does so by depicting comparing the distribution of changes in occupational earnings effects among all occupation movers to the CDF of the cross-sectional earnings growth distribution. These findings suggest that the more important factor for occupational change is the idiosyncratic motive rather than average occupation-

wide earnings differences when explaining the earnings growth distribution.¹²

Occupation switching and cyclical earnings growth To highlight the role of occupation mobility in the cyclical change of earnings growth, Figure 4a decomposes the difference in the earnings growth distribution between expansions and recessions by whether workers were occupational movers or stayers conditioning on employer change. As in Figure 1d we subtract the expansion from the recession earnings growth distribution.

Among employer stayers we observe that occupation movers have larger earnings losses during recessions than occupation stayers. However, these losses are modest, increasing only slightly towards the bottom end of the distribution. These workers also exhibit larger and modest earnings gains in expansions relative to occupational stayers. At the very top of the distribution (above the third quartile) we observe the opposite pattern. In recessions those who changed occupations within the same employer receive larger earnings gains relative to occupation stayers. Overall this leads to the earnings growth distribution of employer stayers/occupation movers to exhibit countercyclical variance. For occupational/employer stayers we instead observe a slight level-change between recession and expansions, whereby earnings losses in recessions are of similar magnitude as earnings gains in expansions.

Figure 4: Cyclical earnings growth distribution by occupation mobility



Note: In the left panel the cyclical change in the annual earnings growth distribution is constructed separately for those workers who (i) simultaneously change employers and occupations, (ii) change employers but did not change occupations, (iii) did not change employers and change occupations, (iv) or did not change either of these dimensions. In the right panel the cyclical changes in earnings are computed based on changes in the estimated occupation fixed effects. This is done separately for expansion and recessions. In either case we fix the quantile of the distribution along the horizontal axis and we subtract the associated expansion from the recession earnings growth along the vertical axis.

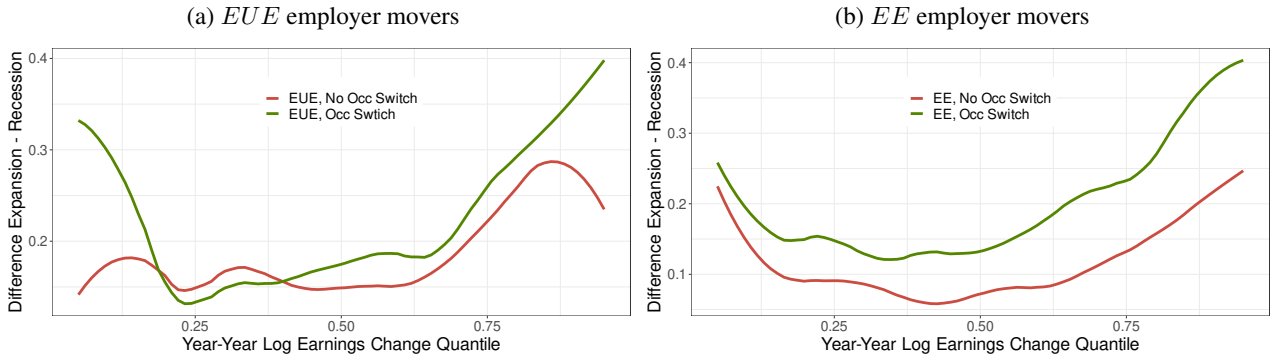
In contrast, employer movers who also changed their occupation exhibit larger losses in recessions relative to occupation stayers and these losses become much larger at the bottom end of the distribution. Further, they also exhibit larger gains in expansions, which also become even larger at the top end of the distribution. These cyclical features create a U-shape pattern that mimics quite closely that

¹²Note that this finding does not contradict that of Groes et al. (2014). They show that (i) the probability of an occupational move is higher among those workers whose wages are at bottom or top end of their current occupation's wage distribution relative to those whose wages are around the median; and (ii) that among those workers who switched occupations, those whose wages are at the bottom (top) end of their pre-separation occupation's wage distribution have a higher probability of moving down (up) the occupation ladder. Although not shown here we also find evidence supporting these patterns in the SIPP when using 2-digit occupations from the 1990 SOC.

of the overall earnings growth distribution (Figure 1d) and implies that the earnings growth distribution of employer/occupation movers is characterised by procyclical skewness. Among those who changed employers but did not change occupations we observe a nearly constant increase in earnings losses during recessions, while an increasing pattern of larger earnings gains during expansions. That is, these workers do not seem to experience an increase in the downside earnings risk of changing employers during recessions.

Figure 4b investigates whether the joint effect of occupation/employer mobility in determining the procyclical skewness of the earnings growth distribution is due to workers moving more often to better or worse occupations at different points of the business cycle. We use the same procedure as described above to generate occupation-wide fixed effects and compute the difference in the fixed effects associated with workers switching occupations. We do this separately for expansions and recessions and then subtract the difference at each quantile. The figure shows that occupation-wide earnings differences do not seem to explain the observed cyclical changes in the earnings growth distribution. Instead, the figure points to earnings change due to idiosyncratic occupational mobility in explaining procyclical skewness. In Online Appendix B, we show that this result is not driven by using the aggregate task-based categories but also holds when using two-digit occupations.

Figure 5: The cyclicity of earnings growth - employer movers and occupation movers/stayers



Note: Among those who change employers and conditional on the type of employer transition (*EUE* or *EE*), the cyclical change in the annual earnings growth distribution is constructed separately for those workers who simultaneously change occupations and for those who did not change occupations. For each case we fix the quantile of the distribution along the horizontal axis and we subtract the associated expansion from the recession earnings growth along the vertical axis. Earnings changes are based on residual earnings after controlling for potential experience, education and month dummies.

Figure 5 decomposes the cyclicity of the earnings growth distribution among employer movers depicted in Figure 4a by the type of transition: *EUE* or *EE*. It is clear from this figure that those who simultaneously changed their occupation and employer either through an *EUE* or *EE* transition exhibit earnings growth distributions characterised by their procyclical skewness. Perhaps unsurprisingly, during recessions the earnings losses among those occupation/*EUE* movers are more pronounced than among occupation/*EE* movers. During expansions, however, both type of transitions lead to similar increases in earnings gains. Among employer movers/occupation stayers only those who experienced an *EE* transition are the ones that exhibit an earnings growth distribution characterised by procyclical skewness. Those who changed employers through an *EUE* transitions exhibit a level change in their earnings losses during recessions, but an increase in earnings gains dur-

ing expansions. Since we observe more *EUE* transitions during recessions, these pattern dominates the cyclical change of the downside earnings risk of the earnings growth distribution among employer movers/occupation stayers as depicted in Figure 4a.

In Online Appendix B we show that this exact same patterns also hold when considering changes hourly wages instead of earnings. This is important as occupation mobility is typically associated with changes in hours worked and the U-shape relationship we find among occupation/employer movers could be mostly driven by changes in hours worked. Our results shows this is not the case. We find that although changes in hours are relative more important among *EUE* occupational movers, changes in hourly wages can still explain a large part of the U-shape relationship characterising the procyclical skewness of the earnings growth distributions among *EUE* and *EE* occupation movers.

To provide a quantitative picture of the importance of occupation and employer movers in shaping the cyclicity of earnings growth, we apply the method developed by Halvorsen et al. (2020) in order to linearly decompose the skewness of the earnings growth distribution over the cycle. We present the results in Online Appendix B. This method shows that employer and occupation movers explain about 60% of the increase in the left-skewness of the earnings growth distribution over the business cycle, even though this group of workers represent no more than 5% of all observations (20,937 person/year observations). *EUE* occupation movers explain 48%, while *EE* occupation movers explain 12%. This is in contrast to *EUE* and *EE* occupation stayers who represent 28% and 8%, respectively, and make up 8% of all observations (36,396 person/year observations).

Taken together the above evidence strongly suggests that the procyclical skewness observed in the overall earnings growth distribution depicted in Figure 1d can be traced back to a combination of larger recessionary earnings losses among those occupation movers who changed employers through *EUE* (and to a lesser extent *EE*) transitions, and the larger earnings gains during expansions among those occupational movers who changed employers through either *EUE* or *EE* transitions. Further, these cyclical changes in earnings are a result of changes in both hourly wages and hours worked among employer/occupation movers. Underlying these transitions we find a prominent role for idiosyncratic occupation/employer mobility and not occupation-wide differences.

3 Theoretical Framework

The patterns documented above highlight the need to investigate occupation in addition to employer mobility to better explain the cyclicity of earnings risk. These patterns also pose two key questions. Do cyclical changes in the earnings risk arise from the occupation dimension or from the employer dimension? Does the importance of each dimension arise from cyclical changes in the returns to mobility or from cyclical changes in worker flows? We now develop a structural model of job search to answer these questions.

3.1 Environment

Time is discrete $t = 0, 1, 2, \dots$. A mass of infinitely-lived, risk-neutral workers with common discount rate β is distributed over a finite number of occupations $o = 1, \dots, O$. At any time t , workers within a given occupation can be either employed or unemployed and can differ in the following dimensions: an idiosyncratic firm-match productivity, ϵ_t , an idiosyncratic occupation-match productivity, z_t , and occupation-specific human capital, x_h .

The idiosyncratic ϵ -productivity determines how well the worker is doing with the employer he is currently working (see Mortensen and Pissarides, 1994). This productivity follows a common and exogenous first-order stationary Markov process, with transition law $\Gamma(\epsilon_{t+1}|\epsilon_t)$. The ϵ -productivity realizations affect a worker only in employment and will allow us to generate employer-to-employer mobility. The idiosyncratic z -productivity represents a worker's "career match" and determines how well he is doing in the current occupation (see Neal, 1999). In the spirit of Braxton and Taska (2022), the z -productivities can be thought of measuring (in reduced form) the evolving distance between the skill requirements of an occupation and a worker's own innate skills. These productivities also follow a common and exogenous first-order stationary Markov process, with transition law $F(z_{t+1}|z_t)$. The z -productivity realizations affect workers both in employment and unemployment and will drive idiosyncratic occupational mobility and any earning growth associated with this type of occupational mobility.

In addition, workers accumulate occupational-specific human capital through a learning-by-doing process. In period t an employed worker with human capital level x_h increases his human capital to x_{h+1} with probability $\chi_e(x_{h+1}|x_h)$, where $h = 1, \dots, H$. A worker's occupational-specific human capital may also depreciate with unemployment. An unemployed worker with human capital level x_h decreases his human capital to x_{h-1} with probability $\chi_u(x_{h-1}|x_h)$ $h = 1, \dots, H$.¹³ Workers can in principle also accumulate firm-specific human capital during employment spells. However, the empirical results of Kambourov and Manovskii (2009b) suggest that the returns to this type of human capital will be zero in the context of our model and will not play a meaningful role in our analysis. Therefore to simplify notation we do not consider this dimension.¹⁴

Business cycles are modelled through fluctuations in the economy-wide productivity, where A_t denotes this aggregate productivity, which follows a first-order stationary auto-regressive process. We also allow some occupations to be more attractive than others in terms of their occupation-wide

¹³The accumulation of occupation-specific human capital, which is shared across all workers with a given occupation, aims at capturing productivity increases that arise from performing the occupation-specific tasks. For simplicity we assume that occupation human capital x_h is orthogonal to the career match z -productivities.

¹⁴Kambourov and Manovskii (2009b) show that returns to seniority within a firm (wage-tenure effects) become negligible once they also include returns to occupational experience in their "Mincerian" style wage regressions after correcting for the well known endogeneity problems arising when estimating wage-tenure profiles (see Altonji and Shakotko, 1987). As we use their estimates to calibrate the occupational-specific human capital process (see Section 4), consistency will require also using their estimates to calibrate any firm-specific human capital process. This will result in the latter only having one level and not playing any meaningful role in our result. The same argument will apply if we were to use alternative estimates such as those in Sullivan (2010) who further consider within-employer occupation mobility as done in here.

productivities. Let $p_{o,t}$ denote the occupation-wide productivity of occupation o at time t and assume it follows an auto-regressive process. Let $\mathcal{P}_{O,t} = \{p_{o,t}\}_{o=1}^O$ denote the vector containing all the occupation-wide productivities at time t .

Production and earnings Firms are passive agents in our model. They only use labour in the production process under a constant return to scale technology and face no capacity constraints when meeting workers. The output of a worker characterised by (z_t, x_h, ϵ_t) in occupation o is given by $y(A_t, p_{o,t}, \epsilon_t, z_t, x_h)$. This production function is increasing and continuous in all of its arguments. To keep the analysis as parsimonious as possible we assume that workers' earnings are also represented by an exogenous function of their current productivities, $\hat{w}(A_t, p_{o,t}, \epsilon_t, z_t, x_h)$, which is strictly increasing and continuous in all of its arguments. This assumption implies that any changes in hours worked and hourly wages are captured (in reduced form) through EE and EUE transitions and changes in workers' productivities.¹⁵ Note, however, that the evolution of worker's output and earnings in the model remain endogenous as workers make employer and occupational transitions decisions that affect their productivities. When unemployed workers receive b each period.

Job destruction Jobs can break-up endogenously as workers may decide to quit to another employer in the same or in a different occupation. This can also occur if workers' productivities fall sufficiently such that they prefer to become unemployed within their occupation. Once unemployed, a worker can decide whether to change occupation or not. In addition, jobs can be destroyed due to exogenous reasons. We consider two sources of exogenous job destruction: one due to worker-firm idiosyncratic reasons and the other due to worker-occupation idiosyncratic reasons. Let $\delta_\epsilon(A_t)$ denote the probability that an employed worker loses his firm-match productivity and is forced to transit into unemployment within his occupation. Similarly, let $\delta_z(A_t)$ denote the probability that a worker loses his occupation-match productivity and is forced to change occupation through unemployment. In the spirit of Huckfeldt (2022) we interpret the latter as "obsolescence shocks".¹⁶

Searching within and across occupations Unemployed and employed workers face a probability $\lambda_u(A_t)$ and $\lambda_e(A_t)$ of meeting a firm when searching for jobs, respectively. Once a meeting takes

¹⁵As discussed in Section 2 and shown in Online Appendix B the cyclical dynamics of the earnings growth distribution and the importance of simultaneous occupation/employer mobility in driving these dynamics arise from changes in both hourly wages and hours worked. Although workers' productivities are typically closely linked with hourly wages, they can also determine workers' choice of hours. Instead of modelling these two dimensions separately we simplify and assume a flexible functional form to match the earnings process observed in the data without assuming a particular wage/hours determination process.

¹⁶This type of job destruction shock allows the model to capture the large, negative earnings changes observed across individuals who were previously employed in high paying jobs, experienced an EUE transition with an occupational switch and ended up re-employed in a lower paying job. Huckfeldt (2022) uses "obsolescence shocks" to match a similar feature in his data. The key difference is that in our setting the obsolescence shock occurs during employment, rendering the worker unemployed and forcing him to change occupations; while in Huckfeldt (2022) this shock hits the worker during unemployment. Further, we allow the obsolescence shock to vary with the business cycle to capture that in recessions we observe more frequent and larger earnings losses accompanied by occupational switching. Guvenen et al. (2021) also document large negative earnings shocks among those with high earnings and labels them "disaster" shocks. Anecdotal evidence suggests that this type of shock has been somewhat common during the Covid-19 pandemic recession, where some individuals, for example, musicians were forced into unemployment and take jobs as postal workers or delivery drivers (see <https://mbspmag.com/arts-and-culture/har-mar-superstar-becomes-mailman/>).

place, a worker draws an initial firm-match productivity $\tilde{\epsilon}$ from $\Gamma(\epsilon, A)$. We allow $\Gamma(\cdot)$ to shift according to the state of the business cycle to capture the possible cyclical changes in the quality of the worker-firm relationships and hence the returns to employer mobility.¹⁷ If the worker finds the firm-match productivity sufficiently attractive, production takes place until the match is destroyed. Otherwise, the worker remains in his current employment state and waits until another meeting takes place. A worker who is searching in an occupation different from his current occupation faces job arrival rates $\lambda_u^c(A_t), \lambda_e^c(A_t)$. These are potentially different from $\lambda_u(A_t), \lambda_e(A_t)$ to capture that job finding in a different occupation can be relatively more difficult (compared to those who are already in the occupation) at particular phases of the business cycle.

Search across occupations is modelled following an imperfect directed search technology in the spirit of Fallick (1993). Occupation mobility entails the benefit of re-starting a worker's z -productivity process, but it comes at the loss of the accumulated occupational-specific human capital upon reallocation. Workers draw their initial career match in any occupation from $F(z, A)$, which we also allow to shift according to the state of the business cycle in order to capture possible cyclical changes in the quality of the worker-occupation relationships. Given differences in occupation-wide labour market conditions $p_{o,t}$, workers are not indifferent about which occupation to draw the new z -productivity. We assume that within each period a worker can receive at most one z -productivity. The worker is endowed with a unit measure of search intensity which he must divide across occupations in order to maximise his chances of receiving such a z . In particular, a worker leaving occupation o has to decide which proportion $s_{\tilde{o}}^i$ of his time to devote to obtain a z -productivity from occupation $\tilde{o} \neq o$, where $i = U, E$ denotes the worker's employment status. Let S^i denote a vector of $s_{\tilde{o}}^i$ for all $\tilde{o} \in O^-$, where O^- denotes the set of remaining occupations such that $\sum_{\tilde{o} \in O^-} s_{\tilde{o}}^i = 1$. The probability that a worker, currently in occupation o , receives a new z from an occupation \tilde{o} is given by $\alpha(s_{\tilde{o}}^i, o)$, where $\alpha(\cdot, o)$ is a continuous, weakly increasing and weakly concave function with $\sum_{\tilde{o} \in O^-} \alpha(s_{\tilde{o}}^i, o) \leq 1$, for all $o \in O$. The latter implies that the probability of not receiving a new z is given by $1 - \sum_{\tilde{o} \in O^-} \alpha(s_{\tilde{o}}^i, o)$.

In the case of receiving the new z -productivity the workers changes occupation. If this worker is unemployed he starts with human capital x_1 and searches for jobs the same period in which reallocation takes place. If no meeting takes place, the worker remains unemployed in the new occupation. If the worker is employed he also starts with human capital x_1 and is able to search for jobs straight away. Follow Jolivet et al. (2006) and Bagger and Lentz (2017) once this employed worker meets a firm, with probability γ he is able to decide whether to accept the new job or not; and with probability $1 - \gamma$ he is forced to accept the new job, as long as it offers a payoff above the value of unemployment. If no meeting takes place this worker stays with his current employer but in the new occupation. The worker's idiosyncratic productivities and occupational human capital then evolve as described above. Workers can decide to reallocate once again, but they must sit out one period in their new employ-

¹⁷The notion that there are cyclical changes in the quality of worker-firm or worker-occupation matches is not new. Moscarini (2001), for example, motivates his analysis on excess reallocation by arguing that "Depressed labour markets perform less well and produce noisier allocations: ceteris paribus, worker-firm matches formed at times of high unemployment are of relatively low quality". He then presents evidence in favour of this hypothesis.

ment state before doing so. In case the worker does not receive a z -productivity he retains his current occupation and employment status for the rest of the period and then starts the occupation mobility process once again.

Timing and state space The timing of the events is summarised as follows. At the beginning of the period the new values of A , \mathcal{P}_O , z , x and ϵ are realised. After these realisations, the period is subdivided into four stages: separation, reallocation, search and matching, and production. At this point it is useful to define the vector $\Omega_t = \{A_t, \mathcal{P}_{O,t}\}$. To simplify notation further we leave implicit the time subscripts, denoting the following period with a prime. We also leave implicit the dependence of output, wages and the exogenous match break-up and job finding probabilities on productivities as described above.

We now turn to formalise unemployed and employed workers' decision problems and derive the corresponding value functions. To save space the associate flow equations and earnings distribution are relegated to Online Appendix C.

3.2 Worker's problem

Unemployed workers Consider an unemployed worker currently characterised by (z, x_h, o) . The value function of this worker at the beginning of the production stage is given by

$$\begin{aligned} W^U(x, z, o, \Omega) = & b + \beta \mathbb{E}_{x', z', \Omega'} \left[(1 - \delta_z(A')) \max \left\{ R^U(x', z', o, \Omega'), \left[(1 - \lambda_U(A')) W^U(x', z', o, \Omega') \right. \right. \right. \\ & + \lambda_U(A') \int_{\underline{\epsilon}}^{\bar{\epsilon}} \max \left\{ W^E(\tilde{\epsilon}, x', z', o, \Omega'), W^U(x', z', o, \Omega') \right\} d\Gamma(\tilde{\epsilon}, A') \Big] \Big\} \\ & \left. + \delta_z(A') R^U(x', \underline{z}, o, \Omega') \right]. \end{aligned} \quad (1)$$

The value of unemployment consists of the flow benefit of unemployment b , plus the discounted expected value of being unemployed at the beginning of next period's reallocation stage. With probability $(1 - \delta_z(A))$ the unemployed worker decides to search across occupations or not. This decision is captured by the choice between the expected net gains from drawing a new \tilde{z} in another occupation and the expected payoff of remaining in the current occupation. The latter entails meeting a firm with probability $\lambda_U(A)$, drawing a firm-match productivity $\tilde{\epsilon}$ and deciding whether to accept it or not. However, with probability $\delta_z(A)$ the worker draws the lowest z -productivity, \underline{z} , and searches across occupations with probability one.

The expected net value for an unemployed worker of searching across occupations, $R^U(\cdot)$, is given by

$$\begin{aligned} R^U(x, z, o, \Omega) = & \max_{s_o^U} \sum_{\tilde{o} \in O^-} \alpha^U(s_o^U) \int_{\underline{z}}^{\bar{z}} \left[\lambda_U^c(A) \int_{\underline{\epsilon}}^{\bar{\epsilon}} \max \left\{ W^E(\tilde{\epsilon}, x_1, \tilde{z}, \tilde{o}, \Omega), W^U(x_1, \tilde{z}, \tilde{o}, \Omega) \right\} d\Gamma(\tilde{\epsilon}, A) \right. \\ & \left. + (1 - \lambda_U^c(A)) W^U(x_1, \tilde{z}, \tilde{o}, \Omega) \right] dF(\tilde{z}, A) + \left(1 - \sum_{\tilde{o} \in O^-} \alpha^U(s_o^U) \right) W^U(x, z, o, \Omega), \end{aligned} \quad (2)$$

where the maximization is subject to $s_o^U \in [0, 1]$ and $\sum_{\tilde{o} \in O^-} s_o^U = 1$. Note that this expression incorporates the value of immediately searching for an employer in the new occupation under the assumption that once a worker receives a new z in a new occupation, he will move to that occupation and enter the search and matching stage during the current period. It is only when the worker does not receive a new z that he remains unemployed in the old occupation and goes directly to the production stage.

Employed workers Now consider an employed worker currently characterised by (z, x, ϵ, o) . The expected value of employment at the beginning of the production stage is described by

$$W^E(\epsilon, x, z, o, \Omega) = w + \beta \mathbb{E} \left[\delta_z(A') R^U(x', \underline{z}, o, \Omega') + \delta_\epsilon(A') W^U(x', z', o, \Omega') + (1 - \delta_z(A') - \delta_\epsilon(A')) \right. \\ \left. \times \max \left\{ W^U(x', z', o, \Omega'), \max \left\{ R^E(\epsilon', x', z', o, \Omega'), \hat{W}^E(\epsilon', x', z', o, \Omega') \right\} \right\} \right]. \quad (3)$$

The value of employment consists of the earnings, plus the discounted value of being employed at the beginning of next period's separation stage, where the worker faces exogenous job loss with probability $\delta_\epsilon + \delta_z$. Otherwise, the worker must decide to separate into unemployment or stayed employed. If the worker remains employed, he enters the reallocation stage and must decide whether to search for jobs in a different occupation or not.

In the case of no reallocation, the expected value of employment at the beginning of the search and matching stage in the current occupation is

$$\hat{W}^E(\epsilon, x, z, o, \Omega) = \int_{\underline{\epsilon}}^{\bar{\epsilon}} \gamma \lambda_E(A) \max \left\{ W^E(\tilde{\epsilon}, x, z, o, \Omega), W^E(\epsilon, x, z, o, \Omega) \right\} d\Gamma(\tilde{\epsilon}, A) \\ + \int_{\underline{\epsilon}}^{\bar{\epsilon}} (1 - \gamma) \lambda_E(A) \max \left\{ W^E(\tilde{\epsilon}, x, z, o, \Omega), W^U(x, z, o, \Omega) \right\} d\Gamma(\tilde{\epsilon}, A) \\ + (1 - \lambda_E(A)) W^E(\epsilon, x, z, o, \Omega), \quad (4)$$

where, with probability $\gamma \lambda_E$ the worker can decide to accept or reject the new job offer. The job acceptance decision truncates the distribution of ϵ such that $W^E(\tilde{\epsilon}, \Omega) > W^E(\epsilon, \Omega)$. Note that in this case the worker's fall back position is keeping his job with the current employer and value of ϵ . With probability $(1 - \gamma) \lambda_E$, however, the worker is forced to take the new ϵ' in a different employer, as long as this gives a value higher than unemployment. This allows us to capture within-occupation employer transitions that are associated with earnings cuts and rises. With probability $1 - \lambda_E$ the worker does not meet a new employer and remains in his current state.

If instead the worker decides to reallocate, the expected net value of searching across occupations is given by

$$R^E(\epsilon, x, z, o, \Omega) = \max_{S^E} \sum_{\tilde{o} \in O^-} \alpha^E(s_{\tilde{o}}^E) \left(\int_{\underline{z}}^{\bar{z}} \left[\int_{\underline{\epsilon}}^{\bar{\epsilon}} \left(\gamma \lambda_E^c(A) \max \left\{ W^E(\tilde{\epsilon}, x_1, \tilde{z}, \tilde{o}, \Omega), W^E(\epsilon, x_1, \tilde{z}, \tilde{o}, \Omega) \right\} \right. \right. \right. \\ \left. \left. \left. + (1 - \gamma) \lambda_E^c(A) \max \left\{ W^E(\tilde{\epsilon}, x_1, \tilde{z}, \tilde{o}, \Omega), W^U(x_1, \tilde{z}, \tilde{o}, \Omega) \right\} \right) d\Gamma(\tilde{\epsilon}, A) \right. \right. \\ \left. \left. + (1 - \lambda_E^c(A)) W^E(\epsilon, x_1, \tilde{z}, \tilde{o}, \Omega) \right] dF(\tilde{z}, A) \right) + \left(1 - \sum_{\tilde{o} \in O^-} \alpha^E(s_{\tilde{o}}^E) \right) W^E(\epsilon, x, z, o, \Omega), \quad (5)$$

where the maximization is subject to $s_o^E \in [0, 1]$ and $\sum_{\tilde{o} \in O-} s_o^E = 1$. Conditional on drawing a z -productivity from another occupation \tilde{o} , the worker meets a new employer with probability $\gamma \lambda_E^c(A)$, draws a new value of ϵ and decides whether to accept it or not. If the offer is rejected, the worker remains with his current employer but switches to occupation \tilde{o} . With probability $(1 - \gamma) \lambda_E^c(A)$, however, the worker is forced to take the new value of ϵ' with a new employer in occupation \tilde{o} , as long as the value of this employment remains above the value of unemployment in the new occupation. Otherwise, the worker transits into unemployment in occupation \tilde{o} . With probability $1 - \lambda_E^c(A)$, the worker remains with his current employer, retains his current value of ϵ , but changes to occupation \tilde{o} .

The formulation in (5) allows us to capture within and across employer occupational mobility in a simple way, with two key properties. First, we do not decouple in which occupation the worker is earning and in which occupation he is searching. This property allows us to ease computational burden by simplifying the state space, perhaps coming at a cost of realism. Second, through EE transitions workers can end up with positive or negative earnings growth when switching occupations as documented in Section 2. In the model this captures the increased risk inherent to occupation mobility: with probabilities $(1 - \gamma) \lambda_E^c$ or $(1 - \lambda_E^c)$ the worker gives up on his fallback option $W^E(\epsilon, x, z, o, \Omega)$. Yet, absent a δ_z shock, this risk is not forced upon the worker. At the earlier (reallocation) stage, workers can opt to not expose themselves to this risk by deciding not to search across occupations, as described in the last maximand of equation (3).

4 Quantitative Analysis

We estimate our model using simulated method of moments. Assume a period to be equal to a month and set the discount rate $\beta = 0.997$. Set $O = 4$, consistent with the task-based categories (Non-Routine Cognitive, Routine Cognitive, Non-Routine Manual and Routine Manual) used to aggregate occupations in Section 2, such that $o \in \{NRC, RC, NRM, RM\}$. The following functional form assumptions yield a set of parameters that we jointly estimate on data from the SIPP based on the annual earnings growth patterns documented in Section 2 as well as on worker labour market flows across employers and occupations computed from the SIPP for our observation period, 1990-2013. We relegate the details of the estimation procedure to Online Appendix D.

4.1 Parametrization

We parametrise the aggregate productivity shock, A , by a two-state Markov process that can be represented by A_I , a cyclical indicator function which takes the value of one when the economy is in an expansion and zero otherwise.¹⁸ The occupation-wide productivity shock, p_o , is such that $p_{o,t+1} = \tilde{p}_o + \rho_p p_{o,t} + v_{o,t+1}$ for all $o \in O$ where \tilde{p}_o is a time-invariant occupation specific productivity and v_o is assumed to be white noise with variance σ_p . Occupation-specific human capital accumulation follows a two state process. Inexperienced workers start in an occupation with productivity x_1

¹⁸This approximation of an AR(1) process simplifies computation and is also consistent with our data work that considered two discrete phases of the business cycle. Expansions, when the aggregate productivity takes its high value, occurs 80% of the time.

and become experienced with probability $\chi(x_2)$ to reach x_2 . To simplify we assume no occupational human capital depreciation while unemployed.

The meeting probabilities are parametrised as $\lambda_i(A_I) = \lambda_{1,i}A_I + (1 - A_I)\lambda_{0,i}$, where $i = U, E$ is the employment status indicator. For those who have just switched occupations, the (potentially) different meeting probabilities are given by $\lambda_i^c(A_I) = \lambda_{1,i}^cA_I + (1 - A_I)\lambda_{0,i}^c$. Recall that, conditional on meeting with a new employer and receiving a job offer, $1 - \eta$ denotes the probability with which the employed worker would not be able to refuse such a job offer (as long as its optimal to remain employed). The exogenous occupation obsolescence shock, δ_z , and firm separation shock, δ_ϵ , are given by $\delta_z(A_I) = \delta_{1,z}A_I + (1 - A_I)\delta_{0,z}$ and $\delta_\epsilon(A_I) = \delta_{1,\epsilon}A_I + (1 - A_I)\delta_{0,\epsilon}$.

The worker-occupation match specific productivity shock is such that $E[z_{t+1}|z_t] = \rho_z z_t + (1 - \rho_z)v_{z,t+1}$, where v_z is distributed following $\tilde{F}(\cdot)$, which we parametrize as a Weibull distribution with shape and scale parameters ν_z and σ_z . Upon a decision to switch occupations new values of z are then drawn from $F(z, A) = A_I\tilde{F}(z) + (1 - A_I)[\omega_z\tilde{F}(z) + (1 - \omega_z)T(z, z_A)]$, where during recessions the reallocating worker draws z from a convolution between $\tilde{F}(\cdot)$ and a tilted uniform distribution $T(z, z_A) = \frac{zz_A+1}{\tilde{z}(\tilde{z}z_A+1) - \tilde{z}(zz_A+1)}$. The worker-firm match specific productivity ϵ remains constant during the duration of the match and only changes upon a voluntary or involuntary employer-to-employer transition. This is motivated by the small variation in annual earning changes observed among employer stayers. Workers draw a new ϵ from $\Gamma(\cdot)$, which is also allowed to shift during recessions in a similar way as $F(\cdot)$. In this case, $\Gamma(\epsilon, A) = A_I\tilde{\Gamma}(\epsilon) + (1 - A_I)[\omega_\epsilon\tilde{\Gamma}(\epsilon) + (1 - \omega_\epsilon)T(\epsilon, \epsilon_A)]$, where $\tilde{\Gamma}(\cdot)$ is parametrized as a Double Exponentially-Modified Gaussian distribution (a normal distribution with exponential tails on both sides) with mean normalised to zero, central variance σ_ϵ and shape parameters lt_ϵ, rt_ϵ governing the exponential distributions describing the left and right tails, respectively. The distributional choice for $\Gamma(\cdot)$ allows for arbitrary skewness and kurtosis, which are both features of the earnings growth distribution in the data. Further, the mixtures $F(\cdot)$ and $\Gamma(\cdot)$ allows us to change the skewness during recessions through z_A and ϵ_A , and to change the variance through ω_z and ω_ϵ . These distributional assumptions are important because they determine the model's earnings growth distributions conditional on previous labour market transition. As we only observe accepted matches in the data, they are crucial for identification.¹⁹

Particularly important is how we design the functions $\alpha^i(\cdot)$, because they determine how workers switch and search across occupations. Our aim is to respect the observed gross and net mobility flows across occupations. We let $\alpha^i(s_{\tilde{o}}) = \alpha_0 e^{\alpha_{\tilde{o}}\alpha_1^i s_{\tilde{o}}^{1-\alpha_1^i}}$, where occupation $\tilde{o} \in O^-$ denotes the search direction, i the worker's employment status and $s_{\tilde{o}}$ denotes search intensity. To show why this functional form is useful consider $i = U$ such that the first order condition of (2) is given by $\alpha'(s_{\tilde{o}}^*) (\Psi^U(\tilde{z}, x_1, \tilde{o}, \Omega) - W^U(z, x_h, o, \Omega)) = \mu$, where μ denotes the multiplier on $\sum_{\tilde{o} \in O^-} s_{\tilde{o}} = 1$.

¹⁹In the estimation we verify that the support of the discretised $F(\cdot)$ and $\Gamma(\cdot)$ distributions are observed in employment for at least some workers. The reason is that we want to investigate whether shifts in the returns to mobility directly affect earnings changes of at least some employed workers. Without this restriction it would be difficult for the model to distinguish, in extreme cases, a higher mass of very bad returns realizations (ϵ or z) that are always rejected from lower job offer arrival rates. This should be kept in mind when considering the flow/return decomposition of earnings changes over the cycle in Section 5.

Together with the latter feasibility constraint, it implies that the search intensity takes a form similar to the Gumbel-distributed random utility model. If the directional terms $\alpha_{\tilde{o}}$ are all equal we obtain that the optimal value of $s_{\tilde{o}}$ is given by

$$s_{\tilde{o}}^* = \frac{e^{\frac{1}{\alpha_1^U} \log(\Psi^U(\tilde{z}, x_1, \tilde{o}, \Omega) - W^U(z, x_h, o, \Omega))}}{\sum_{\tilde{o} \in O^-} e^{\frac{1}{\alpha_1^U} \log(\Psi^U(\tilde{z}, x_1, \tilde{o}, \Omega) - W^U(z, x_h, o, \Omega))}}, \quad (6)$$

and

$$\begin{aligned} \Psi^U(\tilde{z}, x_1, \tilde{o}, \Omega) &= \int_{\underline{z}}^{\tilde{z}} \left[\lambda_U^c \int_{\underline{\epsilon}}^{\tilde{\epsilon}} \max \left\{ W^E(\tilde{\epsilon}, \tilde{z}, x_1, \tilde{o}, \Omega), W^U(\tilde{z}, x_1, \tilde{o}, \Omega) \right\} d\Gamma(\tilde{\epsilon}) \right. \\ &\quad \left. + (1 - \lambda_U^c) W^U(\tilde{z}, x_1, \tilde{o}, \Omega) \right] dF(\tilde{z}). \end{aligned}$$

In Online Appendix D we derive a similar expression for $s_{\tilde{o}}^*$ when considering search across occupations among employed workers. In either case, with $\alpha_{\tilde{o}}$ differing across occupational directions we end up with a multiplier that scales the search direction. Combined with the additive noise which we use to smooth our approximate solution (see Online Appendix D), this means the search direction problem takes a similar form to a nested, multinomial logit discrete choice model. This formulation is useful not only because is simple to estimate but it also allow us to capture net flows across occupations through two distinct channels. The parameter α_1^i determines the extent to which workers adjust their search due to $p_{o,t}$ differences. It inform us whether search is more or less directed. A value closer to zero implies search becomes more directed as differences in the returns across occupations get amplified; while as α_1^i goes to infinity search becomes increasingly random, as workers weight all occupations more equally. The parameter $\alpha_{\tilde{o}}$ scales the probability of moving from occupation o to occupation $\tilde{o} \neq o \in \{NRC, RC, NRM, RM\}$ to capture any skill constraints that restrict worker mobility between occupation pairs.²⁰ The parameter α_0 is a scaling factor to guarantee the proper arrival rates, but does not play a role in the marginal choice of search direction. In Online Appendix D we provide a more detailed derivation of (6) and its relationship with the more used Gumbel-distributed random utility model.

Finally, we assume that per period earnings follow a standard Mincer formulation,

$$\log w_{o,t} = \gamma_w \log w_{o,t-1} + (1 - \gamma_w)(A_t + \tilde{p}_o + p_{o,t} + x_h + z_t + \epsilon_t), \quad (7)$$

where $\gamma_w \in [0, 1]$ captures the extent of stickiness such that with probability γ_w the current earnings remains the same as last period's earnings. Annual earnings are then obtained by summing up monthly earnings, taking into account that during periods of unemployment the worker receives earnings of zero. Note, however, that the job acceptance and occupational mobility decisions of the unemployed depend on the per period payoff when unemployed b .²¹

²⁰For example, it seems unlikely that many workers in *RM* jobs will end up in *NRC* jobs even if the latter occupations offer a much higher average pay.

²¹Given that the formulation of earnings described in (7) implies that past earnings affect workers' job acceptance and reallocation decisions, in the estimation we include earnings in the state space when solving the value functions described in Section 3. As it will be shown, however, we estimate a small value for γ_w , such that past earnings have little effect in practise on workers' mobility decisions.

4.2 Estimation

The above functional forms yield several parameters to estimate comprised of: the set that governs the arrival of job opportunities $\{\lambda_{0,i}, \lambda_{1,i}, \lambda_{0,i}^c, \lambda_{1,i}^c\}_{i=U,E}$, the set $\{\delta_{0,z}, \delta_{1,z}, \rho_z, \nu_z, \sigma_z, \omega_z, z_A, \bar{z}, \underline{z}\}$ that governs the idiosyncratic worker-occupation productivities, the set $\{\delta_{0,\epsilon}, \delta_{1,\epsilon}, \eta, \sigma_\epsilon, lt_\epsilon, rt_\epsilon, \omega_\epsilon, \epsilon_A, \bar{\epsilon}, \underline{\epsilon}\}$ that governs the idiosyncratic worker-employer productivities, the set $\{\rho_p, \sigma_p, \tilde{p}_{NRC}, \tilde{p}_{RC}, \tilde{p}_{NRM}, \tilde{p}_{RM}\}$ that governs the occupation-wide productivities, the set of occupational human capital accumulation $\{x_1, x_2, \chi(x_2)\}$ and the set of directional parameters across occupations $\{\alpha_0, \alpha_1^U, \alpha_1^E, \alpha_{NRC}, \alpha_{NRM}, \alpha_{RM}, \alpha_{NRM}\}$, and finally, the set that governs the aggregate productivity process and payments $\{\gamma_w, b\}$.

To ease computational burden we first normalise x_1 to one, set $\chi(x_2)$ such that human capital accumulation occurs on average after 5 years of occupational tenure and choose x_2 to match the 12% 5-year returns to occupational tenure reported by Kambourov and Manovskii (2009b).²² We also set $b = 0.4$ to match a 40% replacement ratio (see Shimer, 2005). The aggregate productivity process parameters are set to the values of the autocorrelation and unconditional variance of output per worker as observed in the US during the period of study, similar to Shimer (2005), such that its persistence is 0.9580 and variance is 0.0090.²³ We also set $\bar{z} = \bar{\epsilon} = 3$ and $\underline{z} = \underline{\epsilon} = -3$ to generate the idiosyncratic productivity grid.

After fixing the above parameters we recover the remainder ones following a two-step procedure in which we split the parameter set between an inner and outer loop. Given values for the outer loop parameters, we can directly calibrate those in the inner loop such that their values match *exactly* the targeted moments. The inner loop contains the productivity levels \bar{p}_o which we set to match the task-based occupation fixed effects, where the regression to obtain these fixed effects in the simulated data is exactly the same as the one used to residualize earnings in Section 2. This loop also contains the directional parameters of the $\alpha(\cdot)$ function, α_{NRC} , α_{RC} , α_{NRM} and α_{RM} . These are set to match the relative contribution of each occupation o on total net flows.²⁴ We then iterate on the values of the remaining (outer loop) parameters using simulated method of moments, adjusting the inner loop parameters at each iteration. The outer loop parameters are obtained as the solution to

$$\text{Min}(\mathbf{M}^D - \mathbf{M}^S(\cdot))' \mathcal{W}(\mathbf{M}^D - \mathbf{M}^S(\cdot)),$$

where \mathbf{M}^D is a vector of data moments, $\mathbf{M}^S(\cdot)$ is a vector of the same moments obtained from model simulations, which are a function of the parameters to estimate as described above, and \mathcal{W} is a

²²The returns to occupational tenure correspond to the IV estimates of Kambourov and Manovskii (2009). Since we parametrise the returns to occupational human capital outside the simulation minimum distance procedure, we chose to target the IV estimates as they already control for the endogeneity bias present in the OLS returns.

²³Shimer (2005) uses the full post-war time period, which has the advantage of a longer time-series. While the process is different when considering the 80's onwards, using either one does not affect our conclusions. These data is available at a quarterly frequency from FRED, so we convert it to a monthly process, consistent with our simulation procedure.

²⁴Total net flows are calculated as the absolute value of the difference between inflows and outflows per occupation, summed up over all occupations, and divided by two as one person net inflow in some occupation is also counted as a net outflow some other occupation. Hence, the relative contribution of occupation o on total net flows is given by

$$\frac{|Inflow_o - Outflow_o|}{\sum_{o \in O} |Inflow_o - Outflow_o|}.$$

weighting matrix.²⁵

Targeted moments Table 1 and Figure 6 present the set of data moments M^D to estimate the outer loop parameters, which consist of transition flows and productivity moments and percentiles describing 6 (cross-sectional) earnings growth distributions, conditioning on employer and occupation transitions. We now present some heuristic identification arguments that justify our choice of moments, keeping in mind that all parameters need to be estimated jointly. In Online Appendix D we present a graphical representation of the identification of these parameters by evaluating the loss function at the optimal and showing its change in value as we perturbate each parameter in turn.

The parameters governing the arrival probabilities of employment opportunities are informed by the observed transitions probabilities across employers and occupations. In particular, the arrival probability of job offers among occupational stayers, $\lambda_{0,U}$, $\lambda_{0,E}$ and $\delta_{0,\epsilon}$ are informed by the average UE , EE and EU transition rates. The cyclical ratio of the UE , EE and EU rates inform the cyclical components of these offer arrival probabilities, $\lambda_{1,U}$, $\lambda_{1,E}$ and $\delta_{1,\epsilon}$. Similarly, the arrival probabilities of offers among occupational movers, $\lambda_{0,U}^c$, $\lambda_{0,E}^c$, are informed by the average probabilities of an occupational change conditional on the worker's labour market transition, EE , EUE or employer stayer. The parameters $\lambda_{1,U}^c$, $\lambda_{1,E}^c$ are then informed by the cyclical ratios of the occupational change probability through a EE and EUE transition. The ratio of the average unemployment duration among occupational movers and stayers further helps inform unemployed workers offer arrival probabilities.

To inform the curvature parameters α_1^U , α_1^E , we use the variance of the distribution of net flows conditioning on whether the occupation switch occurred through a EUE or EE transition. Since the functional form of $\alpha(\cdot)$ implies that a higher (lower) value of α_1^i leads to workers searching more (less) randomly, one obtains a negative relationship between α_1^i and the differences in each occupation's net flows. We measure dispersion in the latter through the variance of the distribution of these net flows. The shift parameter α_0 helps determine the level of occupational mobility and hence it is informed by the average occupational change probabilities conditional on workers' employment transition.

To inform the idiosyncratic worker-occupation match productivity process we use the earnings growth distributions of occupational movers. These are separated by whether workers made the occupational change through a EE , EUE or within employer transition as depicted in Figure 6. For each of these distributions we target the 10th, 25th, 50th, 75th and 90th percentile. In particular, the earnings losses observed upon an occupational transition through unemployment help inform δ_0^z and δ_1^z , the obsolescence shock parameters governing job loss with a forced occupational move: very large earnings losses among occupational movers imply forced moves because they are driven by high z match individuals who otherwise are unlikely to search outside of their current occupation. The earnings losses of those who changed occupations within their employers or through a EE transition also help to inform ω_z and z_A , the parameters that shift $F(\cdot)$, as it is in recessions where we observe more of these losses. To further inform these parameters we also target the cyclical change

²⁵Our weighting matrix aims to accomplish two key features: (i) normalise to the same scale the transition probabilities and their ratios over the business cycle; (ii) emphasis the role of the earnings change distributions in the minimisation procedure.

Table 1: Targeted moments in the estimation

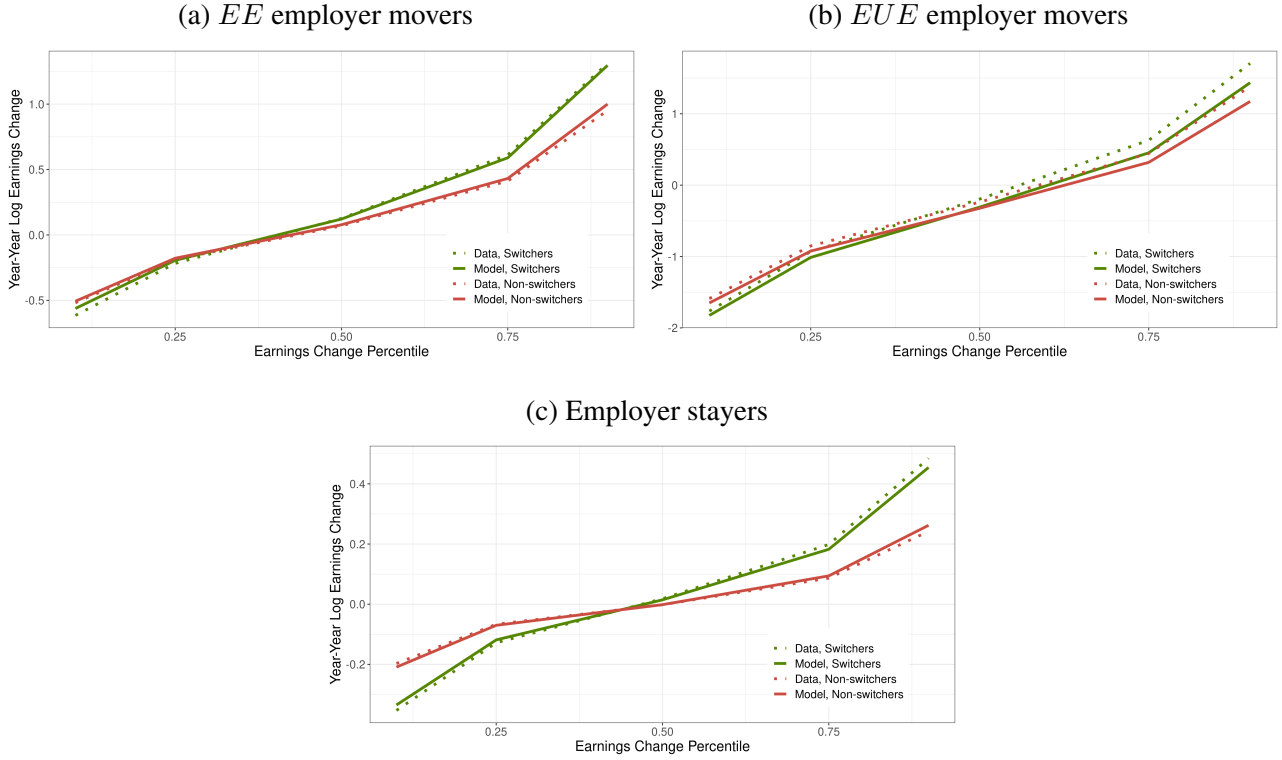
Moment	Model	Data	Moment	Model	Data
Employer Switching					
EE transition prob	0.0296	0.0340 (0.0003)	EE rate - expansion/recession ratio	1.1600	1.1846 (0.0469)
UE transition prob	0.3492	0.3947 (0.0025)	UE rate - expansion/recession ratio	1.0874	1.0876 (0.0244)
EU transition prob	0.0236	0.0223 (0.0002)	EU rate - expansion/recession ratio	0.7437	0.7460 (0.0333)
Occupation Switching					
Prob (Occ. change EE)	0.3107	0.2685 (0.0037)	Prob (Occ. change EE) - exp/rec ratio	1.1068	1.1068 (0.0196)
Prob (Occ. change EUE)	0.2867	0.2892 (0.0034)	Prob (Occ. change EUE) - exp/rec ratio	1.0670	1.0709 (0.0132)
U duration - Occ. movers/stayers ratio	1.2280	1.2709 (0.0215)	Prob (Occ. change Stayer)	0.0101	0.0107 (0.0002)
Variance (Net flows EE)	0.0293	0.0223 (0.0008)	Variance (Net flows EUE)	0.0235	0.0218 (0.0012)
Net flow to NRC	0.1849	0.1851	Net flow to RC	0.3395	0.3432
Net flow to NRM	0.2209	0.2201	Net flow to RM	0.2547	0.2516
Productivities					
NRC wage fixed effect	1.000	1.000	RC wage fixed effect	0.767	0.767
NRM wage fixed effect	0.608	0.608	RM wage fixed effect	0.803	0.803

Note: Bootstrapped standard errors in parenthesis for the moments used in the outer loop.

of the earnings growth distribution as depicted in Figure 1d. In this case we target the 2.5^{th} , 5^{th} , 10^{th} , 25^{th} , 50^{th} , 75^{th} , 90^{th} , 95^{th} , 97.5^{th} percentiles. The shape and scale of the $F(\cdot)$ distribution, ν_z and σ_z , and the persistence of the z -productivity process, ρ_z are informed by the remaining percentiles of the conditional earnings growth distributions of occupational movers. We also use these earnings growth distributions to recover the parameters of the occupation-wide productivity processes, ρ_p and σ_p . In particular we use the earnings growth distribution of those workers who changed occupations within their employers.

For the parameters governing the idiosyncratic worker-firm match productivities we use the earnings growth distributions of occupational stayers. The earnings losses of EE employer movers inform lt_ϵ , the left tail parameter of $\Gamma(\cdot)$, and the rate at which employed workers are forced to move employers within their occupations, η . The earnings losses among EUE employer movers inform ω_ϵ and ϵ_A , which shift $\Gamma(\cdot)$. As above we also use the cyclical change of the earnings growth distribution to inform these parameters. The earnings gains of EE and EUE employer movers then inform σ_ϵ and rt_ϵ , the variance and the shape of the upper tail of $\Gamma(\cdot)$. To inform the wage stickiness parameter, γ_w , we rely on the earnings growth distribution of those workers who did not change occupations or employers, especially the centre of this distribution.

Figure 6: Model fit for the conditional cross-sectional earnings growth distributions (cdf)



Note: The earnings growth distribution is computed separately for *EE* and *EUE* employer movers and employer stayers conditional on whether the worker switched occupation or not. Each graph present these distribution by showing the annual earnings growth value and the corresponding percentile. Our estimation targets the 10th, 25th, 50th, 75th and 90th percentiles of each of these distributions. In Online Appendix D we present the bootstrapped standard errors for these moments.

Parameter estimates Table 2 reports the estimated parameter values. As implied by their simulated standard errors, these parameters are precisely estimated (see also Online Appendix D). The estimated job offer arrival probabilities among workers searching for jobs in their current occupation, $\lambda_{0,U}$, $\lambda_{1,U}$, $\lambda_{0,E}$ and $\lambda_{1,E}$, together with their job acceptance decisions, capture the behaviour of the *UE* and *EE* transitions probabilities among occupational stayers. If a worker decides to search across occupations, he faces a job arrival probability with parameters $\lambda_{0,U}^c$, $\lambda_{1,U}^c$ or $\lambda_{0,E}^c$, $\lambda_{1,E}^c$ until a job is found. These values together with workers' search direction and job acceptance choices determine the *UE* and *EE* transitions probabilities among occupational movers. Indeed, the reason why we find a stronger procyclicality of occupational movers' arrival probabilities is because they have to be scaled by the parameters of the α function, which are less than one, to match the observed transition probabilities. Overall, the model is consistent with the observed procyclical behaviour of the aggregate *UE* and *EE* transition probabilities and the occupational mobility probabilities observed along the four task-based occupational categories we consider.²⁶

²⁶The estimated value of λ_1^U implies that the offer arrival probability of unemployed workers decreases in expansions among occupational stayers. This occurs as the estimation procedure uses this parameter to help match not only the cyclicity of the aggregate *UE* rate, but also to reproduce the cyclical change in the 2.5 percentile of the aggregate earnings growth distributions, which is largely affected by the cyclicity of the earnings distribution for *EUE* transitions. By setting this parameter to zero (i.e. no cyclical change in the job finding probability of occupational stayers), for

Table 2: Estimated parameter values

Job offer arrival				Employer-match productivities				Occupation-match productivities			
$\lambda_{0,U}$	0.8701 (0.0002)	$\lambda_{0,E}$	0.0935 (5.33E-05)	$\delta_{0,\epsilon}$	0.0025 (4.25E-06)	lt_ϵ	3.4350 (0.0010)	$\delta_{0,z}$	0.0084 (4.25E-06)	ν_z	7.7603 (0.0019)
$\lambda_{1,U}$	0.7051 (0.0002)	$\lambda_{1,E}$	0.1854 (7.44E-05)	$\delta_{1,\epsilon}$	0.0002 (2.38E-06)	rt_ϵ	1.4356 (0.0007)	$\delta_{1,z}$	0.0030 (4.23E-06)	σ_z	6.593 (0.0013)
$\lambda_{0,U}^c$	0.1669 (0.0002)	$\lambda_{0,E}^c$	0.0171 (5.33E-05)	η	0.2763 (0.0002)	ω_ϵ	0.9952 (0.0001)	ρ_z	0.0070 (1.95E-05)	ω_z	0.7573 (0.0002)
$\lambda_{1,U}^c$	0.5746 (0.0002)	$\lambda_{1,E}^c$	0.1716 (7.44E-05)	σ_ϵ	0.0347 (2.11E-05)	ϵ_A	-0.3076 (0.0004)			z_A	-0.7180 (0.0004)
Search direction across occupations				Occupation-wide productivities				Payments			
α_0	0.0403 (4.20E-05)	α_{NRC}	-0.4696	ρ_p	0.6168 (0.0001)	\tilde{p}_{NRC}	0 (normalize)	γ_w	0.0949 (2.11E-05)		
α_1^U	0.1398 (0.0002)	α_{RC}	0.5541	σ_p	0.0016 (5.22E-06)	\tilde{p}_{RC}	-0.2658				
α_1^E	0.2990 (0.0002)	α_{NRM}	-0.1796			\tilde{p}_{NRM}	-0.4976				
		α_{RM}	-0.0844			\tilde{p}_{RM}	-0.2189				

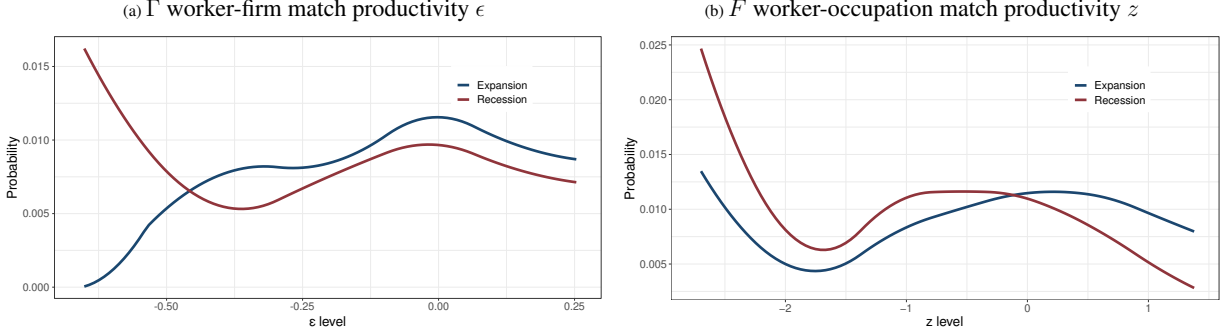
Note: Standard errors in parenthesis only correspond to the outer loop parameters. See Appendix D for details.

Figure 7 depicts the density of the $\Gamma(\cdot)$ and $F(\cdot)$ distributions implied by the estimated parameters governing the worker-firm match productivity ϵ and worker-occupation match productivity z in expansions and recessions. During recessions workers are more likely to draw from worse match productivities when starting in new employers and occupations. During expansion workers are more likely to draw better match productivities. This result is consistent with evidence that shows that during recession new worker-job matches worsen relative to expansions (see Moscarini, 2001). As we will demonstrate later, these cyclical changes in the productivities of new jobs contribute to the cyclical changes in the earnings growth distributions of employer/occupational movers as documented in Section 2. Another important characteristic of these distributions is that worker-occupation match productivities are estimated to be more dispersed than worker-firm match productivities, with variance of 1.041 and 0.585 expansions and of 1.085 and 0.540 in recessions, respectively, and an equal mean normalised to zero. This feature captures the increased earnings risk among occupational movers documented in Section 2. The parameters γ and ρ_z control how often workers move along the $\Gamma(\cdot)$ and $F(\cdot)$ distributions, such that with probability $1 - \gamma$ the worker is forced to change employer and draw a new ϵ and with probability ρ_z the worker draws a new z . Table 2 shows that re-draws of ϵ are about four times more likely than re-draws of z , suggesting that involuntary EE transitions are far more likely than changes in occupation match productivities.

The probability of falling from the $\Gamma(\cdot)$ and $F(\cdot)$ distributions are controlled by $\delta_{0,\epsilon}$, $\delta_{1,\epsilon}$ and $\delta_{0,z}$, $\delta_{1,z}$. As in the data, the probability of an EU transition increases in recessions because workers lose either their firm match or occupation match productivities more often than in expansions. The estimated values imply that workers fall about four times more frequently from $F(\cdot)$ than from $\Gamma(\cdot)$

example, we worsen the fit of the model in this dimension by a factor of five, generating much larger earnings losses than in the data during recessions.

Figure 7: The estimated Γ and F densities



during recessions. As demonstrated later, these shocks generate large earnings losses, particularly among forced occupational movers, as workers at the top of the earnings distribution lose their positions and must start searching for jobs from unemployment.

The $\alpha(\cdot)$ function parameters imply that workers' search across occupations is not fully directed in response to differences in occupation-wide productivities. As discussed in the previous sections, the value of α_1 controls the degree of directness such that a higher value implies that search is less directed. Table 2 shows that employed workers direct their search more than unemployed workers as $\alpha_1^U > \alpha_1^E$. An alternative way to evaluate the degree of directness is by using the effort exerted in searching for jobs in a given occupations, as this choice takes into account the frequency at which unemployed and employed workers encounter job opportunities. In this case, fully random search would arise when a worker targets all remaining occupations with an equal effort of $1/3$; while fully directed search occurs when the worker exert all his effort in targeting one occupation. A simple measure of directness based on search effort is therefore $(\max s_{\tilde{o}} - 1/3)/(1 - 1/3)$, where $\tilde{o} \in O^-$.²⁷ Using the latter we find that unemployed workers have a directness measure of 20.6%, once again lower than the one for employed workers of 21.9%. We also find that during recessions employed workers decrease their degree of search directness, while unemployed workers increase their search directness.²⁸

These measures suggest a mild degree of directness of search across occupations. This occurs due to the importance of the idiosyncratic worker-occupation and worker-firm match components in the overall payoff of new jobs. Since these productivities are randomly drawn from $F(\cdot)$ and $\Gamma(\cdot)$, workers do not differentiate as much between occupations. However, the extent to which individual occupations are targeted heavily depends on the occupation a worker is moving away from. In the estimated model both employed and unemployed workers that leave *RM* occupations have a directness measure of about 31.3% and 29.7%, while those that leave *NRM* occupations a measure of 28.8% and 25.8%. This occurs because *RM* and *NRM* workers mainly target *RC* occupations, which offer the highest chance of drawing a z -productivity due to the high estimated value of α_{RC} and offer a

²⁷For unemployed workers we measure $\max s_{\tilde{o}}$ through the entire unemployment episode, while for employed workers we measure it just before the *EE* transition. In both cases we aggregate all data across the business cycle.

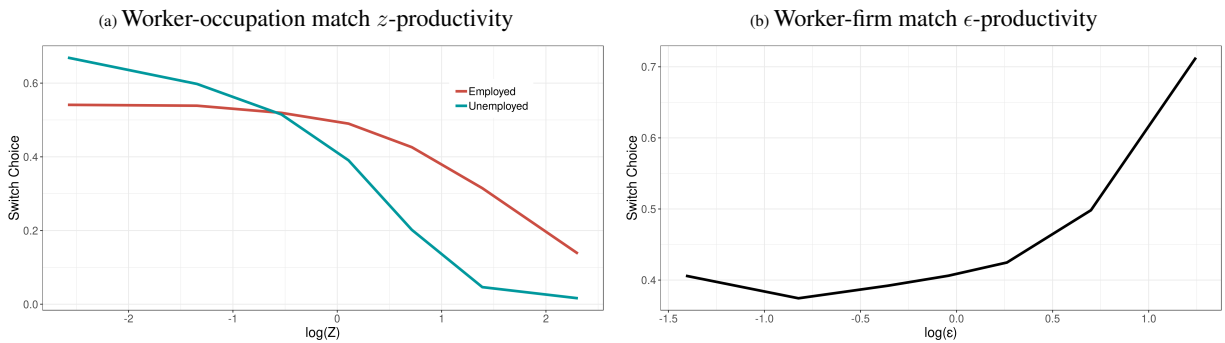
²⁸However, the cyclicity on this measure is very mild. Employed workers increase their search directness from 21.8% to 22.0% from recessions to expansions, while unemployed worker decrease their search directness from 20.5% to 20.7%.

sufficiently high \tilde{p} . On the other hand, workers leaving RC and NRC do face a stronger trade-off between the α_o and \tilde{p}_o and hence they tend to spread more the search, particularly those (employed and unemployed) workers leaving RC occupations who show a degree of directness of about 8%.

Finally, we find that occupation-wide productivity shocks are less persistent but more volatile than aggregate productivity shocks. We estimate (real) earnings also to be rather flexible with the stickiness parameter suggesting that earnings remain at their previous level with only a 5% probability every month.

Occupation and employer switching probabilities Workers in our model make three key decisions: (i) when to change occupations, (ii) when to change employers, and (iii) when to quit into unemployment. The estimated parameters shape the probability of these events and hence how likely are workers to move along the occupation and employer dimensions of the job ladder. Figure 8a shows that the probability of an occupational change monotonically decrease with workers' z -productivities. Employed workers in the lower half of the occupation match productivity range face around a 50% probability of an occupational switch every month. For higher values of z this probability decreases faster reaching around 15% at the highest value of z . Among the unemployed the probability of an occupational change starts even higher at around 60%, but decreases faster than for employed workers essentially disappearing at the highest productivity level. The reason for the latter is that the worker-firm match productivity also plays an important role in occupational switching among employed workers. Figure 8b shows the relationship between ϵ and the probability of occupational switching for employed workers. The positive relationship suggests that z and ϵ can substitute each other as workers move along these two dimensions of the job ladder. That is, workers with higher ϵ -productivities exhibit a higher probability of occupational switching when they also face a lower value of z and vice versa. This highlights that occupational mobility does not necessarily occur more often among those workers lower in the ϵ rank.²⁹ By aggregating across z and ϵ we recover the average occupation and employer switching probabilities shown in Table 1. Endogenous separations into unemployment account for 50.6% of all EU transitions and occur at lower levels of z and ϵ , where matches have a lower probability of survival when a negative productivity shock hits the job.

Figure 8: Probabilities of changing occupations

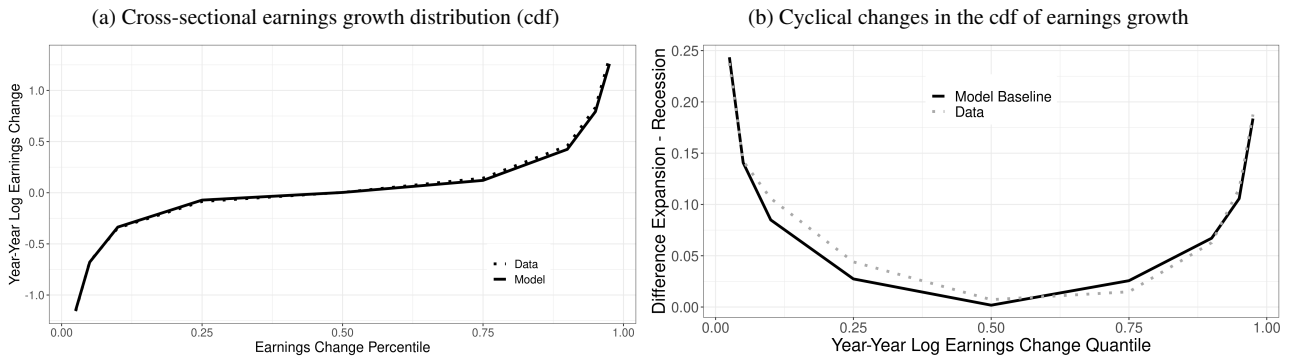


²⁹ Although not shown here, a similar pattern occurs when relating the probability of an employer change to ϵ and z . We obtain that this probability falls monotonically with ϵ , but increases with z .

4.3 Model fit

The model fits the data very well, given the amount of over-identification. Table 1 shows that it replicates the worker flow patterns in the cross-section and over the cycle among occupational movers and stayers as well as their search direction across occupations. Figure 6 shows quantile-by-quantile how well our minimum distance estimator does at bringing together the model and data on the cross-sectional cdf of the earnings growth distributions. Figures 6a and 6c show the model matches nearly perfectly the earnings growth distributions of occupational movers and stayers among *EE* employer movers and employer stayers. Figure 6b shows the fit of the earnings growth distribution among *EUE* employer movers is also very good, having slightly more trouble replicating these workers' earnings gains than their earnings losses. Note that aggregating within type of employer transitions recovers the earnings growth distributions depicted in Figure 1c in Section 2. Figure 9a then shows that when aggregating these conditional earnings growth distributions the model matches nearly perfectly the cdf of the observed cross-sectional earnings growth distribution, as depicted in Figure 1b in Section 2, even below the 10th and above the 90th percentiles which were not targeted.

Figure 9: Earnings growth distribution in the cross-section and over the cycle - model and data

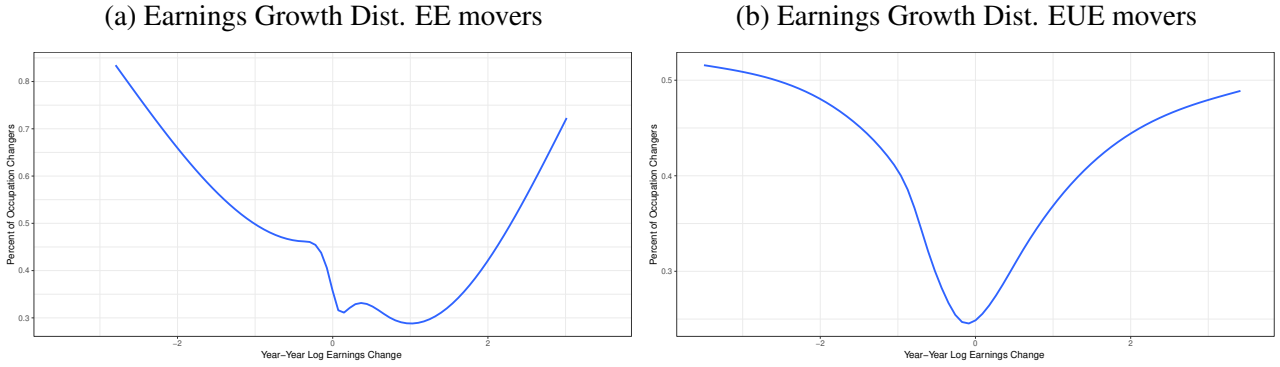


Note: The observed annual earnings growth distribution is constructed for the sample period 1990-2013. The left panel shows the cross-sectional earnings growth distribution (cdf) by plotting the associated earnings change to each selected percentile of the distribution. The percentiles used are 2.5th, 5th, 10th, 25th, 50th, 75th, 90th, 95th, 97.5th. The simulated distribution are constructed in the same way as in the data. The right panel shows the difference between the expansion and recessions earnings growth distributions as documented in Section 2, where we targeted the 2.5th, 5th, 10th, 25th, 50th, 75th, 90th, 95th, 97.5th percentiles.

The importance of occupation mobility in the tails of the cross-sectional earnings growth distribution can be observed when considering the model counterpart to the untargeted relationship between earnings growth and occupation mobility as depicted in Figure 2, Section 2. This figure documents that occupational movers contribute disproportionately to the largest changes in earnings, both positive and negative. In the model the relationship between earnings growth and occupation mobility among all workers is partially a result of the relatively fixed wages of the majority of employer/occupation stayers. It is less clear, however, that the model will be able to obtain the higher probability of occupational switching among those with the largest earnings changes when conditioning on *EE* or *EUE* transitions. There, the effect of drawing a new z -productivity has to dominate the other shock processes that drive extreme earnings changes, such as job loss from high on the ϵ dimension of the job ladder. Figure 10 shows that this is indeed what happens. Conditional on *EUE*

transitions the relationship between earnings growth and occupational mobility reaches its bottom near the mean of earnings growth, just below zero, as in the data. The same holds for the relationship between earnings growth and occupational mobility conditional on EE transitions, where the average worker makes an earnings gain. As the earnings growth associated with a particular transition moves away from the middle of the distribution, the chance it involves an occupation switch increases.

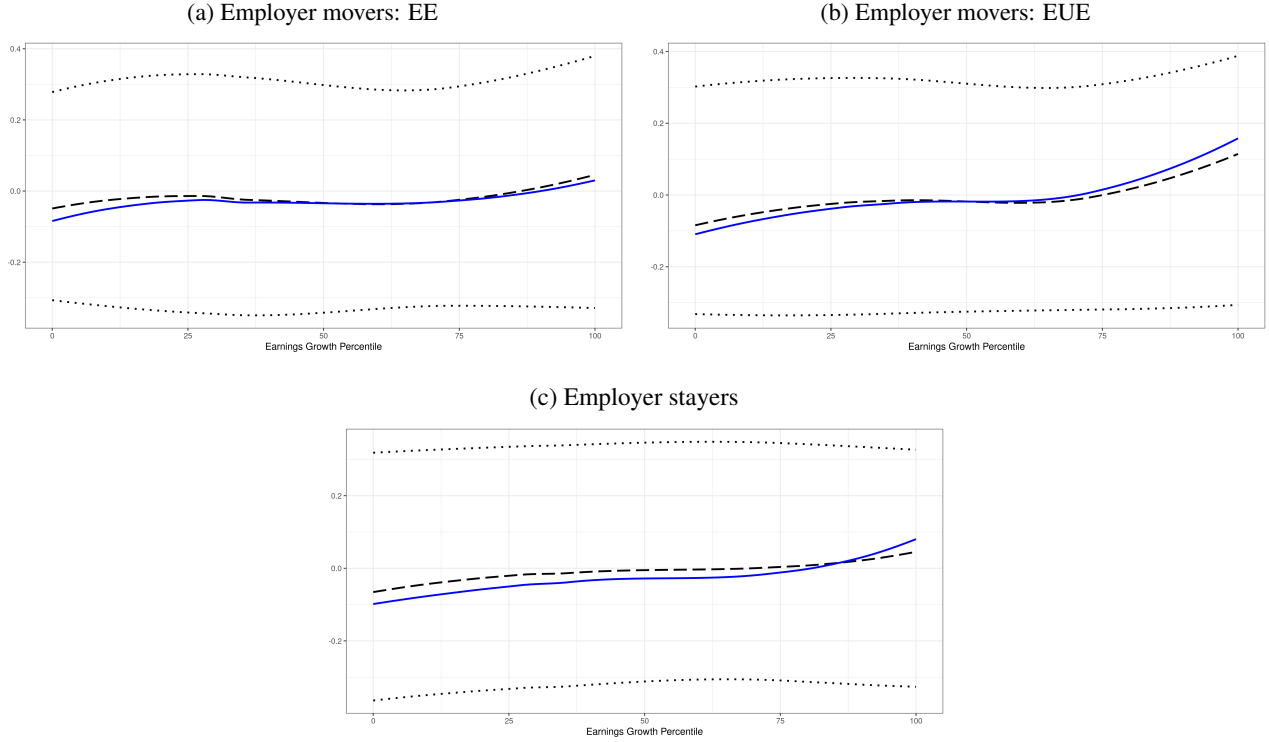
Figure 10: Model counter-part of earnings growth distributions



We can observe the importance of the idiosyncratic z -productivity relative to occupation-wide productivities in the model by computing the untargeted relationship between occupation fixed effects (occupational ladder) and earnings changes. Figure 11 shows that the model is able to reproduce the large earnings changes associated with larger movements both up or down the occupational ladder, as documented in Figure 3, Section 2. As in our empirical analysis, we compute conditional occupation earnings averages by estimating an earnings regression on a set of dummies for the task-based occupational categories. We treat the coefficients on these dummies as the occupation-specific earnings effect and compute the distribution of the difference between these coefficients at their source and destination occupation at each percentile earnings growth. Occupation-wide productivity differences imply that the mean, median, 90th and 10th percentiles are upward sloping and that workers with higher earnings growth are more likely to move to higher paying occupations. However, as in the data, these slopes are very weak and imply that workers main motive to switch occupation arises from the idiosyncratic worker-occupation match productivity. Moreover, this pattern is present when we focus on employer movers, EE or EUE , or employer stayers.

Over the business cycle, Figure 9b compares the targeted difference between the expansion and recession earnings growth distributions in the model and data. As in Figure 1d in Section 2, we condition on a quantile of the earnings growth distribution and then subtract expansion from recession values. Once again the model produces a nearly perfect fit across the distribution. Thus, our model is fully consistent with both the skewness and leptokurtosis of the cross-sectional earnings growth distribution and its prominent procyclical skewness, as well as with the observed worker mobility flows. Further, in Online Appendix D we present the equivalent of Figure 4a in Section 2 but using simulated data from the estimated model. As in the data, it shows that the procyclical skewness of the earnings growth distribution arises from individuals who change occupation and employer at

Figure 11: Model implied occupational ladder



Note: Model-based earnings changes of occupational switchers. Switchers are ranked by their earnings growth in the horizontal axis. For each rank the vertical axis depicts the mean (blue line), median (black line), 90th and 10th percentiles (dotted lines) of the distribution of the differences in occupational earnings effects.

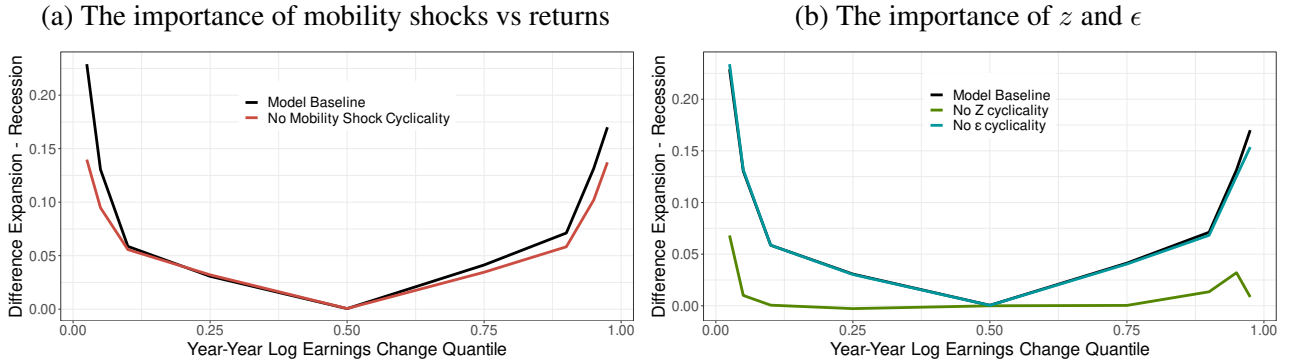
the same time. Also consistent with the data the cyclical change in the earnings growth distribution among employer movers/occupation stayers lies below that of employer/occupation movers and does not exhibit procyclical skewness. The cyclical change in the earnings growth distribution among employer stayers lies further below that of employer movers and also does not exhibit procyclical skewness either for occupation movers or stayers, as in the data.

5 Decomposition of Earnings Growth Dispersion Over The Cycle

What are the forces behind earnings growth and what makes it fluctuate over the cycle? To answer these questions we use our estimated model and decompose the cyclical changes in the earnings growth distribution, depicted in Figure 9b. The aim of this decomposition is to understand whether the importance of the occupation and employer components of a job arise from cyclical changes in the returns to mobility as captured by $F(\cdot)$ and $\Gamma(\cdot)$, or from cyclical changes in the mobility shocks captured by the two unemployment (or job loss) probabilities $\delta_z(A_I)$ and $\delta_\epsilon(A_I)$ and the four employment (or job finding) probabilities $\lambda_i(A_I)$ and $\lambda_i^c(A_I)$. This decomposition will in turn allows us to evaluate the relative importance of the occupation and employer components of a job in explaining the cyclicity of the earnings growth distribution, taking into account workers' employer and occupation reallocation decisions.

To evaluate the importance of these mobility shocks in relation to the returns to mobility, we perform a counterfactual exercise in which all mobility shocks are set to their expansion levels while returns to mobility are allowed to change between expansions and recessions periods. Figure 12a shows the result of this exercise, where the horizontal axis shows a set of quantiles of the earnings growth distribution and the vertical axis shows the difference in earnings growth at each quantile between expansion and recession as described in Section 2. The red line shows the contribution of (worse) recession returns to the cyclical change in the earnings growth distribution. The residual, the gap between black and red, is then attributed to cyclical changes of the unemployment/employment mobility shock processes. Consider first earnings growth that lie between the 10th and 90th percentiles of the distribution. Here we observe that nearly all of the changes in these percentiles arise from cyclical changes in returns to mobility. Further out into the bottom and top tails, however, the decomposition shows that mobility shocks have a more prominent role. Below the 10th percentile about 50% of the increase in earnings losses during recessions is due to cyclical changes in the mobility shock process. Above the 90th percentile up to 25% of the decline in earnings gains during recessions can be also attributed to cyclical changes in the mobility shocks.

Figure 12: Cyclical change in the earnings growth distribution



Note: Each panel shows the difference between the expansion and recessions earnings growth distributions using the 2.5th, 5th, 10th, 25th, 50th, 75th, 90th, 95th, 97.5th percentiles. In the left panel we show the impact of holding all the mobility shocks $\delta_z(A_I)$, $\delta_\epsilon(A_I)$, $\lambda_i(A_I)$ and $\lambda_i^c(A_I)$ constant at their expansion levels. In the right panel we show the impact of holding the $F(\cdot)$ or Γ constant at their expansion levels.

5.1 Idiosyncratic productivity shocks

Given the importance of changes in the returns to mobility in explaining the the earnings growth distribution over the cycle, we evaluate whether changes in $F(\cdot)$ or $\Gamma(\cdot)$ drive this result. Figure 12b shows the model's baseline fit and two counterfactuals, where either $F(\cdot)$ or $\Gamma(\cdot)$ are held at their expansions levels. For each counterfactual we let the remainder distribution and mobility shocks vary over the cycle as in the baseline estimation.

This counterfactual exercises demonstrate that worker-occupation match productivities are much more important in accounting for the cyclical change of the earnings growth distribution than worker-firm match productivities. Restricting the z -productivity distribution to be constant over the cycle,

significantly decreases the difference between the expansion and recessions earnings changes. Figure 12b shows that returns to idiosyncratic occupational mobility explains all the difference between expansions and recessions in the interquartile range. These returns also explain nearly all of what can be attributed to overall returns in both tails. In contrast, when restricting the ϵ -productivity distribution to be constant over the cycle, the model hardly suffers in its ability to generate the difference between the expansion and recessions earnings changes.

5.2 Job loss and job finding

The above analysis demonstrates that the ability of mobility shocks to generate cyclical changes in earnings growth arises from its impact on the left and right tails of the earnings growth distribution, particularly the left tail. We now consider this feature further by investigating what type of mobility shocks are the major contributors in each of these tails. If workers' choices played no role in determining the measure of workers across employers and occupations, we could simply turn off the cyclical elasticity parameters of the mobility shocks one at a time and compare the earnings growth distributions with and without them. Instead, however, since different workers in our model respond to the different unemployment/employment shocks differently, the outcomes of the model are not linear in the cyclicalities of these shocks, but instead there are non-trivial interactions between each of them and the associated returns. The solution to this problem is to use a Shapley-Owen decomposition, in which we take every combination of turning on and off the cyclical elasticities and then attribute to each shock its average effect. In this decomposition we still allow for cyclical changes in $F(\cdot)$ and $\Gamma(\cdot)$ along with the aggregate and occupation-wide shocks.

With the six cyclical elasticities of the mobility shocks, there are a total of 61 combinations of counterfactuals we need to evaluate in which all but one element of the mobility shocks are set constant at their expansion levels and the remainder is allowed to change between expansion and recessions periods as in the baseline estimation. To make these experiments explicit, define $\Lambda(\mathbb{I})$ as a function that takes a vector of indicators and gives a vector of mobility shocks. If all the indicators of the vector's elements are set to zero, the function $\Lambda(\mathbb{I})$ gives the values of the mobility shocks at their expansion values, thus not allowing the mobility shocks to change over the cycle. However, if the i^{th} element of the vector of indicators is set to one, then $\Lambda(\mathbb{I})$ implies that the mobility shock associated with i is allowed to vary with the cycle. Let $\tilde{G}_R(\Delta w|\Lambda(\mathbb{I}))$ denote the distribution of earnings growth in all recession periods given $\Lambda(\mathbb{I})$ and $\tilde{G}_E(\Delta w|\Lambda(\mathbb{I}))$ denote the distribution of earnings growth in all expansion periods given $\Lambda(\mathbb{I})$. We then can evaluate counterfactuals with $\Lambda(\mathbb{I})$ for every combination of indicators, and each time storing $\tilde{G}_R(\Delta w|\Lambda)$, $\tilde{G}_E(\Delta w|\Lambda)$ and computing their difference. Notice that these counterfactuals not only affect the mobility shocks, but also the value functions and hence the workers' mobility choices.

The top row of Table 3 presents the baseline model's difference between the expansion and recession earnings growth distributions. The final row, labeled "Total" is the portion of this that is attributed to the cyclicalities of all mobility shocks. That is, it is the wedge between the black and the red lines

Table 3: Contribution of cyclical mobility shocks in the tails

Quantile	0.025	0.05	0.10	0.90	0.95	0.975
Expansion - Recession (model)	0.229	0.131	0.059	0.071	0.132	0.170
<i>Contribution of cyclical mobility shocks</i>						
Occupation loss, δ_z	0.122	0.061	0.020	-0.003	-0.007	-0.029
Employer loss, δ_ϵ	0.012	0.010	0.001	0.000	-0.001	-0.002
Unemp. occ. movers job finding, λ_U^c	-0.007	-0.005	-0.002	-0.001	-0.002	-0.008
Unemp. occ. stayers job finding, λ_U	0.001	0.001	0.001	-0.001	-0.003	-0.006
Emp. occ. movers job finding, λ_E^c	-0.009	-0.007	-0.003	0.011	0.023	0.044
Emp. occ. stayers job finding, λ_E	-0.029	-0.024	-0.013	0.008	0.019	0.034
Total	0.089	0.036	0.003	0.013	0.029	0.033

on Figure 12a, focusing on the left and right tails. The remainder rows decompose this difference at various quantiles by the impact each individual mobility shock has in creating the wedge such that all rows add to the final row's total.

The decomposition shows that cyclical changes in the probability of job loss makes the mobility shocks account for about half of the cyclical change of the left tail. The job finding probabilities among employed workers, in contrast, contribute in a negative way to earnings losses in this tail, and to some extent offsets the small effect of the unemployed job finding probabilities. This occurs as lower job finding probabilities among the employed during recessions lead to lower EE transitions. The smaller contribution of unemployed workers' job finding probabilities can be explained by two factors. First, since recessions are relatively short, the timing of re-employment matters. While the increase in separations into unemployment unambiguously increases downside earnings risk, as most of these separators will have been working prior to the recession, the decline in job finding probabilities during recessions might affect these workers (as well as those already unemployed) by increasing their jobless spell.³⁰ Second, since we are targeting the cyclical change in the average UE transition rate, we are not fully matching the long tail of the unemployment duration as this is not properly captured by the Poisson process assumed here. This implies that the model is attributing some of the effects of long-term unemployment to cyclical changes in $F(\cdot)$ and $\Gamma(\cdot)$.

Table 3 shows that among the unemployment shocks, the obsolescence δ_z shock, which forces employed workers to undergo an occupation switch through unemployment, is by far the most important one. It can explain the entire contribution of the mobility shocks on the left tail of the distribution. In comparison, cyclical changes in job loss due to the destruction of worker-firm match productivity, δ_ϵ , explains a very small portion of the increase in large earnings losses during recessions. Thus, these losses arise mainly from occupations "shutting down" on some workers (particularly those with high earnings), forcing them into unemployment and to switch occupation. The decomposition depicted in Figure 12b shows that upon changing occupations these workers face a worse $F(\cdot)$ from which to draw z , contributing the remainder 50% in explaining the large earnings losses among EUE

³⁰Note also that λ_U, λ_U^c go in opposite directions because λ_U actually increases in recessions, a feature described in Section 4.2

transitions observed in the data during recessions.

At the right tail, Table 3 shows a different picture. In this case the importance of mobility shocks arises from cyclical changes in the job finding probability among employed workers who decided to change occupations. With lower job finding probabilities during recessions, workers have less opportunities to switch employers and occupations at the same time, contributing to the collapse of the right tail of the earnings growth distribution.

6 No Occupational Mobility

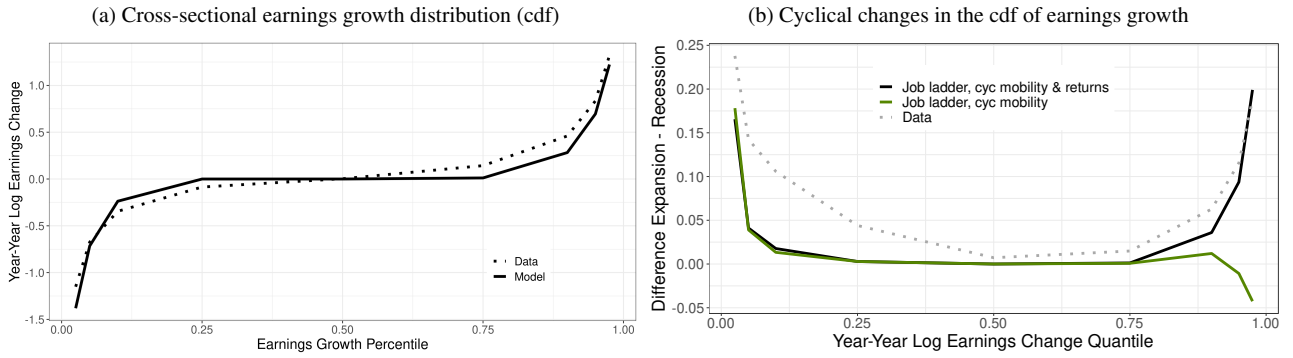
The analysis presented until now shows the importance of occupational mobility in explaining the cyclical changes in the earnings growth distribution. An important question is whether we need a model that encompasses employer and occupational mobility to simultaneously explain the cross-sectional and cyclical properties of earnings growth distribution *and* worker transition rates. Previous work has shown that a suitable extended version of the canonical on-the-job search model is consistent with the main properties of the cross-sectional earnings growth distribution as depicted in Figure 1 and workers transition rates (see Hubmer, 2018 and Karahan et al., 2020). However, this work does not tackle the cyclical properties of earnings changes or worker transition rates. We now show that without occupational mobility, our job ladder model is not able to reproduce well the *cyclical* properties of the earnings growth distribution as depicted in Figure 1d. By also investigating whether this result is driven by cyclical changes in the returns to employer mobility or in the job loss and job finding probabilities, we show that this job ladder model has very different policy implications relative to our baseline model.

To formalize this exercise we shutdown occupational mobility in the baseline model by eliminating any differences in the z -productivities across workers and occupation-wide productivities. In this version workers lose their jobs and become unemployed with probability $\delta(A_I)$. They encounter job offers from unemployment with probability $\lambda_U(A_I)$ and draw a match-specific productivity ϵ from $\Gamma(\epsilon, A)$, which also can change over the cycle. The ϵ -productivity maps one-to-one to the earnings associated with such a job offer. When employed, they draw a new ϵ with probability $\lambda_E(A_I)$. With probability γ the worker decides whether to accept the new job or not; and with probability $1 - \gamma$ he is forced to accept the new job, as long as it offers a payoff above the value of unemployment.

We structurally estimate such a model using the same moments presented in Section 4, except for those pertaining to occupational mobility. In Online Appendix D we show that this version of our model is able to replicate very well the targeted average EE , EU and UE transition probabilities as well as their expansion/recessions ratios. Figure 13a show that the model is able also to replicate the cross-sectional earnings growth distribution (as depicted in Figure 1b in Section 2), capturing well its skewness and leptokurtosis, consistent with the results in Hubmer (2018) and Karahan et al. (2020). Underlying this fit, however, our estimation reveals that although the model is able to capture well the targeted earnings growth distribution conditional on EE employer movers, it fails to capture the targeted earnings growth distributions conditional on EU employer movers and

employer stayers (see Online Appendix D). In particular, among *EUE* employer movers the model generates not only larger earnings losses relative to the data, but it hardly generates any earnings gains. This is a consequence of the shape of the estimated $\Gamma(\cdot)$, which implies employed workers can climb the job ladder and achieve high values of ϵ relatively fast. This feature makes the model consistent with the earnings growth distribution of *EE* movers. However, when these workers fall into unemployment, the estimated $\Gamma(\cdot)$ implies that at re-employment workers are more likely to draw low values of ϵ . Even though these workers might not accept the lowest ϵ draws (particularly during recessions), they will still face a higher probability of becoming re-employed in jobs associated with a low ϵ . The initial large drop in earnings due to job loss coupled with low re-employment earnings then leads to larger earnings losses and smaller earnings gains among *EUE* workers relative to the data.

Figure 13: Earnings growth distribution implied by the model without occupations



Note: The observed annual earnings growth distribution is constructed for the sample period 1990-2013. The left panel shows the cross-sectional earnings growth cumulative distribution by plotting the associated earnings change to each selected percentile of the distribution. The percentile used are: 2.5, 5, 10, 25, 50, 75, 90, 95, 97.5. We construct the simulated distribution in the same way as done in the data. The right panel shows the difference between the expansion and recessions earnings growth distributions as documented in Section 2. In this case we use the same percentiles as in the cross-sectional distribution.

Figure 13b presents the implications of this model for the targeted cyclical change of the earnings growth distribution. The black curve shows the cyclical changes in the earnings growth distribution when allowing $\Gamma(\cdot)$ to vary between expansions and recessions as estimated above. The green curve shows the implications when $\Gamma(\cdot)$ is kept constant at its expansion level; i.e. no cyclical returns to mobility. The first key result of this exercise is that allowing the returns to mobility to change over the cycle is crucial to obtain procyclical skewness in the earnings growth distribution. Otherwise, this model exhibits countercyclical variance, which is counterfactual. The role of a cyclical $\Gamma(\cdot)$ is important as it resolves two opposing effects. On the one hand, longer unemployment spells in recessions imply annual earnings fall. As workers become re-employed their earnings then increase from a lower base and generate larger earnings growth. On the other hand, more opportunities to make *EE* transitions imply larger earnings growth in expansions. With a constant $\Gamma(\cdot)$ the effect of workers moving out of unemployment dominates and the model predicts both larger gains and losses in recessions. Allowing $\Gamma(\cdot)$ to change over the cycle implies the second effect dominates.

The second key result is that even allowing $\Gamma(\cdot)$ to change over the cycle the model remains far

from the data. Figure 13b shows that it consistently under-predicts the difference in the earnings losses/gains between expansions and recessions, except at the extremes. The lack of fit on this dimension arises as the model is unable to fully resolve a second related trade-off: reproducing the observed worker transitions flows and the earnings growth distribution over the cycle. The main tension arises from the relationship between unemployment and earnings changes. By reproducing the observed UE transition probability in expansions and recessions, the model already implies a degree of cyclical change in the earnings losses through unemployment. However, the very large gap between the simulated and the data distributions below the median, implies that these recessionary earnings losses are too similar to the ones in expansions. In order to generate larger earnings losses during recessions the model either has to have (i) a steeper ϵ ladder, so that when workers fall they do so from higher up in the ladder. However, a steeper ϵ ladder comes into tension with the observed earnings changes associated with EE transitions, which the model is able to reproduce quite well. Alternatively the model has to have (ii) counterfactually long unemployment durations, which comes into tension with the cyclicity of UE transitions that it also reproduces quite well. As the analysis of Sections 4 and 5 show, the aforementioned trade-off can be fully resolved when we account for occupational mobility in addition to employer mobility. This is because occupational mobility creates an additional source of downside earnings risk which drastically increases in recessions.

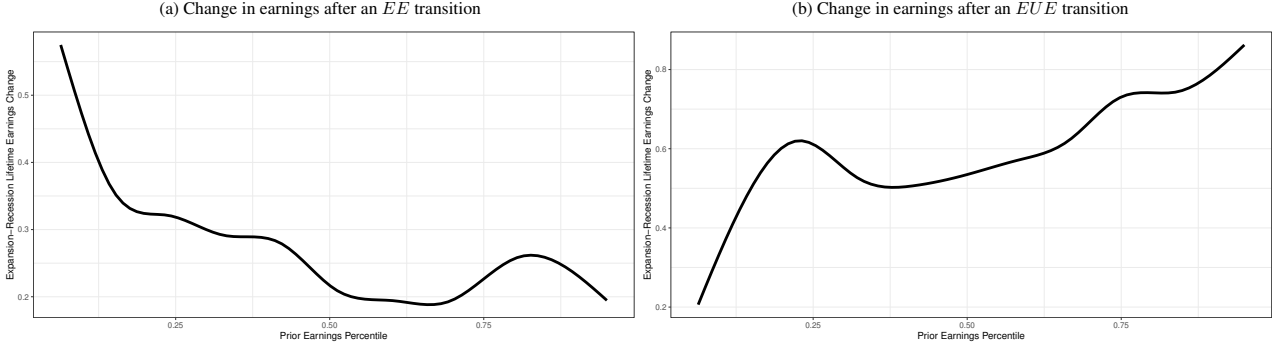
The third key result is that from the viewpoint of the no occupation model the larger earnings losses it generates in recessions arise purely from cyclical changes in the mobility shocks. A similar conclusion can be drawn when considering the (moderately) larger earnings gains it generates in expansions. Figure 13b shows that the cyclical difference in the earnings growth distribution remains nearly unaffected below the 75th percentile when not allowing the returns to mobility to change over the cycle. This implication is in contrast with the occupation mobility model where cyclical changes in the returns to mobility explain a large part of the cyclical pattern of the earnings growth distribution. Each model therefore suggest emphasising two very different policy tools for reducing the larger earnings losses during recessions. The no occupation model would favour policies that aim at getting unemployed workers back to work quickly through, for example, reductions in the generosity or length of unemployment insurance payments, like the ones implemented in the UK during the Covid-19 pandemic and in Germany through the Hartz reforms. The occupation mobility model instead would favour policies that aim at improving the job quality individuals can expect at re-employment, through, for example, re-training schemes.

7 Sullyng Effects of Recessions

Our model implies that cyclical changes in the returns to mobility and the mobility shocks affect workers differently. For example, by affecting their ability to climb the job ladder during expansions and recessions, workers might suffer differently from the sullyng effects of recessions. The nature of the job ladder in our model implies that those workers with lower earnings are the more likely to make an EE transition (with or without an occupational move). During recessions this probability

decreases (as λ_E and λ_E^c fall), not allowing them to climb the job ladder as fast as in expansions. In addition, worse returns to mobility also imply that if they manage to change jobs they will expect to draw lower ϵ and z productivities, leading to lower earnings gains relative to expansions.

Figure 14: Changes in lifetime earnings after an EE and EUE transition

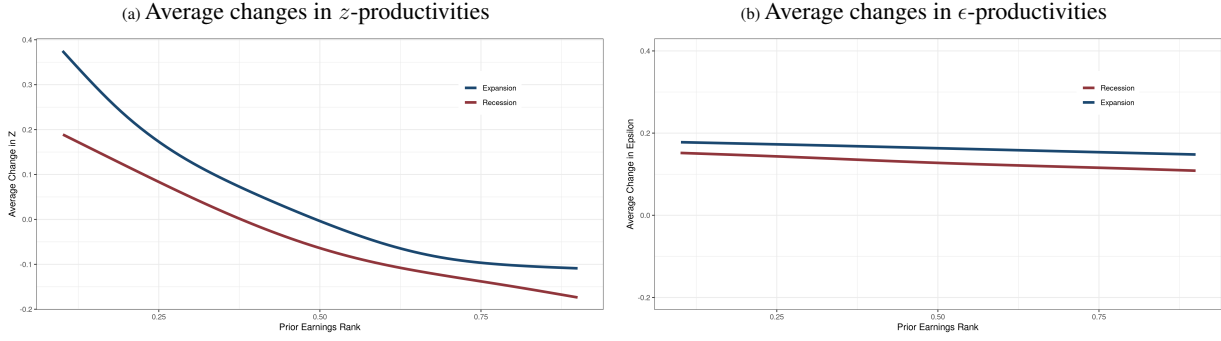


Note: The left panel depicts the expansions and recessions difference of workers' average lifetime earnings prior to an EE transition and their average lifetime earnings after this transition conditional on a worker's earnings immediately prior the transition. A positive value implies that during expansions workers typically have larger increases in average lifetime earnings after an EE transition. The right panel present the same exercise but using EUE transitions instead. A positive value implies that during recessions workers typically have larger losses in average earnings after an EUE transition.

Our model further implies that this sully effect persists over time. To show this implication we compare a worker's average lifetime earnings prior an EE transition to his average lifetime earnings after the EE transition. We do this separately for expansions and recessions periods and subtract the two values. Figure 14a shows the result of this exercise conditional on the worker's rank in the earnings distribution immediately prior the EE transition. The fact that the difference between expansion and recessions is positive across prior earnings percentiles shows that workers who made an EE transition increased their lifetime earnings more in expansions than in recessions. The fact that this difference decreases with earnings percentiles implies that low-paid workers suffer disproportionately more in terms of lifetime earnings from the sully effects of recessions than do high-paid workers.

Figure 15 shows that this occurs mainly through the occupational dimension of the job ladder. Figure 15a presents the average change in z -productivities due to an EE transition conditioning on a worker's rank in the earnings distribution immediately prior the EE transition. Figure 15b shows instead the average change in ϵ -productivities. As low-paid (typically low z) workers are more willing to switch occupations than higher paid (typically higher z) workers (see Figure 8), they are also more sensitive to cyclical changes in the returns to mobility and the mobility shocks, particularly changes in $F(\cdot)$ and λ_E^c . This implies that during recessions low-paid workers who make an EE transition will face an average increase in their z -productivity that is about 15 to 10 percentage points lower than in expansions, while workers with earnings at the top of the distribution face a reduction of about 5 percentage points. This strong differential effect is not present when considering ϵ -productivities. Figure 15b shows that although during a recession all workers face a lower average changes in their ϵ -productivities after an EE transition (due to a worsening in $\Gamma(\cdot)$ and λ_E), this decrease is small in comparison and there is no meaningful differential effect across the percentiles of the earnings distribution.

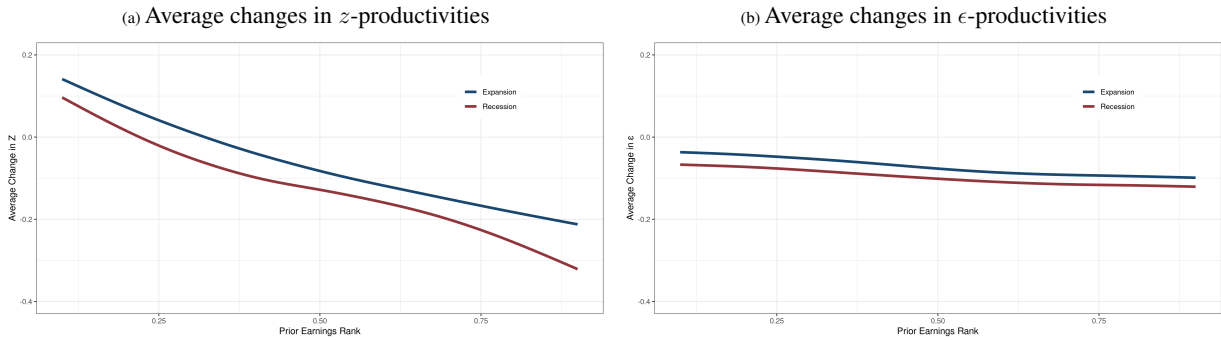
Figure 15: Climbing the job ladder



Note: The left panel depicts the average change in z -productivities worker obtain after an EE transition conditional on a worker's earnings immediately prior the transition in recessions and expansions periods. The right panel instead depicts the average change in e -productivities worker obtain after an EE transition conditional on a worker's earnings immediately prior the transition in recessions and expansions periods.

Cyclical changes in the returns to mobility and the mobility shocks also affect workers' re-employment outcomes after falling from the job ladder. Figure 14b shows that now it is high-paid workers who suffer disproportionately more in recessions from an EUE transition than low-paid workers. In this case we plot the difference between expansion and recessions average lifetime earnings losses after changing employers through an unemployment spell. The fact that the difference is positive show that lifetime earnings losses are larger in recessions relative to expansions in line with the results of Davis and Von Wachter (2011) and Huckfeldt (2021). The fact that this difference increases with workers' prior earnings percentiles indicates that losses are even stronger during recessions among the high-paid workers.

Figure 16: Falling from the job ladder



Note: The left panel depicts the average change in z -productivities worker obtain after an EUE transition conditional on a worker's earnings immediately prior the transition in recessions and expansions periods. The right panel depicts the average change in e -productivities worker obtain after an EUE transition conditional on a worker's earnings immediately prior the transition in recessions and expansions periods.

Figure 16 shows these effects once again occur along the occupational dimension of the job ladder. This figure shows the average change in z and e productivities after and EUE transition by a worker's rank in the earnings distribution. The disproportionately worse outcomes of high-paid workers come as a result of forced occupational mobility due to a higher value of δ_z and a worsening of $F(\cdot)$ during recessions. The same argument as above applies: as high-paid (typically high z) workers are less likely to change occupations voluntarily, they are more sensitive to a higher prevalence of the ob-

solescence shock and a higher likelihood of drawing a lower value of z during recessions. Figure 16b shows that this differential effect is not present along the employer dimension, where the worsening in ϵ productivities is much smaller and very similar across the percentiles of the earnings distribution.

8 Conclusion

In this paper we have studied the structural causes of cyclical earnings risk through the lens of a job ladder model with endogenous occupation and employer mobility. A series of important research has highlighted that higher-order moments of the earnings growth distribution are the most cyclical, that is, the tails of the distribution move the most over the cycle. Thus, crucial to understanding a recession is to understand why largest earnings losses become even worse and top earnings gains become smaller. We document the importance of career changes, seen through occupational mobility, in explaining the tails of the earnings growth distribution. Workers who switch occupations and employers simultaneously have considerably more disparate outcomes in terms of earnings growth, both to the upside and downside and whether through unemployment or directly through an EE transition.

Recessions bring with them a reduction in the opportunity to find a new employer or to switch careers while simultaneously changing the returns to such transitions. Particularly salient, they bring a larger chance of job loss and a larger chance of job loss that displaces one from his career. These transitions also have different outcomes, for example, making the cost of job loss different in expansions or recessions. Our estimated model shows that changes in the returns to occupation mobility are the main source cyclical earnings risk. Beyond this force, job loss that also comes with career displacement plays an important role in explaining the increase in the larger earnings losses during recessions. During expansion, the increased chances of improving the occupational dimension of a job through direct job-to-job transitions contributes to increasing the largest earnings gains.

Our framework has focused on ex-post worker heterogeneity mostly through differences in idiosyncratic occupation and firm match-specific productivities. Recent work has emphasised instead the role of ex-ante heterogeneity in the probabilities of job loss and job finding among workers. Occupational components, whose importance is the message of this paper, can meaningfully interact with these ex-ante differences. For instance, the interaction can elucidate for which types of workers the cyclical collapse of the career and employer ladder are especially relevant and likewise can clarify the role of career paths to type-level differences. These extensions could be particularly useful to understand the great amount of heterogeneity in workers' experiences of the business cycle, seen both as cyclical changes in worker transition rates and earnings growth.

Online Appendix

<https://github.com/CTVproject/CTVW/raw/main/CVWOnlineAppendix.pdf>

References

- [1] Altonji, J. and R. Shakotko. 1987. “Do wages rise with seniority? *Review of Economic Studies*, 54(3): 278-297.
- [2] Arellano, M, R. Blundell and S. Bonhomme. 2017. “Earnings and Consumption Dynamics: A Non-linear Panel Data Framework”, *Econometrica*, 85(3): 693-734.
- [3] Bagger, J. and R. Lentz. 2019. “An Empirical Model of Wage Dispersion with Sorting”, *Review of Economic Studies*, 86(1): 153-190.
- [4] Barlevy, G. 2002. “The Sullyng Effect of Recessions,” *Review of Economic Studies*, 69(1): 65-96.
- [5] Braxton, J.C. and B. Taska. 2022. “Technological Change and the Consequences of Job Loss”, Mimeo, Department of Economics, University of Michigan.
- [6] Burdett, K. 1978. “A Theory of Employee Search and Quits”, *American Economic Review*, 68(1): 212-220.
- [7] Busch, C. 2020. “Occupational Switching, Tasks, and Wage Dynamics”. Mimeo, Department of Economics, Ludwig Maximilian University (LMU) Munich.
- [8] Busch, C., D. Domeij, F. Guvenen, and R. Madera. 2021. “Skewed Idiosyncratic Income Risk over the Business Cycle: Sources and Insurance”, *American Economic Journal: Macroeconomics*, forthcoming.
- [9] Card, D., J. Heining, J. and P. Kline. 2013. “Workplace Heterogeneity and the Rise of West German Wage Inequality.” *The Quarterly Journal of Economics*, 128(3): 967-1015.
- [10] Carrillo-Tudela, C. and L. Visschers. 2021. “Unemployment and Endogenous Reallocations Over the Business Cycle”. IZA Working Papers No. 7124.
- [11] Davis, S. J. and Von Wachter, T. 2011. “Recessions and the Costs of Job Loss”, *Brookings Papers on Economic Activity*, 43: 1-72.
- [12] De Nardi, M., G. Fella, M. G. Knoeff, G. Paz-Pardo, and R. Van Ooijen. 2021. “Family and Government Insurance: Wage, Earnings, and Income Risks in the Netherlands and the US,” *Journal of Public Economics*, 193, 104-327.
- [13] Fallick, B. C. 1993. “The Industrial Mobility of Displaced Workers”. *Journal of Labor Economics*, 11(2): 302-323.
- [14] Fujita, S. and G. Moscarini. 2017. “Recall and Unemployment”. *American Economic Review*, 102(7): 3875-3916.

- [15] Gregory, V., G. Menzio and D. Wiczer. 2021. “The Alpha, Beta and Gamma of the Labor Market”. Mimeo, Department of Economics, New York University.
- [16] Groes, F., P. Kircher, and I. Manovskii. 2015. “The U-Shapes of Occupational Mobility,” *Review of Economic Studies*, 82(2): 659-692.
- [17] Guvenen, F., F. Karahan, S. Ozkan, and J. Song. 2021. “What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Risk?,” *Econometrica*, 89(5): 2303-2339.
- [18] Guvenen, F., S. Ozkan, and J. Song. 2014. “The Nature of Countercyclical Income Risk,” *Journal of Political Economy*, 122(3): 621-660.
- [19] Halvorsen, E., H. A. Holter, S. Ozkan, and K. Storesletten. 2020. “Dissecting Idiosyncratic Income Risk,” Available at SSRN: <https://ssrn.com/abstract=3669146>.
- [20] Harmenberg, K. 2021. “The Labor-Market Origins of Cyclical Skewness”. Mimeo, Copenhagen Business School, Denmark.
- [21] Huckfeldt, C. 2022. “Understanding the Scarring Effects of Recessions”. *American Economic Review*, forthcoming.
- [22] Hubmer, J. 2018. “The Job Ladder and Its Implications for Earnings Risk,” *Review of Economic Dynamics*, 29, 172-194.
- [23] Jolivet, G., Postel-Vinay, F., Robin, J-M., 2006. “The Empirical Content of the Job Search Model: Labor Mobility and Wage Distributions in Europe and the US”. *European Economic Review*, 50 (4): 877–907
- [24] Kambourov, G. and I. Manovskii. 2009a. “Occupational Mobility and Wage Inequality”. *Review of Economic Studies*, 76(2): 731-759.
- [25] Kambourov, G. and I. Manovskii. 2009b. “Occupational Specificity of Human Capital”. *International Economic Review*, 50(1): 63-115.
- [26] Kambourov, G. and I. Manovskii. 2008. “Rising Occupational and Industry Mobility in the United States: 1968-97”. *International Economic Review*, 49(1): 41-79.
- [27] Karahan, F., S. Ozkan, and J. Song. 2020. “Anatomy of Lifetime Earnings Inequality: Heterogeneity in Job Ladder Risk vs. Human Capital,” FRB of New York Staff Report (908).
- [28] Kramer, J. 2022. “The Cyclicity of Earnings Growth along the Distribution - Causes and Consequences”. Mimeo, Institute for International Economic Studies, Stockholm, Sweden.
- [29] Kudlyak, M. And R. Hall. 2019. “Job-Finding and Job-Losing: A Comprehensive Model of Heterogeneous Individual Labor-Market Dynamics”, FRB San Francisco WP No. 19-05.

- [30] Lise, J. 2012. "On-the-Job Search and Precautionary Savings," *Review of Economic Studies*, 80(3): 1086-1113.
- [31] Lucas, R. and E. Prescott. 1974. "Equilibrium Search and Unemployment". *Journal of Economic Theory*, 7: 188-209.
- [32] Mortensen, D. and C. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment". *Review of Economic Studies*, 61(3): 397-415.
- [33] Moscarini, G. 2001. "Excess Worker Reallocation". *Review of Economic Studies*, 68(3): 593-612.
- [34] Neal, D. 1999. "The Complexity of Job Mobility Among Young Men". *Journal of Labor Economics*, 17(2): 237-524.
- [35] Pilosoph, L. 2021. "A Multisector Equilibrium Search Model of Labor Reallocation". Mimeo.
- [36] Postel-Vinay, F. And H. Turon. 2010. "On-the-job Search, Productivity Shocks, and the Individual Earnings Process". *International Economic Review*, 51(3): 599-629.
- [37] Shimer, Robert. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies". *American Economic Review*, 95(1): 25-49.
- [38] Storesletten, K., C. I. Telmer, and A. Yaron. 2004. "Consumption and Risk Sharing Over the Life Cycle," *Journal of Monetary Economics*, 51(3): 609-633.
- [39] Sullivan, P. 2010. "Empirical evidence on occupation and industry specific human capital". *Labour Economics*, 17: 567-580.
- [40] Wiczer, D. 2015. "Long-term Unemployment: Attached and Mismatched?" St. Louis FED, WP 2015-042A.