STAFF REPORTS

NO. 1206
SEPTEMBER 2026

Cyclical Earnings, Career and Employment Transitions

Carlos Carrillo-Tudela | Ludo Visschers | David Wiczer

FEDERAL RESERVE BANK *of* NEW YORK

Cyclical Earnings, Career and Employment Transitions
Carlos Carrillo-Tudela, Ludo Visschers, and David Wiczer
Federal Reserve Bank of New York Staff Reports, no. 1206
September 2026
https://doi.org/10.59576/sr.1206

Abstract

Career changes across occupations are central to the cyclical behavior of workers' earnings growth. Using SIPP data, we show that workers who switch both employers and occupations account for 58 percent of the increase in the left-skewness of earnings growth during recessions. This is driven mainly by cyclical variation in the earnings changes upon a transition, rather than in transition flows. In a business cycle model with on-the-job search over employer and occupation ladders, cyclical shifts in the quality of worker-occupation matches are the main source of cyclical earnings risk, but they generate procyclical skewness only through workers' optimal mobility decisions: without these, procyclical skewness becomes countercyclical variance. Our results give a structural foundation to reduced-form income processes with higher-order risk, and point to policies that target occupational reallocation directly.

JEL classification: E24, E30, J62, J63, J64
Key words: earnings, unemployment, business cycle, search, occupational mobility

---

Wiczer: Federal Reserve Bank of New York (email: david.wiczer@ny.frb.org). Carrillo-Tudela: University of Essex (email: cocarr@essex.ac.uk). Visschers: U. Carlos III Madrid (email: lvissche@eco.uc3m.es). The authors thank the editor and referees for their constructive comments, which have helped to improve the paper greatly. They also thank Joe Altonji, Serdar Ozkan, Serdar Birinci, Richard Rogerson, Matthew Shapiro, Ija Trapeznikova, Gianluca Violante, Haomin Wang, Christian Hellwig, and Moritz Kuhn, as well as seminar participants at the University of Bristol, Copenhagen, Florida, Edinburgh, Essex, Kansas State, Konstanz, Carlos III, the St. Louis Fed, Kansas City Fed, Bank of Portugal, the Chilean Central Bank, Southampton, Tilburg, Alicante, Mannheim, Copenhagen Business School, Umea, and Berlin, 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, the Vienna Macro workshop 2022, SAET Paris 2023, Human Capital and Labor Market Dynamics (Le Mans) 2023, Spanish Macro Network 2024, Simposio AEe 2024, Lisbon Macro 2025, Winter Workshop on Search and Matching (Aussois) 2026, and Berlin Macro 2026. Carrillo-Tudela acknowledges financial support from the UK Economic and Social Research Council (ESRC), award reference ES/V016970/1. Visschers acknowledges financial support from the UK Economic and Social Research Council (ESRC), award references ES/V016970/1, and ES/L009633/1, Fundación BBVA (Ayudas a Equipos de Investigación Científica Sars-Cov-2 y Covid-19), Maria de Maeztu grant (award reference CEX2021-001181-M), Programa ATRAE of the Spanish Ministry of Science, Innovation and Universities, (award reference ATR2023-145734), and Project Grant PID2024-160896OB-I00, the latter three funded by MICIU/AEI /10.13039/501100011033.

1 Introduction

Understanding the nature of earnings risk is important because it impacts individuals' career decisions, their consumption, savings and investment decisions, and ultimately inequality. In this paper, we study the joint behavior 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 evidence showing that the distribution of earnings growth exhibits procyclical skewness (see Guvenen, et al., 2014 and Busch, et al., 2022).[1] Less is known, however, about the type of cyclical earnings risk, whether this risk is mainly caused by cyclical changes in the frequency of job loss and job finding or cyclical changes in the "returns" to mobility (i.e. earnings changes conditional on a transition), and how workers' mobility decisions interact with them both. Studying these features is important as they help determine the sources of idiosyncratic earnings risk over the cycle. Without this understanding, it remains difficult to evaluate, for example, whether governments should emphasize labor 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 demonstrates that career changes, which we observe as occupational mobility, are the main driver behind the cyclical patterns of the annual earnings growth distribution, particularly through the impact on its tails. Employer mobility on its own does not contribute as much as occupational mobility in shaping cyclical earnings growth. We show that the worsening in the quality of occupation-worker matches during recessions is the main risk that underlies the observed procyclical skewness of the earnings growth distribution. However, workers' occupation mobility decisions are crucial for this result. When workers do not react to how earnings risk changes over the cycle, our analysis predicts that recessions are instead characterized by larger earnings losses and larger earnings gains, eliminating the procyclical skewness observed in the data and instead generating an earnings growth distribution that is characterized by countercyclical variance. Our analysis also demonstrates that targeted policies that ease workers' access to certain occupations, e.g., through re-training schemes, appear more effective than untargeted ones in improving workers' lifetime expected utility, but only when these ease access to the most desirable occupations.

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. We present novel evidence showing the importance of occupational mobility driving the cyclicality of the earnings growth distribution and its procyclical skewness. Our empirical analysis establishes three novel facts. First, the procyclical skewness of the earnings growth distribution is overwhelmingly ac-

[1] Procyclical skewness implies that the right tail of the distribution compresses and the left tail expands in recessions and vice versa in expansions (Busch et al., 2022).

[2] For example, during the recent Covid-19 pandemic the UK Government implemented unemployment benefit cuts for 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, which imposed severe penalties on the level of unemployment benefits individuals could claim if they rejected a suitable job offer irrespective of the industry/occupation.

counted for by workers who simultaneously change employers and occupations. We find that they can explain close to 60% of the cyclical change in skewness. The importance of occupation mobility does not appear driven by occupation-wide variation in pay or composition effects due to worker characteristics. Second, the cyclical change in skewness reflects cyclical movements in earnings changes conditional on a transition, not only in the flows of workers across occupations. Third, the large recessionary earnings losses driving procyclical skewness partially revert at the three-year horizon, more so for occupation movers; while the large cyclical earnings gains persist for both occupation movers and stayers.

To interpret these data patterns, we propose a model of on-the-job search in which a job consists of two dimensions: the employer and the occupation dimension. These dimensions create two job ladders that workers climb and can fall from. We use this model to investigate (i) whether cyclical changes in earnings risk arise from the occupation or from the employer dimension, (ii) the importance of workers' occupation mobility decisions in shaping observed earnings growth cyclically in lieu of changing earnings risk, and (iii) the implications for policies that target access to specific occupations in order to help workers obtain better employment.

We untangle these forces by extending 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 match quality (offer) distributions. In our model, workers' 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 worker-specific match quality and specific human capital as well as occupation-wide productivity differences. Once the worker has decided to reallocate, he decides how to direct his search effort across occupations. The latter is captured through an imperfect directed search function that allows us to estimate the degree of search directness across occupations for employed and unemployed workers as well as the ease of access of each occupation.

We use simulated method of moments to structurally estimate our model and recover the employer and occupation ladder from earnings growth data. In particular, we show that the employer ladder dimension of the model does not provide sufficient variation to capture large cyclical shifts in the earnings growth distribution tails. This is because the one-job ladder model creates a close link between the size of the earnings losses and gains it can generate, which limits the asymmetric cyclical changes of the tails of the implied earnings growth distribution. The moments that recover the employer ladder thus leave it with too little cyclically-adjustable risk.

Our empirical findings suggest that the occupation dimension of the ladder is crucial to replicate the data as it creates a separate channel through which downside and upside earnings risk can vary over the cycle at a different rate from the employer ladder. Two features of the estimated occupation ladder drive this. First, it is far more dispersed than the employer ladder, so occupation moves span a much larger range of gains and losses. Second, in recessions both distributions shift down, but the chances for a large upside occupational match erode. For those who do move, the best draws become scarce. A worker stuck in a poor occupation match therefore faces a double blow in a downturn with fewer opportunities to draw a new match and a distribution from which the most valuable draws have largely disappeared.

Quantitatively, this model is quite successful at capturing the cyclical behavior of the earnings growth distribution, the cyclical behavior of worker mobility, and a range of cross-sectional patterns that characterize earnings growth dynamics and worker mobility. Specifically, our model reproduces (i) the job loss, re-employment and direct employer transition probabilities cross-sectionally and over the business cycle; (ii) gross and net occupational mobility patterns among employer stayers and movers, including the increasing relationship between the size of the earnings change and the probability of an occupation switch; (iii) the expansion and recession earnings growth distributions of workers who changed employer with and without occupation switch; and (iv) the cross-sectional and the cyclical change in the overall earnings growth distribution, characterized by procyclical skewness. Beyond these targeted moments, the model also delivers the same degree of dispersion in earning levels as in the data, and generates an earnings process inline with the data, where earnings gains are largely preserved at the three-year horizon while large losses recover only partially, and the implied stochastic income process is similar to estimates obtained on the PSID (Guvenen, 2009; Krueger et al., 2016). The model generates these features through the state-dependent persistence highlighted by Arellano et al. (2017), providing a structural interpretation of their econometric findings.

The model disentangles the effects of firm-worker and occupation-worker shocks in explaining the cyclicality of earnings growth. We find that cyclical changes in worker-occupation match quality can explain the vast majority of the difference between the expansion and recession earnings growth distributions, particularly between the 10^{th} and 90^{th} percentiles. Outside this range, two major conclusions emerge. At the bottom tail, a significant proportion of the largest earnings losses in recessions arise from those individuals who lost their jobs *and* had to change occupations. These types of events are similar in spirit to the "disaster" shocks in Guvenen et al. (2021) or "obsolescence" shocks in Huckfeldt (2022). At the top tail, the frequency of opportunities to improve the occupational dimension of a job becomes more important to explain the largest earnings gains observed in expansions. Underlying all these channels, however, are workers' occupation mobility decisions. We show that by not allowing workers to respond to these sources of cyclical risk, the model fails to generate procyclical skewness altogether, even when all productivity distributions and transition probabilities vary as estimated.

With this framework, we investigate the cost of business cycles and its policy implications. We find that the sullying effects of recessions persist over time: workers who change employers increase their lifetime earnings by less in recessions than in expansions, and this gap is largest for low-paid workers and operates mainly through the occupation dimension of the job ladder. Motivated by this result, we use the model to compare labor market policies that target specific occupations against untargeted ones. Improving workers' access to the most productive occupations raises match quality and lifetime expected utility, particularly for unemployed workers and particularly in recessions, whereas increasing the availability of jobs in the less productive occupations reduces welfare by drawing search away from better careers. Such targeted policies also compare favorably with untargeted increases in the generosity of unemployment benefits, which improve match quality only through workers' selectivity, at the cost of considerably longer unemployment spells.

Related Literature This paper contributes to the literature investigating the properties of idiosyncratic earnings risk. While measures of earnings risk have typically been obtained using log-normal earnings processes, recent papers show that idiosyncratic earnings risk presents non-Gaussian features: using administrative data for the US, Guvenen et al. (2021) find that the earnings growth distribution is left-skewed and has excess kurtosis (see also Arellano et al., 2017), and De Nardi et al. (2021), Halvorsen, et al. (2024) and Busch (2020) find similar features in data from the Netherlands, Norway and Germany. Over the business cycle, Guvenen et al. (2014), Busch et al. (2022), Harmenberg (2021) and Kramer (2022), among others, show that earnings risk is mainly characterized by its procyclical skewness.[3] We show that procyclical skewness arises primarily from those who simultaneously changed employers and occupations, and mainly occurs through variation in the distribution of earnings growth obtained upon a transition, rather than variation in the transition flows.[4]

We also contribute to the literature that uses job search models to explain earnings growth dynamics. Quantitative versions of the canonical job ladder model have been shown to replicate salient properties of cross-sectional earnings dynamics, highlighting the role of workers' job acceptance decisions (see Low et al., 2010; Postel-Vinay and Turon, 2010; Altonji et al., 2013; Hubmer, 2018; Ozkan et al., 2023). We build on Low et al. (2020) by introducing an occupation ladder alongside an employer ladder and study their cyclical behavior. We show that cyclical changes in workers' occupation mobility decisions *and* the earnings risk created by cyclical shifts in the quality of worker-occupation matches is what delivers the procyclical skewness of the earnings growth distribution. A version of our model without occupational mobility, akin to the canonical job ladder model, is not able to replicate this feature of the data.[5]

Kambourov and Manovskii (2009a,b) and Busch (2020) have highlighted the role of occupational mobility in explaining earnings dynamics through the returns to occupational tenure and long-run changes in earnings inequality. Carrillo-Tudela and Visschers (2023) show its importance in determining the joint behavior of cyclical unemployment and its duration distribution. Huckfeldt (2022) shows its role in shaping the long-term costs of job loss during recessions. The novelty here is to show that occupational mobility is the main driver behind cyclical earnings growth.

The rest of this paper is organized as follows. In Section 2, we present empirical work highlighting the role of occupational mobility in the procyclical skewness of the earnings growth distribution. Sections 3 and 4 present our two job ladder model, its estimation, identification and fit to the data.

[3]This is in contrast to Storesletten, et al. (2004) who find that earnings risk is characterized by countercyclical variance.

[4]As in De Nardi et al. (2021) and Halvorsen et al. (2024) we find that changes in hours are important for cyclical earnings risk, but the importance of simultaneous occupation and employer mobility remains strongly visible when considering hourly wages.

[5]Harmenberg (2021) investigates 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.

Section 5 evaluates the roles of the occupation ladder, workers' mobility decisions and each source of employment risk. Section 6 considers the implications for the sullying effects of recessions and for occupation-targeted labor market policies. Section 7 concludes. The Online Appendix contains details of our data construction, estimation procedure, identification and robustness exercises.

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 is interviewed in waves of four months. Individuals are asked to report information on their current and previous employment status, occupations, industries, and earnings (hourly wages and hours worked). Using this information, we now describe how we construct year-to-year earnings growth and link these to labor market transitions across employment status, employers and occupations, where we define employer, occupation and earnings based on a worker's main job for each month. In Online Appendix B, we provide a more detailed discussion of our sample construction, and our approach to deal with measurement error in occupation and earnings information.

Labor market transitions Direct employer changes are labeled EE transitions, while employer changes that occurred with an intervening spell of non-employment that lasted for at least a month are labeled EUE transitions. The latter begin with an employer to non-employment EU transition, and end with a non-employment to employment UE transition. The jobless spell in between these transitions takes into account both unemployment and non-participation periods.[6] Thus, in a given month, a labor market transition will be categorized as either an EE, EU, or UE transition.

To measure occupation changes, we harmonize the occupation classification across SIPP panels using the crosswalk translation scheme created by IPUMS based on the 1990 Standard Occupational Classification (1990 SOC). For our benchmark analysis, we aggregate using four task-based occupation 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 considerable change in the main tasks performed. We then compare the task-based occupation for a given individual across employers or within the same employer, depending on whether an employer transition occurred. We also verify that the results we document below hold when using less aggregated occupation codes.

A well-known issue in the SIPP is that labor market transitions are over-proportionally reported between the last month of a wave and the first month of the next wave. To minimize the bias introduced by these "seam effects", we aggregate transitions at a wave level. We consider a wave to contain an employer change when we observe an EE, EU or a UE transition in any month during

[6] We label as "temporary recalls" those workers whose EUE transition ends with a return to the employer they left immediately before the non-employment spell. We categorize these cases as continuing employment spells without an employer change, on the grounds that these workers very likely remained attached to their previous employers and did not face the same reallocation and search frictions as those who did not expect to return to their previous employers.

that wave. If any of these transitions occurred across waves, we count them as belonging to the wave in which the transition originated. If we observe multiple transitions of the same individual within the same wave, we count them separately.[7]

A worker is labeled as an "occupation and employer mover" when he/she has an employer and occupation change and as an "occupation stayer and employer mover" when only an employer change is observed. An individual is labeled as an "occupation mover and employer stayer" when we only observe a change in the occupation dimension of a job within the wave. For an individual to be labeled an "employer and occupation stayer" he/she should not have changed either of these dimensions during the wave. Under this categorization, we find that 86.6% of all wave-level observations are made up of employer stayers, 4.3% involved an EE transition, and 9.0% involved an EUE transition. Within the four task-based categories, 1.3% of all wave-level observations involved a simultaneous EE and occupation transition (34.1% of all EE transitions), 3.1% involved an EUE transition with an occupation change (37.8% of all EUE transitions), and 1.2% involved an employer stay with an occupation transition (1.4% of all employer stays).

Earnings To construct our earnings measure, we first deflate nominal monthly earnings by the Personal Consumption Expenditure price index. We then residualize log earnings on worker characteristics, so that the earnings growth distribution minimizes variation driven by composition — consistent with our subsequent modeling assumptions. Specifically, letting $w_{i,t}$ denote the observed real earnings of an employed individual i in month t, we estimate $\log w_{i,t} = x'_{i,t}\beta + \mu_t + \epsilon_{i,t}$, where $x_{i,t}$ is a vector consisting of a quadratic in potential experience, education, gender, and race, and μ_t denotes calendar month dummies. Our earnings measure for individual i is then $\exp(\tilde{w}_t + \hat{\epsilon}_{i,t})$, where $\tilde{w}_t$ is the cross-sectional mean of $\log w_{i,t}$ in month t and $\hat{\epsilon}_{i,t}$ is the estimated residual. This measure combines a component common to all workers in a given month $(\tilde{w}_t)$ with an individual-specific component $(\hat{\epsilon}_{i,t})$. Exponentiating ensures that earnings are strictly positive, so that the distinction between employment and non-employment is preserved.

To link annual earnings changes to labor market transitions, we aggregate monthly earnings to the wave level. For a given *reference wave*, we define the annual earnings change as the difference between the sum of monthly earnings in the reference wave together with the two following waves (twelve months in total) and the sum of monthly earnings in the three preceding waves (also twelve months). Starting from the fourth wave of each panel and ending two waves before its conclusion, we advance the reference wave one wave at a time and attribute the resulting earnings change to the EE, EU, or UE transition occurring in the reference wave, distinguishing transitions with and without an occupation change. Reference waves containing no employer transition are attributed to employer stayers. The procedure amounts to a rolling twelve-month-versus-twelve-month comparison, with each window labeled by the transition (or non-transition) at its midpoint.

Because some EUE transitions involve re-employment after long non-employment spells, the associated annual earnings can be very low. We therefore compute earnings changes using inverse

[7] Our approach to treat multiple transitions of the same individual within the same wave is one of many we could have followed. However, since multiple transitions only represent about 4% of all transitions or 0.4% of all observations, we do not expect our treatment choice to have meaningful effects on our results.

Figure 1: Earnings Growth Distributions of Employer Changers by Occupational Mobility

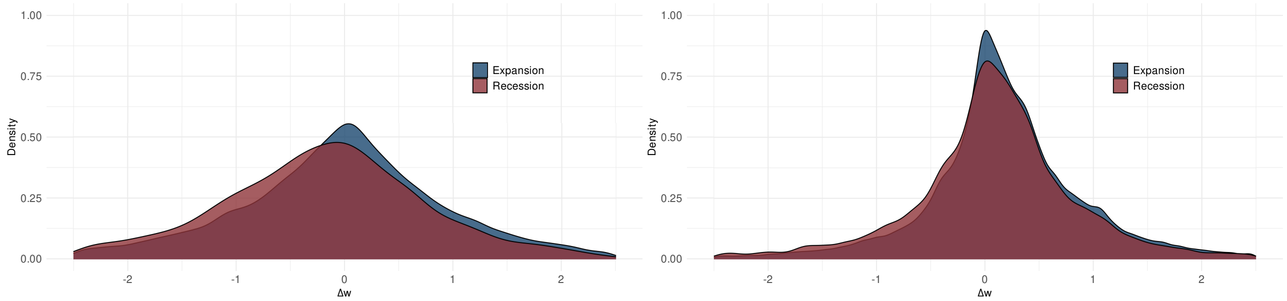

(a) Occupational Switchers (b) Non-Switchers

Note: Kernel density estimates of the distribution of annual earnings growth for workers changing employers with or without an occupation switch in expansion and recession periods. Annual earnings growth is based on residualized earnings, and recessions are defined as periods in which the HP-filtered unemployment rate lies in the top 20% of realizations as described in Section 2.1.

hyperbolic sine (IHS) differences rather than log differences.[8] Let $\Delta_{i,t}$ denote the year-to-year IHS difference in earnings for individual i measured at reference wave t, and let Δ denote the pooled distribution of these changes across individuals and time.

Business Cycle We define "recessions" as periods in which the HP-filtered unemployment rate lies in the top 20% of realizations, while "expansions" are defined as periods when unemployment is below this threshold. This approach allows us to analyze cyclical changes in the earnings growth distribution according to the state of the aggregate labor market instead of NBER recessions/expansion periods, which typically encompass only the rise in unemployment and not its slow recovery. Further, this approach does not require making any interpolation assumptions on earnings growth to cover SIPP data gaps.

2.2 The cyclicality of occupational movers' earnings growth

A substantial body of empirical research has already established that employer transitions account for a disproportionate share of large earnings changes in the labor market.[9] Building upon this observation, we use the distribution of annual earnings growth $G(\Delta)$ to show that occupation mobility exhibits an even stronger association with large earnings changes.

Among the group of workers who changed employers, either EUE or EE, Figure 1 depicts the kernel density estimates of this distribution separately for occupation switchers and non-switchers, for expansion and recession periods.[10] It shows that across the business cycle occupation switchers exhibit substantially greater earnings growth dispersion than non-switchers, as evidenced by the lower and flatter peak of the density functions. The greater dispersion reflects that the largest earnings gains and losses observed in the data are significantly more likely to coincide with occupation switches.

[8]IHS differences approximate log differences except at very low or zero earnings, where they avoid implausibly large changes. We further restrict the sample to workers with annual earnings of at least $1,040 — the threshold corresponding to ten hours per week at the national minimum wage over a quarter, following Guvenen et al. (2014). As a result, our IHS-based estimates are very close to what log differences would yield.

[9]In Online Appendix C we verify that the preponderance of substantial earnings changes, whether large positive or negative, occurs in conjunction with an EE or EUE transition, rather than within continuing employment relationships.

[10]We exclude workers who remain with the same employer from these figures because, as shown in Online Appendix C, their substantial concentration of mass at the center of the distribution, reflecting minimal annual earnings changes, would obscure the cyclical movements in the tails that are the primary focus of our analysis.

In Online Appendix C we confirm this feature. We show that the proportion of occupation switches associated with large (positive or negative) earnings changes is larger than the proportion associated with smaller earnings changes. To quantify the importance of occupation switchers to earnings growth dispersion, we calculate the variance of the aggregate earnings growth distribution. We divide the sample into $k = 4$ groups: (i) occupation and employer movers, (ii) occupation stayers and employer movers, (iii) occupation movers and employer stayers, and (iv) occupation and employer stayers. We find that the 5.7% of observations associated with occupation switching contribute about 50% of the overall variance. Among employer movers the contribution of occupation switchers is even bigger, where switchers contribute twice as much to the variance as occupation stayers.[11]

Figure 1 also shows a leftward shift of the densities during recessions, indicating a deterioration in typical earnings growth outcomes, which is considerably more pronounced among occupation switchers. During recessions, the density for occupation switchers becomes flatter, reflecting an increase in dispersion, while the mode shifts from a slightly positive value to a negative one. However, the most dramatic cyclical variation arises in the tails of the distribution, where substantial reallocation of mass occurs. The density at large negative earnings changes rises substantially in recessions, and rises by more than twice as much for occupation switchers as for non-switchers. For example, at $\Delta \approx -1$ (a 63% decline in annual earnings), the recession-minus-expansion gap in the density is 5.7 percentage points for switchers and 2.4 percentage points for non-switchers. At the same time, positive earnings changes are lower in recessions than expansions, and the corresponding cyclical gap is markedly larger for occupation switchers than non-switchers.

2.3 Skewness over the business cycle

The cyclical movement of the earnings growth distribution among occupation switchers depicted in Figure 1a constitutes a primary source behind the procyclical skewness of $G(\Delta)$. We quantify the importance of occupational mobility to procyclical skewness by evaluating the contributions of $k = 6$ groups separately: occupation switchers and non-switchers, each subdivided by the type of employer transition — EE, EUE, and employer stays. Relative to the four groups used in the variance calculation in Section 2.2, we now split employer movers into EE and EUE transitions to emphasize their importance in explaining the cyclical change in left-skewness. The overall change in skewness is given by

$$\Delta Skew = \frac{1}{std(\Delta|exp)^3} \cdot \frac{1}{N_{exp}} \sum_k \sum_{i \in S_{k,exp}} n_i \left[(\Delta_i - E[\Delta|exp])^3 \right] - \frac{1}{std(\Delta|rec)^3} \cdot \frac{1}{N_{rec}} \sum_k \sum_{i \in S_{k,rec}} n_i \left[(\Delta_i - E[\Delta|rec])^3 \right],$$

where $S_{k,exp}$ ($S_{k,rec}$) refers to all observations in group k during expansions (recessions), n_i refers to the sample weight of an observation i, N_{exp} (N_{rec}) to the weighted sum of all earnings change observations during expansions (recessions) and the time subscript is left implicit. The first term

[11] The variance contribution of each group is as follows: (i) occupation and employer movers, 41.7%, (ii) occupation stayers and employer movers, 18.7%, (iii) occupation movers and employer stayers, 5.8% and (iv) occupation and employer stayers, 33.8%.

is the skewness of $G(\Delta|exp)$ and is calculated to be -1.89, while the second term is the skewness of $G(\Delta|rec)$ and equals -2.96 (see Table 2 below). Hence in our data, left-skewness increases in recessions by 56.6%. As a comparison, the variance of $G(\Delta)$ only increases by 18.2% from 0.22 in expansions to 0.26 in recessions.

Keeping constant the overall mean and standard deviation of earnings growth in expansions and recessions, $E[\Delta|c]$ and $std(\Delta|c)$ with $c \in \{exp, rec\}$, the contribution of each of the k groups to the cyclical change of skewness can be obtained as the ratio between

$$\frac{1}{std(\Delta|exp)^3} \sum_{i \in S_{k,exp}} \frac{n_i}{N_{exp}} \Big(\Delta_i - E[\Delta|exp]\Big)^3 - \frac{1}{std(\Delta|rec)^3} \sum_{i \in S_{k,rec}} \frac{n_i}{N_{rec}} \Big(\Delta_i - E[\Delta|rec]\Big)^3$$

and $\Delta Skew$. Table 1 reports the results of this decomposition and shows that those workers who changed occupations are the major contributors to the observed increase in the left-skewness during recessions. This group of workers contribute 58.4% to the increase in left-skewness, while the contribution of workers who did not switch occupations is around two-thirds of that of occupation movers. This happens even though the latter group represent a small proportion of all observations, 5.7%. Breaking this down further, EUE transitions accompanied by an occupation switch alone account for 50.1% of the cyclical change in left-skewness, while EE transitions with an occupation switch contribute 7.8%. Excluding occupation switchers from the sample would therefore mechanically reduce $\Delta Skew$ by more than half.[12]

Table 1: Decomposition of Cyclical Change in Skewness

	No Occ. Switch			Occ. Switch		
	Emp. Stay	EE	EUE	Emp. Stay	EE	EUE
Contribution to $\Delta Skew$	2.0%	5.8%	33.8%	0.5%	7.8%	50.1%
Fraction of Observations	0.865	0.027	0.051	0.012	0.014	0.031

Note: Decomposition of cyclical change in third central moment across occupation switchers and non-switchers. Annual earnings growth is based on residualized earnings, and recessions are defined as periods in which the HP-filtered unemployment rate lies in the top 20% of realizations as described in Section 2.1.

Graphical representation of cyclical skewness To visualize the importance of simultaneous occupation and employer transitions in driving the procyclical skewness of $G(\Delta)$, it is useful to consider a quantile function representation. Figure 2a depicts these quantile functions separately for expansion and recession periods, where the horizontal axis denotes quantiles of $G(\Delta)$ (ranging from 0 to 1), and the vertical axis measures the corresponding earnings growth rate at each quantile. For instance, the value at a given quantile gives the earnings growth in either expansions or recessions, and the vertical distance between the two curves indicates the extent to which earnings growth at this quantile is higher during expansions than recessions.

[12] Our decomposition relates to the one proposed by Halvorsen et al. (2024), Lemma 1. However, there are two important differences. They use the average earnings change of each group k instead of the average earnings change in the sample to compute the skewness of each group. Further, directly applying their decomposition leaves a co-skewness term that is ambiguous where to attribute, while in our case the co-skewness terms are attributed to the group whose mean is far from the unconditional mean. Nevertheless, in Online Appendix C we show that, conditional on the co-skewness terms, our results do not change when using their decomposition. Occupation switchers contribute about 68.5% to the cyclical skewness of the earnings growth distribution.

Figure 2: Cyclical Variation of the Earnings Growth Distribution

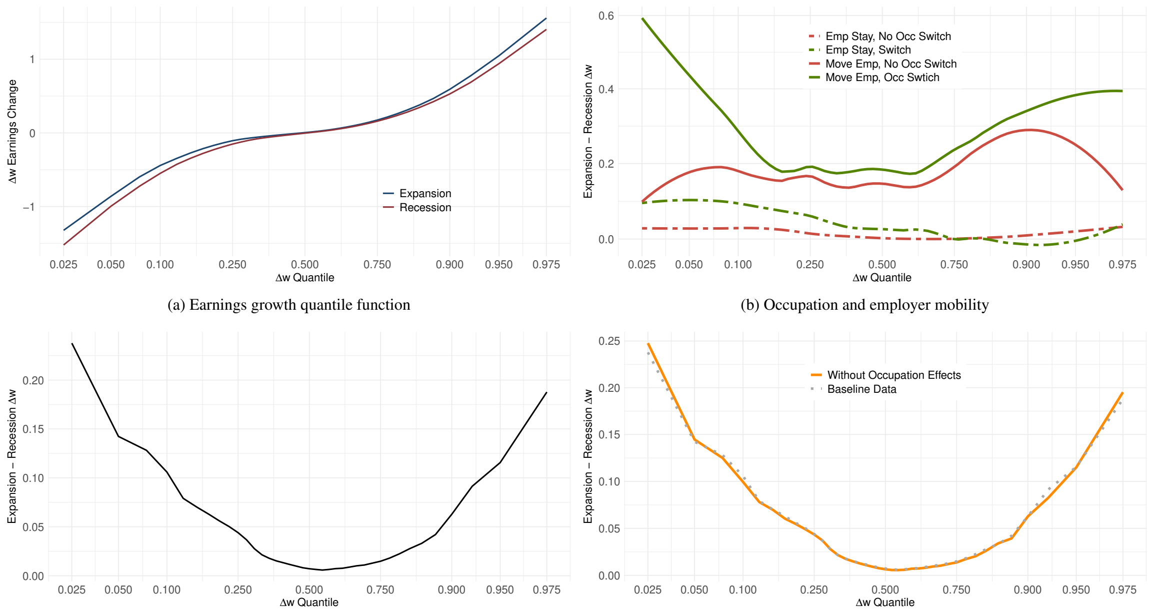

(a) Earnings growth quantile function

(b) Occupation and employer mobility

(c) Quantile function cyclical difference - baseline earnings

(d) Quantile function cyclical difference without occupation effects

Note: The top-left panel shows the CDF of annual earnings growth in expansion and recession periods, while the bottom-left panel shows the difference between the two. The top-right panel shows the difference between annual earnings growth in recession and expansion periods, separately for occupation and employer movers, occupation stayers and employer movers, occupation movers and employer stayers, and occupation and employer stayers. In this case, we use the quantiles obtained from the unconditional sample. The bottom-right panel shows this difference based on residualized earnings without occupation-fixed effects, where the latter are computed using 2-digit occupations.

Figure 2c depicts this (vertical) difference between the expansion and recession quantile functions to show more clearly the widening gap between the expansion and recession curves as we move away from the median toward the tails. The figure shows that the cyclical difference in earnings growth is characterized by a U-shaped pattern, indicating an asymmetric response: recessions disproportionately increase the frequency and magnitude of large negative earnings changes while simultaneously reducing the frequency and magnitude of large positive earnings changes. For example, workers at the 10^{th} percentile of the earnings growth distribution have Δ values about 0.10 lower in recessions than in expansions, while workers at the 90^{th} percentile have Δ values about 0.07 higher in expansions than in recessions. This U-shaped pattern visually represents the procyclical skewness of the earnings growth distribution documented above and by Guvenen et al. (2014).[13] This U-shape highlights how much more cyclical the tails of $G(\Delta)$ are than the median, despite the fact that changes in median earnings over time are closely followed macroeconomic indicators of earnings cyclicality.

Figure 2b depicts the same cyclical difference between the quantile functions but now separately for (i) occupation and employer movers, (ii) occupation stayers and employer movers, (iii) occupation movers and employer stayers, and (iv) occupation and employer stayers. The results re-confirm the importance of occupation and employer movers in driving the procyclical skewness of the earnings growth distribution. This group of workers exhibits larger losses in recessions and these losses become

[13]The interpretation of this U-shaped pattern becomes clearer when we consider alternative distributional shifts. For example, if cyclical variations in earnings growth consisted solely of a mean shift — that is, if the entire distribution shifted uniformly without any change in shape — then the curve in Figure 2c would be a horizontal line at a constant positive value. Conversely, if recessions were characterized by an increase in variance with no change in skewness, such that both positive and negative earnings changes became more extreme proportionally, the difference curve would exhibit a monotonically downward-sloping pattern, with larger differences at higher quantiles.

much larger at the bottom end of the distribution. They also exhibit larger gains in expansions, which become even larger at the top end of the distribution. These cyclical features then create a U-shaped pattern that mimics quite closely that of Figure 2c, and implies that the earnings growth distribution of occupation and employer movers is also characterized by procyclical skewness. Further, in Online Appendix C we show that a similar U-shaped pattern emerges regardless of whether the occupation and employer change occurred either through an EUE or EE transition. Occupation movers that stayed with their employers also have larger earnings losses during recessions. However, these losses are modest, increasing only slightly towards the bottom end of the distribution.

Among those who changed employers but did not switch occupations, we only observe an increasing pattern for 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. Note that the difference between expansion and recession quantile functions among occupation and employer movers lies above that of occupation stayers and employer movers. This indicates that the former group experiences larger earnings changes across the cycle relative to the latter group. For occupation and employer stayers we observe a slight level-change between recessions and expansions, whereby earnings losses in recessions are of a similar magnitude as earnings gains in expansions.

Wages and hours The above U-shaped pattern still holds when considering changes in hourly wages instead of earnings. This is important as occupation mobility can be associated with changes in hours worked and the U-shape we find among occupation and employer movers could be mostly driven by changes in hours worked. Our results show that hours are not the only drivers. In Online Appendix C, we use hourly wages while employed instead of total earnings and show that changes in hourly wages can explain a large part of the U-shape characterizing the procyclical skewness of $G(\Delta)$ and also among EUE and EE occupation movers.

Worker characteristics Worker composition effects can also play an important role for procyclical skewness, and certainly the cyclicality of worker earnings differs across observable groups. Importantly, however, most of our key patterns hold across groups. Table 2 shows that all demographic groups experience more left-skewness during recessions and groups that generally have more transitions and more cyclical unemployment contribute relatively more to the overall increase in left-skewness. We find that males contribute slightly more than females, while young workers (less than 30 years old) contribute more than twice as much to cyclical skewness as older workers, and that the increase in left-skewness mostly arises from white individuals (relative to non-whites). We divide education into three mutually exclusive categories: college, high-school diploma, and high-school dropouts; and find that those with a high-school diploma are the largest group and contribute more to cyclical skewness than the other two. Within all these categories the earnings growth associated with occupational mobility remain more important than the earnings growth of occupational stayers.

Table 2 also shows that the increase in left-skewness during recessions originates in the group of short-tenured workers. These are workers who typically exhibit a high level of employer and occupation churn and are akin to the γ-types in Gregory et al. (2025). A key finding, however, is the very large left-skewness of long-tenured workers. These workers are instead typically higher earners who

Table 2: Central Moment Skewness by Business Cycle Phase

	Expansion	Recession	Contribution to $\Delta Skew$	Proportion in the Pop
Female	-1.96	-2.89	0.45	0.47
Male	-1.78	-3.03	0.55	0.53
White	-1.86	-3.13	0.93	0.79
Non-white	-1.93	-2.38	0.07	0.21
College Graduate	-1.85	-3.52	0.27	0.30
High School Diploma	-1.92	-2.87	0.49	0.61
No High School	-1.59	-2.39	0.24	0.09
Young (age <30)	-0.93	-1.97	0.77	0.23
Older (age $\geq$30)	-2.96	-3.70	0.23	0.77
Occ. short-tenure ($\leq$ 88 months)	-1.48	-2.35	0.87	0.52
Occ. long-tenure ($>$ 88 months)	-3.33	-4.24	0.13	0.48
Overall	-1.89	-2.96	1.00	100

Note: Decomposition of cyclical change in third central moment. Annual earnings growth is based on residualized earnings, and recessions are defined as periods in which the HP-filtered unemployment rate lies in the top 20% of realizations as described in Section 2.1. The short and long tenure classification is in relation to the median tenure observed in our data.

experience low separation probabilities and are akin to the α-types in Gregory et al. (2025). The large left-skewness of these workers' earnings growth distributions arises from (infrequent) transitions into non-employment that lead to an occupational change and very large earnings losses. Table 2 shows that the left-skewness accounted for by these high-tenured workers further increases in recessions. This increase arises from higher separations into unemployment and worsening re-employment pay conditions during recessions. In Online Appendix C we show that higher-tenured workers indeed face an increased transition probability into unemployment during recessions. In Section 2.5 (below) we present evidence suggesting the importance of cyclical changes in the pay conditions or earnings "returns" to occupation mobility.

Online Appendix C also shows the contribution to cyclical skewness among those who made a voluntary employer transition relative to those who made an involuntary one.[14] Among EE and EUE movers, the proportion of involuntary transitions is smaller than that of voluntary ones, involuntary mobility is less common among occupation switchers, and increases in recessions, with the ratio of involuntary to voluntary transitions among employer and occupation movers increasing from 0.30 in expansions to 0.40 in recessions.[15] Voluntary mobility accounts for about three-quarters of the total cyclical change in left-skewness, while involuntary mobility accounts for about one-fifth. The increase of involuntary EUE mobility and their associated earnings losses during recessions are important drivers of the increase in the left tail of $G(\Delta)$, and these effects are more pronounced

[14]Transitions are classified as voluntary when workers declared that the main reason for stopping work with their employer was (i) retirement or old age, (ii) school/training, (iii) temporary job and ended, (iv) quit to take another job, (v) unsatisfactory work arrangements (hours, pay, etc.) and (vi) quit for some other reason. Transitions are involuntary when workers declared that they separated because (i) of layoff, (ii) childcare problems, (iii) other family/personal obligations, (iv) illness or injury, (v) discharged or fired, (vi) employer bankrupt, (vii) employer sold business, and (viii) slack work or business conditions. Noting that there is significant non-response to this question (37.4% of transition observations in our sample), our analysis conditions on a response making it difficult to compare with the results in Table 2.

[15]The finding that EUE involuntary mobility is less common than voluntary mobility is consistent with Flaaen et al. (2019). Among EE occupation switchers, the proportion of involuntary transitions increases from 11.7% in expansions to 17.1% in recessions, while for EE non-switchers it increases from 15.6% to 19.1%. Among EUE occupation switchers, the proportion of involuntary transitions increases from 23.2% in expansions to 31.1% in recessions, while for EUE non-switchers it increases from 31.1% to 39.6%.

among occupation switchers. The importance of these transitions during recessions reflects the larger increase in the earnings losses of long-occupation tenured workers who transit into unemployment involuntarily as discussed above.

2.4 Occupation-wide effects

The annual earnings changes associated with occupational mobility can arise from workers moving from occupations with lower average earnings to occupations with higher average earnings, and vice versa. We now turn to investigate the extent to which occupation-wide effects might be driving the procyclical skewness of the earnings growth distribution. We do so by first examining whether the cyclical patterns documented in Section 2.3 hold within each of the four task-based occupation categories. We then remove occupation-wide variation from the earnings measure and recompute the difference between expansion and recession earnings growth quantile functions.

We start by evaluating the increase in left-skewness during recessions of all the earnings changes that originate from each of the four task-based occupation categories. We find that the earnings growth distribution arising from each of these four occupational categories exhibits a pronounced increase in its left-skewness during recessions, with manual occupations having a somewhat smaller percentage increase in their left-skewness relative to cognitive occupations.[16] As in the aggregate, the increase in left-skewness in each of these occupation categories arises from occupation switchers. To show this, we re-estimate the kernel densities of the earnings growth distribution for occupation switchers and non-switchers, as depicted in Figure 1, but now separately for those originating in each of the four task-based occupations. In Online Appendix C, we show that all four task-based occupation categories share a very similar pattern to the one depicted in Figure 1. Occupation switchers from each source occupation exhibit an earnings growth distribution that is more dispersed with a stronger leftward shift during recessions than those of occupation stayers. These results suggest that the mechanism behind cyclical skewness –occupation switchers driving the leftward shift of the earnings growth distribution in recessions– operates similarly across all four task-based occupational categories.

To further investigate the importance of occupation-wide effects, we re-estimate our earnings equation by explicitly controlling for occupation-specific fixed effects υ_o such that $\log w_{i,t} = x'_{i,t}\beta + \mu_t + \upsilon_o \mathbb{1}(o = o(i,t)) + \epsilon^{\text{nofx}}_{i,t}$, where the last term $\epsilon^{\text{nofx}}_{i,t}$ is now net of occupation fixed effects. Following our previous earnings construction, we derive monthly log earnings net of occupation-wide effects by replacing $x'_{i,t}\beta + \mu_t + \upsilon_o \mathbb{1}(o = o(i,t))$ with the cross-sectional mean of $\log w_{i,t}$ in month t. We construct annual earnings in an analogous way as before and use the IHS transformation of these annual earnings to derive the year-to-year earnings change net of occupation-wide effects, $\Delta^{\text{nofx}}_{i,t}$. By calculating the recession and expansion quantile functions using $\Delta^{\text{nofx}}_{i,t}$, we can evaluate the role of occupation-wide effects in determining the U-shaped pattern that characterizes the procyclical skewness of $G(\Delta)$. To minimize composition effects arising from differences across the occupations

[16] The left-skewness of the earnings growth distribution arising from non-routine cognitive occupations increases from -1.81 to -3.21, while that of routine cognitive increases from -1.74 to -2.94. A similar pattern is observed among the manual occupations, where the left-skewness of the earnings growth distribution arising from non-routine manual occupations increases from -1.58 to -2.47, and that of routine manual occupations increases from -1.96 to -2.67.

that make up each of the four (coarser) task-based categories, we perform this exercise using the 22 occupation categories of the two-digit 1990 SOC. If occupation-wide effects are important in determining the procyclical skewness of $G(\Delta)$, we should observe a dampening or disappearance of the U-shape. Figure 2d shows the results of this exercise. It suggests that controlling for occupation-wide effects when constructing our earnings measure has little effect on the original U-shape.

In Online Appendix C, we investigate the role of occupation-wide effects further by computing the distribution of changes in occupation fixed effects conditional on earnings growth quantiles, $\tilde{G}(\Delta v_o|q)$, where the quantiles q are obtained using our baseline earnings measure. We also compute the unconditional distribution of changes in occupation fixed effects, $\tilde{G}(\Delta v_o)$. We show that $\tilde{G}(\Delta v_o|q)$ varies little across earnings growth quantiles, so large positive and large negative earnings changes are associated with similar conditional distributions of Δv_o. We also show that the $var_q(E[\Delta v_o|q])$ increases in recessions and contributes to the countercyclical increase in the variance we also find for $\tilde{G}(\Delta v_o)$, where $var(\Delta v_o) = E_q[var(\Delta v_o|q)] + var_q(E[\Delta v_o|q])$. This evidence thus suggest that occupation-wide effects create a force towards the increase in the variance of the $G(\Delta)$ during recessions. If they were the main drivers of cyclical changes in $G(\Delta)$, we would then observe a countercyclical increase in its variance instead of the documented increase in left-skewness.

2.5 The role of transition flows on cyclical earnings growth

The increase in left-skewness during recessions can be naturally linked to the cyclical changes in workers' job separations and job finding probabilities. It is well documented that the probabilities of an EE or UE transition increase during expansions, while the probability of an EU transition increases in recessions. In our data the expansion-to-recession ratios of these probabilities are 1.18, 1.09, 0.75, respectively, verifying their documented cyclicality. Further, the probabilities of an occupational change given an EE or an EUE transition also increase in expansions, such that their expansion-to-recession ratios are 1.11 and 1.07, respectively. We now show evidence suggesting that these transition flows do contribute to the *cyclicality* of the tails of $G(\Delta)$, but do not appear to be their major driver. Instead, we show that the earnings growth distribution of occupation switchers *conditional* on these transitions flows is able to explain most of the cyclical variation in $G(\Delta)$.

For this purpose, we implement a reweighting decomposition in the spirit of DiNardo, et al. (1996) and Machado and Mata (2005) that separates the contribution of cyclical changes in transition flows from cyclical changes in earnings growth conditional on those flows. We decompose $G(\Delta)$ in expansions and in recessions by considering the weighted sum of 10 conditional distributions for each state of the business cycle. Let $G(\Delta|\mathbb{I}_k, \mathbb{I}_s, \mathbb{I}_d)$ denote the earnings growth distribution conditional on whether workers switched occupations or not $k \in \{sw, nsw\}$, the state of the business cycle $s \in \{exp, rec\}$, and the duration of the transition, d, where $d = 0$ implies an employer stay or an EE transition and $d \in \{1, 2\text{-}3, 4\text{-}6, > 6\}$ implies an EUE transition with a non-employment duration of 1 month, 2-3 months, 4-6 months and more than 6 months.[17] To recover the full distribution, each

[17] The indicator function $\mathbb{I}_k$ takes the value of one when $k = sw$ and zero otherwise; $\mathbb{I}_s$ takes the value of one when $s = exp$ and zero otherwise; and $\mathbb{I}_d$ takes the value of one when the duration of the transition lies within a corresponding category and zero otherwise.

of the conditional distributions is weighted by the fraction of the population belonging to each of the 10 groups in the respective phase of the business cycle. We then construct a counterfactual recession earnings growth distribution, $G^c(\Delta|rec)$, by applying the recession weights to each of the associated expansion distributions, $G(\Delta|\mathbb{I}_k, exp, d)$. The difference $G(\Delta|exp) - G^c(\Delta|rec)$ gives us the extent to which cyclical changes in the transitions flows explain cyclical changes in $G(\Delta)$ as captured by $G(\Delta|exp) - G(\Delta|rec)$.[18]

Table 3: The Contribution of Transition Flows to Cyclical Changes in $G(\Delta)$

Percentiles	$[0, .05]$	$(.05, .1]$	$(.1, .25]$	$(.25, .75)$	$[.75, .9)$	$[.9, .95)$	$[.95, 1]$
Transition flows	0.193	0.141	0.085	0.082	0.100	0.055	0.023
Flows & Occ. Δ	0.889	0.874	0.892	0.854	0.733	0.775	0.796

Note: Decomposition of cyclical change in third central moment. Annual earnings change based on residualized earnings after controlling for potential experience, education and month dummies. Recessions are defined as periods in which the HP-filtered unemployment rate is in the top 20% of realizations.

Table 3 shows the result of this decomposition. The first row, reports the proportion of the cyclical difference in earnings growth that is explained by changes in transition flows, $[G(\Delta|exp) - G^c(\Delta|rec)]/[G(\Delta|exp) - G(\Delta|rec)]$. The relative contribution of these flows is stronger at the bottom end due to the increase in flows into non-employment — by far the most cyclical of the transition probabilities, with an expansion-to-recession ratio of 0.75 compared to 1.09 and 1.18 for UE and EE transitions. However, the importance of flows diminishes nearly monotonically as we move towards the top end of the distribution. Hence, cyclical changes in transition flows have a non-trivial but small impact: they account for about a fifth of the cyclical difference at the bottom of the distribution but only about one-tenth by the lower quartile.

Building on this decomposition, we construct a second counterfactual $\tilde{G}^c(\Delta|rec)$ which now considers not only the recession transition weights but also the recession distributions associated with occupation switchers, $G(\Delta|sw, rec, d)$, and keeps the distribution of occupation stayers at their expansion levels. This counterfactual allows us to investigate the extent to which the observed earnings changes associated with occupational switching, conditional on transition flows, contribute to the difference between the expansion and recession earnings growth distributions. The second row of Table 3 presents the results. It implies that cyclical changes in earnings growth associated with occupation switching explain about 70% to 80% of the difference between $G(\Delta|exp)$ and $G(\Delta|rec)$. This represents roughly four times the contribution of cyclical changes in flows alone reported in the first row. That is, cyclical changes in the "earnings returns" to occupational mobility —the earnings change conditional on a given transition— appear to explain the bulk of the U-shape, while cyclical changes in transition flows seem to play a smaller role.

[18] Our decomposition combines the reweighting approach of DiNardo, et al. (1996) with the conditional-distribution approach of Machado and Mata (2005). Because our "covariates" are categorical (the occupation-switch indicator and a coarse non-employment duration bin), the conditional distributions $G(\Delta|\mathbb{I}_k, \mathbb{I}_s, \mathbb{I}_d)$ are estimated directly as empirical conditional CDFs — without the quantile-regression step that Machado and Mata (2005) use to handle continuous covariates. The counterfactual recession distribution then mixes the expansion conditional distributions with the recession population weights of each cell.

Figure 3: The Persistence of Earnings Changes

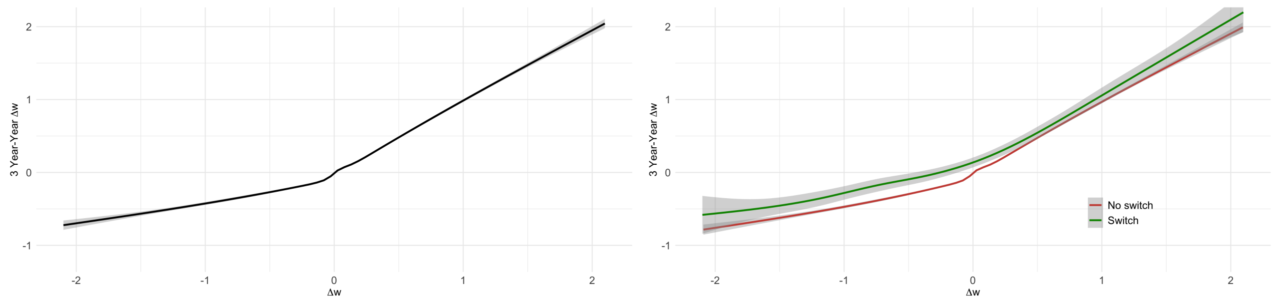

(a) All workers

(b) Occupation switchers vs non-switchers

Note: The left panel shows the relationship between workers' annual earnings changes and these workers' earnings changes at 3-year horizons. The right panel shows these same relationships separately for occupational switchers and non-switchers, using the task-based occupational categories.

The substantial deterioration of these "earnings returns" during recessions for occupation switchers, combined with the limited role of transition flows, suggests that the quality of matches involving an occupational switch tends to be lower in recessions than in expansions. This pattern aligns with the broader literature on cyclical match quality: Bowlus (1995) provides empirical evidence that worker-firm matches formed during recessions are of relatively lower quality, and Moscarini (2001) advances a related theoretical argument that depressed labor markets "perform less well and produce noisier allocations." Indeed, we find that workers who end an unemployment spell with an occupation switch during a recession have a higher separation rate from their re-employment job than occupation stayers.[19]

Our decomposition results suggest that cyclical changes in earnings changes conditional on transition flows, account for the bulk of the U-shape, while cyclical changes in transition flows play a smaller role. This decomposition, however, cannot disentangle workers' choices from the cyclical environment: observed transition flows reflect worker choices in the face of changing returns and opportunities, and observed earnings changes reflect acceptance decisions shaped by both flows and alternative returns elsewhere. We address this in Section 3 by developing a structural model of workers' employer and occupational mobility choices that allows for cyclical variation in both flows and the match quality of jobs.

2.6 Persistence of earnings changes

Finally, we investigate whether the cyclical asymmetries and the occupation-switcher patterns documented in earlier subsections persist at longer horizons. Given that each SIPP panel follows individuals for at most four years, we construct for each worker their three-year earnings changes and compare them with their one-year earnings changes.

Figure 3a considers these relationships by plotting workers' initial annual earnings change against

[19] In particular, we regress the spell duration of the re-employment job on a constant, a recession indicator, an indicator of whether the job started in a recession, an indicator for occupation change, and the interaction for the last two variables. We find that all variables are statistically significant at least at a 5% level, with the coefficient estimates of 1.04, -2.62, 2.46, -0.14, and -0.11. These findings complement Bowlus (1995), who presents empirical evidence showing that the quality of newly formed employment matches worsens during recessions relative to expansions and that starting wages reflect this cyclical deterioration.

their subsequent three-year earnings changes. It shows that those who did not change their annual earnings also did not change their three-year earnings. As we move towards the tails of $G(\Delta)$, we observe a pronounced asymmetric pattern. Earnings *gains* are very persistent at the three-year horizon, with a relationship that closely traces the 45-degree line. For example, a worker who experienced an annual earnings change of 1.5 also experienced an earnings change of 1.5 at the three-year horizon. In contrast, earnings *losses* are less persistent, as shown by the convex relationship between one-year and three-year losses. The figure shows that the larger earnings losses recover by more than half by the third year. These large losses are concentrated among workers who experienced longer non-employment spells, particularly the long-tenured involuntary EUE movers highlighted in Section 2.3. These findings then suggest that the large earnings losses driving the procyclical skewness of $G(\Delta)$ are not fully transitory: a substantial portion remains unrecovered at the three-year horizon.

Figure 3b shows that the same asymmetric pattern is present separately for occupation switchers and non-switchers. A key feature is that occupation switchers' annual earnings losses recover more at a three-year horizon than occupation stayers with the same annual earnings losses. This suggests that those occupation and employer movers who experienced earnings losses have lower earnings initially but have a steeper earnings rise in the future.[20]

3 Theoretical Framework

These empirical patterns motivate a structural model with on-the-job search and endogenous occupation mobility, in which workers face a two-dimensional job ladder along both the employer and the occupation dimensions of a job. We use this model to evaluate the contributions of workers' endogenous mobility decisions and cyclical labor market risk in shaping cyclical changes in $G(\Delta)$. In turn, this will allows us to assess whether the cyclical changes in $G(\Delta)$ arise from the occupation or the employer dimension of the ladder.

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 $o = 1, \dots, O$. Business cycles are modeled through fluctuations in the economy-wide productivity A_t, which follows a first-order stationary autoregressive process. We also allow some occupations to be more attractive than others in terms of their occupation-wide productivity. Let $p_{o,t}$ denote the occupation-wide productivity of occupation o at time t and assume it also 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.

Employer ladder Within an occupation, unemployed and employed workers search for jobs and face a probability $\lambda_U(A_t)$ and $\lambda_E(A_t)$ of meeting a firm in a given period. Once a meeting takes place, a worker draws a productivity ϵ from an exogenous distribution $\Gamma(.; A)$. We interpret ϵ as

[20] Note that our findings refer to earnings *changes*, while other work has focused on the persistent cost of occupation displacement based on earnings *levels* (see Huckfeldt, 2022). The two are fully consistent: a faster recovery in earnings changes among occupation switchers can coexist with a persistent gap in their long-run earnings levels relative to non-switchers.

capturing the idiosyncratic firm-worker productivity of an employment match, akin to Low et al. (2010), and assume that it follows a common and exogenous first-order stationary Markov process once the match is formed. If the worker is unemployed, he decides whether to accept employment at that firm or wait in unemployment for a better offer. To capture voluntary and involuntary employer mobility, we follow Jolivet et al. (2006) and assume that with probability $1-\gamma$ the employed worker is able to decide whether to accept the job offer if the ϵ-productivity is sufficiently attractive; and with probability γ he is forced to accept it, as long as it offers a payoff above the value of unemployment. The $\Gamma(.;A)$ distribution thus describes the employer dimension of the job ladder and by allowing it to shift according to the state of the business cycle we capture the possible cyclical changes in the quality of the worker-firm matches.

Occupation ladder Along the occupation dimension, the worker's position in the job ladder is captured by three components, p_o, z and x_h. The idiosyncratic occupation-worker productivity z represents a worker's "career match" and determines how well he is doing in the current occupation, akin to Neal (1999). Once an occupation-worker match is formed, this productivity follows a common and exogenous first-order stationary Markov process and their realizations affect workers both in employment and unemployment as long as he stays within the same occupation. Differences in z reflect that not all workers in the same occupation are equally productive, even if they have similar observable characteristics, as they might be exposed to different occupation match-specific, industry and location factors that shape their career prospects within the occupation.[21]

In addition, workers accumulate occupational-specific human capital x_h 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 $h = 1, \ldots, H$. 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 $\chi_u(x_{h-1}|x_h)$ $h = 1, \ldots, H$.[22]

Search across occupations Workers can always decide to change (p_o, z, x_h) by searching in a different occupation. Changing occupations comes at the cost of losing the accumulated occupational-specific human capital upon reallocation, but entails the benefit of re-starting a worker's z-productivity process. Search across occupations is modeled following an imperfect directed search technology that

[21] The idiosyncratic z-productivity can be viewed as reflecting the productivity a worker has in a sub-market, defined as the cross-product of an occupation/industry/location tuple, where each aggregate occupation o is composed of a large number of these sub-markets. For example, the productivity of a chemical engineer working for the photo industry in Rochester, New York, can differ from the productivity of another, equally experienced, chemical engineer working in the oil industry in Houston, Texas. The chemical engineer in Houston could be more productive than the one in Rochester, not only because he might be better suited to be a chemical engineer, but also because the oil industry may be more productive than the photo industry, and location-specific public and private investments could make chemical engineers in Houston more productive than those in Rochester.

[22] The accumulation of occupation-specific human capital, which is shared across all workers with a given occupation o, aims at capturing productivity increases that arise from performing the occupation-specific tasks. For simplicity, we assume that occupational human capital x_h is orthogonal to the career match z-productivities and occupation-wide p_o-productivities. Workers can, in principle, also accumulate firm-specific human capital during employment spells. However, the empirical results of Kambourov and Manovskii (2009b) and Sullivan (2010), who in addition to between-employer occupation mobility also consider within-employer occupation mobility as done here, 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.

involve two stages. The first stage entails the decision to search in a different occupation, and if so, how to allocate search effort across occupations. In particular, once a worker has decided to reallocate, he is endowed with a unit measure of search effort which must be divided across the remaining occupations in order to maximize the expected value of searching across occupations. We assume that within each period a worker can receive at most one z-productivity. A worker leaving occupation o has to decide which proportion $s^i_{\tilde{o}}$ of his search effort to devote to obtain a z from occupation $\tilde{o} \neq o$, where $i = U, E$ denotes the worker's employment status. Let $\mathcal{S}^i$ denote a vector of $s^i_{\tilde{o}}$ for all $\tilde{o} \in O^-$, where O^- denotes the set of remaining occupations such that $\sum_{\tilde{o}\in O^-} s^i_{\tilde{o}} = 1$. The probability that a worker, currently in occupation o, receives a new z from an occupation $\tilde{o}$ is given by $\alpha(s^i_{\tilde{o}}, o)$, where $\alpha(., o)$ is a continuous, weakly increasing and weakly concave function with $\sum_{\tilde{o}\in O^-} \alpha(s^i_{\tilde{o}}, o) \leq 1$, for all $o \in O$. The latter implies that the probability of not receiving a new z is given by $1 - \sum_{\tilde{o}\in O^-} \alpha(s^i_{\tilde{o}}, o)$.

Initial z-productivities in any occupation are drawn from an exogenous distribution $F(z; A)$. We allow this distribution to also shift according to the state of the business cycle to capture possible cyclical changes in the quality of the worker-occupation matches.[23] Cyclical variations in both $F(z; A)$ and $\Gamma(\epsilon; A)$ are motivated by the statistical decomposition exercise of Section 2.5, which suggests that both cyclical variation in the transition probabilities and in the earnings returns to occupation and employer mobility could play a role in shaping the cyclical properties of the aggregate earnings growth distribution. Those earnings returns can be approximated through the match quality distributions.

Search and matching in the new occupation The second stage of the job search process across occupations entails finding a job in the new occupation. Workers start with human capital x_1 and are able to search for jobs the same period in which reallocation takes place. Unemployed and employed workers meet firms with probabilities $\lambda^c_U(A_t), \lambda^c_E(A_t)$, respectively. 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. Once a meeting takes place, employed workers face the same γ probability that determines if the transition to a new employer is voluntary or not. If no meeting takes place and the worker is unemployed, the worker remains unemployed in the new occupation. If the worker is employed and no meeting takes place, this worker stays with his current employer but in the new occupation. Workers can decide to reallocate once again, but they must sit out one period in their new employment 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 again.

Production and earnings The output of an employed worker characterized by (z_t, x_h, ϵ_t) in occupation o is given by $y(A_t, p_{o,t}, \epsilon_t, z_t, x_h)$, where $y(.)$ is increasing and continuous in all of its arguments. Unemployed workers receive b each period. Individual earnings w are determined by worker and firm bargaining such that w is strictly increasing and continuous in all of the arguments of $y(.)$. We use the bargaining protocol proposed by Nagypal (2008) such that earnings each period are obtained through Nash Bargaining but with delay instead of a permanent severance of the firm-worker

[23]Cyclical shifts in $F(z; A)$ can arise from cyclical shifts in the occupation match-specific, industry, and location components captured by z. In the example above, a negative aggregate shock affecting Rochester's photo industry more than Houston's oil industry would reduce the relative quality of the Rochester chemical engineer's occupation match.

relationship determining the disagreement payoffs (see also Hall and Milgrom, 2008, and Elsby and Gottfries, 2022). When disagreement entails the worker receives b and the firm receives zero as per-period payoffs, but otherwise the worker is allowed to search, bargained earnings are given by $\Lambda y(A_t, p_{o,t}, \epsilon_t, z_t, x_h) + (1 - \Lambda)b$, where Λ denotes the worker's exogenous bargaining power.[24] Note that the evolution of workers' output in the model remains endogenous as they make employer and occupational transition decisions that affect their productivities.

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. In addition, workers can fall from the job ladder through exogenous reasons. In keeping with the two dimensions of the job ladder, 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-worker 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-worker match productivity and is forced to change occupation through unemployment. In the spirit of Huckfeldt (2022) we interpret the latter as "obsolescence shocks".[25]

Inspection vs. experience goods The two ladders differ not only in the productivity components but also in the information available to the worker at the moment of moving. Movement along the *employer* ladder is treated as an "inspection" good: when meeting a new firm, the worker observes the value of the firm-match ϵ before deciding to accept. Movement along the *occupation* ladder is instead modeled as an "experience" good in the following sense: the worker must commit to changing occupations *before* observing the realizations of z and the new p_o. While he knows $F(z; A)$ and the process governing p_o, the specific realizations are revealed only after he begins searching in the new occupation. Crucially, this commitment is not fully reversible: a worker who finds the resulting (z, p_o) disappointing can search again the following period, but he has already lost his occupation-specific human capital and his original z is replaced by a new draw. The employer ladder differs in both respects: ϵ is observed before the acceptance decision, and an unfavorable draw is simply rejected at no cost. This asymmetry between the two ladders captures the risk of occupational mobility underlying the patterns we document in Section 2.

Sources of labor market risk The model thus has several sources of cyclical labor market risk. Similar to Low et al. (2010), the risk to employment mobility is given by (i) the probabilities of job loss and forced reallocation $\delta_z(A_t)$, $\delta_\epsilon(A_t)$, γ; (ii) the job finding probabilities $\lambda_U(A_t)$, $\lambda_E(A_t)$,

[24] See Online Appendix D for a detail discussion of the bargaining protocol. In Section 4 we map the negotiated w to monthly earnings, which one could further decompose into hours worked and hourly wages. As discussed in Section 2 and shown in Online Appendix C, the cyclical dynamics of the earnings growth distribution and the importance of simultaneous occupation and employer mobility in driving these dynamics arise from changes in both hourly wages and hours worked. For simplicity, we capture their joint effect through changes in the components of $y(.)$.

[25] 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.

$\lambda^c_U(A_t)$, $\lambda^c_E(A_t)$ and (iii) the match quality distributions $F(z;A)$ and $\Gamma(\epsilon;A)$. The risk to earnings arise from the stochastic evolution of productivities A, $\mathcal{P}_O$, z, ϵ and occupation human capital x. Workers react to these risks by deciding whether to accept new employment opportunities and whether to quit into unemployment, and whether to change occupations and how to direct their search across occupations.

Timing and state space The timing of the events is summarized as follows. At the beginning of the period the new values of A, $\mathcal{P}_O$, z, x and ϵ are realized. After these realizations, 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 earnings on productivities as described above.

3.2 Workers' problem

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

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

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 value for an unemployed worker of searching across occupations is given by

$$\begin{aligned} R^U(x,z,o,\Omega) = \max_{\mathcal{S}^U}\sum_{\tilde{o}\in O^-}\alpha^U(s^U_{\tilde{o}})\int_{\underline{z}}^{\overline{z}}\Bigg[&\lambda^c_U(A)\int_{\underline{\epsilon}}^{\overline{\epsilon}}\max\Big\{W^E(\tilde{\epsilon},x_1,\tilde{z},\tilde{o},\Omega), W^U(x_1,\tilde{z},\tilde{o},\Omega)\Big\}d\Gamma(\tilde{\epsilon};A) \\ + \quad &\big(1-\lambda^c_U(A)\big)W^U(x_1,\tilde{z},\tilde{o},\Omega)\Bigg]dF(\tilde{z};A) + \Big(1-\sum_{\tilde{o}\in O^-}\alpha^U(s^U_{\tilde{o}})\Big)W^U(x,z,o,\Omega), \end{aligned} \tag{2}$$

where the maximization is subject to $s^U_{\tilde{o}} \in [0,1]$ and $\sum_{\tilde{o}\in O^-} s^U_{\tilde{o}} = 1$. Note that this expression incorporates the value of immediately searching 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 characterized by (ϵ, x, z, o). The expected value of employment at the beginning of the production stage is described by

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

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 stay 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

$$
\begin{aligned}
\hat{W}^E(\epsilon, x, z, o, \Omega) = & \int_{\underline{\epsilon}}^{\overline{\epsilon}} (1-\gamma)\lambda_E(A) \max\Big\{W^E(\tilde{\epsilon}, x, z, o, \Omega), W^E(\epsilon, x, z, o, \Omega)\Big\} d\Gamma(\tilde{\epsilon}; A) \\
& + \int_{\underline{\epsilon}}^{\overline{\epsilon}} \gamma\lambda_E(A) \max\Big\{W^E(\tilde{\epsilon}, x, z, o, \Omega), W^U(x, z, o, \Omega)\Big\} d\Gamma(\tilde{\epsilon}; A) \\
& + (1-\lambda_E(A))W^E(\epsilon, x, z, o, \Omega), \qquad (4)
\end{aligned}
$$

where, with probability $(1-\gamma)\lambda_E$, 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 $\gamma\lambda_E$, however, the worker is forced to take the new ϵ' in a different employer, as long as this gives a value higher than unemployment. 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 value of searching across occupations is

$$
\begin{aligned}
R^E(\epsilon, x, z, o, \Omega) = & \max_{\mathcal{S}^E} \sum_{\tilde{o} \in O^-} \alpha^E(s^E_{\tilde{o}}) \Bigg(\int_{\underline{z}}^{\overline{z}} \Bigg[\int_{\underline{\epsilon}}^{\overline{\epsilon}} \Big((1-\gamma)\lambda^c_E(A) \max\Big\{W^E(\tilde{\epsilon}, x_1, \tilde{z}, \tilde{o}, \Omega), W^E(\epsilon, x_1, \tilde{z}, \tilde{o}, \Omega)\Big\} \\
& + \; \gamma\lambda^c_E(A) \max\Big\{W^E(\tilde{\epsilon}, x_1, \tilde{z}, \tilde{o}, \Omega), W^U(x_1, \tilde{z}, \tilde{o}, \Omega)\Big\}\Big) d\Gamma(\tilde{\epsilon}; A) \qquad (5) \\
& + \; (1-\lambda^c_E(A))W^E(\epsilon, x_1, \tilde{z}, \tilde{o}, \Omega) \Bigg] dF(\tilde{z}; A) \Bigg) + \Big(1 - \sum_{\tilde{o} \in O^-} \alpha^E(s^E_{\tilde{o}})\Big) W^E(\epsilon, x, z, o, \Omega),
\end{aligned}
$$

where the maximization is subject to $s^E_{\tilde{o}} \in [0, 1]$ and $\sum_{\tilde{o} \in O^-} s^E_{\tilde{o}} = 1$. Conditional on drawing a z-productivity from another occupation $\tilde{o}$, the worker meets a new employer with probability $\lambda^c_E(A)$, draws a new value of ϵ and with probability $1-\gamma$ 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 $\gamma\lambda^c_E(A)$, 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 $1-\lambda^c_E(A)$, 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 $\gamma\lambda_E^c$ 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

To investigate the quantitative properties of our model, we estimate it 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 labor 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 parametrize 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. This is consistent with our empirical analysis in Section 2 which considered two discrete phases of the business cycle. We approximate the evolution of A_I by an AR(1) process with persistence ρ_A and white noise variance, σ_A. The occupation-wide productivity shock, p_o, is also modeled as an AR(1) process such that $p_{o,t+1} = \tilde{p}_o + \rho_p p_{o,t} + \upsilon_{o,t+1}$ for all $o \in O$, where $\tilde{p}_o$ is a time-invariant occupation specific productivity and υ_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 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 parametrized 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 $\lambda_i^c(A_I) = \lambda_{1,i}^c A_I + (1-A_I)\lambda_{0,i}^c$. 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 $\delta_\epsilon(A_I) = \delta_{1,\epsilon}A_I + (1-A_I)\delta_{0,\epsilon}$.

Upon a decision to switch occupations, new values of z are drawn from $F(.; A_I)$, depending on the state of the business cycle. Each distribution is estimated non-parametrically as described below. The worker-occupation match specific productivity then evolves such that $z_{t+1} = z_t$ with probability $1-\rho_z$ and updates to a new value $z_{t+1} = z'$ with probability ρ_z, where z' is drawn from $F(.; A_I)$.[26] The worker-firm match-specific productivity ϵ remains constant during the duration of the match and only changes upon an EE transition. This is motivated by the small variation in annual earning changes observed among employer stayers (see Online Appendix C). Workers draw a new ϵ from $\Gamma(.; A_I)$, which is also estimated non-parametrically separately for expansion and recession periods.[27]

In parameterizing $\alpha^i(\cdot)$ we aim to respect the observed gross and net mobility flows across occupations. We let $\alpha^i(s_{\tilde{o}}) = \alpha_0\, e^{\alpha_{\tilde{o}}(1-\alpha_1^i)}\, s_{\tilde{o}}^{\alpha_1^i}$, where occupation $\tilde{o} \in O^-$ denotes the search direction, i the worker's employment status, and $s_{\tilde{o}}$ search intensity. The parameter $\alpha_1^i \in [0, 1]$ governs the extent to which workers adjust their search direction in response to differences in $p_{\tilde{o},t}$ across occupations. The parameters $\alpha_{\tilde{o}}$ scale the probability of moving into each occupation $\tilde{o} \in O^-$ and capture differences in ease of access across occupations.[28] The parameter α_0 is a scaling factor that ensures the proper arrival probabilities but plays no role in the marginal choice of search direction. This functional form is also computationally convenient: when $i = U$, for example, optimal search intensity takes a form similar to that implied by a 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

$$s^*_{\tilde{o}} = \frac{e^{\frac{1}{\alpha_1^U}\log\left(\Psi^U(\tilde{z},x_1,\tilde{o},\Omega)-W^U(z,x_h,o,\Omega)\right)}}{\sum_{o'\in O^-} e^{\frac{1}{\alpha_1^U}\log(\Psi^U(\tilde{z},x_1,o',\Omega)-W^U(z,x_h,o,\Omega))}}, \tag{6}$$

where $\Psi^U(\tilde{z}, x_1, \tilde{o}, \Omega)$ is the expected value of search in occupation $\tilde{o}$ as in equation (2). In Online Appendix D, we discuss the relationship with the random utility model, derive a similar expression for employed workers and show that with $\alpha_{\tilde{o}}$ differing across occupations we end up with a multiplier that scales the search direction. Combined with the additive noise used to smooth our approximate solution, workers' two-stage occupation mobility decision (first deciding whether to reallocate and then deciding how to allocate search effort across occupations) takes a similar form to a nested multinomial logit discrete choice model.

Finally, to capture any stickiness in the earnings process that does not arise from the persistence

[26]Note that this within-match evolution of z means that the cyclical match-quality distribution $F(.; A_I)$ shapes both employment-mobility risk (through new draws upon a transition) and within-match earnings risk (through stochastic re-draws at rate ρ_z). This parametrization avoids introducing a separate cyclical shock process for the evolution of z within a match.

[27]In the estimation, we verify that the support of the discretized $F(.; A_I)$ and $\Gamma(.; A_I)$ 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.

[28]The functional form for $\alpha^i(\cdot)$ highlights two reasons why search can be more or less directed. In particular, $\log\alpha^i(\cdot) = \log\alpha_0 + (1-\alpha_1^i)\,\alpha_{\tilde{o}} + \alpha_1^i \log s_{\tilde{o}}$ shows that the probability of obtaining a new z from occupation $\tilde{o}$ is a convex combination of baseline accessibility and log search intensity. As α_1^i approaches one, search becomes more directed, with differences in productivities across occupations increasingly amplified; as α_1^i goes to zero, search ceases to respond to productivity differences and is governed instead by baseline ease of access $\alpha_{\tilde{o}}$. For tractability we assume this ease of access is independent of the origin occupation. Allowing $\alpha_{\tilde{o}}$ to also vary with the origin o (e.g. $\alpha_{o,\tilde{o}}$) would increase the number of search-effort parameters significantly.

of the above productivity dynamics, we approximate bargained earnings as

$$w_{o,t} = \begin{cases} w_{o,t-1} & \text{Prob} = \eta \\ \Lambda(A_t + \tilde{p}_o + p_{o,t} + x_h + z_t + \epsilon_t) + (1-\Lambda)b & \text{Prob} = 1-\eta, \end{cases} \tag{7}$$

where η is the probability that in period t the previous period earnings are carried over. Given 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.

In Online Appendix D we extend Nagypal's (2008) bargaining protocol by combining it with Gottfries' (2025) bargaining with infrequent renegotiation. There we show that equation (7) is the exact solution to the earnings that arise from our extended bargaining protocol when productivity y_t is constant and is an approximation with error of order $\eta(1-\rho)$ when y_t follows an AR(1) productivity process with persistence ρ. As we will show, we estimate a small value for η and persistent productivity processes, such that our approximation is tight and past earnings having little effect in practice on workers' mobility decisions. In Online Appendix D we also investigate allowing labor market tightness to determine earnings dynamics, as would happen in the DMP framework. We find very similar cyclical changes in the earnings growth distribution under both approaches.

4.2 Estimation procedure

The above functional forms yield several parameters to estimate. These are comprised of the set that governs the arrival of job opportunities $\{\lambda_{r,i}, \lambda^c_{r,i}\}_{i=U,E,\ r=0,1}$, the set $\{\delta_{r,z}, \rho_z, \check{F}_{j,r}\}_{r=0,1}$ that governs the idiosyncratic worker-occupation productivities, the set $\{\delta_{r,\epsilon}, \gamma, \check{\Gamma}_{j,r}\}_{r=0,1}$ that governs the idiosyncratic worker-employer productivities. We use the discrete set of j points in recession and expansion, $\{\check{F}_j\}_{r=0,1}, \{\check{\Gamma}_j\}_{r=0,1}$, to define histograms that approximate the two distributions, $F(.; A_I)$ and $\Gamma(.; A_I)$.[29] 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 $\{\alpha_0,\ \alpha^U_1,\ \alpha^E_1,\ \alpha_{NRC},\ \alpha_{NRM},\ \alpha_{RM},\ \alpha_{NRM}\}$, and finally, the set that governs the aggregate productivity process $\{\rho_A, \sigma_A\}$ and the set that governs payments $\{\Lambda, \eta,\ b\}$.

To ease computational burden we first normalize 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).[30] We also set $b = 0.4\bar{w}$ to match a 40% replacement ratio (see Shimer, 2005), where $\bar{w}$ is the average wage to scale these benefits. $\Lambda = 0.5$ as is typically done in the literature (see Petrongolo and Pissarides, 2001). The AR(1) aggregate productivity process parameters are set to the values of the autocorrelation and

[29] For each of these histograms, we define five bins centered at the middle of each quintile. We fix the mass at the median, then choose relative weights for each of these other quantiles. We then interpolate and extrapolate between these points to define the $F(.; A_I)$ and $\Gamma(.; A_I)$ distributions with a large number of bins to approximate a continuous distribution from which agents can draw, which helps smooth out agents' decisions, approximating a continuous distribution with cutoff rules.

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

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 0.958 and variance is 0.009.[31] After fixing the above parameters, we recover the remaining ones following a two-step procedure in which we split the parameter set between an inner and outer loop. 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 of the main text. This loop also contains the parameters 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. The remaining (outer loop) parameters are estimated using simulated method of moments (see Online Appendix D).

4.3 Targeted moments and parameter estimates

Table 4 and Figures 4 and 5 present the set of data moments M^{D} to estimate the outer loop parameters as well as the model fit to these moments. We now present some identification arguments that justify our choice of moments, keeping in mind that all parameters need to be estimated jointly. We also discuss the fit of the model and the value of the parameter estimates.

Table 4: Targeted moments in the estimation

Moment	Model	Data	Moment	Model	Data
Employer Switching					
EE transition prob	0.0291	0.0341	EE prob - expansion/recession ratio	1.1829	1.1846
		(0.0003)			(0.0469)
UE transition prob	0.3809	0.4010	UE prob - expansion/recession ratio	1.0943	1.0876
		(0.0025)			(0.0244)
EU transition prob	0.0168	0.0163	EU prob - expansion/recession ratio	0.7439	0.7460
		(0.0002)			(0.0333)
Occupation Switching					
Prob (Occ. change | EE)	0.2471	0.2685	Prob (Occ. change | EE) - exp/rec ratio	1.0549	1.0536
		(0.0037)			(0.0196)
Prob (Occ. change | EUE)	0.3088	0.2892	Prob (Occ. change | EUE) - exp/rec ratio	1.0800	1.0809
		(0.0034)			(0.0132)
U duration - Occ. movers/stayers ratio	1.1750	1.1637	Prob (Occ. change | Stayer)	0.0102	0.0107
		(0.0215)			(0.0002)
Variance (Net flows | EE)	0.0235	0.0223	Variance (Net flows | EUE)	0.0167	0.0218
		(0.0008)			(0.0012)
Net flow to NRC	0.1878	0.1851	Net flow to RC	0.3290	0.3432
Net flow to NRM	0.2256	0.2201	Net flow to RM	0.2576	0.2516
Earnings Growth Distribution Employer Stayers					
Occ stayers - $90/10$ percentiles	0.5446	0.4384	Occ movers - $90/10$ percentiles	0.8546	0.8146
		(0.0008)			(0.0119)
Occ stayers - Skewness	0.0661	0.1110	Occ movers - Skewness	0.2013	0.0988
		(0.0022)			(0.0119)
Productivities					
NRC fixed effect	1.000	1.000	RC fixed effect	0.767	0.767
NRM fixed effect	0.608	0.608	RM fixed effect	0.803	0.803

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

[31] Shimer (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. We convert the original quarterly data from FRED to a monthly process, consistent with our simulation procedure.

Table 5: Estimated parameter values

Job offer arrival unemployed and employed				Employer- and occupation-match productivities			
$\lambda_{0,U}$	0.9497 (0.0152)	$\lambda_{0,E}$	0.0435 (0.0010)	$\delta_{0,\epsilon}$	0.0028 (0.0001)	$\delta_{0,z}$	0.0080 (0.0001)
$\lambda_{1,U}$	0.7393 (0.0105)	$\lambda_{1,E}$	0.0774 (0.0007)	$\delta_{1,\epsilon}$	0.0003 (0.0000)	$\delta_{1,z}$	0.0036 (0.0000)
$\lambda^c_{0,U}$	0.1537 (0.0095)	$\lambda^c_{0,E}$	0.0124 (0.0027)	γ	0.3500 (0.0003)	ρ_z	0.0733 (0.0086)
$\lambda^c_{1,U}$	0.5288 (0.0364)	$\lambda^c_{1,E}$	0.0912 (0.0010)	**Payments**			
				η	0.0930 (0.0004)		
Search direction across occupations				**Occupation-wide productivities**			
α_0	0.0438 (0.0003)	α_{NRC}	-0.4712	ρ_p	0.9876 (0.0000)	$\tilde{p}_{NRC}$	0 (normalize)
α_1^U	0.2141 (0.0087)	α_{RC}	0.5548	σ_p	0.0299 (0.0001)	$\tilde{p}_{RC}$	-0.2658
α_1^E	0.3952 (0.0043)	α_{NRM}	-0.1991			$\tilde{p}_{NRM}$	-0.4976
		α_{RM}	-0.0608			$\tilde{p}_{RM}$	-0.2189

Note: Standard errors in parentheses only for estimated parameters.

Transition parameters The parameters governing the arrival probabilities of employment opportunities are informed by the observed transition probabilities across employers and occupations. In particular, the arrival probabilities of job offers among occupation stayers, $\lambda_{0,U}, \lambda_{0,E}, \lambda_{1,U}, \lambda_{1,E}$, are informed by the average UE and EE transition rates and their expansion to recession ratio. The arrival probabilities of offers among occupation movers, $\lambda^c_{0,U}, \lambda^c_{0,E}, \lambda^c_{1,U}, \lambda^c_{1,E}$, are informed by the average probabilities of an occupation change conditional on an EE, EUE transition and their expansion to recession ratio, as well as the probability of an occupational change among employer stayers. The ratio of the average unemployment duration among occupational movers and stayers further helps inform unemployed workers' offer arrival probabilities. The job displacement parameters $\delta_{0,\epsilon}, \delta_{0,z}, \delta_{1,\epsilon}, \delta_{1,z}$, are informed by the average EU probability and its cyclical ratio, where $\delta_{0,z}$ and $\delta_{1,z}$ are also informed by the average and cyclical ratio of the occupational change EUE probability. These parameters are further informed by the observed frequency of very large earnings losses in the cross-section and over the business cycle. In particular, the 2.5th, 5th and 10th percentiles of the earnings growth quantile function and the U-shape characterizing its cyclical variation, as shown in Figures 5a and 5b, provide extra information to recover these parameters. Since γ determines the extent of negative earnings changes among EE movers who did not switch occupations, we target the lower tails of the earnings growth quantile functions among those EE movers who did not switch occupations as shown in Figures 4c and 4d.

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 a higher (lower) value of α_1^i leads to workers searching in a more (less) directed way across occupations, 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 occupation change probabilities conditional on workers' employment transition.

Overall, the model is able to reproduce very well the moments characterizing the employer mobility and gross and net occupational mobility patterns.[32] The model also does well in matching the untargeted unemployment duration distribution.[33] Table 5 reports the estimated parameter values associated with these transition probabilities. These parameters imply that workers in the model are more likely to make voluntary EE transitions than involuntary ones and more likely to make involuntary EU transitions than voluntary ones. We define voluntary transitions as those in which the worker decided to move, while involuntary transitions are those that are imposed on the worker.

Among those workers who made an EE transition with an occupation switch in a given month, we find that 84.7% of them decided to make this transition; while among those who made an EE transition without an occupation switch, 84.5% of them decided to make this transition. In contrast, among EUE transitions that did not involve an occupation switch, 97.4% started with an involuntary separations into unemployment; while among EUE transitions that did involve an occupation switch, 98.1% started with an involuntary separations into unemployment. These involuntary EU separations arise from the exogenous δ_ϵ and δ_z shocks. However, in both cases, the unemployment spell ended with workers deciding to accept job offers. Taken together, these results imply voluntary employer changes that are accompanied by an occupational switch are 5.6 times more likely than involuntary ones. The estimated parameters also imply that the differences between voluntary and involuntary transitions becomes even stronger during expansions. The fraction of voluntary EE transitions that had (did not have) an occupation change increases from 77.5% (76.8%) in recessions to 85.0% (85.3%) in expansions. Among EU separations these (voluntary) shares decrease from 4.3% (3.7%) in recessions to 0.4% (1.91%) in expansions. These patterns are broadly inline with the ones we document in Section 2.[34]

The estimated parameters further 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 more directed. Table 5 shows that employed workers are able to direct their search much more than unemployed workers as α_1^E is nearly twice as large as α_1^U. The parameters $\alpha_{\tilde{o}}$ imply it is easiest to access routine cognitive occupations and hardest to access non-routine cognitive ones, as suggested

[32] The estimated value of $\lambda_{1,U}$ implies that the offer arrival probability of unemployed workers in the same occupation decreases in expansions. Note, however, the aggregate offer arrival probability, including the one when changing occupations, increases in expansions making the UE transition probability procyclical. This differences could be capturing (in reduce form) that meeting jobs outside a worker's current occupation crowded out meeting jobs in the worker's current occupation.

[33] Using SIPP wave-level intervals, the model's shares of completed spells whose duration fall in each interval are 45%, 26%, 14%, while in the data these values are 57%, 23%, 10%. The remaining share of spells have a duration greater than 12 months. The unemployment duration hazard generated by the model is essentially flat and close to the average UE probability. Given that in our sample we minimize the presence of temporary layoffs, the empirical unemployment hazard also becomes much flatter, but still downward sloping (see Carrillo-Tudela and Visschers, 2023), and closer to the one implied by our model.

[34] The main caveat is that in the data we observe more voluntary than involuntary separations into unemployment. However, the model's relationship between voluntary and involuntary transitions is not fully comparable with the data for two main reasons. In the data, the definition is somewhat ad-hoc and depends on whether the worker responded truthfully. Further, there is a significant number of missing responses in the SIPP which potentially introduces biases to the our empirical results.

by the relative contributions of each of the task-based occupation to net flows.

Productivity and payments The persistent parameter of the z-productivity process, ρ_z, determine the frequency of redraws from $F(.; A_I)$. Redraws that lead to low values of z make it more likely for workers to decide to switch occupations, while high values make it more likely for them not to switch. This parameter is therefore informed by the probabilities of an occupation switch reported in Table 4. Since the frequency of redraws also determines the shape of the earnings growth distribution of occupation movers, ρ_z is further informed by the earnings growth distributions of EE occupation switchers in Figures 4a and 4b and the dispersion (as measured by the 90-10 percentile ratio) and skewness of the earnings growth distribution of those who switched occupations within their employers, reported in Table 4. Table 5 shows that z-redraws are estimated to occur with a 7% probability every month, which is the same as the probability of redrawing ϵ among those employed workers who decided to search in their occupation.[35] The earnings growth distribution of employer and occupation stayers together with the net mobility moments also inform ρ_p and σ_p. We find that occupation-wide productivity shocks are more persistent and volatile than aggregate productivity shocks.

The stickiness parameter η in the earnings process described in equation (7) is informed by the dispersion and skewness of the earnings growth distribution among employer and occupation stayers as well as by the cross-sectional and cyclical change of earnings growth quantile function, particularly the percentiles around the median, where employer and occupation stayers are typically located. Since we estimate a small value for $\eta = 0.0930$ and y_t is persistent, the earnings dynamics in equation (7) appears well approximated by our extended bargaining protocol, where renegotiations do not necessarily occur every period.

4.4 The worker-occupation and worker-firm match distributions

We now present the identification arguments for the match quality distributions $F(z; A)$ and $\Gamma(\epsilon; A)$. These distributions are the main drivers of the occupation and employer ladders and play a crucial role in our analysis. Conditional on the observed transition probabilities, $F(z; A)$ and $\Gamma(\epsilon; A)$ are shaped by the observed cross-sectional earnings growth quantile function, the difference between the expansion and recession earnings growth quantile functions and the expansion and recession earnings growth quantile functions of EE occupation switchers and of EE non-switchers.

Figure 4 shows the fit of the model relative to these four EE earnings growth quantile functions, where we target the 10^{th}, 25^{th}, 50^{th}, 75^{th}, and 90^{th} percentiles. The earnings growth quantile functions of occupation switchers are steeper, indicating more earning growth dispersion than for occupation stayers. Indeed, the dispersion of EE earnings growth of occupation switchers seen here by the difference between 90th and 10th percentiles is 50% more than that of occupation stayers. Figure 5a shows the fit of the model in relation to the overall cross-sectional earnings growth quantile function, while Figure 5b shows the fit relative to the U-shape that describes the difference between expansion and recession earnings growth quantile functions. In these cases, we target the 2.5^{th}, 5^{th}, 10^{th}, 25^{th}, 50^{th}, 75^{th}, 90^{th}, 95^{th}, 97.5^{th} percentiles.

[35] Redraws of ϵ among these workers occur with probability $\lambda_{1,E}$ in expansions and with probability $\lambda_{0,E}$ in recessions. Since recessions occur 20% of the time in our simulations, the average value of λ_E across the business cycle is 0.07.

Figure 4: Earnings growth quantile functions for EE movers - model and data

(a) Occupational switchers in expansion

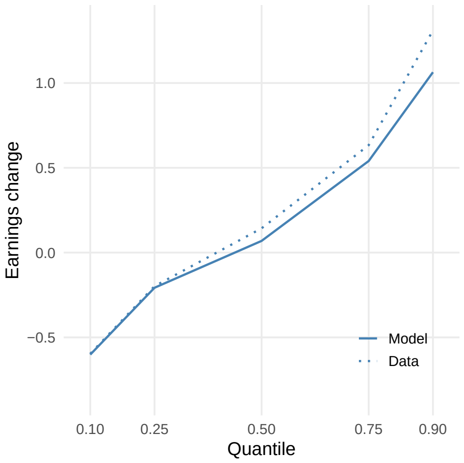

(b) Occupational switchers in recession

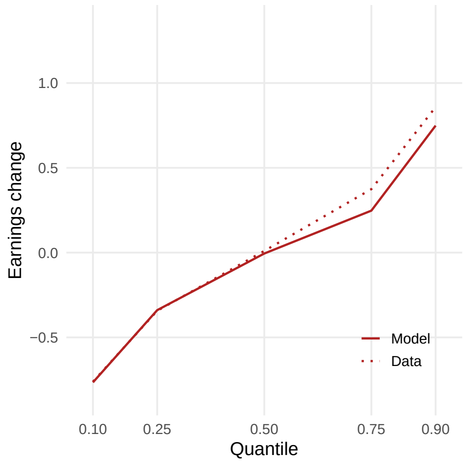

(c) Occupational stayers in expansion

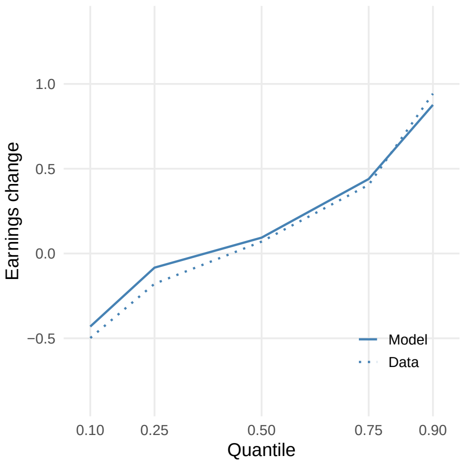

(d) Occupational stayers in recession

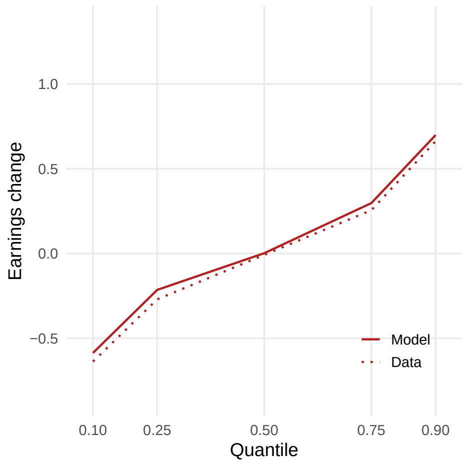

Note: The earnings growth distribution is computed for EE employer movers stayers conditional on whether the worker switched occupation or not. Each graph present these distribution be showing the annual earnings growth value and the corresponding percentile. Our estimation targets quintiles of each of these distributions. In Online Appendix D we present the bootstrapped standard errors for these moments.

Figure 5: Earnings growth quantile functions - model and data

(a) Cross-section quantile function

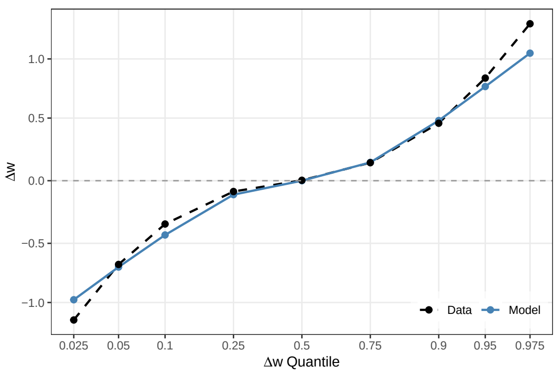

(b) Cyclical changes in the quantile functions

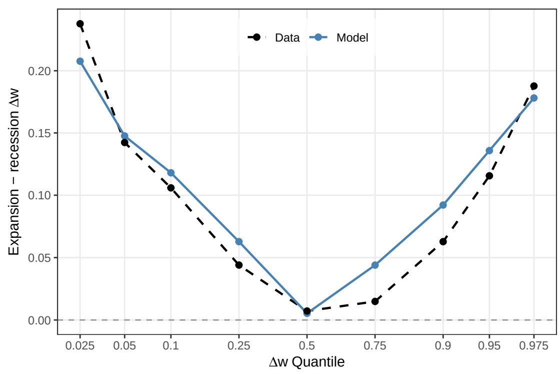

Note: The observed annual earnings growth distribution is constructed for the sample period 1990-2013. The left panel shows the cross-sectional earnings growth distribution (cdf) by plotting the associated earnings change to each selected percentile of the distribution. The percentiles used are 2.5^{th},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

In Online Appendix A we provide a formal proof that shows how to identify a *wage* offer distribution based on the distribution of *wage changes*. We do this using the Burdett-Mortensen framework extended with involuntary transition (godfather) shocks. This result is important as it shows how we can link the empirical earnings growth distribution of employer movers tightly to the underlying offer distribution, $\Gamma(.)$. We establish that the EE wage growth distribution is sufficient to isolate two candidate wage offer distributions, with the true one among them. The "incorrect" candidate can then generically be ruled out by comparing the mean wages (or higher wage moments) of the two candidate distributions, completing the identification of the wage offer distribution.

In Online Appendix A and in Figure 8a (below) we show that in the one job ladder setting without the occupation dimension, the identified $\Gamma(\epsilon; A)$ is *not able* to generate a good fit to the aforementioned U-shaped pattern, even if targeted.[36] The intuition for this negative result is that with only the employer ladder and matching the observed transition probabilities, the model exhibits a too-close link between large earnings losses and large earnings gains. Relative to expansions, the forces that help match an increased prevalence of large EE earnings losses in recessions –while the middle of the earnings growth distribution remains unchanged– also imply an increased prevalence of large earnings gains. Online Appendix A shows that this force is strong enough to be present even as we allow λ_E, the godfather shocks γ *and* the offer distribution to vary freely over the cycle.

We overcome this limitation of the employer ladder model by introducing a second job ladder that slows down more in recessions and generates proportionally more large earnings changes than the employer ladder. The evidence documented in Section 2.5 motivates the importance of an occupation ladder as the missing component. The $F(z; A)$ distribution then performs this function and allows the model to generate the observed cyclical change in the earnings growth distribution.

We use the same logic as with $\Gamma(\epsilon; A)$ to recover the distributions $F(z; A)$, and target the earnings growth quantile functions of EE occupation *switchers* in expansions and recessions, and the U-shape that characterizes the cyclical change in the earnings growth distribution. With two job ladders, the increasing share of occupation switchers observed as one moves towards either tail of the earnings growth distribution (see Section 2.2 and Online Appendix C), drives how much lower and higher the quantile function of EE occupation switchers extends and provides information about the range of $F(.)$ relative to $\Gamma(.)$.[37] Figure 5b shows the model is now able to replicate very well the U-shaped pattern. In Online Appendix D we also show that it reproduces the aforementioned relationship between the share of occupation switcher and the earnings growth quantiles. As in the data, occupation mobility contributes disproportionately to the largest positive and negative earnings changes.

[36]We show this by applying the same non-parametric estimation as describe in Section 4.2, to give the job ladder model enough flexibility to match the data.

[37]If the $F(.)$ distribution were to have a very narrow range, earnings growth of occupation switchers would nearly coincide with those of occupation stayers. As we increase the range of $F(.)$ keeping sufficient mass in its tails, large earnings changes will more often entail a double change of z and ϵ. This logic implies that the importance of occupation switchers observed at the the tails of the earnings growth distribution in Figure 1 (Section 2.2), and through the steepness of the quantiles functions of EE occupation switchers in Figure 4, generates a relatively wide $F(.)$ distribution in our estimation, as shown in Figure 6 by comparing the x-axis of the two panels.

Figure 6: The estimated Γ and F densities

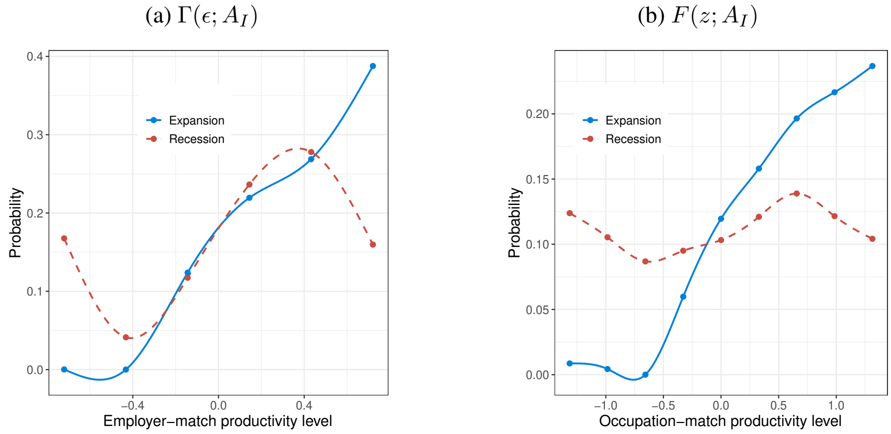

4.5 Two job ladder dynamics

Our model is able to reproduce all the targeted moments shown in Table 4 and Figures 4 and 5 through a novel mechanism. Workers climb and fall along an occupation and an employer ladder that shift between expansion and recessions. The larger dispersion of the occupation switchers earnings growth distribution (see Section 2.2 and Figure 4), maps into a large dispersion of the worker-occupation match quality distribution $F(z; A)$, larger than the dispersion of the firm-worker match distribution $\Gamma(\epsilon; A)$ as shown in Figure 6. At the same time, the estimated parameters and distributions imply that the speed at which workers move along the occupation ladder is slower than along the employer ladder. These two features create a dispersion-mobility differential between the occupation and employed dimensions of the job ladder that enables the model to generate the earnings dynamics we discuss in the next section.

In expansions, the occupation ladder becomes an important source of mobility and large earnings gains. The larger dispersion of $F(.)$ relative to $\Gamma(.)$, in particular its right tail, and the higher values of $\lambda^c_{1,U}, \lambda^c_{1,E}$ contribute to make workers more eager to switch occupations and try to draw a higher new z-productivity. In recessions, the two ladders work very differently. The lower values of $\lambda^c_{0,U}, \lambda^c_{0,E}$ make other occupations less accessible, while the distribution for new worker-occupation match qualities worsens. This implies that drawing a low z becomes more likely, and more difficult to recover from. As a consequence workers will more often refrain from searching across occupations. Since this in itself limits the magnitude of earnings changes that can arise from voluntary occupation switching, the estimated model also generates large earnings losses through an increased δ_z which pulls high paid workers into a forced occupation switch through unemployment.

Even though mobility along the job ladders slows down in recessions, the estimated model predicts that EUE occupation switchers separate from their re-employment jobs faster in recessions than in expansions. This occurs as the lower quality of the re-employment jobs makes these forced occupation movers more likely to quit to another employer and start climbing back the occupation ladder. This result is consistent with the empirical evidence we present in Section 2.5.[38]

[38] Regressing the spell duration of the re-employment job on the same specification as described in Section 2.5 on simulated data delivers an estimated coefficient of -0.022 on the indicator for occupation change, and a coefficient of -0.052 for the interaction between the latter and starting the job in a recession. These values a lower than in the SIPP but

Figure 7: Multi-year persistence of earnings changes - model and data

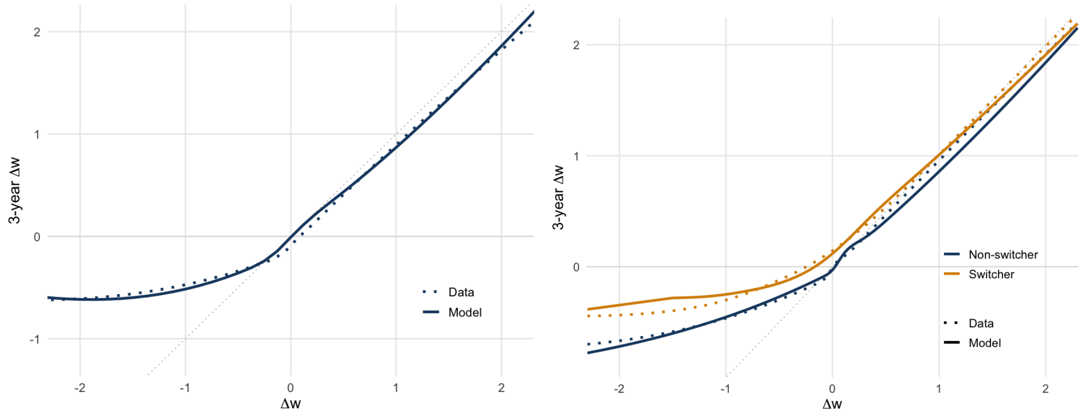

(a) All workers

(b) Occupation switchers vs non-switchers

Note: Model and data comparison on the relationship between workers' annual earnings changes and these workers' earnings changes at a 3-year horizon. The left panel consider all workers, while the right panel conditions on occupations switchers and non-switchers, using the task-based occupational categories.

The cyclical shifts of $F(.)$ and $\Gamma(.)$ also have a direct implication on the dispersion of earnings levels the model generates. In order to investigate whether the dispersion of the match quality distributions is reasonable, we compare the cross-sectional distribution of earnings *levels* of the model with the SIPP data. We find that a close untargeted fit, with the variance in the data and model equal to 0.665 and 0.603, and the 90/10 ratios equal to 6.9 and 7.2, respectively.

4.6 Persistence of earnings changes and occupation mobility

The above discussion shows that our estimated model is able to reproduce the cross-sectional and cyclical patterns of the aggregate earnings growth distribution. We now investigate the untargeted multi-year persistence of earnings changes and show it naturally arises from our two job ladder model. Section 2.6 documents an asymmetric pattern, where at a three-year horizon, large one-year earnings gains are highly persistent. Large one-year losses, by contrast, recover asymmetrically and convexly, with occupation switchers losses recovering faster than non-switchers and a substantial fraction of the largest losses remaining unrecovered at the three-year horizon.

Figure 7 shows the model reproduces very closely this asymmetric pattern. In the model, large earnings gains can be realised by climbing either the employer and/or occupation ladder. Lower-earning workers realize large earnings gains more often because they face a higher acceptance probability. Once a worker has realised a large gain, they are less likely to accept further moves. The largest negative earnings changes are concentrated among long-tenured higher-earning workers who fall off these ladders through an involuntary δ_ϵ and δ_z shocks. Their recovery proceeds through subsequent EUE and EE climbs at a slower rate than the fall itself, because climbing requires sequential meetings and acceptances while falling is instantaneous. For occupation switchers this recovery happens through climbing both of the ladders, explaining why the model generates more recovery for this group than for occupation stayers as in the data. Thus, the asymmetric recovery is the joint product of the climb-and-fall dynamics on the occupation and employer ladders and the persistence of the

deliver the same conclusion.

productivity processes.

This nonlinear pattern appears to capture, at a multi-year horizon, the state-dependent persistence documented by Arellano et al. (2017). They show that the persistence of earnings depends jointly on where workers sit in the earnings distribution and on the direction and size of the shock they receive. The high-persistence situations correspond in our model to workers whose position on the ladders is consistent with the shock they receive: high earners receiving positive shocks remain high; low earners receiving negative shocks remain low, or drop into unemployment and return near the bottom. In both cases a worker's rank in the earnings distribution does not change much. The low-persistence situations correspond to large movements along the ladders: high earners receiving negative shocks fall off via EUE or EE down, while low earners receiving positive shocks climb up. In both cases the new earnings position depends largely on the realization of the shock rather than on the starting position.The same asymmetry also shapes the unconditional distribution of earnings changes, making it negatively skewed and leptokurtic at a one-year horizon as implied by the quantile function depicted in Figure 5a and at a two-year horizon as documented by Guvenen et al. (2021).[39]

This state-dependent structure has a direct implication when estimating the well known Restricted Income Profiles (RIP) process specification on model simulated data. In the latter the persistence parameter will be a weighted average of the local persistence across these different situations, with the weights determined by how common each situation is in the population. To verify this, we estimate a RIP specification and obtain a persistence value of 0.907. This value is somewhat below the ones typically estimated in the PSID data; for example, Krueger et al. (2016) report 0.9695 and Guvenen (2009) reports 0.988. Our model RIP process further generates a very similar unconditional dispersion of the persistent component to that implied by these studies, with the stationary variance in the model equal to 0.658, while the above two studies imply 0.639 and 0.629, respectively.

4.7 Discussion

The above analysis shows that the two job ladder model is overall successful in replicating a wide range of targeted and untargeted patterns. In terms of untargeted patterns, we highlight the ability of the model in generating an empirically reasonable (i) persistence and unconditional dispersion of a standard income process, (ii) the negative skewness and excess kurtosis of longer-horizon earnings changes, (iii) dispersion in earnings levels, (iv) unemployment duration distribution, and (v) faster separate rates from re-employment jobs among occupation switchers in recessions. Although not shown above, in the estimated model, as in the empirical patterns reported in Section 2.3, workers with low occupation tenure contribute more than high tenured workers to the cyclical skewness of the overall earnings growth distribution. However, the latter group exhibits much larger left-skewness than short-tenured workers as they generate the largest earnings losses upon displacement into unem-

[39] At a two-year horizon, the model reproduces the negative skewness of earnings changes with a moment skewness of -0.26, against -0.26 in the PSID documented by Guvenen et al. (2021), see their Table II. By quantiles the change distribution is close to symmetric in both the model and data, with the Kelley skewness being -0.0001 in the model and -0.02 in the data, making the negative skewness a tail feature rather than a quantile one. The model likewise matches robust kurtosis (Crow-Siddiqui 6.47 vs 6.83), while moment kurtosis is lower (6.66 vs. 12.26): the model captures the bulk of the excess kurtosis but generates somewhat thinner extreme tails.

ployment, particularly in recessions, in line with the data.[40]

5 Decomposing Cyclical Earnings Growth

5.1 The importance of occupation mobility

We now evaluate further the role of occupation mobility in determining the cyclical skewness of the earnings growth distribution. We do this by first considering a version without occupations and hence eliminating the occupation ladder from our baseline model. This version also allows us to show that the main implications drawn from the identification arguments of Γ hold in a more general environment. We estimate this one job ladder model using the same moments as in Section 4, except for those pertaining to occupational mobility. We find that it replicates the targeted average EE, EU and UE transition probabilities and their expansion/recession ratios, the cross-sectional earnings growth quantile function, and the earnings growth quantile functions of EE movers in both expansions and recessions.[41] We then return to the two job ladder model, adding back the occupation dimension and investigating the role of workers' occupation mobility decisions in shaping cyclical changes in the earnings growth distribution.

No occupation ladder Figure 8a plots the cyclical change in the aggregate earnings growth distribution generated by this model against its empirical counterpart. It shows that the one job ladder model produces only a slightly larger left tail in recessions and essentially no cyclical difference in the large earnings gains. It therefore significantly under-predicts the U-shaped difference between the earnings losses and gains of expansions and recessions, complementing the discussion in Section 4.4. This under-prediction reflects a tension the model cannot resolve. To match the larger recession losses observed in the data, the model would need either (i) a steeper ϵ-ladder, so that separating workers fall from higher up, or (ii) counterfactually long unemployment durations, so that the unemployed accumulate larger losses before re-employment. Both margins, however, are already disciplined by moments the model matches: a steeper ϵ ladder would overshoot the earnings changes of EE movers, while longer durations would break the cyclicality of UE transitions, which it also reproduces. The moments that pin down the employer ladder thus leave it with too little cyclically-adjustable risk. Absent an occupation ladder, the model has no separate channel through which downside and upside risk can vary over the cycle. The two job ladder model resolves this tension because occupation mobility supplies exactly such a channel.

Occupation mobility decisions We now show that workers' occupation mobility decisions create the key economic force that allows the model to endogenously generate the observed cyclicality of the earnings growth distribution (as shown in Figure 5b). Once again, consider the baseline model with an occupation and an employer ladder, but now we do not let workers adjust their occupation mobility decisions over the business cycle. In particular, we hold these decisions constant at their *recession*

[40]Thus, the model can capture to an important degree a latent feature differentiating worker types as described in Gregory et al. (2025).

[41]In this version of the model 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 $\Gamma(\epsilon; A)$. When employed, they draw a new ϵ with probability $\lambda_E(A_I)$ and are forced to accept it with probability γ.

Figure 8: The role of occupation for the cyclicality of the aggregate earnings growth distribution

(a) One job ladder - No occupations (b) Acyclical occupation mobility decisions

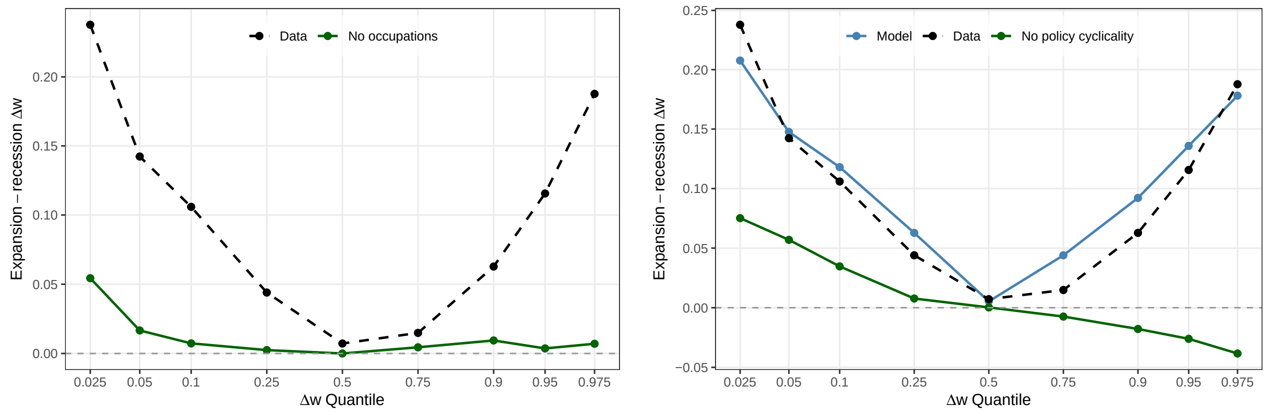

Note: The left panel shows the difference between the expansion and recessions quantile functions of the earnings growth distributions under no occupations. The right panel shows the same difference in the baseline model when occ. switching decisions do not change with the business cycle.

levels, while all exogenous earnings risk and productivities adjust as in the baseline model. That is, we do not allow workers to adjust their occupation mobility decisions as the transitions probabilities, $F(z; A)$ or $\Gamma(\epsilon; A)$ change from recessions to expansions.

Figure 8b shows the difference between the expansion and recession earnings growth quantile functions implied by this counterfactual. If cyclically changing mobility decisions did not matter, we should find no meaningful change in the ability of the baseline model to generate the U-shape that characterizes the cyclicality of the earnings growth distribution. Instead, we find that this ability disappears. By not allowing workers to reoptimize their occupation mobility decisions over the business cycle in lieu of cyclically changing employment risk, the model now generates a earnings growth distribution characterized by countercyclical variance, with recessions exhibiting larger earnings losses and larger earnings gains than expansions.

Summary Taken together, the above analysis shows that we not only need to have an occupation ladder in addition to the employer ladder to explain the cyclicality of earnings changes, but also that workers' occupation mobility decisions are crucial and hence the need of a structural model that encompasses both aspects.

5.2 The contribution of different employment risks

We now investigate the contributions of the different exogenous employment risk we introduced in Section 3.1 in shaping the cyclicality in the earning growth distribution. Recall that these exogenous risks are the probability of job loss and forced employer mobility $\delta_z(A)$, $\delta_\epsilon(A)$, γ; the job finding probabilities and the match quality distributions, $\lambda_u(A)$, $\lambda_e(A)$, $\lambda_u^c(A)$, $\lambda_e^c(A)$; and the match quality distributions $F(z; A)$ and $\Gamma(\epsilon; A)$. The aim is to understand to what extent cyclically changing job loss and job finding probabilities and match quality distributions determine the difference between the expansion and recession earnings growth quantile functions. Crucially, in all of the following counterfactual exercises all workers are free to re-optimize their job acceptance and occupation mobility decisions rules between expansion and recessions periods.

Job loss and job finding To evaluate the importance of the job loss and job finding probabilities, we perform a counterfactual exercise in which all these probabilities, $\delta_z(A)$, $\delta_\epsilon(A)$, $\lambda_u(A)$, $\lambda_e(A)$, $\lambda_u^c(A)$, and $\lambda_e^c(A)$, are set to their recession levels; while $F(z;A)$ and $\Gamma(\epsilon;A)$ are allowed to change between expansions and recessions periods as in the baseline model. The importance of cyclically changing job loss and job finding probabilities can be evaluated by the worsening in the ability of the model to generate the observed difference between the expansion and recession earnings growth quantile functions. Figure 9a shows the result of this exercise and makes clear that these probabilities are important but only at the lowest quantiles of the earnings growth distribution.

Figure 9b investigates this result further by differentiating between the job loss and the job finding probabilities. We find that holding the job finding probabilities constant does not explain the cyclical behavior of earnings losses, but can explain some of the cyclical behavior of earnings gains, pulling down the right-arm of the U-shape. In contrast, holding the job loss probabilities constant leads to a significant drop in the difference between expansion and recession earnings losses and hardly affect the difference between expansion and recession earnings gains. Although both job loss probabilities pull down the left-arm of the U-shape pattern, the stronger effect arises from cyclical changes in the probability of falling from the occupation ladder due to δ_z. Its effect on the larger earnings losses is about twice as large as that of δ_ϵ. The large earnings losses that δ_z generate mainly arise from high-paid, high-tenured workers who fall from the occupation ladder and are forced to reallocate. Since δ_z increases in recessions, these losses become more frequent, creating a larger difference between the expansions and recession large earnings losses. However, these figures also suggest that cyclical changes in the match quality distributions, $F(z;A)$ and $\Gamma(\epsilon;A)$, are more important sources of earnings risk.

Match quality distributions Figure 9c shows the result of evaluating the relative importance of the employer ladder through cyclically changing in $\Gamma(\epsilon;A)$ in driving the previous result. For this exercise we allow $F(z;A)$ and job loss and job finding probabilities to adjust with the business cycle as in the benchmark model. However, we hold $\Gamma(\epsilon;A)$ at its recession level and re-compute the difference between the expansion and recession earnings growth quantiles. It is clear that a cyclically changes in the employer ladder does not meaningfully alter the model's ability to replicate the cyclicality of the earnings growth distribution.

Figure 9d instead shows that the main source of cyclical earnings risk is the occupation match quality distribution, $F(z;A)$. Holding this distribution at its recession levels effectively eliminates any difference between expansion and recession large earnings and significantly reduces the difference between the larger earnings gain in expansion and recession, pulling down (but not collapsing) the right arm of the U-shape pattern.

Summary The above results thus show that cyclical changes of the occupation ladder through $F(z;A)$ and $\delta_z(A)$ are the main sources of cyclical employment risk. This is in line with the results drawn from the empirical decomposition presented in Section 2.5. However, our model shows these sources of employment risk on their own cannot explain the cyclical change in the earnings growth distribution. To fully explain the cyclical dynamics we need to take into account how work-

Figure 9: The importance of cyclical employment risk on the earnings growth distribution

(a) All transition probabilities (b) Job loss and job finding

(c) $\Gamma(\epsilon, A)$ (d) $F(z, A)$

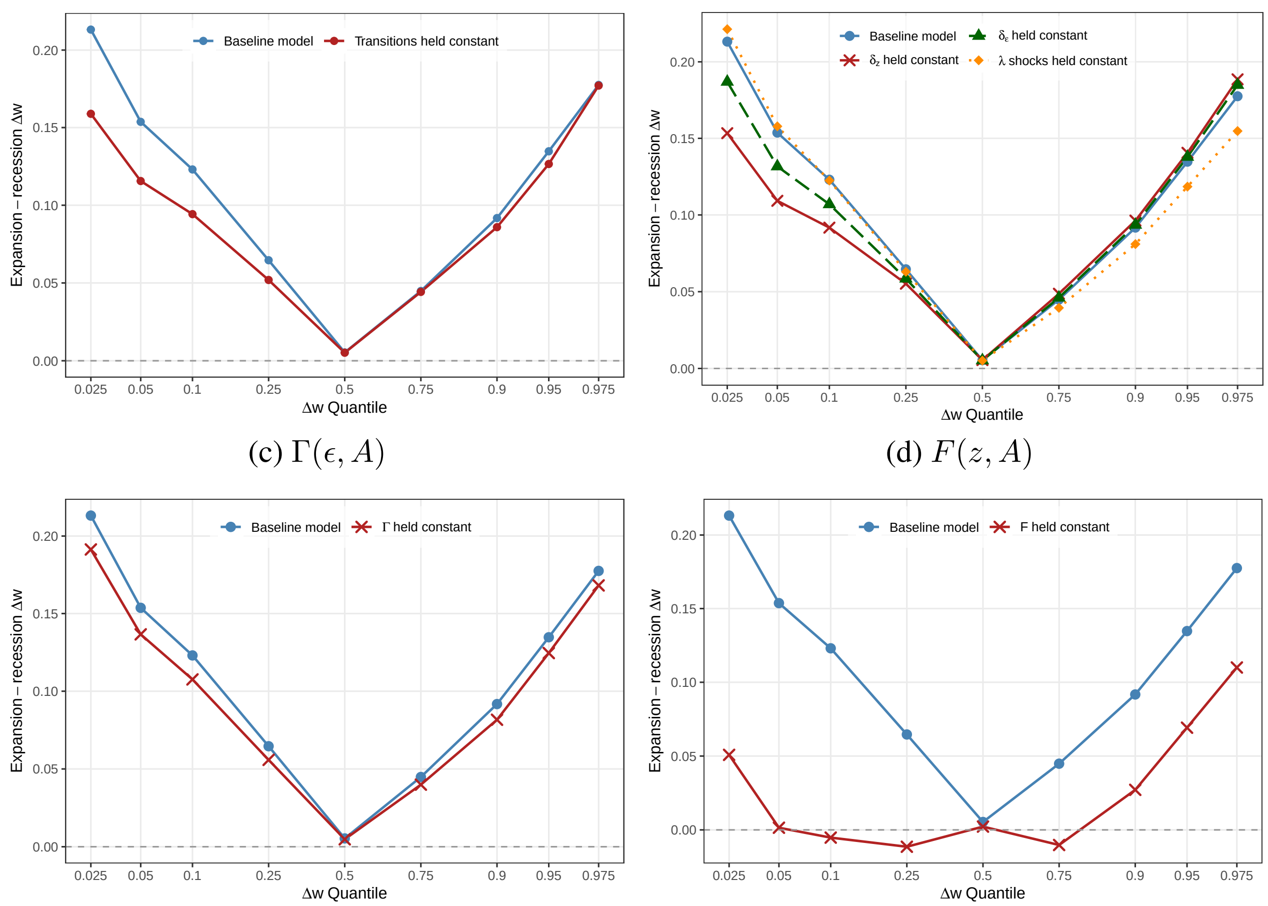

Note: Each panel shows the difference between the expansion and recessions earnings growth distributions. In the top left panel we show the impact of holding all the mobility shocks $\delta_z(A_I)$, $\delta_\epsilon(A_I)$, $\lambda_i(A_I)$ and $\lambda_i^c(A_I)$ constant at their recession levels. In the top right panel we hold constant $\delta_z(A_I)$, $\delta_\epsilon(A_I)$ and the job arrival probabilities at their recession levels, one counterfactual at a time. In the bottom panels we show the impact of holding the $\Gamma(\epsilon; A)$ and $F(z; A)$ constant at their recession levels, respectively.

ers' occupation mobility decisions adjust to these changing risks.

6 Implications

We now turn to investigate the cyclical implications of our model for lifetime earnings and utility measures. We first consider how the worsening of meeting probabilities and match quality distributions during recessions that correspond to the collapse of both job ladders affect lifetime earnings among EE movers. We then consider the (partial equilibrium) effects of targeted occupation policies across the business cycle.

6.1 Sullying Effects of Recessions

Cyclical changes in the match-quality distributions and the mobility shocks affect workers differently. By affecting their ability to climb the occupation and/or the employer job ladders during expansions and recessions, workers located at different ranks of the earnings distribution might suffer differently from the collapse of the jobs ladders during recessions. During recessions the probability of climbing these job ladders decreases (as λ_E and λ_E^c fall), not allowing lower paid workers to climb the job ladder as fast as in expansions. In addition, worse match-quality distributions 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 10: Changes in lifetime earnings after an EE transition

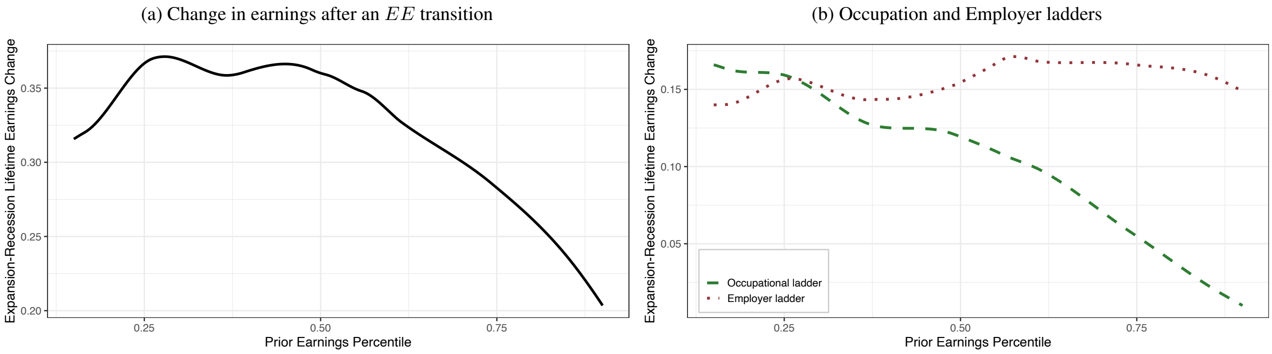

Note: The left panel depicts the expansions and recessions difference of workers' average lifetime earnings prior to an EE transition and their average lifetime earnings after this transition conditional on a worker's earnings immediately prior the transition. The right panel decomposes the separate roles of the occupation and employer ladders.

We show that this sullying effect persists over time through their effect on lifetime earnings. In particular, 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 10a 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 sullying effects of recessions than do high-paid workers.

Figure 10b shows that this occurs mainly through the occupational dimension of the job ladder. It decomposes the change in average lifetime earnings after an EE transition into the employer and the occupation job ladders, by constructing the average lifetime earnings pre- and post- transition based on the ϵ and $p_o + z$ components, respectively. We observe that the differential effect of the employer job ladder across workers' prior earnings percentiles is essentially flat. In contrast, the differential effect is strongly decreasing across prior earnings percentiles when we consider the occupation job ladder. This implies that low-paid workers are disproportionately affected relative to higher paid workers by the sullying effects of recession due to a worse occupation match-quality distribution and less opportunities to climb the occupation ladder. This result motivates our next section which investigates the implications of labor market policies that help workers access better occupations.

6.2 Occupation targeting

In recent years there has been renewed interest in evaluating the effects of policies that help workers, particularly the unemployed, broaden their search and access more occupations (see e.g. Belot et al., 2019). These policies are varied and can involve providing more information about relevant job opportunities in alternative occupations, re-training to narrow occupation-specific skill requirements gaps, among others aspects. Their aim is to provide better job opportunities and improve workers' re-employment chances through reallocation. In our model, we can approximate these policies through the parameters α_{NRC}, α_{RC}, α_{NRM} and α_{RM} as they control the ease with which workers can access a given occupation. Changing accessibility has two key effects on workers' decisions in our model. The first is that workers will re-optimize whether they would like to switch occupation or not, and how they target their search across the different occupations. The second is that workers will re-optimize their job acceptance decisions. For unemployed workers these two margins jointly shape unemployment duration through opposing forces: easier access raises the meeting rate, shortening spells, while the higher reservation value makes workers choosier, lengthening them. These type of policies contrast with untargeted policies that aim to bring individuals back to work quicker. In the model this type of policies can be approximated through the per-period value of unemployment b, which captures the generosity of unemployment benefits. In contrast to the occupation targeting levers, changing b mainly affects workers' job acceptance decisions and for the unemployed the duration of their unemployment spell, and will little impact on how workers divide their search across occupations. We now show that target policies appear more effective than untargeted ones in improving workers' lifetime expected utility, but only when these ease worker access to the most desirable occupations.

To evaluate these different type of policies levers, we aggregate across all states and occupations $(A, \mathcal{P}_O, x_h, z, \epsilon)$ to obtain workers' average lifetime expected utility $\bar{W}^i$ separately for each employment status $i = U, E$. This provides a measure of total welfare for the employed and unemployed that capture per period payments and option values. To evaluate the targeted policy leavers, we calculate the elasticity of $\bar{W}^i$ with respect to a change in the job meeting probability in a given occupation o through a change in $\alpha^i(s_o) = \alpha_0\, e^{\alpha_o(1-\alpha_1^i)}\, s_o^{\alpha_1^i}$ as we vary the values of α_o for a given occupation $o = \{NRC, RC, NRM, RM\}$. Note that when computing these elasticities, $\mathcal{E}_o^i$, we allow workers to re-optimize on the two margins we described above. This implies that relative changes in the direction of search intensities s_o are fully taken into account both in the numerator and denominator. These elasticities are helpful to compare changes in $\bar{W}^i$ between employed and unemployed workers who exhibit different degrees of search directedness, $\alpha_1^E \neq \alpha_1^U$, and job acceptance strategies.

Figures 11a and 11b shows the change in $\bar{W}^U$ and $\bar{W}^E$ relative to their baseline values as we increase $\alpha^i(s_o)$ due to changes in the accessibility to each occupation. The slope of each curve gives the values of $\mathcal{E}_o^i$ and shows that expected lifetime utilities increase as NRC occupations become more accessible, but lifetime utilities decrease when the other three occupations become more accessible. To understand these results, note that NRC occupations not only have the highest time-invariant occupation productivity $\tilde{p}_o$, but also the lowest α_o, making NRC occupations the most productive and desirable but the hardest to access (see Table 5). Raising α_{NRC} then increases the option value of a better job match in NRC occupations and hence raises the value of reallocating away from RC, NRM and RM occupations. Conditional on deciding to reallocate, workers consequently re-direct their search away from these other occupations towards NRC occupations. The better prospects from reallocating towards NRC increase the values of unemployment and employment in the remaining

Figure 11: Changes in lifetime expected utility as occupations become more accessible

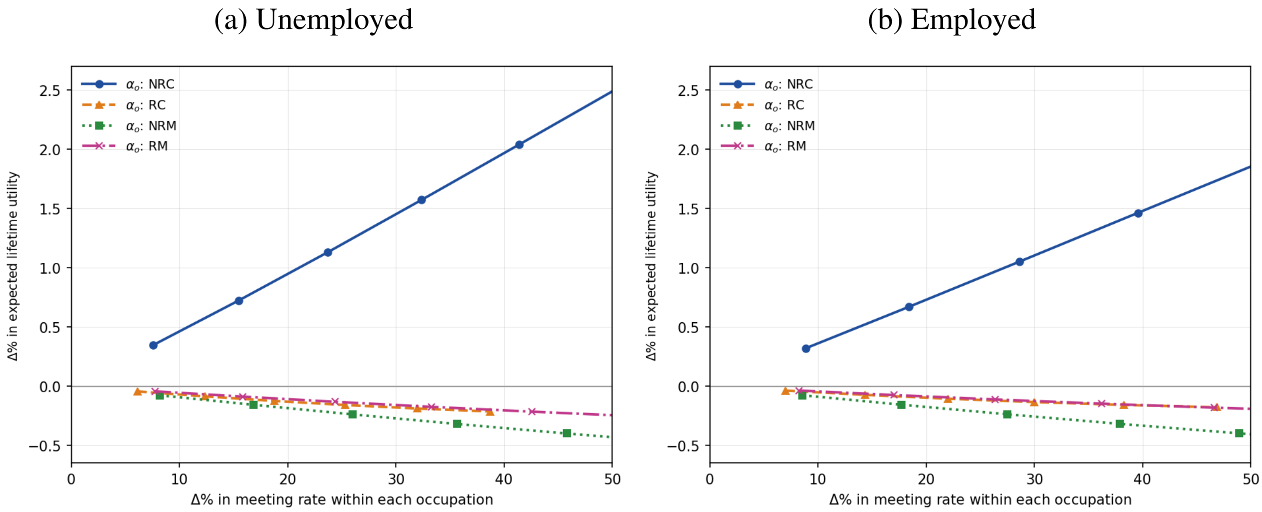

Note: The left panel changes in $\bar{W}^U$ as $\alpha^U(s_o)$ changes due to increases of α_o for a given occupation $o = \{NRC, RC, NRM, RM\}$. The right panel presents the same exercise for employed workers.

occupations. This makes workers pickier in their job acceptance decisions, particularly unemployed workers in RC, NRM and RM who reject more job offers in any of these occupations but achieve better matches. Employed workers also gain, but predominantly by reallocating into the more productive NRC occupations rather than by becoming more selective. Because unemployed workers benefit from both margins while the employed benefit mainly through reallocation, the elasticities are larger for the unemployed, $\mathcal{E}_o^U > \mathcal{E}_o^E$. In contrast, improving accessibility in RC, NRM or RM has the opposite effect: it draws workers' search and reallocation toward these lower-productivity occupations and away from NRC. Workers therefore end up in less desirable occupations and form relative worse matches on average, which lowers $\bar{W}^U$ and $\bar{W}^E$ relative to baseline.

To highlight the effects of improving the accessibility to NRC occupations, consider the increase in α_{NRC} such that the NRC meeting rate increases by 50% as shown the above figures. Among employed workers this leads to an improvement of 4.2% in average job quality, measured as $p_o + z + \epsilon$, across the business cycle, where 75% of this increase arises from better jobs in NRC occupations and the remainder from being more selective in accepting jobs in the other occupations. Among the unemployed the effect is stronger. They experience an 8.1% increase in job quality across the cycle, where 60% of it arises from better jobs in NRC occupations and the remainder from the increases selectivity in the other occupations. We also find that this type of policy is more effective during recessions. The same increase in α_{NRC} during recessions implies an increase in average job quality of 6.9% among employed workers and 9.4% among unemployed workers, where now the selectivity margin plays a much more important role. During expansions job quality increases are 2.8% for all workers and occupation upgrading becomes the dominate force.

An alternative way to improve job quality purely through the selectivity margin is by increasing b, particularly during recessions. For example, doubling the replacement rate from 0.4 to 0.8 during recessions leads to an increase in average job quality of 5.5% and 3.3% among unemployed and employed workers, respectively. Hence this policy lever does not seem as effective as improving access to NRC occupations to increase the productivity of jobs among new hires. Moreover, the increase in selectivity arising from the higher b also leads to a 14% longer average unemployment duration, while the above increase in α_{NRC} leads only to 8% longer duration during recessions.

These exercises illustrate that there is no logical tension between the importance of worker-occupation match effects in the cyclical dynamics of the earnings growth distribution and the potential benefits of policies to facilitate particular occupational transition. Reallocation frictions in the (partially directed) cross-occupation search process stand in the way of searching for better jobs in occupations that pay better on average. We show that policy interventions that reduce these reallocation frictions can be beneficial.

7 Conclusion

In this paper we have studied the structural causes of cyclical earnings risk through the lens of a two job ladder model with endogenous occupation and employer mobility. 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. Our model shows that the main source of risk is the cyclical change in the distribution of worker-occupation match productivities and that endogenous occupation reallocation decisions are crucial to understand the increase in the left-skewness of the earnings growth distribution.

We show these results in an environment with only ex-post worker heterogeneity, arising from firm and occupation match-specific productivities, occupation-wide productivities and human capital accumulation. This approach complements the work of Gregory, et al. (2025) and Ozkan, et al. (2023), who emphasize instead ex-ante worker heterogeneity, particularly in the probabilities of job loss and job finding. An abiding difficulty in this literature 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 employment risk. Nevertheless, more generally, occupational components can meaningfully interact with ex-ante worker differences. Tackling these interactions can 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. We leave this analysis for future research.

References

[1] Altonji, J., A. Smith and I. Vidangos. 2013. "Modeling Earnings Dynamics", *Econometrica*, 81(4): 1395-1454.

[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] Belot, Michele, P. Kircher and P. Muller. 2019. "Providing Advice to Job Seekers at Low Cost: An Experimental Study on On-Line Advice", *Review of Economic Studies*, 86(4): 1411-1447.

[4] Bowlus, A. J. 1995. “Matching Workers and Jobs: Cyclical Fluctuations in Match Quality”, *Journal of Labor Economics*, 13(2): 335-350.

[5] Burdett, K. 1978. “A Theory of Employee Search and Quits”, *American Economic Review*, 68(1): 212-220.

[6] Busch, C. 2020. “Occupational Switching, Tasks, and Wage Dynamics”. Mimeo, Department of Economics, Ludwig Maximilian University (LMU) Munich.

[7] Busch, C., D. Domeij, F. Guvenen, and R. Madera. 2022. “Skewed Idiosyncratic Income Risk over the Business Cycle: Sources and Insurance”, *American Economic Journal: Macroeconomics*, 14(2): 207-242.

[8] Carrillo-Tudela, C. and L. Visschers. 2023. “Unemployment and Endogenous Reallocations Over the Business Cycle”. *Econometrica*, 91(3): 1119-1153.

[9] 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.

[10] DiNardo, J., N. M. Fortin and T. Lemieux. 1996. “Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach”, *Econometrica*, 64(5): 1001-1044.

[11] Elsby, M. W. L. and A. Gottfries. 2022. “Firm Dynamics, On-the-Job Search, and Labor Market Fluctuations”, *Review of Economic Studies*, 89(3): 1370-1419.

[12] Flaaen, A., M. D. Shapiro and I. Sorkin. 2019. “Reconsidering the Consequences of Worker Displacements: Firm versus Worker Perspective”, *American Economic Journal: Macroeconomics*, 11(2): 193-227.

[13] Gottfries, A. 2025. “Bargaining with Renegotiation in Models with On-the-Job Search”, *Review of Economic Dynamics*, forthcoming.

[14] Gregory, V., G. Menzio and D. Wiczer. 2025. “The Alpha Beta Gamma of the Labor Market”, *Journal of Monetary Economics*, 150.

[15] Guvenen, F. 2009. “An Empirical Investigation of Labor Income Processes”, *Review of Economic Dynamics*, 12(1): 58-79.

[16] 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.

[17] Guvenen, F., S. Ozkan, and J. Song. 2014. “The Nature of Countercyclical Income Risk,” *Journal of Political Economy*, 122(3): 621-660.

[18] Hall, R. E. and P. R. Milgrom. 2008. “The Limited Influence of Unemployment on the Wage Bargain”, *American Economic Review*, 98(4): 1653-1674.

[19] Halvorsen, E., H. A. Holter, S. Ozkan, and K. Storesletten. 2024. “Dissecting Idiosyncratic Earnings Risk,” *Journal of the European Economic Association*, 22(2): 617-668.

[20] Harmenberg, K. 2021. “The Labor-Market Origins of Cyclical Skewness”. Mimeo, BI Norwegian Business School, Norway.

[21] Huckfeldt, C. 2022. “Understanding the Scarring Effects of Recessions”. *American Economic Review*, 112(4): 1273-1310.

[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] Kramer, J. 2022. “The Cyclicality of Earnings Growth along the Distribution - Causes and Consequences”. Mimeo, Institute for International Economic Studies, Stockholm, Sweden.

[27] Krueger, D., K. Mitman and F. Perri. 2016. “Macroeconomics and Household Heterogeneity”, in *Handbook of Macroeconomics*, Vol. 2, 843-921. Elsevier.

[28] Low, H., C. Meghir and L. Pistaferri. 2010. “Wage Risk and Employment Risk over the Life Cycle”, *American Economic Review*, 100(4): 1432-1467.

[29] Machado, J. A. F. and J. Mata. 2005. “Counterfactual Decomposition of Changes in Wage Distributions Using Quantile Regression”, *Journal of Applied Econometrics*, 20(4): 445-465.

[30] Moscarini, G. 2001. “Excess Worker Reallocation”. *Review of Economic Studies*, 68(3): 593-612.

[31] Nagypál, É. 2008. “Worker Reallocation over the Business Cycle: The Importance of Employer-to-Employer Transitions”. Mimeo, Department of Economics, Northwestern University.

[32] Neal, D. 1999. “The Complexity of Job Mobility Among Young Men”. *Journal of Labor Economics*, 17(2): 237-261.

[33] Ozkan, S., J. Song and F. Karahan. 2023. “Anatomy of Lifetime Earnings Inequality: Heterogeneity in Job-Ladder Risk versus Human Capital”, *Journal of Political Economy Macroeconomics*, 1(3): 506-550.

[34] Petrongolo, B. and C. A. Pissarides. 2001. “Looking into the Black Box: A Survey of the Matching Function”, *Journal of Economic Literature*, 39(2): 390-431.

[35] 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.

[36] Shimer, Robert. 2005. “The Cyclical Behavior of Equilibrium Unemployment and Vacancies”. *American Economic Review*, 95(1): 25-49.

[37] 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.

[38] Sullivan, P. 2010. “Empirical evidence on occupation and industry specific human capital”. *Labour Economics*, 17: 567-580.

ONLINE APPENDIX

Appendix A — Identification and Estimation based on Wage Changes

In this appendix, we develop the identification and estimation of a Burdett and Mortensen (1998) model with godfather shocks based on labor market flows and distributions of wage ***changes***.

Identification of the Burdett and Mortensen model with Godfather Shocks

Environment Here, we lay out in brief terms the key equations of the Burdett and Mortensen ("BM") model (1998) of wage posting with on-the-job search augmented with 'godfather' shocks. The latter captures randomly drawn wage offers that workers cannot refuse: they have to move to it. Note that in this appendix we adopt the standard Burdett and Mortensen notation, different from the main text. Our focus in this appendix is on linking the distribution of EE *wage* changes (which is related to, but mathematically different from, year-to-year earnings changes) uniquely to the underlying wage offer distribution.

We define a cumulative worker offer distribution F over a discrete number N of wage levels on an evenly spaced wage grid, $\{w_1, \ldots, w_N\}$, with probability mass function f_n (to which we sometimes loosely refer as wage offer density). We restrict ourselves (without loss of generality) to the case that all these wages are acceptable for workers coming out of unemployment. Given the wage offer distribution, the workers are distributed over these wages in the steady state according to g_n and associated cdf G_n.

With rate λ_1 an employed worker receives an offer randomly drawn from firm wage offer cdf F and can decide to switch employer, which they optimally will do if and only if the wage in the new offer is (weakly) higher. Godfather shocks arrive instead with rate γ and imply that a worker is forced to accept an offer in another firm, again drawn randomly according to firm wage offer cdf F. As a consequence of receiving job offers or godfather shocks, workers change employers and wages. The cumulative *wage change* distribution is denoted by H; it lives on a $2N-1$ evenly spaced grid, between $w_1 - w_N$ and $w_N - w_1$. For simplicity, we can define H over a grid $\{-(N-1), \ldots, 0, 1, \ldots N-1\}$ where $-(N-1)$ corresponds to $w_1 - w_N = -d \cdot (N-1)$ (and d is the distance between wage grid points).

This environment is contained in the model laid out in the main paper, when we shut down the occupational dimension, aggregate and occupation-wide productivity shocks and human capital accumulation. A further nonconsequential difference is that in the main text the arrival rate of standard job offers is phrased as $\lambda_{.,E}(1-\gamma)$, which here is λ_1; while the godfather arrival rate is $\lambda_{.,E}\gamma$ in the main text, and γ here.

Given this structure, we want to uniquely map a set of observables including the wage change distribution into fundamentals including the firm wage offer distribution. Specifically, given an observed job finding rate, separation rate from employment in unemployment, and EE mobility rate, the wage *change* distribution and an additional 'simple' wage moment (such as the mean-min ratio of wages taken out of unemployment), we can generically uncover uniquely $\lambda_0, \lambda_1, \delta, \gamma$ and $f(w)$, as we will show below.

Labor market flows Suppose we have a mass 1 of workers and a mass 1 of firms. In steady state inflows and outflows must balance. Hence, for unemployment we find $\lambda_0 u = \delta(1-u)$ and (substituting this into the flow balance) for firms and employed workers, we have

$$\big(\delta + \gamma + \lambda_1 G_{n-1}\big) f_n \;=\; \big(\delta + \gamma + \lambda_1 (1 - F_n)\big) g_n. \tag{8}$$

We can also derive flow balance for employment in wages at or below w_n,

$$(\delta + \gamma) F_n = \Big(\delta + \gamma + \lambda_1 (1 - F_n)\Big) G_n; \tag{9}$$

multiplying the left-hand side by $(1 - G_n + G_n)$ and bringing the terms involving the second G_n to the left-hand side, yields $(\delta + \gamma) F_n (1 - G_n) = (\delta + \gamma + \lambda_1)(1 - F_n) G_n$. Therefore, the following key distribution relation holds

$$\frac{G_n}{1 - G_n} = \lambda \, \frac{F_n}{1 - F_n}, \text{ with } \lambda \overset{def}{=} \frac{\delta + \gamma}{\delta + \gamma + \lambda_1}. \tag{10}$$

This defines a tight and direct relationship between the firm offer cdf F_n and the worker wage cdf G_n, linked through the ratio of exogenous mobility ($\delta + \gamma$), which includes mobility through unemployment, to overall mobility ($\delta + \gamma + \lambda_1$). The more important exogenous mobility is, the closer the worker distribution is to the wage offer distribution. The fact that G_n is a direct function of only F_n and labor market flow parameters δ, γ and λ_1 will be helpful below.

Identification of arrival rates

It turns out that in this setting the identification problem dichotomizes and we can infer flow rates separately from the firm wage offer distribution. For the flow rates we need to know the aggregate employer-to-employer switching rate, m_{EE}, and the distribution of wage changes associated with it, H. Since we can read off δ directly from the EU-rate, we take it to be known. The following proposition summarizes this with the proof immediately afterwards.

Proposition 1. *Given δ, the observed aggregate EE flow rate m_{EE} and the distribution H of wage changes of EE movers, we can identify γ, λ_1.*

Proof. Define $m(k)$ to be the employer-to-employer flow with wage change of k steps on the wage grid ($k > 0$ captures moving k gridpoints up on the wage grid; $k = 0$ a move to a firm with the same wage; $k < 0$ moving k gridpoints down). Then let $m^+_{EE} = \sum_{k>0} m(k)$ the flow of workers who change employers with a wage gain; $m^0_{EE} = m(0)$ the flow of employer changers without a wage change, and m^-_{EE} the flow of employer changers with a wage loss. Underlying this $m(k) = (\gamma + \lambda_1) \sum_{l=1-k}^{N} g_l f_{l+k}$ for $k > 0$; for $k = 0$, under the assumption that when meeting a firm with the same wage the worker flips a coin, $(\gamma + 0.5\lambda_1) \sum_{l=1}^{N} g_l f_l$ (this proof and the proofs below go through under alternative tie-breaking rules); and $m(k) = \gamma \sum_{l=1}^{N-k} g_{l+k} f_l$ for $k < 0$. Since $\sum_k \sum_l f_l g_k = 1$, the following needs to hold:

$$\frac{m^+_{EE}}{\gamma + \lambda_1} + \frac{m^-_{EE}}{\gamma} + \frac{m^0_{EE}}{\gamma + 0.5\lambda_1} = 1 \tag{F1}$$

The above condition (F1) directly involves the observed flow rates. A second condition relates λ to the ratio of the sum of all wage gains to the sum of all wage losses. In particular, we can rewrite equation (10) and sum over all n, to find $\sum G_n(1-F_n) = \lambda \sum F_n(1-G_n)$. Take $\sum_{k\geq 1} k \sum_i g_i f_{i+k}$ and rewrite it as $\sum_{j>i}(j-i)\, g_i f_j$. Since $(j-i)$ can be written as $\mathbf{1}\{i < l \leq j\}$, this equals $\left(\sum_{i<l} g_i\right)\left(\sum_{j\geq l} f_j\right)$; we therefore have, for the sum of observed individual positive wage gains

$$U_{obs} = \sum_{k=1}^{N-1} k\, m(k) = (\gamma+\lambda_1)\sum_l \Big(\sum_{i<l} g_i\Big)\Big(\sum_{j\geq l} f_j\Big) = (\gamma+\lambda_1)\sum_l G_l(1-F_l). \tag{11}$$

Similarly, for the sum of observed individual wage losses

$$D_{obs} = \sum_{k=1}^{N-1} k\, m(-k) = \gamma \sum_l F_l(1-G_l). \tag{12}$$

Putting (11) and (12) together with (10), we have

$$\frac{U_{\text{obs}}}{D_{\text{obs}}} = \lambda\,\frac{\gamma+\lambda_1}{\gamma} = \frac{(\gamma+\lambda_1)(\delta+\gamma)}{\gamma\,(\delta+\gamma+\lambda_1)}. \tag{F2}$$

Now consider (F1) and (F2) in (γ, λ_1)-space. It is straightforward to establish that (F1) implies a strictly decreasing relation of λ_1 as a function of γ. Intuitively, to get the same EE flows, if γ is larger, than λ_1 needs to be smaller. On the other hand, equation (F2) implies a strictly upward-sloping relationship. The model restricts $U_{obs}/D_{obs} > 1$, as can be seen directly from equation (F2). Using the implicit function theorem, we find that $d\lambda_1/d\gamma$ has the sign of $(\lambda_1(\delta+2\gamma+\lambda_1))/(\gamma(\gamma+\delta)) > 0$. Intuitively, if γ is larger, to get to the same ratio of sum of positive wage changes to negative wage changes, λ_1 needs to be larger as well. Hence, there is only one (γ, λ_1) that satisfies (F1) and (F2). □

Note that in particular this implies that, even if more than one distributions of wage offers were to consistent with the same observed m_{EE} and H, given the same δ, these all imply the same γ, λ_1.

The spectral factorization problem

As a result of the above, we have pinned down $\lambda_1, \delta, \gamma$. (With our assumption that all wages are acceptable, we can also directly observe λ_0, which means that all arrival rates are identified.) We now use this to make progress towards identification of f. From the EE flows with a given wage change, we can now pin down the values of $c(k)$ by the left-most equality relationship in equation (13).

$$c(k) = \frac{m(k)}{\gamma+\lambda_1} = \sum_{i=1}^{N-k} g_i f_{i+k} \qquad \text{if } k \in \{1,\dots,N-1\}$$

$$c(k) = \frac{m(k)}{\gamma+0.5\lambda_1} = \sum_{i=1}^{N} g_i f_{i+k} \qquad \text{if } k = 0$$

$$c(k) = \frac{m(k)}{\gamma} = \sum_{i=1-k}^{N} g_i f_{i+k} \qquad \text{if } k \in \{-(N-1),\dots,-1\} \tag{13}$$

By the structure of the BM model with godfather shocks, the terms $c(k)$ are, as the chosen notation already hint at, cross-correlations with known values (the middle terms involving $m(k), \gamma, \lambda_1$ in (13)). Since g is a function only of f and known parameters, via equation (8) or (10), a mapping we further denote by $g = \mathcal{G}(f)$, we have that every $c(k)$ is a function of f (and parameters) alone. Thus let us define the mapping from f to $c(k)$ as

$$\Phi(f) := \big(c(-(N-1)), \ldots, c(0), \ldots, c(N-1)\big) \in \mathbb{R}^{2N-1}.$$

Our identification problem is to go the other way around: from observed $(c(k))_{k=-(N-1)}^{N-1}$ to f. To abbreviate notation, we refer to the vector of cross-correlations as $\vec{c}$.

Towards this, we associate to f polynomial $\widehat{F}(z) = \sum_{j=1}^{N} f_j z^{j-1}$, of degree $N-1$, and to g $\widehat{G}(z) = \sum_j g_j z^{j-1}$; then $\widehat{F}(1) = \widehat{G}(1) = 1$. Define the *reversal* $\widehat{\widehat{G}}(z) := z^{N-1}\widehat{G}(1/z)$, whose coefficient of z^m is g_{N-m}. Then, we have

$$P(z) := \widehat{F}(z)\,\widehat{\widehat{G}}(z) \;=\; \sum_{k=-(N-1)}^{N-1} c(k)\, z^{N-1+k}, \tag{P}$$

a polynomial of degree $2N-2$ whose coefficients are the wage-change frequencies, as given in equation (13). We have $P(1) = \widehat{F}(1)\widehat{\widehat{G}}(1) = 1$ and the extreme coefficients $c(N-1) = g_1 f_N$ and $c(-(N-1)) = g_N f_1$ are single products (relating intuitively to the only possible combinations of current and new wage that can yield the largest possible wage gain and loss). Recovering an offer distribution consistent with the wage change distribution means factoring P into the degree-$(N-1)$ pieces $\widehat{F}$ and $\widehat{\widehat{G}}$, with $\widehat{F}$ and $\widehat{G}$ tied together through the distribution relation (8).

Wage offer distribution solutions for a given wage change distribution typically come in *pairs*

Consider a solution f to the problem $\Psi(f) = \vec{c}$. Now construct an alternative wage offer distribution, $\tilde{f}$, in the following way: from f construct $g = \mathcal{G}(f)$ as explained above (deriving g from f using the flow balance equations or distributional relation (10)). Then, define $\tilde{f}_i = g_{N+1-i}$, for $i = 1, \ldots, N$. Formally, let us define the *dual* of an offer distribution f by a mapping T

$$Tf \overset{def}{=} \text{reverse}\big(\mathcal{G}(f)\big), \qquad (\text{reverse}\, x)_i \overset{def}{=} x_{N+1-i}. \tag{14}$$

Proposition 2. *Take as given arrival rates, $\delta, \gamma, \lambda_1$. If f is a wage offer that produces a wage change distribution H, then Tf (its 'dual' or 'twin') defined in equation* (14) *also produces H. Hence, given wage change distribution H, f is only partially identified: generically, at best up to a dual pair $\{f, Tf\}$. However, information on wage levels allows us to distinguish between this pair:*
(i) if $f \neq Tf$, generically the mean wage in the population, $\mu(\mathcal{G}(f)) \neq \mu(\mathcal{G}(Tf))$ differs;
(ii) if $f \neq Tf$, there is always a level wage moment that separates the twins.
The set of offer distributions that are self-dual, i.e. $Tf = f$ is of Lesbesgue measure zero.

Importantly, at this stage, we still cannot rule out that there exist f and other $f' \neq Tf$ that produce the same wage change distribution. Instead, the proposition says that if f is a solution, Tf is as well, and helps us to distinguish within this dual pair, using (some) wage level data. Furthermore, on a small set of wage offers f is self-dual, $Tf = f$. Call this set Σ. This set has Lebesgue measure zero for every $N \geq 2$, hence generically, at this stage we have at least one pair of solutions $\{f, Tf\}$ to (P), with $f \neq Tf$. Let us now proceed to the proof.

Proof. Let us start by showing that T is an involution that preserves the data. That is, T maps probability distributions to probability distributions, satisfies $T(Tf) = f$, and $\Phi(Tf) = \Phi(f)$. First, we show that the worker distribution $\tilde{g}$ associated with $\tilde{f} = Tf$, is $\tilde{g}_i = f_{N+1-i}, \text{ for } i = 1, \dots, N$. Consider the distribution relation in (10). Given the specification of $\tilde{f} = Tf$, we have that for the cdf associated with $\tilde{f}$, $\tilde{F}$, it holds that $\tilde{F}_n = 1 - G_{N-n}$, hence ratio $\lambda\tilde{F}_n/(1-\tilde{F}_n) = \lambda(1-G_{N-n})/G_{N-n}$. By the distribution relation applied to f, g, we have that the latter equals $(1-F_{N-n})/F_{N-n}$, and using this in the distribution relation for $\tilde{f}$ and $\tilde{g}$, we find $\tilde{G}_n/(1-\tilde{G}_n) = (1-F_{N-n})/F_{N-n}$ and therefore $\tilde{G}_n = 1 - F_{N-n}$. This implies that $\tilde{g}_n = f_{N+1-n}$. A number of conclusions can now be drawn: (1) $T(Tf) = f$; (2) Phrased in n polynomial terms, cf. equation (P), the wage offer polynomial of Tf is $\widehat{Tf}(z) = \sum_i g_{N+1-i} z^{i-1} = \widehat{\widehat{G}}(z)$, and the reversed-worker polynomial associated with $\mathcal{G}(Tf)$ becomes $\widehat{F}(z)$; so T sends the ordered pair $(\widehat{F}, \widehat{\widehat{G}})$ to $(\widehat{\widehat{G}}, \widehat{F})$, leaving the product P and every $c(k)$ unchanged. Therefore $\Phi(Tf) = \Phi(f) = \vec{c}$ and Tf is another solution to $\Phi^{-1}(\vec{c})$.

We now tackle (ii) before (i), considering the case that there is indeed a 'twin' issue of partial identification, $f \neq Tf$. Let $\tilde{g} := \mathcal{G}(Tf)$ be the twins' implied worker distributions. Then $g \neq \tilde{g}$, and there is an integer $p \in \{1, \dots, N-1\}$ with $\sum_n w_n^p g_n \neq \sum_n w_n^p \tilde{g}_n$, where w_n^p is the n-th power of wage level w_n. Hence the wage-change data together with the level moments of the cross-sectional wage distribution up to order $N-1$ point-identify the offer distribution. By the involution T, $\tilde{g}$ is the reversal of f: indeed $f = T(Tf) = \text{reverse}\big(\mathcal{G}(Tf)\big)$, and reversal is its own inverse. Consider $e := g - \tilde{g}$. First, we have $e \neq 0$. If not, towards contradiction, $e = 0$ would say $\mathcal{G}(f) = \text{reverse}(f)$, i.e. $Tf = \text{reverse}(\mathcal{G}(f)) = f$, contradicting $f \neq Tf$. Second, $\sum_n e_n = 0$ since g and $\tilde{g}$ are both probability vectors, i.e. $\sum_n w_n^0 e_n = 0$. Suppose, for another contradiction, that $\sum_n w_n^p e_n = 0$ also for every $p = 1, \dots, N-1$. Then $Ve = 0$, where $V = \big(w_n^p\big)_{p=0,\dots,N-1;\ n=1,\dots,N}$ is the Vandermonde matrix of the distinct nodes $w_1, \dots, w_N$, which is invertible; so $e = 0$—a contradiction. Hence some $p \in \{1, \dots, N-1\}$ gives unequal worker moments.

Let us now consider (i): for almost every f a single cross-sectional mean resolves the dual pair. Towards showing this and to save on notation, redefine $\mu(x) := \sum_n n\, x_n$ for the mean in grid units, rather than wage levels. We want to show that the offer distribution and its dual, and the cross-sectional wage distribution g and its dual $g' = \mathcal{G}(Tf)$ have means that only coincide at a zero measure set. First, consider that by the reversal property of T, we have $\mu(Tf) = N + 1 - \mu(g)$ and (following the beginning of this proof) $\mu(g') = N + 1 - \mu(f)$. Reordering this, and using summation by parts, we have $\mu(x) = N - \sum_{n<N} X_n$ (where $X_n = \sum i = 1^n x_i$). Therefore,

$$\mu(g) - \mu(g') \ = \ \mu(f) - \mu(Tf) \ = \ \mu(f) + \mu(g) - (N+1) \ = \ (N-1) \ - \ \sum_{n<N} (F_n + G_n) \quad (15)$$

Now, from the distribution balance in equation (10), we know that G_n is a function of F_n, in particular $G_n = \lambda F_n/(1-F_n) + \lambda F)$. If we multiply 15 with $\prod_{n<N}((1-F_n) + \lambda F_n)$, which is strictly bigger than zero, we have a polynomial equation such that it equals 0 if and only if $\mu(g) = \mu(g')$ and $\mu(f) = \mu(Tf)$. This polynomial is not the zero polynomial: as $f_N \to 1$, the value of the RHS of (15) goes to a strictly positive value bounded away from zero (since $\sum_{n<N}(F_n + G_n) \to 0$); while if $f_1 \to 1$ to a strictly negative value bounded away from zero ($\sum_{n<N}(F_n + G_n) \to 2(N-1)$). Then

from Okamoto's lemma (the zero set of a polynomial that is not the zero polynomial has Lebesgue measure zero, Okamoto, 1973)

Finally, one can ask how large the set of offer distributions are that are self-dual, ie. $f = Tf$.Self-duality $f = \tilde{f} = Tf$ implies that $\tilde{F} = F$. By the distribution flow balance of equation (10), we have that $G_n = \lambda F_n/((1 - F_n) + \lambda F_n$. By the the definition of T, it follows that $1 - \tilde{F}_{N-n} = \lambda F_n/((1 - F_n) + \lambda F_n)$. By self-duality, we thus have

$$1 - F_{N-n} = \frac{\lambda F_n}{(1 - F_n) + \lambda F_n}, \; \forall\, n = 1, \ldots, (N-1)/2 \tag{16}$$

Each of these constraints defines a bilinear polynomial on $\mathbb{R}^{N-1}$, hence the solution set Σ is of zero measure inside it. In fact, from standard arguments it follows that the dimension of Σ is weakly smaller than $(N-1)/2$. □

While the above means that typically we only need a simple wage level moment to discard one of the pair $\{f, Tf\}$, at this stage we cannot rule out that there are potentially multiple pairs of wage offer distributions that all lead to the same wage change distribution. This is addressed in the next section.

Generic Global Identification

We now want to count the number solutions $\{f', \ldots\}$ of $\Phi(f') = c$, where given c is the vector of cross-correlations that can result from the BM model, over the space of *complex* numbers, for cases that are not exceptional (in a Lebesgue zero-measure sense). The reason for this is that, different from the reals, the count of solutions of a polynomial system is the *same* number for almost all right-hand sides (this constant is called the *degree* of the map). We will show, below, the degree is *at most two*, by finding one explicit, fully computable data point whose complex solution set has exactly two well-behaved elements. Since this is a statement over complex numbers, this than needs to be translated back to "at most two over $\mathbb{R}$ for almost every offer distribution."

To move the analysis to the complex side, we extend the set of offer 'distributions' f to complex values: let

$$X := \left\{ f \in \mathbb{C}^N : \textstyle\sum_n f_n = 1 \right\} \cong \mathbb{C}^{N-1}. \tag{17}$$

We want to employ a key result from algebraic geometry, but to do so we need to introduce some further concepts and explain why our case at hand and, in particular, our mapping Φ fits these concepts.

- *Variety*: Let $m \geq 1$ and let $p_1, \ldots, p_s \in \mathbb{C}[t_1, \ldots, t_m]$ be finitely many polynomials ($s = 0$ allowed). The (affine) *variety* they define is their common zero set

$$V(p_1, \ldots, p_s) \;:=\; \{\, y \in \mathbb{C}^m \;:\; p_1(y) = 0, \;\ldots,\; p_s(y) = 0 \,\},$$

and a subset $V \subseteq \mathbb{C}^m$ is a variety if and only if it equals $V(p_1, \ldots, p_s)$ for some such finite list.

- *Irreducible Variety*: variety V is *reducible* if it splits as $V = V_1 \cup V_2$ with each V_i a variety strictly smaller than V; *irreducible* means no such split exist.

- *Proper Subvariety* of Y. A *subvariety* of Y is a variety contained in Y; *proper* means not all of Y. On an *irreducible* variety Y, every proper subvariety $B \subsetneq Y$ has $\dim\ B \leq \dim\ Y - 1$.

- *Rational Map*: A map with m output slots is *rational* if each output slot is one rational function, i.e.

a ratio of two polynomials, of the inputs, defined at exactly the input points where no denominator on the list vanishes. Below, the dashed arrow $X \dashrightarrow Y$ signifies that it need not be defined everywhere.
- *Dominant*: A map φ is *dominant onto* Y if its image is dense in Y: not trapped inside any proper subvariety of Y.

The source $X = \{f \in \mathbb{C}^N : \sum_n f_n = 1\} \cong \mathbb{C}^{N-1}$ is an affine space, the simplest irreducible variety. Φ is rational: the worker map $\mathcal{G}$, the correlations $c(k)$, and hence Φ are *rational* in f (ratios of polynomials): their only denominators are the factors $\delta + \gamma + \lambda_1(1 - F_n)$ from (8). The denominators vanish only on the proper subvariety $\{F_n = (\delta + \gamma + \lambda_1)/\lambda_1$ for some $n\}$. Write $X_0 \subset X$ for the (open, dense) complement of the latter, on which Φ is defined and holomorphic.[42] On the real simplex these denominators lie in $[\delta + \gamma, \delta + \gamma + \lambda_1]$, so the simplex sits inside X_0. Note further that the real simplex, being a full-dimensional subset of the real points of X, is Zariski dense in X: a polynomial vanishing on the simplex vanishes on all of X: this can help us go between real and complex statements below. The closure of the $\Phi(X_0)$, $\bar{Y} := \overline{\Phi(X_0)}$, is irreducible (the closure of the image of an irreducible variety is irreducible), and dominance follows directly from its construction: a polynomial vanishing on the image vanishes on the image's closure. Now consider the dimension of $\bar{Y}$, it will be shown formally regarding a particular f_ϵ, the concentrated offer, in a proof below (but does not require any of the results in between). From the reasoning there, it follows $\dim\, \bar{Y} \geq N - 1$. Since images do not gain dimensions and closures do not raise it, we also have $\dim\, \bar{Y} \leq N - 1$, therefore $\dim\, \bar{Y} = N - 1$.

Given this, we can now use the following result from algebraic geometry:

Result 1. *Let $\varphi : X \dashrightarrow Y$ be a dominant rational map of irreducible complex varieties of the same dimension. Then there is an integer $d = \deg \varphi \geq 1$ and a* proper subvariety $B \subsetneq Y$ *such that every $y \in Y \setminus B$ in the image has* exactly d *preimages.*

This is proposition 7.16 of Harris (1992) for rational maps, applied to $\mathbb{C}$ and equal dimension.[43] Given the last sentence of the Result, we also have that the property that complement of a proper subvariety is dense (in the ordinary topology) and it has full Lebesgue measure. Therefore, in our case, generic wage change distributions, admissible in the sense that they are generated by at least one $f \in X$, according to the BM model with godfather shocks, have exactly d preimages in X. Towards uncovering the value of d, we collect two more results from algebraic geometry.

Result 2 (Continuity of roots: Hurwitz's theorem)**.** *If polynomials p_ε of degree d converge coefficient-wise, as $\varepsilon \downarrow 0$, to a polynomial p_0 of degree d with d* simple *roots, then for all small ε the polynomial p_ε has exactly d simple roots, converging to those of p_0.*

Result 3. *If a holomorphic map on an open subset of $\mathbb{C}^m$ has an injective derivative at a point, it is injective on a neighborhood of that point.*

[42] A map into $\mathbb{C}^{2N-1}$ is holomorphic exactly when each coordinate function is; each coordinate of Φ is a ratio of polynomials in f, polynomials are holomorphic, and a quotient of holomorphic functions is holomorphic wherever its denominator is nonzero.

[43] See also Corollary III.10.7 of R. Hartshorne (1977).

Result 2 is Hurwitz' theorem, while result 3 comes from the *holomorphic inverse function theorem* (if F is holomorphic on an open set in $\mathbb{C}^m$ with values in $\mathbb{C}^m$ and $\det DF(z_0) \neq 0$, then F is a biholomorphism of a neighborhood of z_0 onto a neighborhood of $F(z_0)$, cf. P. Griffiths and J. Harris (1994), Ch. 0).

We now want to establish that we can read the degree off a well-chosen $(f_\epsilon, \Phi(f_\epsilon))$. Well-chosen, in this context, means that it is informative about the degree d. Lemma 1 provides sufficient conditions for an y_0 and associated preimages $\{x_1, \ldots, x_m\}$ to informative about the degree.

Lemma 1 (Counting bound)**.** *Let $\varphi : X \dashrightarrow Y$ be dominant between irreducible complex varieties of the same dimension, holomorphic on the open set X_0 where it is defined, and let y_0 be a point of $\varphi(X_0)$ such that:*

(i) the preimage $\varphi^{-1}(y_0) \cap X_0$ is a finite set $\{x_1, \ldots, x_m\}$;

(ii) φ is locally one-to-one at each x_i (e.g. by Result 3, $D\varphi_{x_i}$ injective);

(iii) φ is proper over y_0*: every sequence $x'_n \in X_0$ with $\varphi(x'_n) \to y_0$ has a subsequence converging to a point of X_0.*

Then $\deg \varphi \leq m$.

Proof. Let $d := \deg \varphi$ and let $B \subsetneq \bar{Y}$ be the bad set of Result 1 , so points of $\varphi(X_0) \setminus B$ have exactly d preimages. Such points come arbitrarily close to y_0: the image $\varphi(X_0)$ contains a dense open subset of the irreducible variety $\bar{Y}$, removing the proper subvariety B keeps it dense, and $y_0 \in \bar{Y}$. Choose $y_n \to y_0$ with $y_n \in \varphi(X_0) \setminus B$, and for each n list the d *distinct* preimages $x_n^{(1)}, \ldots, x_n^{(d)}$.

Suppose $d > m$. For each j, $\varphi(x_n^{(j)}) = y_n \to y_0$, so by properness (iii) we may pass to a common subsequence along which every $x_n^{(j)} \to x^{*(j)} \in X_0$; by continuity $\varphi(x^{*(j)}) = y_0$, hence each $x^{*(j)} \in \{x_1, \ldots, x_m\}$. Since $d > m$, two indices $j \neq j'$ share a limit: $x^{*(j)} = x^{*(j')} = x_i$. But then, for large n, $x_n^{(j)} \neq x_n^{(j')}$ are two *distinct* points, both inside an arbitrarily small neighborhood of x_i, mapping to the *same* value y_n—contradicting local injectivity (ii) at x_i. Hence $d \leq m$. □

Local Identification Given the above, in particular (ii) of Lemma 1, it is worthwhile delving a bit deeper into local injectivity, the classical rank condition for local identification. Let us briefly go back to an f that is in the interior of the simplex. The derivative $D\Phi_f$ is the Jacobian of partial derivatives $\partial c(k)/\partial f_n$ on the tangent space $\{\delta f : \sum_n \delta f_n = 0\}$ of dimension $N - 1$.

Lemma 2. *If $f_N > 0$, interior to the simplex and $\widehat{F}$ and $\widehat{\widehat{G}}$ share no common roots, then $D\Phi_f$ is injective.*

Proof. A small change of the offer distribution δf induces, through $g = \mathcal{G}(f)$, a perturbation $\delta\widehat{\widehat{G}}$. and lies in $\ker D\Phi_f$ iff every coefficient of P is unchanged, i.e. $\delta P \equiv 0$. Differentiating $P = \widehat{F}\,\widehat{\widehat{G}}$ by the product rule, $\delta P = \delta\widehat{F} \cdot \widehat{\widehat{G}} + \widehat{F} \cdot \delta\widehat{\widehat{G}}$, so $\delta P = 0$ implies $\delta\widehat{F} \cdot \widehat{\widehat{G}} = -\widehat{F} \cdot \delta\widehat{\widehat{G}}$.

The right side is divisible by $\widehat{F}$, so the left-hand side should be as well. If $\widehat{F}$ divides a product but shares no root with $\widehat{G}$, it must divide the other, $\delta\widehat{F}$. Each root ρ of $\widehat{F}$, with multiplicity m, must appear as a root (with full multiplicity) on $\delta\widehat{F}$. The divisibility of $\delta\widehat{F}$ by $\widehat{F}$ means that $\delta\widehat{F} = \hat{Q}\widehat{F}$ for some polynomial $\hat{Q}$, and degrees add: $\deg \widehat{F} = \deg \hat{Q} + \deg \widehat{F}$. By construction $\deg \delta\widehat{F} \leq N-1$, while $\deg \widehat{F} = N-1$ (given $f_N > 0$). It follows that $\deg Q = 0$ and Q is a constant, c_Q, hence $\delta\widehat{F} = c_Q\widehat{F}$. A small change of the offer distribution δf must remain on the simplex. Hence $\delta\widehat{F}(1) = \sum_n \delta f_n = 0$, while $\widehat{F}(1) = 1$. Hence $c = 0$. $\square$

The result of Lemma 2 gives *local* uniqueness of the solution, over $\mathbb{C}$ (as long as $f_N > 0$) as over $\mathbb{R}$.

The complex preimage is the dual pair.

Proposition 3 (Generic degree two)**.** *For every $N \geq 2$, $\deg_{\mathbb{C}} \Phi = 2$. Equivalently, generic wage-change data are produced by exactly two complex offer distributions, the dual pair.*

We prove $\deg \Phi \leq 2$ by applying Lemma 1 at a wage change distribution whose complex preimage has exactly two points. This point, the 'witness' (since it testifies to the degree) is the *concentrated wage offer distribution*

$$f_\varepsilon = (\varepsilon, \dots, \varepsilon,\ 1-(N-1)\varepsilon), \qquad 0 < \varepsilon \text{ small},$$

with data $y_0 := \Phi(f_\varepsilon)$ and factorization $P = \widehat{F}\,\widehat{\widehat{G}}$. It is perhaps intuitive that the wage offer concentrated at the top can provide insights: for any worker not hired at the top, there will be nearly always be one upward move afterwards: to the top wage. Conversely, all moves down nearly all (non-zero wage change) start from the top wage. Other than the dual that completely flips this (nearly all wage offers at the bottom), it appears hard to come up with a third offer distribution that could give rise to the wage change patterns associated with the concentrated offer distribution. The results below formalize this.

Below, we qualify the size in orders of magnitude of ε, writing $h = \Theta(\varepsilon^r)$ if $c_1\varepsilon^r \leq |h| \leq c_2\varepsilon^r$ for constants $0 < c_1 \leq c_2$ as $\varepsilon \downarrow 0$. With this in hand, we first establish two auxiliary results on the size of the roots. First, we show that the $2N-2$ roots of P are distinct and dichotomize in large and small roots, each set of $N-1$ elements. Then we show that

Lemma 3 (Root sizes)**.** *There is $\varepsilon_0 > 0$ such that for all $0 < \varepsilon < \varepsilon_0$, the $2N-2$ roots of P are all distinct and split into*

- $N-1$ small *roots, the roots of $\widehat{F}$, each of modulus $\Theta(\varepsilon^{1/(N-1)})$; and*
- $N-1$ large *roots, the roots of $\widehat{\widehat{G}}$, each of modulus $\Theta(\varepsilon^{-1/(N-1)})$.*

In particular $\widehat{F}$ and $\widehat{\widehat{G}}$ do not share any common root, and no root of P lies within a fixed distance of $z = 1$.

Proof. Small roots. Since $f_j = \varepsilon$ for $j < N$ and $f_N = 1-(N-1)\varepsilon$,

$$\widehat{F}(z) = \big(1-(N-1)\varepsilon\big)\, z^{N-1} + \varepsilon\big(1 + z + \cdots + z^{N-2}\big),$$

with leading coefficient $1-(N-1)\varepsilon \neq 0$, so $\widehat{F}$ has exactly $N-1$ roots. To locate them, substitute $z = \varepsilon^{1/(N-1)}\zeta$ and divide by ε:

$$\varepsilon^{-1}\,\widehat{F}\big(\varepsilon^{1/(N-1)}\zeta\big) = \big(1-(N-1)\varepsilon\big)\,\zeta^{N-1} + 1 + \sum_{j=1}^{N-2} \varepsilon^{j/(N-1)}\zeta^j \;\xrightarrow[\varepsilon\downarrow 0]{}\; \zeta^{N-1} + 1$$

uniformly on compact ζ-sets, since each intermediate power $\varepsilon^{j/(N-1)} \to 0$. The limit $\zeta^{N-1}+1$ has $N-1$ distinct roots, the $(N-1)$-st roots of -1, all on $|\zeta|=1$. By Result 2, for small ε the rescaled polynomial has exactly $N-1$ simple roots near those, so $\widehat{F}$ has $N-1$ simple roots of modulus $\Theta(\varepsilon^{1/(N-1)})$.

Large roots. By (10), $G_n = \lambda F_n/((1-F_n)+\lambda F_n) = \lambda n\varepsilon + O(\varepsilon^2)$ for $n \le N-1$, so $g_n = \lambda\varepsilon + O(\varepsilon^2)$ for $n \le N-1$ and $g_N = 1 - O(\varepsilon)$; hence

$$\widehat{\widehat{G}}(z) = g_N + \sum_{m=1}^{N-2} g_{N-m}\, z^m + g_1\, z^{N-1},$$

with leading coefficient $g_1 = \lambda\varepsilon + O(\varepsilon^2) \neq 0$, so $\widehat{\widehat{G}}$ has exactly $N-1$ roots. Substitute $z = \varepsilon^{-1/(N-1)}\xi$ and set $H_\varepsilon(\xi) := \widehat{\widehat{G}}\big(\varepsilon^{-1/(N-1)}\xi\big)$. Its constant term is $g_N \to 1$; its top term is $g_1\,\varepsilon^{-1}\xi^{N-1} = \lambda\,\xi^{N-1}\big(1+O(\varepsilon)\big)$; and each intermediate term ($1 \le m \le N-2$) is

$$g_{N-m}\,\varepsilon^{-m/(N-1)}\,\xi^m \;=\; \lambda\,\varepsilon^{1-m/(N-1)}\,\xi^m\big(1+O(\varepsilon)\big) \;\longrightarrow\; 0,$$

since $m < N-1$ makes the exponent $1 - m/(N-1)$ positive. Hence

$$H_\varepsilon(\xi) \underset{\varepsilon\downarrow 0}{\longrightarrow} 1 + \lambda\,\xi^{N-1}$$

uniformly on compact ξ-sets. The limit has $N-1$ distinct roots, of modulus $\lambda^{-1/(N-1)} > 1$; by Result 2, $\widehat{\widehat{G}}$ has $N-1$ distinct roots $z = \varepsilon^{-1/(N-1)}\xi$ of modulus $\Theta(\varepsilon^{-1/(N-1)})$.

For small ε the small family ($\to 0$) and large family ($\to \infty$) are disjoint, so the $2N-2$ roots are distinct; and $P(1) = \widehat{F}(1)\widehat{\widehat{G}}(1) = 1 \neq 0$ with every root converging to 0 or to ∞, so no root is near 1. $\square$

Lemma 4 (Top of the Wage Offer Distribution)**.** *Every $f^\dagger \in X_0$ with $\Phi(f^\dagger) = y_0$, where $y_0 = \Phi(f_\epsilon)$ has top mass $f^\dagger_N \in \{f_N, g_1\}$. Thus, either $f^\dagger_N = 1-(N-1)\varepsilon = \Theta(1)$ or $f^\dagger_N = \lambda\varepsilon + O(\varepsilon^2) = \Theta(\varepsilon)$, and moreover, the offer mass at the highest wage either equals f_N or the top mass associated with its dual Tf, g_1.*

Proof. The largest positive wage changes can only arise from one product $c(N-1) = g_1 f_N$. Likewise $c(-(N-1)) = g_N f_1$. Given the distribution balance equation (10), we can find the solution for g_1 as a function of f_1, and g_N of f_N. These are:

$$g_1 = \frac{\lambda f_1}{(1-f_1)+\lambda f_1}, \quad g_N = \frac{f_N}{f_N + \lambda(1-f_N)} \tag{18}$$

Substituting these into the cross-correlation equation (13) at the extreme wage changes, the lowest wage mass x and top wage mass y that hit these wage changes need to satisfy

$$\begin{aligned} c(N-1)\Big(1-(1-\lambda)x\Big) &= \lambda xy \\ c(-(N-1))\Big(\lambda+(1-\lambda)y\Big) &= xy. \end{aligned} \tag{19}$$

We can now ask which other values x, y are consistent with the observed frequency of the largest positive and negative wage changes. Note that, substituting out y or x yields a quadratic equation in the other variable. Since the leading coefficient is strictly positive (as long as $c(N-1) > 0$, $c(-(N-1)) < 0$ and $0 < \lambda < 1$), it has at most two solutions. By construction $x = f_{\epsilon,1}, y = f_{\epsilon,N}$ is a solutions. However, $x = g_{\epsilon,N}, y = g_{\epsilon,1}$, the mass of the reversed worker distribution of f_ϵ (Tf_ϵ) is also a solution, as can be readily algebraically verified. There are no further solutions. $\square$

Proof of Proposition 3. Fix $\varepsilon < \varepsilon_0$ as in Lemma 3 (small enough that the conclusions of Lemmas 3 and 4 hold).

Let $f^\dagger \in X_0$ be any complex preimage of y_0. Then $P = \widehat{F}^\dagger\,\widehat{G}^{\dagger *}$. Because P has full degree $2N-2$ and nonzero constant term ($c(\pm(N-1)) = \Theta(\varepsilon) \neq 0$), and $\widehat{F}^\dagger, \widehat{G}^{\dagger *}$ both have degree $\leq N-1$, both must have degree exactly $N-1$; in particular $f_N^\dagger \neq 0$. So $\widehat{F}^\dagger$ is a degree-$(N-1)$ factor of P: writing $\kappa := f_N^\dagger$ for its leading coefficient,

$$\widehat{F}^\dagger(z) = \kappa \prod_{i\in I}(z - r_i),$$

where I is an $(N-1)$-element subset of the (distinct) roots of P, and $\widehat{F}^\dagger(1) = 1$ fixes $\kappa = 1/\prod_{i\in I}(1-r_i)$.

Let q be the number of *large* roots in I (so $N-1-q$ small ones). By Lemma 3, each small root gives $1 - r_i = 1 + O(\varepsilon^{1/(N-1)})$, hence a product $\Theta(1)$, while each large root gives $|1 - r_i| = \Theta(\varepsilon^{-1/(N-1)})$. Therefore

$$\Big|\prod_{i\in I}(1-r_i)\Big| = \Theta\big(\varepsilon^{-q/(N-1)}\big), \qquad |f_N^\dagger| = |\kappa| = \Theta\big(\varepsilon^{q/(N-1)}\big),$$

with Θ-constants depending only on the roots of P—fixed once ε is fixed—not on which preimage is being examined. By Lemma 4, $|f_N^\dagger|$ is $\Theta(1)$ or $\Theta(\varepsilon)$. For small ε the quantity $\varepsilon^{q/(N-1)}$ is $\Theta(1)$ only when $q = 0$ and $\Theta(\varepsilon)$ only when $q = N-1$; for $1 \leq q \leq N-2$ it satisfies $\varepsilon \ll \varepsilon^{q/(N-1)} \ll 1$, matching neither (there are finitely many values of q, so a single ε_0 works for all of them at once). Hence $q \in \{0,\ N-1\}$.

Case $q = 0$. I consists of all $N-1$ small roots, which are exactly the roots of $\widehat{F}$ (Lemma 3). Two degree-$(N-1)$ polynomials with the same roots and the same value 1 at $z = 1$ are equal, so $\widehat{F}^\dagger = \widehat{F}$ and $f^\dagger = f_\varepsilon$.

Case $q = N-1$. I consists of all the large roots, the roots of $\widehat{\widehat{G}}$, so $\widehat{F}^\dagger = \widehat{\widehat{G}} = \widehat{Tf_\varepsilon}$ (Proposition 2), and $f^\dagger = Tf_\varepsilon$.

Thus every preimage is f_ε or Tf_ε; both are genuine preimages ($\Phi(Tf_\varepsilon) = \Phi(f_\varepsilon)$ by Proposition 2), and they are distinct (one concentrated at the top wage, the other at the bottom). So $\Phi^{-1}(y_0) \cap X_0 = \{f_\varepsilon,\ Tf_\varepsilon\}$: two points.

Step 2: checking that hypotheses of Lemma 1 at y_0 are satisfied. (i) we have already. *Local injectivity (ii):* at f_ε and Tf_ε the two polynomial factors are the pair $\{\widehat{F}, \widehat{\widehat{G}}\}$ (Lemma 3), that have no common roots, with nonvanishing leading coefficients ($1-(N-1)\varepsilon$ and g_1 respectively); so by the complex version of Lemma 2, $D\Phi$ has full rank $N-1$ at both points, and Result **??** makes Φ locally one-to-one there.

Properness (iii): Let $Y_0 := \{y :\ c(N-1) \neq 0,\ c(-(N-1)) \neq 0\} \ni y_0$, and fix a compact neighborhood $K \subset Y_0$ of y_0 inside $\bar{Y}$ (on which $P(1) = 1$). For $y \in K$ the data polynomial P_y has degree $2N-2$, with all coefficients bounded and the leading and constant coefficients bounded away from 0. Two consequences: *(a)* its roots lie in a fixed annulus $\{c \leq |z| \leq C\}$; *(b)* they stay a fixed distance $\delta > 0$ from $z = 1$—writing $1 = P_y(1) = c(N-1)\prod_i(1-r_i)$, the product $\prod_i |1-r_i| = 1/|c(N-1)| \geq 1/M$ is bounded below while every factor is at most $1+C$, so if one factor were smaller than δ the product would be smaller than $\delta\,(1+C)^{2N-3}$; hence $\delta := \left[M(1+C)^{2N-3}\right]^{-1}$ works.

Now take any sequence $f'_n \in X_0$ with $y_n := \Phi(f'_n) \to y_0$; eventually $y_n \in K$. As in Step 1, each offer polynomial $\widehat{F}'_n$ is a degree-$(N-1)$ factor of P_{y_n} with $\widehat{F}'_n(1) = 1$, i.e. $\widehat{F}'_n = \kappa_n \prod_{i \in I_n}(z - r_i)$ with $\kappa_n = 1/\prod_{i \in I_n}(1-r_i)$; since every $|1-r_i| \in [\delta,\, 1+C]$, we get $|\kappa_n| \leq \delta^{-(N-1)}$, and the remaining coefficients of $\widehat{F}'_n$ (κ_n times elementary symmetric functions of roots bounded by C) are bounded. The *complementary* factor $\widehat{G}'^{*}_n = P_{y_n}/\widehat{F}'_n$ satisfies the same normalization $\widehat{G}'^{*}_n(1) = \widehat{G}'_n(1) = 1$ (worker mass sums to one) and is the analogous product over the complementary root set, so its coefficients are bounded by the same argument. Hence the pairs (f'_n, g'_n) lie in a fixed compact set; pass to a subsequence with $f'_n \to f^*$ and $g'_n \to g^*$ (note $f^* \in X$, the hyperplane $\sum_n f_n = 1$ being closed). It remains to check that f^* is *not a pole* of Φ. Along the sequence the distributional link from (10) holds: $G'_n\left(\lambda + (1-\lambda)(1-F'_n)\right) = \lambda F'_n$ at every index (once the denominator is cleared this is a polynomial identity in the pair (f', g'), so it passes to the limit):

$$G^*_n\, D^*_n = \lambda F^*_n, \qquad D^*_n := \lambda + (1-\lambda)(1-F^*_n), \qquad n = 1, \ldots, N.$$

If some $D^*_n = 0$, then $\lambda F^*_n = 0$, so $F^*_n = 0$, for where it follows that $D^*_n = \lambda + (1-\lambda) \neq 0$, contradiction. So every $D^*_n \neq 0$: the limit f^* lies in X_0, the limit relations give $g^* = \mathcal{G}(f^*)$ (solve them recursively, exactly as in (8)), and Φ is continuous at f^* with $\Phi(f^*) = \lim_n y_n = y_0$. This is hypothesis (iii).

To wrap up: we can now apply Lemma 1 at the 'witness' concentrated offer; this yields $\deg \Phi \leq 2$. □

Estimation of the BM model with Godfather Shocks from Earnings Change Distributions

Given this, we now proceed to estimate the above simple version of the BM model with godfather shocks. We do this across the cycle, for the recession and expansion data, with two versions. (I) we allow λ_1, γ to vary freely, within the bounds dictated by the model; (II) we take the proportion of the EE arrival rate that is a godfather shock as constant across the cycle, based on data pooled across expansion and recession, and let the overall EE arrival rate vary over the cycle. In both versions, the wage offer distribution to be estimated is approximated by 50-point discrete distribution.

To do the estimation, we proceed in two steps. (step 1): we can directly calibrate the implied γ and λ_1 from the two moment conditions from proposition 1 (δ is taken straight from the data). In (step 2), we use these to calculate c_k, and use these to estimate f by quasi maximum likelihood on the $c(k)$ data (incorporating the relation with the steady-state worker distribution $g = \mathcal{G}(f)$ given the parameters from (step 1).)

The approach here is first and foremost illustrative: there are two aspects that are important abstractions here but are addressed in the main model. First, in this section we treat the expansion-recession comparison as one across steady states; the full model incorporates the transitional dynamics of workers as the economy goes through the business cycle. Second, rather than treating the earnings growth data as year-by-year data with potentially compounded transitions as in the main model, here we simply take these as *wage* changes in the BM model for EE flows. However, there are interesting lessons to learn from this exercise.

				I. (λ_1, γ) free			II. $\lambda_1/(\lambda_1+\gamma)$ constant		
	δ	m_{EE}	R_{data}	γ	λ_1	mode	γ	λ_1	mode
pooled	0.016	0.034	1.29	0.024	0.031	exact	0.022	0.028	exact m_{EE}
expansion	0.015	0.035	1.35	0.023	0.043	exact	0.023	0.029	exact m_{EE}
recession	0.020	0.030	0.86	0.028	0	boundary	0.019	0.024	exact m_{EE}

Table 6: "Boundary": recession data violate $R > 1$, so no exact solution exists; the boundary calibration sets $\lambda_1 = 0$ (pure godfather) "exact m_{EE}" calibration to match observed EE flows

Based on the EE and EU moments of Table 4 in the main text, the statistics $R_{data} = U_{obs}/D_{obs}$, m^+_{EE}, m^0_{EE}, m^-_{EE} from earnings growth distribution, and the relations between these from proposition 1, we calculate the resulting parameters (under columns I.) in Table (6). Note that the dominance of the sum of negative changes in recession is inconsistent with the BM model with godfather shocks, so we take the closest the model the model can come to this situation in steady state, which occurs when all EE flows are driven by godfather shocks. Under II, we show the parameters for the case that we keep the proportion of godfather shocks constant over the cycle but adjust the levels to match m_{EE}. In the second step, we estimate an offer distribution across 50 earnings grid points.[44] Interestingly, in the numerics, the second-best cluster is exactly the dual $T\hat{f}$ of the best, at minimal numerical distance ($\sim 10^{-13}$) after independent discovery.

Results We now show the fit to the earnings growth distribution and the resulting offer distributions, for version I. The first-order differences across versions I and II are not large. Note that, because we have not used an additional level moment, we will show both of the "twin" offer distributions.

Figure 1: Fit of the EE earnings change distribution

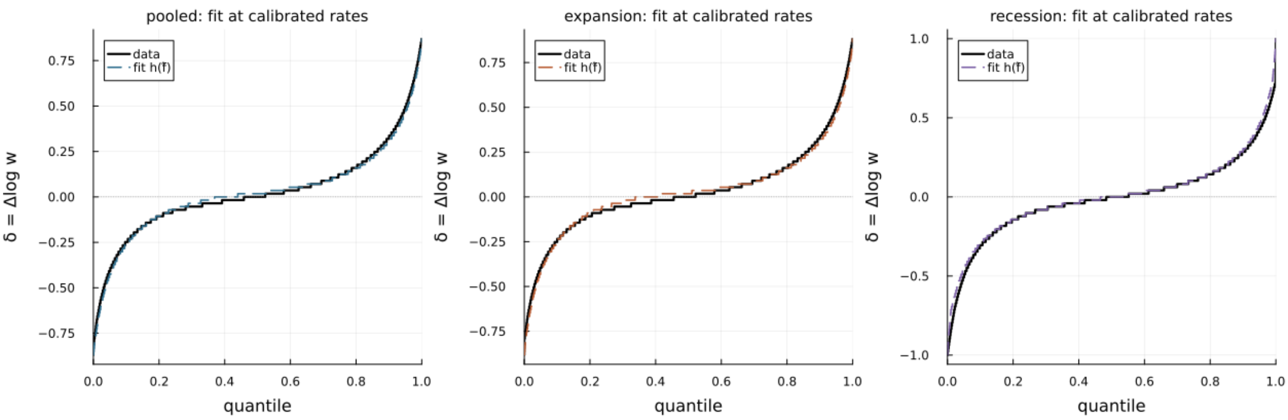

[44]We use a TikTak multistart architecture: 10^6 Sobol points, 2048 CMA-ES local searches with start points blending toward the incumbent, L-BFGS polish of every endpoint, and dual seeding.

In Figure 1, we show the fit of the model with an equidistant wage grid of 50 points. The fit is good relative to the largest observed EE earnings changes.[45] Figure 2 then shows the estimated wage offer distributions. Across the cycle, these are heavily skewed wage offer distributions, which reconcile the coexistence of large wage gains and losses with frequent small wage changes (through EE transitions) by the existence of a lot of small wage changes either at the very top or bottom of the distribution.

Figure 2: Fit of the EE earnings change distribution

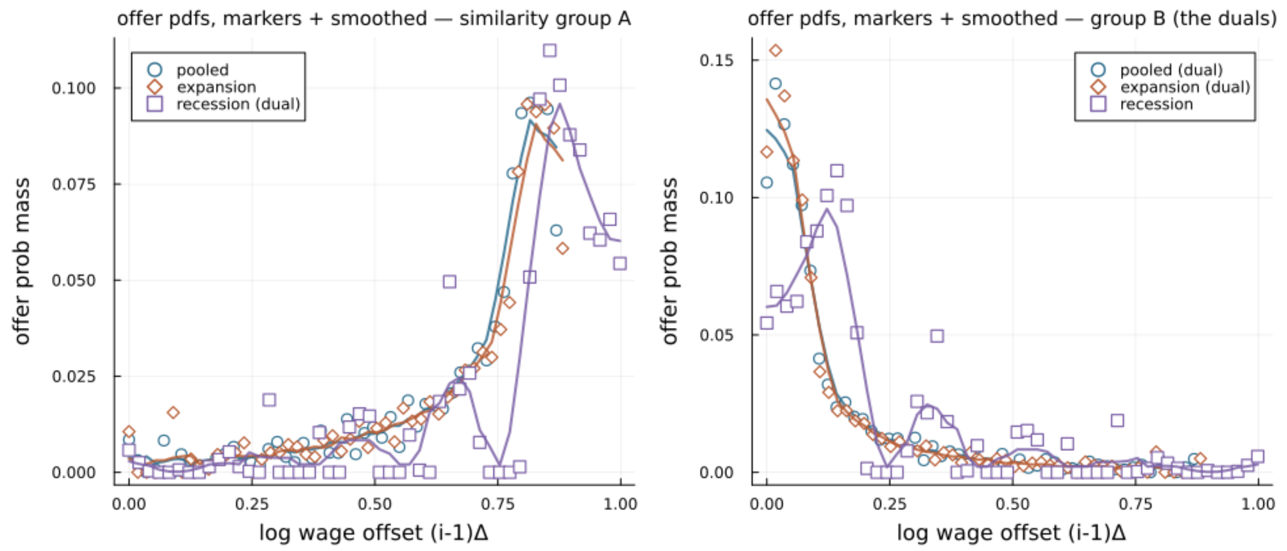

When we investigate how the earnings growth distribution of EE movers changes over the cycle, we observe a roughly similar shape (but not scale) as for the overall cyclical shift of the earnings change distribution, see the solid black line in Figure 3. It turns out that the simple BM model with godfather shocks has trouble matching U-shape in this figure. The scale of the very largest earnings changes makes it harder to observe the earlier pictures on the fit of the earnings change distribution across expansion and recession, but Figure 3 brings this more directly. To give the simple BM model with godfather shocks the best chance of actually replicating the U-shaped cyclical shift, we add this shift as a target and shift weights towards matching the U-shaped cyclical pattern rather than the earnings growth distributions themselves. The dotted baseline in the graph captures the case when we only target the earnings growth distributions of Figure 1; as weights ω increase, we increase the weight on hitting the cyclical shift of the EE earnings growth distribution itself. We do this for both version I and II of the estimation, the right-hand panel refers to the cyclically constant proportion of godfather shocks.

We clearly trade in some fit of the earning change distribution when putting more weight on matching the cyclical shift pattern. More importantly, we see that the model suffers from a tension: when it incorporates that there is a stronger shift towards more frequent large earnings losses through EE moves in recessions, it also unleashes a force towards more frequent large earnings gains, while data

[45]Notice that in the two-ladder model, the fatter tails will be estimated to relate more to the occupational ladder than the employer ladder; here they are generated by the BM offer distribution.

Figure 3: Fit of the *Cyclical Shift* of EE earnings change distribution

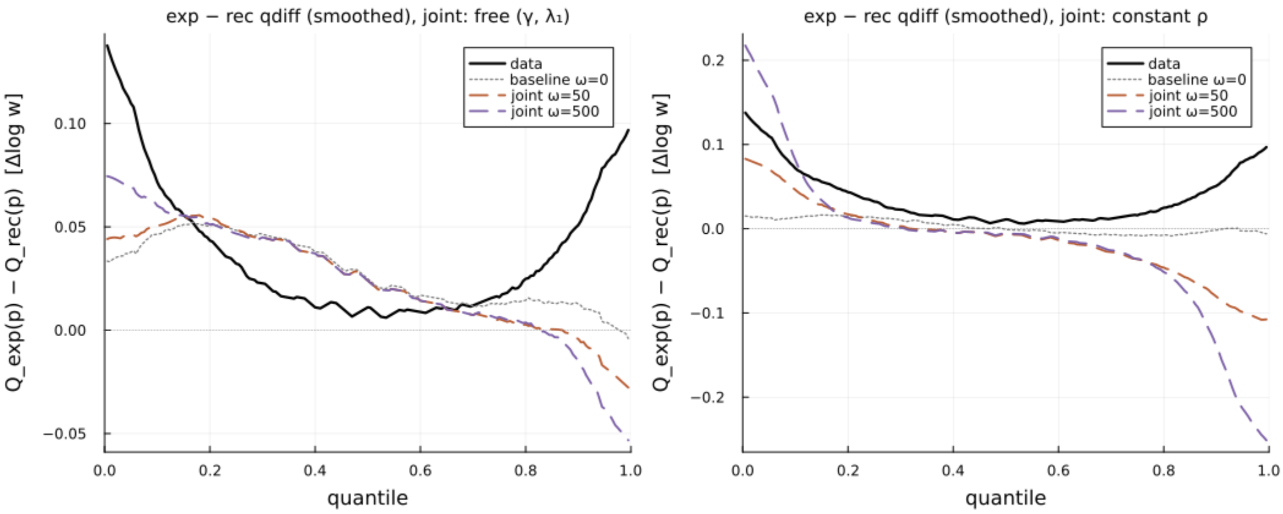

Figure 4: Fit of the EE earnings change distribution, when Cyclical Shift is also targeted

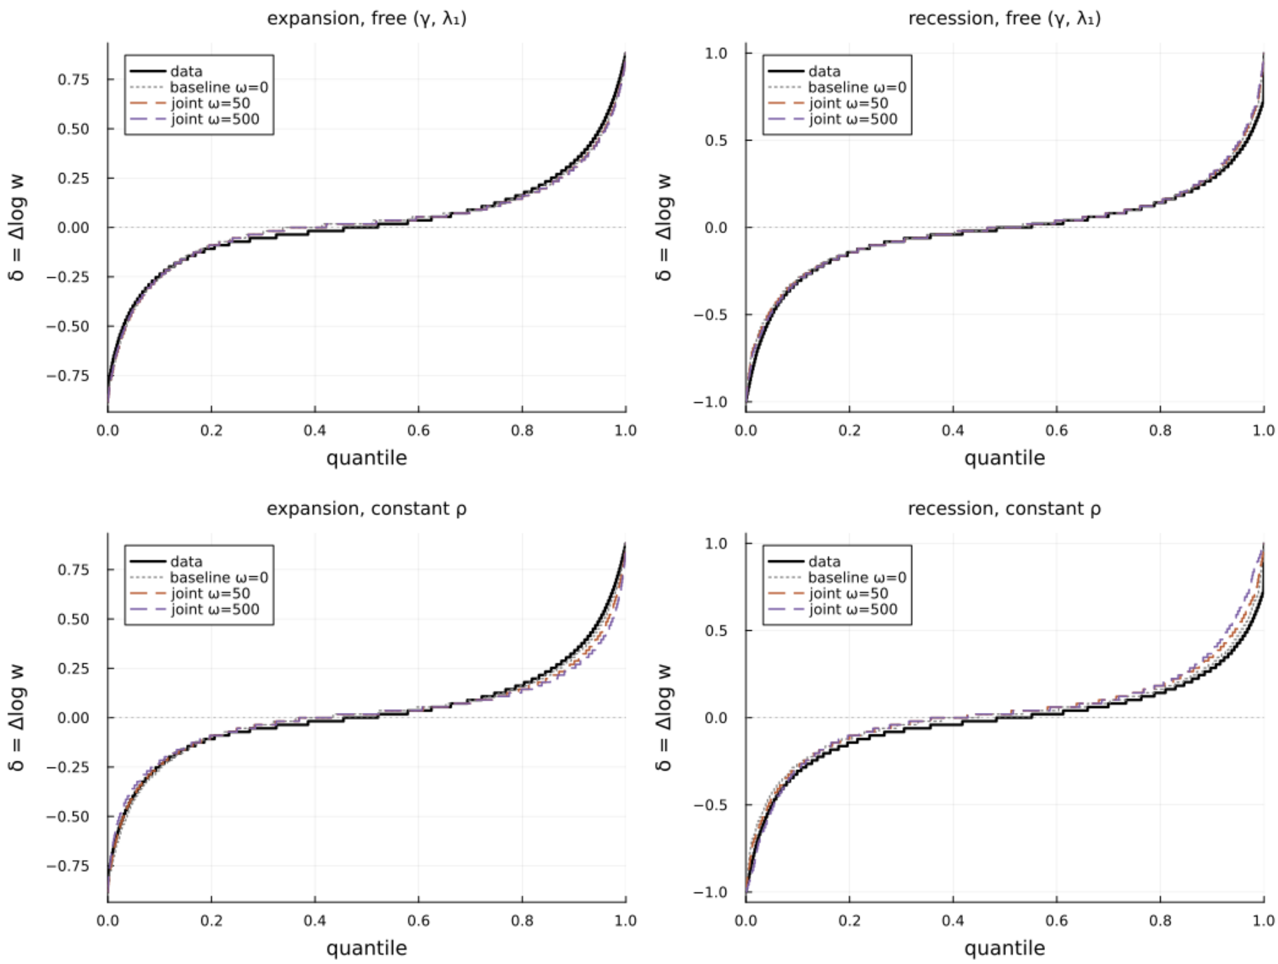

demands diminished upside earnings transitions in a recession. This occurs even when the estimation has the freedom to choose a recession-specific offer distribution, and adjust the λ_1 and γ (respecting the flows). Behind this, intuitively, is that to have frequent large earnings losses, one needs enough workers with high initial wages and frequent enough bad godfather draws, whose co-existence requires workers are able to reach the high initial wages often enough.[46] We expect some of this result also to carry over to the overall distribution as the periods of zero income by themselves at further earnings dispersion but in a rather symmetric way.[47]

Appendix B — Data

The Survey of Income and Program Participation (SIPP) is a longitudinal dataset based on a representative sample of the US civilian non-institutionalized population. It is divided into multi-year panels, each comprising a new sample of individuals subdivided into four rotation groups. Individuals in a given rotation group are interviewed every four months, so that information is collected for each month. At each interview, individuals are asked, among other things, about their employment status, occupations, industries, and monthly earnings over the previous four months.[48] The SIPP's high-frequency interview schedule and panel dimension allow us to follow workers over time, construct uninterrupted unemployment (or, more broadly, non-employment) spells beginning with a separation from employment and ending with a return to employment, and analyze workers' occupational mobility patterns conditional on unemployment duration, as outlined in Section 2 of the main text.

Survey design and use of data

We consider the period 1990–2013, using the 1990–1993, 1996, 2001, 2004, and 2008 panels. For the 1990–1993 panels we use the Full Panel files as the basic data sets, appending the monthly weights obtained from the individual waves (sometimes referred to as core wave data). For these panels we also use the occupational information from the core waves rather than the Full Panel files, for two reasons: (i) the Full Panel files do not always carry an imputation flag for occupations; and (ii) between the 1990 and 1993 panels firm identities were retrospectively recoded based on core wave firm identifiers. For the 1996, 2001, 2004, and 2008 panels, no Full Panel file is produced; we append the individual wave information using the identifier `lgtkey` and merge in the person weights of workers with information for the entire panel (or an entire year). Job identifier information is also clearly specified in these later panels. Since 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.

The SIPP's sample design implies that in *all* panels the first and last three months have fewer than four rotation groups and therefore a smaller sample size, so we only consider months with information for all four rotation groups. Matching occupations after a separation requires that we observe re-employment; if a separation occurs too close to the end of the panel, we would non-randomly select short unemployment spells. We therefore also exclude periods with fewer than two waves remaining before the panel ends, a restriction also necessary to compute annual earnings growth. We weight

[46]The steady state nature of the comparison here can be important, nevertheless, we find a similar tension in the fully dynamic model.

[47]Given the identification of the offer distributions, up to a twin distribution, and all flows except outflow rate λ_0 from unemployment, it is straightforward to extent the current model.

[48]See `http://www.census.gov/sipp/` for a detailed description of the dataset.

the remaining observations using the SIPP's person weights, `wpfinwgt`. The data also exhibit seam effects: transitions are more likely to be reported at a seam (i.e., between waves, and therefore at 4, 8, 12, ... months) than within a wave. As discussed in the main text, we therefore conduct our analysis at the wave level.

Sample selection, labor market status, and transitions

For the 1990–2008 panels, we consider all workers between 18 and 65 years of age who are not in self-employment or in the armed forces. We measure employment transitions using two sources of information. The first is the monthly employment status recode. Using the SIPP 2001 wording as an example, we consider a worker to be employed during a month if the individual reports one of the following: "with a job entire month, worked all weeks"; "with a job all month, absent from work without pay 1+ weeks, absences not due to layoff"; or "with a job all month, absent from work without pay 1+ weeks, absences due to layoff". When a worker is employed for part of the month and unemployed for the other part, we classify the month as unemployment only if the worker is unemployed in week 2 *and* has been unemployed for at least four weeks in total; shorter unemployment spells within the month are classified as employment. A worker who reports "no job/business — looking for work or on layoff" during one of these weeks is classified as unemployed. We adopt this classification because we want unemployment entry to capture a meaningful weakening of attachment to the previous employer, rather than short absences after which the worker returns to the same firm. The four-week threshold further limits the role of short-term absences and temporary layoffs, motivated by Fujita and Moscarini (2017), who document that many workers with very short unemployment spells return to their previous employer. We want to focus on the unemployed who at least *consider* employment in other firms and possibly other occupations.

To measure employer-to-employer transitions, we use the second source of information: start and end dates and job numbers. While each employer is assigned a unique number, the month in which the job number changes does not always correspond to the actual timing of the employer transition, and some job number changes are spurious. We therefore corroborate job number changes with employment end dates: we require that the individual is fully employed in both adjoining months, that one employment relationship ends, and that the job number switches within the wave.

Assigning occupations to non-employed workers

The SIPP collects information on a maximum of two jobs that an individual might hold simultaneously. For each job we have information on, among other things, hours worked, total earnings, three-digit occupation, and three-digit industry codes. We drop observations with imputed occupations (and industries). When an individual holds two jobs simultaneously, we define the main job as the one with more hours; ties are broken by total earnings, though such ties are very rare. Once the main job is identified, the worker is assigned the corresponding two-, three-, or four-digit occupation.

Each unemployment spell that begins and ends within the panel can be assigned a "source" occupation (immediately before the spell) and a "destination" occupation (immediately after re-employment). If the source code is missing (e.g., due to imputation) and an occupation is reported in an earlier wave with continuous employment in between, we carry the earlier occupation forward. A worker is classified as an occupation mover if source and destination differ. We conservatively classify as occupational stays cases in which an individual employed at two firms simultaneously at the moment of separation subsequently re-employs in an occupation matching either prior job, even the one with fewer hours; the effect on the mobility statistics is small.

We construct the occupational mobility statistics from transitions of the form: at least one month of employment (with a non-imputed occupational code), followed by an unemployment spell of at least one month, followed by at least one month of employment (with a non-imputed occupational code). We label these as EUE transitions. We also consider transitions in which the jobless spell, while including at least one month of unemployment, also contains non-participation months; we label these as E-NUN-E transitions, or NUN spells of non-employment. To span the intermediate cases, we further consider spells that begin with an EU transition (with the worker possibly later reporting they stopped looking for work) and those that end with a UE transition, labelled E-UN-E, E-NU-E, and — when both restrictions apply — E-UNU-E. We also consider a version in which the entire jobless spell is classified as non-employment (ENE).

Occupational classifications The SIPP uses the Census of Population Occupational System, closely related to the Standard Occupational Classification (SOC). The 1984–1991 panels use the 1980 Census Occupational Classification, the 1992–1996 and 2001 panels use the 1990 classification (which differs only slightly), and the 2004 and 2008 panels use the 2000 classification (which differs more substantially). To obtain a uniform three-digit coding system, we recode the 1980 and 2000 classifications into the 1990 classification using the IPUMS crosswalk, very similar to the one developed by Dorn (2009) and Autor and Dorn (2013).[49] From these three-digit codes we further aggregate into two coarser groupings: a two-digit version corresponding to the 22 Standard Occupational Codes (the federal statistical standard), and a one-digit version combining occupations into four task-based categories defined by the routine/non-routine and cognitive/manual dimensions: Non-routine Cognitive, Routine Cognitive, Non-routine Manual, and Routine Manual.

It is well known that measurement error in occupational codes may give rise to spurious transitions, as discussed for example by Kambourov and Manovskii (2008) and Moscarini and Thomsson (2007). Since 1986, the SIPP interviewing procedure carries the previous interview's occupational code forward whenever the worker reports no change in job type or employer. This "dependent interviewing" reduces spurious occupational transitions among employer stayers, but coding errors remain among employer movers. Because each earnings change observation requires a specific occupational code, we do not adopt probabilistic correction methods of the kind used by Kambourov and Manovskii (2008); instead we rely on a high degree of aggregation. Carrillo-Tudela and Visschers (2023) show that, among employer movers, correcting for coding errors when using the four task-based categories reduces the observed gross occupational mobility rate by about 5 percentage points, so observed mo-

[49]We exclude the Armed Forces from all classifications. The 1980 and 1990 classifications are available at `https://www.census.gov/people/io/files/techpaper2000.pdf`, the 2000 classification at `http://www.bls.gov/soc/socguide.htm`, and further information at `http://www.census.gov/hhes/www/ioindex/faqs.html`.

bility remains substantial after correction. Moreover, to the extent that misclassification bias does not vary much over the cycle (as documented in Carrillo-Tudela and Visschers, 2023), coding error introduces a downward bias to the role of occupational mobility in explaining the procyclical skewness of the earnings growth distribution: observed stayers (very likely true stayers) exhibit weak procyclical skewness relative to observed movers, and since the observed-mover group likely contains some true stayers due to misclassification, true procyclical skewness among genuine movers should be even greater than what we observe.

Earnings construction

Monthly earnings We deflate nominal monthly earnings in the SIPP by the Personal Consumption Expenditure price index. Letting $w_{i,t}$ denote the resulting real earnings of individual i in month t, we estimate

$$\log w_{i,t} = x'_{i,t}\beta + \mu_t + \epsilon_{i,t}$$

on the subsample of months with positive earnings ($w_{i,t} > 0$), where $x_{i,t}$ contains a quadratic in potential experience, education, gender, and race, and μ_t are calendar-month dummies. To abstract from variation in earnings changes driven by observable worker characteristics, we construct a monthly earnings measure in *levels* that combines a component common to all workers in a given month with an individual-specific residual:

$$\hat{w}_{i,t} = \begin{cases} \exp(\tilde{w}_t + \hat{\epsilon}_{i,t}) & \text{if } w_{i,t} > 0, \\ 0 & \text{if } w_{i,t} = 0, \end{cases}$$

where $\tilde{w}_t$ is the cross-sectional mean of $\log w_{i,t}$ in month t (computed over employed individuals) and $\hat{\epsilon}_{i,t}$ is the estimated residual. The residualisation is therefore applied in logs to positive-earnings months only; zero-earnings months enter $\hat{w}_{i,t}$ directly as zeros in levels. This preserves the distinction between employment and non-employment that is essential for our subsequent linkage of earnings changes to labor market transitions.

Cleaning reporting errors Survey earnings data are subject to reporting error, and occasionally misreported earnings bias the variance of earnings changes upwards — particularly for employer and occupation stayers, whose true earnings changes are smaller and for whom measurement error is therefore relatively larger. A common cleaning method applies time-series break-detection, as in Gottschalk (2005), to reject small transitory changes. We do not adopt this approach for two reasons. First, it would itself impart leptokurtosis to the earnings change distribution we study. Second, evidence from administrative data (Kurmann and McEntarfer, 2019) suggests that small earnings changes are not solely measurement error, and Hudomiet (2015) finds that annual earnings changes for workers with interrupted careers are in fact *larger* in administrative data than in survey data. The latter implies that, if anything, our SIPP-based estimates understate the true scale of earnings changes among EUE movers.

Following Busch et al. (2022), we instead clean the residual earnings data by dropping the bottom and top 2% of the wave-frequency earnings distribution and excluding imputed earnings. We further drop a small number of observations (less than 1% of the sample) that exhibit a one-period change exceeding 200% which reverts such that the two-period change is below 10% — a pattern consistent with a shifted decimal or entry error. We retain such changes when they coincide with an employment status change at either monthly or wave frequency.

Annual earnings and earnings changes Because the final month of each SIPP wave records an artificially high number of transitions, we link earnings to transitions at the wave level. Letting t index the first month of a reference wave, we define annual earnings preceding and following the reference wave as

$$\hat{w}^{a,-}_{i,t} = \sum_{\tau=1}^{12} \hat{w}_{i,t-\tau}, \qquad \hat{w}^{a,+}_{i,t} = \sum_{\tau=0}^{11} \hat{w}_{i,t+\tau},$$

where the sums include zero-earnings months arising from non-employment spells. Both windows therefore reflect the extensive margin directly.

Because some EUE transitions involve re-employment after long unemployment spells, the associated $\hat{w}^{a,\pm}_{i,t}$ can be very low. To handle this uniformly across all workers, we use the inverse hyperbolic sine (IHS) transformation rather than logs when computing earnings changes. Letting $\tilde{w}^{a,\pm}_{i,t}$ denote the IHS of $\hat{w}^{a,\pm}_{i,t}$, our year-to-year earnings change measure is

$$\Delta_{i,t} = \tilde{w}^{a,+}_{i,t} - \tilde{w}^{a,-}_{i,t}.$$

The IHS transformation coincides with the log for moderate and large earnings and remains well-defined at zero, so the resulting $\Delta_{i,t}$ behaves like a log change away from the lower tail while accommodating long non-employment spells. To avoid the small range over which IHS and log differences diverge meaningfully, we restrict the sample to workers with annual earnings of at least \$1,040 — corresponding to ten hours per week at the national minimum wage over a quarter, following Guvenen et al. (2014). In practice, our IHS-based estimates are very close to what log differences would yield.

Appendix C — Earnings growth and occupational mobility

Cross-sectional earnings growth distribution

In this section we present additional features of our SIPP data in relation the earnings growth distribution that we referred to in Section 2 of the main text. Guvenen et al. (2021) shows that key features of the earnings growth distribution are that it is left skewed and very leptokurtic, with approximately Pareto-distributed tails. Figure 5a plots the log density of the earnings growth distribution to better visualize this latter property. Figure 5 presents the density in levels to highlight its leptokurtosis. We present these densities by stacking the distributions associated with EUE, EE transitions and employer stayers on top of each other to show their role in shaping earnings growth distribution.

Figure 5c instead presents the quantile function of the earnings growth distribution, where we graph the earnings changes in the y-axis and the quantiles of the distribution in the x-axis. This figure further emphasizes that large earnings changes lie at the bottom and top tails of the distribution. Figure 5d conditions this quantile function by whether the earnings change was associated with the worker

Figure 5: Earnings growth distribution in the cross-section

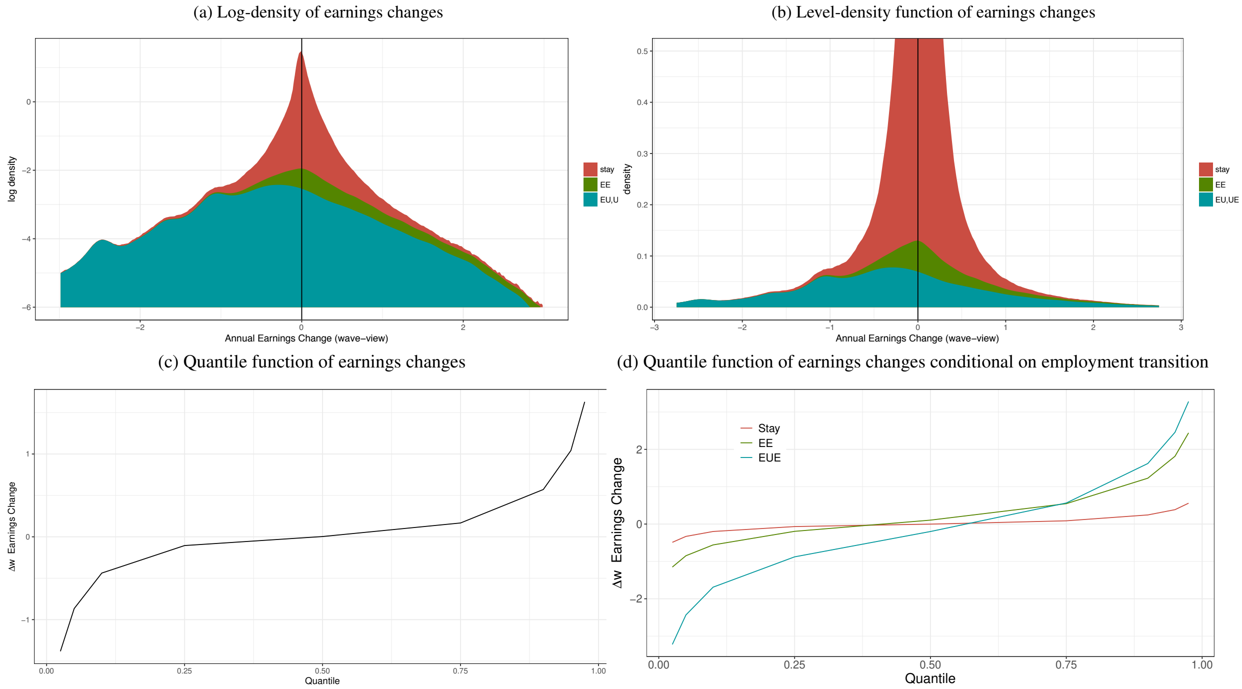

Note: The annual earnings growth distribution is constructed for the sample period 1990-2013. It is based on residual earnings after controlling for potential experience, education and month dummies. Online Appendix A presents the details of the definition of earnings and worker transitions. Recessions are defined as periods in which the HP-filtered unemployment rate is in the top 20% of realizations. For the bottom graphs we use the quantiles obtained from the unconditional sample.

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 5a. 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.

Figure 6 further shows that we arrive to a similar conclusion by analyzing earnings growth conditional on previous earnings. Figure 6a shows that the SIPP data is also consistent with the relationship between earnings growth and previous earnings documented by Guvenen et al. (2014), who used the previous five-year earnings percentile based on SSA data for all US. Although here we only use the previous year earnings, we find that workers with the larger earnings changes are also those who had the lowest earnings, while those with progressively higher previous year earnings are associated with smaller changes. One of the advantages of the SIPP relative to the SSA data is that the former provides better information about individuals' labor market histories and demographic characteristics. These characteristics are the ones we exploit in this paper and it is reassuring that, despite the much smaller number of observations, the SIPP and the SSA data present consistent pictures of the earnings change distribution. The role of the underlying labor market flows can be gauged from Figure 6b, which presents the same relation, but for a sample restricted to only employer stayers. The earnings growth of employer stayers are not only much less dispersed, but the probability of an earnings change (positive or negative) is much less sensitive to these workers' previous year earnings. In fact, the earnings growth distribution does not change much across the previous year's earnings distribution, apart from the probability of relatively larger improvements among the very low-earners, and the probability of earnings losses among very high earners.

Figure 6: Earnings growth distribution conditional on previous earnings

(a) Conditional on previous year earnings

(b) Previous year earnings and employer staying

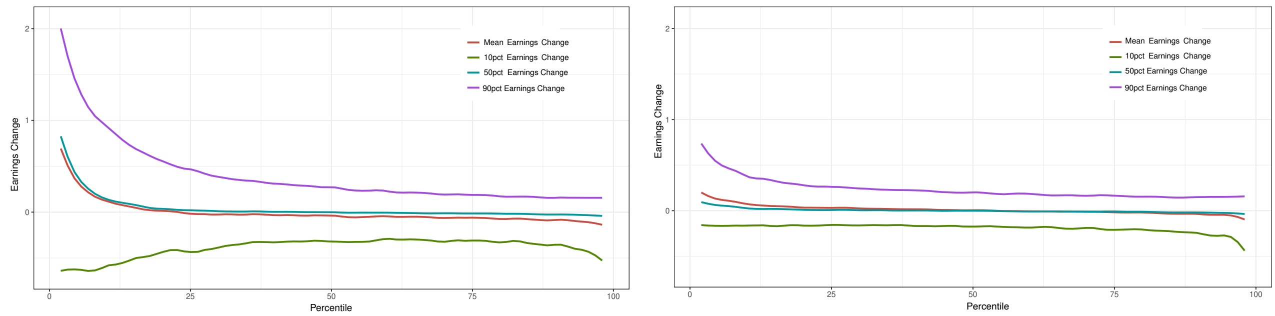

Note: The annual earnings growth distribution is constructed for the sample period 1990-2013. It is based on residual earnings after controlling for potential experience, education and month dummies. Section 2.1 presents the details of the definition of earnings and worker transitions. Recessions are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

This evidence shows that those workers who experienced large earnings changes are also those who had low previous year earnings *and* changed employers. Further, the 90th percentile curve in Figure 6a shows that workers who were at the bottom of the earnings distribution climb the most, while the slow decline of the curve reflects that the larger are previous earnings, the lower is the gain from changing employer. These patterns are broadly inline with the implications of standard job ladder models.

Probability of occupation switching and earnings growth Figure 7 depicts the probability of an occupational change associated with a given value of earnings growth quantiles. Figure 7a shows that when pooling together employer movers and stayers the fat tails of the annual earnings growth distribution depicted in Figure 5a are associated with a high probability of an occupation change (see also Guvenen et al. 2021). Small earnings changes centered around zero, however, are associated with a much smaller probability of an occupation change. Although the large difference in the probability of an occupation move can be accounted for by the larger propensity to change occupations among employer movers, the same pattern remains when only considering employer movers. Figures 7b and 7c show this pattern among those individuals who changed employers through an EE or EUE transition. 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.

Figure 7d instead presents the share of EUE and EE occupation switchers and non-switchers (left axis) as well as the share of stayers (right axis) for each quantile of the earnings growth distribution. It makes clear that the increase in the probability of an occupation switch at the tails of the earnings growth distribution documented in the previous graphs arises from the convergence between the shares of both EE and EUE occupation switchers and non-switchers, where EUE shares follow a U-shape relationship and EE shares an inverse U-shape relationship. The levels of these curves

Figure 7: Earnings growth distribution and probability of occupation mobility

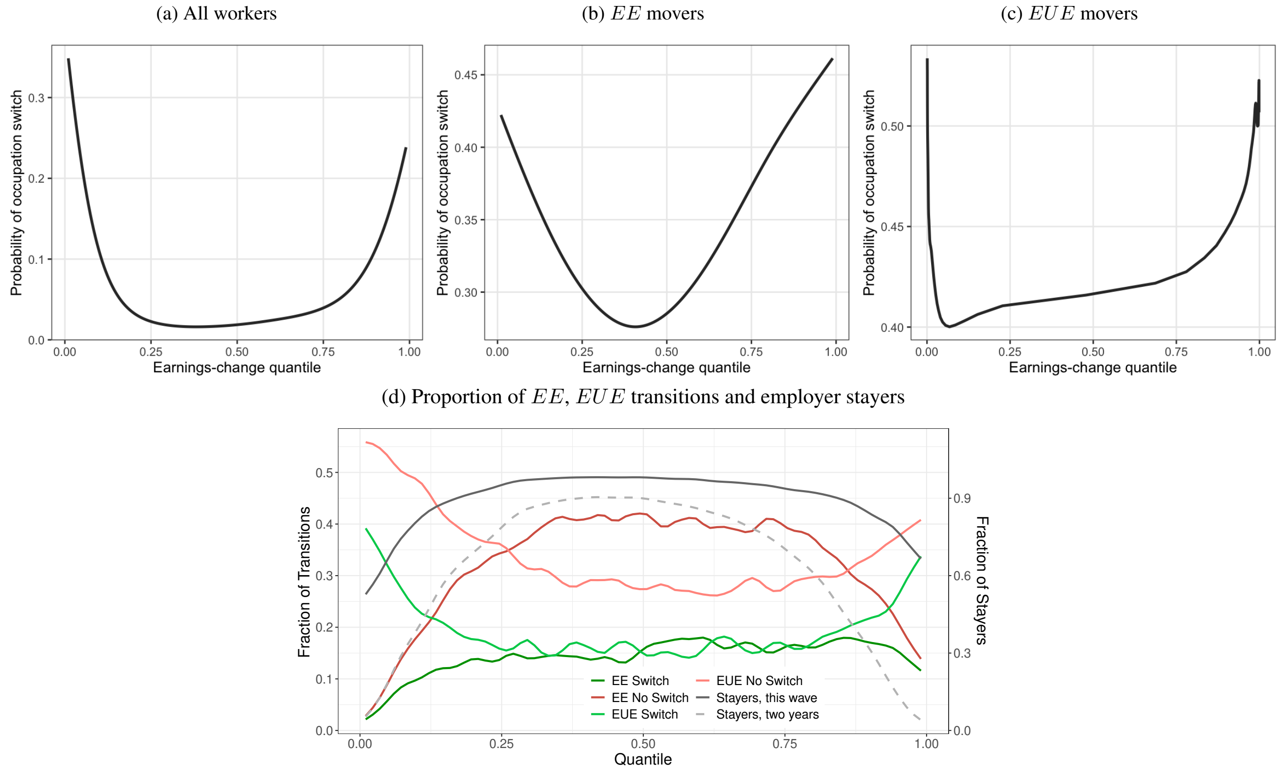

Note: The probability of occupational change is computed as the proportion of workers with a given earnings change who took a job in a different task-based occupation: Non-routine Cognitive, Routine Cognitive, Non-routine Manual and Routine Manual. Earnings change quintiles in the horizontal axis of all figures are based on the same *overall* residual earnings distribution, controlling for potential experience, education and month dummies.

are consistent with the finding that the average occupation mobility rate among employer switchers is about 28% at our level of aggregation (4 task-based occupation categories). This figure also differentiates between individuals who, relative to the reference wave, did not change employers within the two year span for which we compute earnings changes (dash curve), and those who did not change employers within the reference wave but did report an employer change in previous or posterior waves within the two year span (solid curve). A comparison between these two curves shows that the former represent a very small share (close to zero) at the tails of the earning growth distribution, re-affirming that employers movers are those who populate these tails.

Cyclical skewness of the earnings growth distribution

As discussed in Section 2 of the main text, we provided a way to decompose the increase in the left-skewness of the earnings growth distribution we observe in recessions. We build on Halvorsen et al. (2024), whose proposed decomposition implies a co-skewness term,

$$\Delta Skew_k = \left(\frac{std(\Delta w|k, exp)}{std(\Delta w|exp)}\right)^3 skew(\Delta w|k, exp) - \left(\frac{std(\Delta w|k, rec)}{std(\Delta w|rec)}\right)^3 skew(\Delta w|k, rec) + \Delta coskew,$$

where $std(\Delta w|k,.)$ and $skew(\Delta w|k,.)$ denote the standard deviation and skewness of the earnings growth distribution of workers in group k and $std(\Delta w|k,.)$ denotes the standard deviation of the aggregate earnings growth distribution. The first term denotes the contribution of group k in expansions, while the second one the contribution in recessions.

The first row of Table 7 show the outcome of Halvorsen et al. (2024) decomposition (HHOS). As

Table 7: Decomposition of Cyclical Change in Skewness

	No Occ. Switch			Occ. Switch			
	Emp. Stay	EE	EUE	Emp. Stay	EE	EUE	co-skew
HHOS Skewness Contrib	0.080	0.036	-0.024	0.007	0.045	0.148	0.708
CVW Skewness Contrib	0.020	0.058	0.338	0.005	0.078	0.501	
Fraction of Observations	0.865	0.027	0.051	0.012	0.014	0.031	

Note: Decomposition of cyclical change in third central moment following Halvorsen et al. (2024). Weights across groups held constant over the cycle. Annual earnings change based on residualized earnings after controlling for potential experience, education and month dummies. Recessions are defined as periods in which the HP-filtered unemployment rate is in the top 20% of realizations.

mentioned in the main text, conditional on the co-skewness term, we find that occupational switchers explain about 65% of the increase in left skewness of the earnings growth distribution during recessions. Our benchmark analysis instead uses a linear decomposition, where the co-skewness terms are attributed to the group whose mean is far from the unconditional mean. The second row of Table 7 presents our benchmark (CVW) decomposition. The main implication is the same: those workers who change occupations and employers at the same time contribute 57.9% to the observed increase in the left-skewness of the earnings growth distribution during recessions. This happens even though occupation/employer movers represent 4.5% of all observations in our sample. This table also shows that the relatively large contribution of EUE occupation movers across both decompositions.

Table 8: Cyclical Skewness Decomposition: Involuntary vs Voluntary Transitions

	No Occ Switch					Occ Switch				
		Voluntary		Involuntary			Voluntary		Involuntary	
	Stay	EE	EUE	EE	EUE	Stay	EE	EUE	EE	EUE
Contribution	0.042	-0.012	0.318	0.003	-0.062	0.011	0.049	0.380	0.002	0.269
Fraction	0.907	0.012	0.024	0.002	0.014	0.012	0.006	0.015	0.001	0.006

Note: Decomposition of cyclical change in third central moment following Halvorsen et al. (2024). Weights across groups held constant over the cycle. Noting that there is a significant number individuals who do not respond to this question, our analysis conditions on a response.

In the main text we discussed the role of involuntary and voluntary transitions in generating the increased left-skewness of the earnings growth distribution during recessions. Within the SIPP, we use the `ersend` variable for the reason a match ended. We classify transitions as voluntary when workers declared that the main reason for stopping work with their employer was (i) retirement or old age, (ii) school/training, (iii) temporary job and ended, (iv) quit to take another job, (v) unsatisfactory work arrangements (hours, pay, etc.) and (vi) quit for some other reason. A transition is classified as involuntary when workers declared that they separated because (i) of layoff, (ii) childcare problems, (iii) other family/personal obligations, (iv) illness or injury, (v) discharged or fired, (vi) employer bankrupt, (vii) employer sold business, and (viii) slack work or business conditions. An important caveat with this variable is that one might reasonably argue that this information might not be reported accurately and too many report voluntary separations, particularly among those who transit into non-employment. Another caveat with this variable is that there is a large amount of missing data. Table 8 shows the results of the decomposition taking into account the reason for separation. The increase in left-skewness during recessions arising from involuntary EUE transitions with an occupation switch are more important than involuntary EUE transitions without an occupation switch, although the latter represent more than twice as many observations than the former. Voluntary EE and EUE transitions that are accompanied by an occupational switch are also more important in explaining the increase in left-skewness than voluntary EE and EUE transitions without an occupation switch.

As discussed in the main text, the increase in left-skewness can be visually represented through the U-shape pattern obtained from the difference between the expansion and recession quintile functions of the earnings growth distribution. As shown in Table 7 the lion share of the increase in left-skewness arises from occupation switchers who also changed employers. The decompositions also show that both EE and EUE transitions with occupational switching contribute to the increase in left-skewness. This result can be visually verified in Figure 8, where we decomposes the cyclicality of the earnings growth distribution among employer movers 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 characterized by their procyclical skewness.

Figure 8: Cyclical earnings growth: employer/occupation movers and stayers

(a) EUE employer movers

(b) EE employer movers

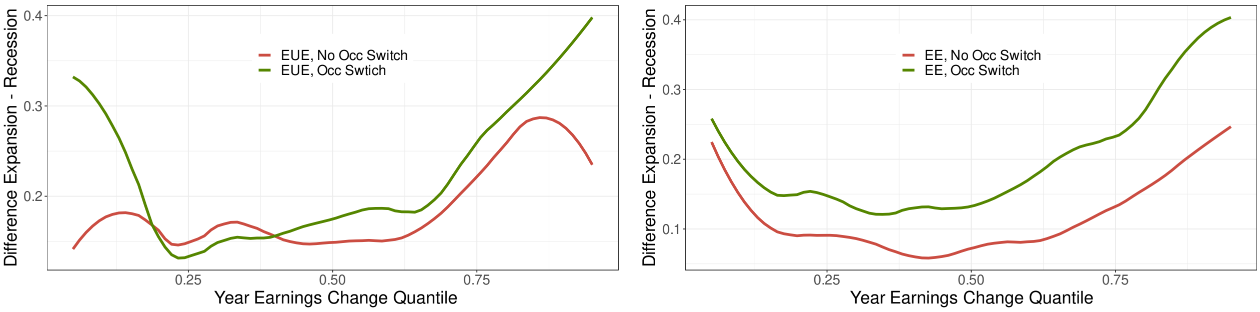

Note: Among those who change employers and conditional on the type of employer transition (EUE or EE), the cyclical change in the annual earnings growth distribution is constructed separately for those workers who simultaneously change occupations and for those who did not change occupations. For each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

Perhaps unsurprisingly, during recessions the earnings losses among those occupation and EUE movers are more pronounced than among occupation and EE movers. During expansions, however, both type of transitions lead to similar increases in earnings gains. Among employer movers and occupation stayers only those who experienced an EE transition are the ones that exhibit an earnings growth distribution characterized 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 during 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 and occupation stayers as depicted in Figure 2a in the main text.

Figure 9 shows that the same patterns also hold when considering changes in hourly wages instead of earnings. To construct annual changes in hourly wages, we accumulate average residualized hourly wages over the prior and posterior year following a transition, just as we did for earnings. When the worker is non-employed, we assign a wage rate of zero. Hence, we isolate out changes along the

Figure 9: Wage growth distribution over the cycle and the importance of occupational movers

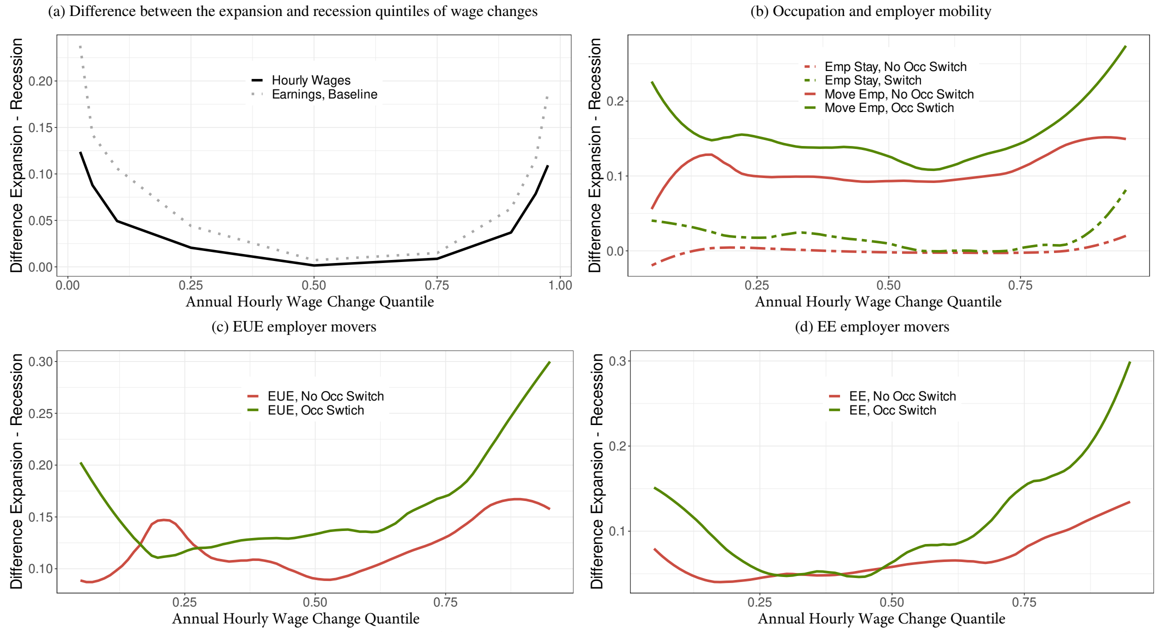

Note: The annual wage growth distribution is constructed for the sample period 1990-2013. It is based on residual wages after controlling for potential experience, education and month dummies. Quantiles are based on the annual wage growth distribution. Recessions are defined as periods in which the HP-filtered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

intensive margin of hours. The motivation for the latter is because we are focusing on workers who have large earnings changes one might be concern that these reflect occupation transitions that are accompanied by movements in and out of part-time work and/or changes in work habits. Including periods of non-employment here is fully consistent with our theoretical framework as our proposed job ladder model also include these periods. We observe that wage changes have remarkably similar U-shape pattern as earnings changes.

Figure 9a shows that the aggregate hourly wage growth distribution has the same U-shape pattern that characterizes the increase in left-skewness in the annual earnings growth distribution. This establishes that the basic cyclical skewness in not only a feature on changes in hours worked. Figure 9b shows that occupation changes are the primary drivers of this skewness again when using wages rather than earnings. Figures 9c and 9d that the countercyclical skewness arises from both EE and EUE transitions that also involve an occupational switch.

Unemployment transitions of high occupation tenured workers

In Section 2.2 of the main text we discussed the role of workers who have high and low occupation tenure in driving the increase in left-skewness of the earnings growth distribution during recessions. We showed that the earnings growth distribution of high-tenured workers has the largest left-skewness relative to other demographic groups, and that this left-skewness significantly increases in recessions. Here we show evidence that high-tenured workers face and increased transition probability into unemployment during recessions, which we argue is one of the key factors lying behind the increase in left-skewness of the earnings growth distribution among these workers. Figure 10a shows a clear negative relationship between the occupational tenure of workers and the probability that they will transition into unemployment. This relation decreases steeply at low tenure levels and tends to flatten

Figure 10: Unemployment separation probabilities of high-tenured workers

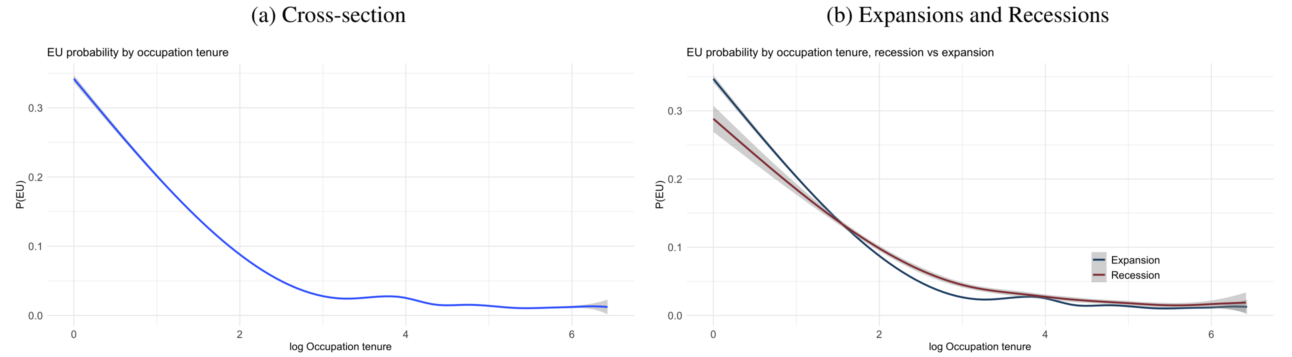

Note: Relationship between the probability of transiting into unemployment by occupational tenure. High-tenured individuals are defined as having tenures above the median, which is calculated to be 88 months. Low-tenured individuals are those who have tenure durations below this threshold. Recessions are defined as periods in which the HP-filtered unemployment rate is in the top 20% of realizations.

out at high tenure levels. Figure 10b shows that the same decreasing pattern between occupational tenure and the probability of transiting into unemployment holds in recession and expansions periods. The key feature, however, is that during recessions high occupational tenured individuals face a higher separation probability into unemployment relative to expansions. This is in contrast to low occupational tenured individuals, who face a lower separation probability into unemployment during recession relative to expansions.

Earnings growth and occupation identities

In Section 2.3 of the main text we argued that occupation-wide effects do not seem to meaningfully contribute to the increase in left-skewness of the year-to-year earnings growth distribution during recessions. We investigate this along a number of dimensions: (i) whether occupations contribute differently to the overall skewness; (ii) whether the earnings changes of occupation switchers are similarly more dispersed across (origin) occupations and whether occupation switchers contribute more to the cyclical tail shift (iii) the role of changes in the (wage) occupation fixed effect when changing occupations, across the cross-sectional earnings growth distribution, and (iv) across the cycle.

Skewness contribution across Occupations We start by evaluating the role of each of the four task-based occupational categories, Non-routine Cognitive, Routine Cognitive, Non-routine Manual and Routine Manual, in contributing to cyclical skewness. Table 9 shows that the left-skewness of the earnings growth distribution arising from non-routine cognitive occupations increases from -1.81 to -3.21, while that of routine cognitive increases from -1.74 to -2.94. Among the manual occupations the left-skewness of the earnings growth distribution arising from non-routine manual occupations increase from -1.58 to -2.47, and that of routine manual occupations increases from -1.96 to -2.67. The largest contribution to overall left-skewness arises routine cognitive occupations.

Similarity of Earnings Change Patterns of Occupation Switchers across Occupations The increase in left-skewness in each of these occupation categories appears to arise from occupation switchers that also changed employers. Figure 11 depicts the kernel density estimates of the earnings growth distribution arising from each of these four task-based categories separately for occupation switch-

Table 9: Central Moment Skewness by Business Cycle Phase: Origin Occupation

	Expansion	Recession	Contribution to $\Delta Skew$	Population in the Pop
Non-routine Cognitive	-1.81	-3.21	0.19	0.28
Routine Cognitive	-1.74	-2.94	0.43	0.31
Non-routine Manual	-1.58	-2.47	0.20	0.13
Routine Manual	-1.96	-2.67	0.15	0.27

Note: Decomposition of cyclical change in third central moment. Annual earnings growth are based on residualized earnings, while recessions defined as periods in which the HP-filtered unemployment rate lies in the top 20% of realizations.

ers and non-switchers, conditional on having experienced either an EE or EUE transition, and for expansion and recession periods. We find that the leftward shift of the density during recessions is more pronounced among occupation switchers within each task-based category. During recessions, the most dramatic cyclical variation arises in the tails of the distribution where the cyclical gap in both the left and right tails is larger for occupational switchers than non-switchers.

Taking out occupation fixed effects from positive monthly earnings, replacing them by average effect To further analyze the role of occupation-wide effects, we considered a second measure of earnings by controlling for occupation fixed effects, υ_o, when residualising earnings as described in the main text. Following from the details of the construction of our benchmark earnings measure (see Appendix B), monthly earnings without occupation fixed effects are such that $\hat{w}_{i,t}^{\text{nofx}} = \exp(\tilde{w}_t + \epsilon_{i,t}^{\text{nofx}})$ if $w_{i,t} > 0$. Importantly, whenever a worker has zero monthly income, we leave it unchanged: $\hat{w}_{i,t}^{\text{nofx}} = 0$ if $w_{i,t} = 0$. We construct annual earnings in an analogous way as described in Appendix B, such that $\hat{w}_{i,t}^{\text{nofx},a,-} = \sum_{\tau=0}^{11} \hat{w}_{i,t-\tau}^{\text{nofx}}$ and $\hat{w}_{i,t}^{\text{nofx},a,+} = \sum_{\tau=1}^{12} \hat{w}_{i,t+\tau}^{\text{nofx}}$. Letting $\tilde{w}_{i,t}^{\text{nofx},a,-}$ and $\tilde{w}_{i,t}^{\text{nofx},a,+}$ denote the inverse hyperbolic sine transformation of these annual earnings, the year-to-year earnings change net of occupation-wide effects is define as $\Delta_{i,t}^{\text{nofx}} = \tilde{w}_{i,t}^{\text{nofx},a,+} - \tilde{w}_{i,t}^{\text{nofx},a,-}$. We can then decompose each earnings growth quantile into an occupation-wide component and an "idiosyncratic" component (the new residualized earnings), where $\Delta_{i,t}$ is the original residualized annual earnings change, $\Delta_{i,t}^{\text{nofx}}$ the new residualized earnings. We perform this exercise using the 22 occupation categories of the two-digit 1990 SOC. Figure 2d in the main text depicts the difference between expansion and recession quantile functions obtained using $\Delta_{i,t}^{\text{nofx}}$. It also shows the difference of these quantile functions using $\Delta_{i,t}$. Hence a comparison between the two allows us to gauge the importance of occupation-fixed effects, as $\Delta_{i,t}\upsilon_o = \Delta_{i,t} - \Delta_{i,t}^{\text{nofx}}$ measures the change in earnings due to the changes in the occupation-wide component. As seen in Figure 2d, the occupation-fixed effects do not appear to meaningfully contribute to the cyclical skewness of the earnings growth distribution.

Change of Occupation Fixed Effecs across the Earnings Growth Distribution To further investigate the role of occupation-wide effects, we now consider the previously estimated occupation fixed effects υ_o, and compute the difference between these fixed effects at the point of a labor market transition. For example, if we observed an EE transition in a given month, we compare the υ_o associated with the immediately preceding occupation to the fixed effect $\upsilon_{o'}$ associated with the occupation immediately proceeding the transition. If we observed an EUE transition we compare the fixed effect associated with the occupation at the point of transition into unemployment with the fixed effect

Figure 11: Earnings Growth Distributions of Employer Changers by Occupational Mobility: Origin Occupation

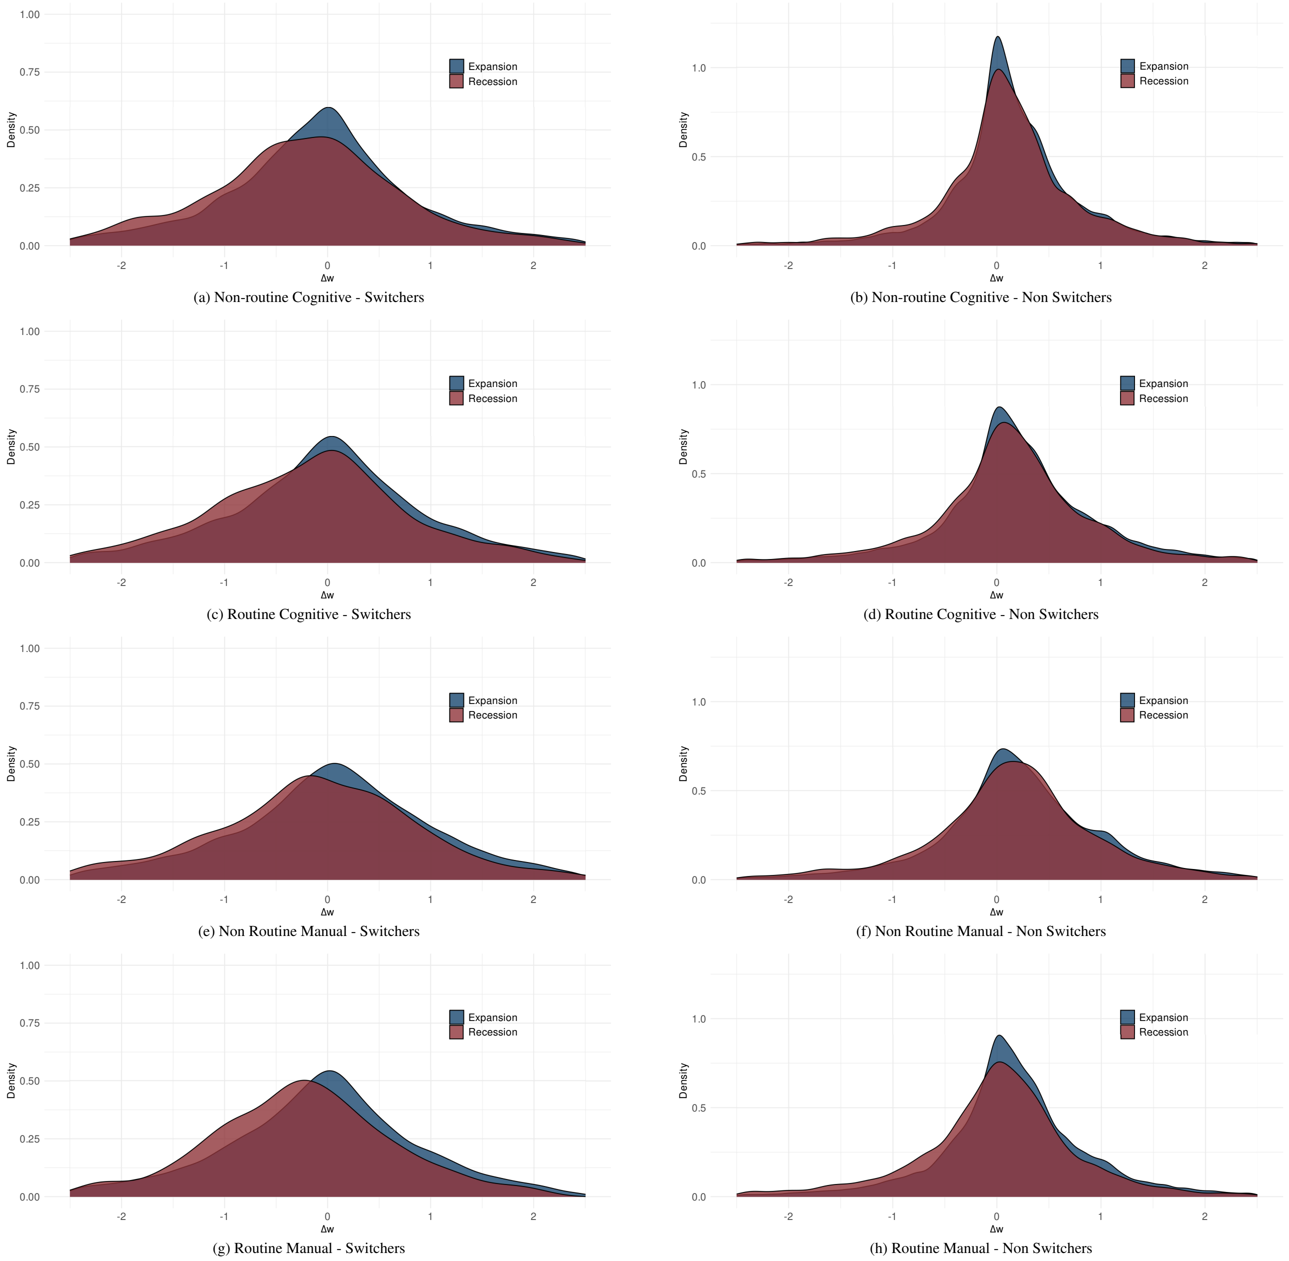

(a) Non-routine Cognitive - Switchers

(b) Non-routine Cognitive - Non Switchers

(c) Routine Cognitive - Switchers

(d) Routine Cognitive - Non Switchers

(e) Non Routine Manual - Switchers

(f) Non Routine Manual - Non Switchers

(g) Routine Manual - Switchers

(h) Routine Manual - Non Switchers

Note: Kernel density estimates of the distribution of annual earnings growth for workers changing employers with or without an occupation switch in expansion and recession periods. Annual earnings growth are based on residualized earnings, while recessions defined as periods in which the HP-filtered unemployment rate lies in the top 20% of realizations.

associated with the occupation at the point of re-employment. Label these changes Δv_o. Note that $\Delta v_o = 0$ when there is no occupation change, $\Delta v_o > 0$ when there is an occupation change and the worker moved to an occupation associated with a higher fixed effect, and $\Delta v_o < 0$ when there is an occupation change and the worker moved to an occupation associated with a lower fixed effect. Let $\tilde{G}(\Delta v_o|q)$ denote the distribution of changes in occupation-wide fixed effects associated with each earnings growth quantile q in the aggregate growth distribution $G(\Delta)$.

Figure 12 ranks occupation switchers by their earnings growth quintile (x-axis) and relates each of these workers' rank to the associated $\tilde{G}(\Delta v_o|q)$. For each rank we show the values of Δv_o at

Figure 12: $\tilde{G}(\Delta \upsilon_o|q)$ for different labor market transitions

(a) Employer movers: EUE

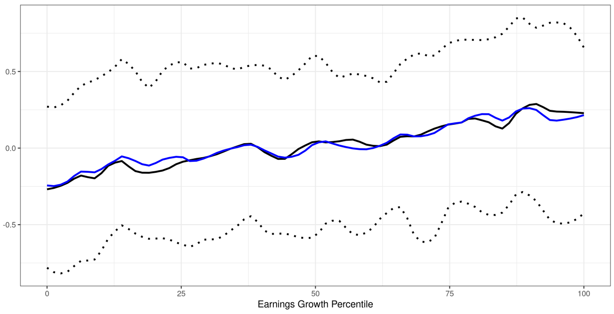

(b) Employer movers: EE

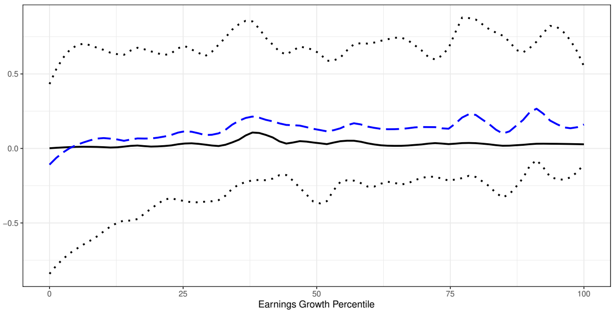

(c) All movers - comparison with the earnings growth CDF

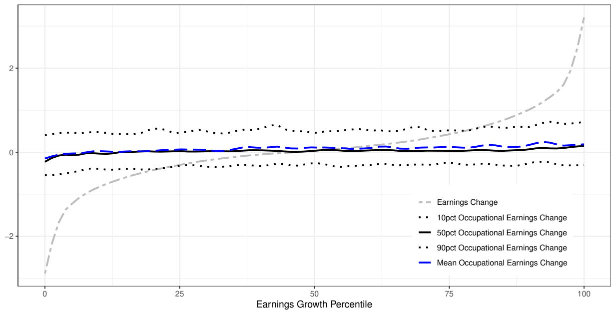

Note: Occupational switchers are ranked by their earnings growth in the horizontal axis. For each rank the vertical axis depicts the mean, median, 90^{th} and 10^{th} percentiles of the distribution of the differences in occupational earnings effects. Occupations defined at the 2-digit level of aggregation.

the mean, median, 90^{th} and 10^{th} percentiles of $\tilde{G}(\Delta \upsilon_o|q)$. The bottom and top curves show the sequence of values for $\Delta \upsilon_o$ associated with the 10^{th} and 90^{th} 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 labor market transition and using the 22 occupation categories of the two-digit 1990 SOC. Figure 12c shows that the magnitude of $\Delta \upsilon_o$ across $\tilde{G}(\Delta \upsilon_o|q)$ is small compared to the magnitude of earnings growth. It does so by comparing the various values of $\Delta \upsilon_o$ among all occupation movers across all q to the earnings changes at each quantile of the earnings growth distribution.[50] Figure 12a zooms in and shows the sequence of values of $\Delta \upsilon_o$ across $\tilde{G}(\Delta \upsilon_o|q)$ for EUE movers at the mean, median, 90^{th} and 10^{th} percentiles. We observe that these are all upward sloping, implying that workers with higher earnings growth are more likely to move to higher paying occupations. For EE movers, however, Figure 12b only show a noticeable upwards sloping pattern for the mean and the 10^{th} percentile. However, a key feature of these figures is that the upward sloping patterns have a very subdued effect on $\tilde{G}(\Delta \upsilon_o|q)$. That is, large earnings changes are associated with similar positive and negative values of $\Delta \upsilon_o$ across all types of labor market transitions.

Cyclical Shift of Changes in Occupation Fixed Effects along Earnings Growth Distribution To assess whether the cyclical pattern of $G(\Delta)$ is mirrored in $\Delta \upsilon_o$, Figure 13a plots again the mean values of $\tilde{G}(\Delta \upsilon_o|q)$ for each quantile q but now for all workers and separately for expansion and recession

[50] Note that this finding does not contradict that of Groes et al. (2015). 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.

Figure 13: Quantile functions of $\tilde{G}(\Delta v_o|q)$ and $\tilde{G}(\Delta v_o)$ in expansion and recession periods

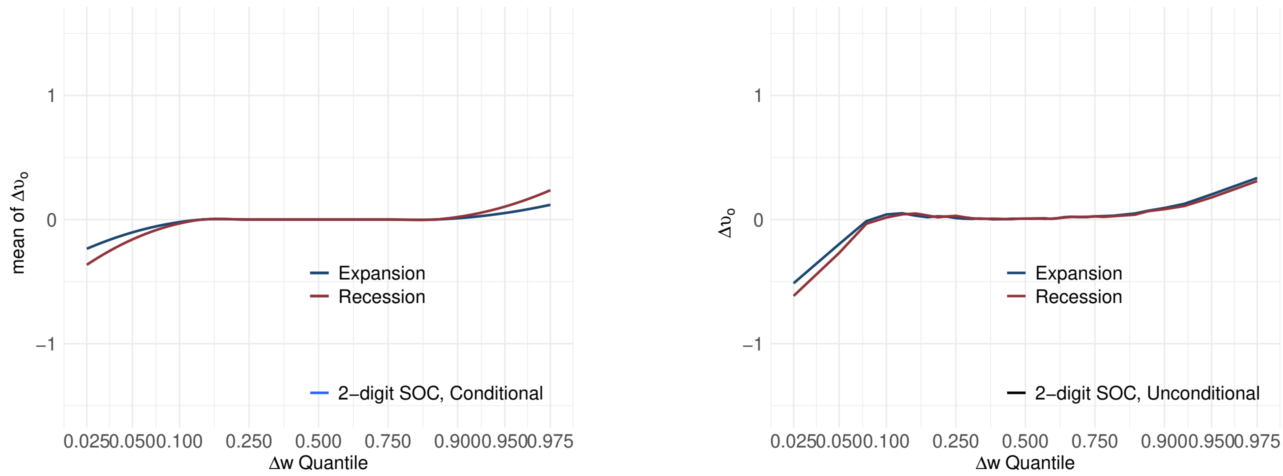

(a) Mean of Δv_o given quantile q of $G(\Delta)$

(b) Quantile function of Distribution Δv_o itself

Note: Panel (a) additionally plots the mean of the occupation-fixed-effect change Δv_o conditional on each quantile q of $G(\Delta)$, separately for expansion and recession. Panel (b) plots the unconditional quantile function of Δv_o separately for expansion and recession. Recessions are defined as periods in which the HP-filtered unemployment rate lies in the top 20% of realizations.

periods. The graph shows that these values hardly change between expansion and recessions periods between the 10th and 90th quantiles of the earnings growth distribution, but the means deteriorate at lower quantiles and *improve* at higher quantiles during recessions. Thus, these patterns imply that $var_q(E[\Delta v_o|q])$ increases noticeably during recessions, more so than a change in skewness.

Now consider the unconditional distribution of Δv_o, $\tilde{G}(\Delta v_o)$. Figure 13b plots the quantile function of $\tilde{G}(\Delta v_o)$ separately for expansions and recessions. Note that in this case, each quantile in the x-axis refers to the quantile of the distribution of occupation fixed effect changes $\tilde{G}(\Delta v_o)$ and *not* to a quantile of the distribution of earnings growth $G(\Delta)$. The figure shows that the quantile functions of $\tilde{G}(\Delta v_o)$ are also characterized by a flat plateau near zero with negative values in the bottom tail and positive values in the top tail. The recession curve lies below the expansion curve in the bottom tail and is very similar (nearly indistinguishable) to the expansion curve in the top tail. Given that the mean of $\tilde{G}(\Delta v_o)$ changes very little between expansion and recessions, the cyclical change in the quantile function of $\tilde{G}(\Delta v_o)$ implies that $var(\Delta v_o)$ increases in recessions. Noting that total variance $var(\Delta v_o) = E_q[var(\Delta v_o|q)] + var_q(E[\Delta v_o|q])$, the countercyclical increase in $var_q(E[\Delta v_o|q])$ thus contributes to the countercyclical increase in $var(\Delta v_o)$. As mentioned in the main text, the key message is not that Δv_o does not change between recessions and expansions, but that its cyclical variation contributes more the increase in the variance of the earnings growth distribution and less to the increase in its left-skewness.

It is important to highlight that one cannot directly compare the magnitude of the differences between the expansion and recessions curves depicted in Figures 13a and 13b with those of the quantile functions of the earnings growth distribution depicted in Figure 2 in the main text. This is because Figures 13a and 13b use Δv_o at the point of a labor market transition, while Figure 2 uses annual earnings growth. In the tails of the distribution of the latter, periods of nonemployment play an important role, during which there is no role of an occupation wage fixed effect. Thus, in addition to the cyclical differences in the quantities on the y-axis of Figure 13a being smaller than those on the y-axis in Figure 2d, differences in occupation wage fixed effects across the cycle will also only partially filter through to annual earnings growth when there are periods of nonemployment involved, hence their unimportance in Figure .

Cyclical worker flows and occupation identities

We complement the previous analysis of earnings changes and occupation identities by investigating the occupation worker flows across the four task-based occupation categories and how they change between expansion and recession periods. For each transition type (EE and EUE, measured at the wave level as described in Section 2 of the main text) and each occupation category, we compute the gross mobility rate defined as the share of transitions originating in the category that involve an occupation change, and the net mobility rate defined as the difference between gross inflows and gross outflows relative to the number of transitions originating in that category.

Figure 14 show the gross mobility rates of each occupation and how they changes between expansion and recession periods. This graph confirms that the procyclicality of the aggregate occupation mobility probability among EUE and EE transitions, reported in Table 4 of the main text, occurs across all four task-based occupations and is not driven by a subset of occupations. This implies that, for example, in expansions there is a higher probability that workers will leave their current (or pre-unemployment) occupation, irrespectively of the identity of the occupation.

Figure 14: Gross occupational mobility rates over the business cycle

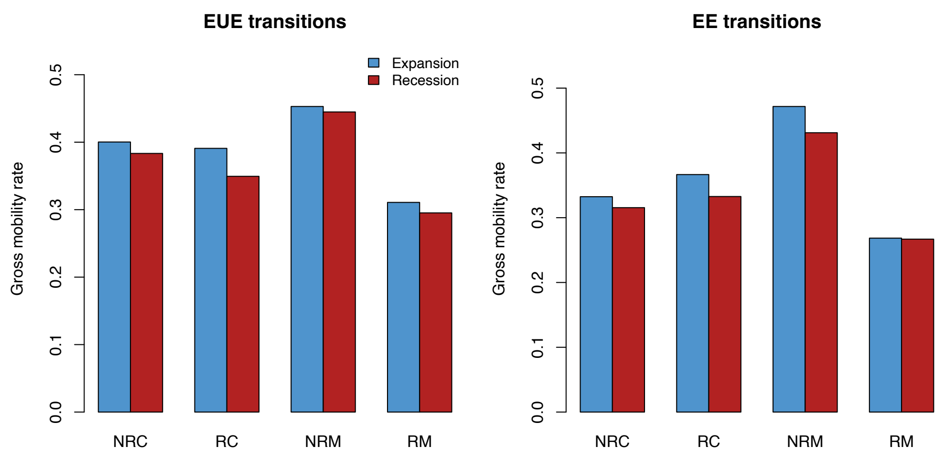

Note: Gross mobility rate of each task-based occupation category, defined as the share of transitions originating in the category that involve an occupation change, computed separately for EUE transitions (left panel) and EE transitions (right panel) and for expansion and recession periods. Transitions are measured at the wave level. Recessions are defined as periods in which the HP-filtered unemployment rate lies in the top 20% of realizations.

Figure 15 instead presents the net mobility rates and shows how they change between expansion and recession episodes. We find that the aggregate net mobility rate among EUE transitions increases in recessions with a recession to expansion ratio of 2.02. During recessions there are more net inflows into routine cognitive and non-routine manual occupations and more outflows among non-routine cognitive and routine manual occupations. Among EE transitions the net mobility rate decreases

Figure 15: Net occupational mobility rates over the business cycle

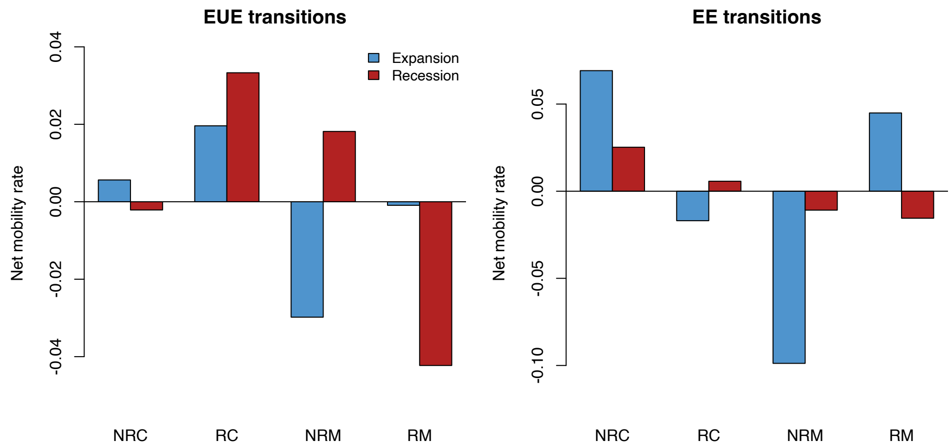

Note: Net mobility rate of each task-based occupation category, defined as gross inflows minus gross outflows relative to the number of transitions originating in the category, computed separately for EUE transitions (left panel) and EE transitions (right panel) and for expansion and recession periods. Transitions are measured at the wave level. Recessions are defined as periods in which the HP-filtered unemployment rate lies in the top 20% of realizations.

in recessions with an recession to expansion ratio of 0.26, where recessions go together with higher inflows into cognitive occupations and higher outflows from manual occupations.

Although we observe that recessions change net worker reallocation across occupations from the EE channel to the unemployment channel, net mobility is an order of magnitude smaller than gross mobility for nearly all categories and both transition types (EE and EUE): gross rates range between 27% and 47%, while most net rates are below 3% in absolute value. The main exception arises among EE transitions, where NRM occupations exhibit a net outflow rate of 9.9% and NRC occupations a net inflow rate of 6.3%. The vast majority of occupational mobility therefore reflects workers flowing simultaneously in both directions between occupation categories rather than directed net movements (excess reallocation). Although this conclusion is obtained based on a coarse occupation aggregation a similar conclusion emerges when considering the two-digit occupations. Together with the analysis of earnings changes and its relation to occupation identities performed in the previous subsection, our evidence suggests that the component of occupation mobility that drives excess reallocation is also the one that drives the cyclical skewness of the earnings growth distribution.

Appendix D — Estimation

Simulation procedure

The parameters to be estimated are those listed in Sections 4.1 and 4.2: the job-arrival set $\{\lambda_{r,i}, \lambda^c_{r,i}\}_{i=U,E,\ r=0,1}$; the worker-occupation set $\{\delta_{r,z}, \rho_z, \check{F}_{j,r}\}_{r=0,1}$; the worker-employer set $\{\delta_{r,\epsilon}, \gamma, \check{\Gamma}_{j,r}\}_{r=0,1}$; the occupation-wide set $\{\rho_p, \sigma_p, \tilde{p}_{NRC}, \tilde{p}_{RC}, \tilde{p}_{NRM}, \tilde{p}_{RM}\}$; the human-capital set $\{x_1, x_2, \chi(x_2)\}$; the search-direction set $\{\alpha_0, \alpha^U_1, \alpha^E_1, \alpha_{NRC}, \alpha_{RC}, \alpha_{NRM}, \alpha_{RM}\}$; the aggregate-productivity set $\{\rho_A, \sigma_A\}$; and the payment set $\{\Lambda, \eta, b\}$. The match-quality distributions $F(\cdot; A_I)$ and $\Gamma(\cdot; A_I)$ are estimated non-parametrically, replacing the parametric families used in earlier drafts: each is a histogram whose relative masses at the 10^{th}, 25^{th}, 75^{th} and 90^{th} percentiles — with the median mass held fixed — are the free parameters $\{\check{F}_{j,r}\}$ and $\{\check{\Gamma}_{j,r}\}$, estimated separately for expansion and recession ($r = 0, 1$) and interpolated onto a fine grid from which agents draw.

As mentioned in the main text we fix $\beta = 0.997$, normalize 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 (2009). 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 $\rho_A = 0.9580$ and $\sigma_A = 0.0090$. The idiosyncratic productivity grids are evenly spaced on $[\underline{z}, \bar{z}] = [\underline{\epsilon}, \bar{\epsilon}] = [-3, 3]$, based on the deciles of F and Γ.

We recover the remaining parameters by a two-step, inner/outer-loop procedure. 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 of the main text. This loop also contains the parameters 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. The latter 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's net inflow in some occupation is also counted as a net outflow in some other occupation. Hence, the relative contribution of occupation o on total net flows is given by $\frac{|Inflow_o - Outflow_o|}{\sum_{o \in O} |Inflow_o - Outflow_o|}$. We then iterate on the values of the remaining outer loop parameters, adjusting the inner loop parameters at each iteration. The outer loop parameters are obtained as the solution to

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

where $\mathbf{M^D}$ is a vector of data moments, $\mathbf{M^S}(.)$ 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 weighting matrix, which we take to be the identity matrix. We also normalize to the same scale the all cross-sectional and business cycle moments.

We solve the model by value function iteration using global methods. Local methods would be unsuitable because they truncate some of the variation in earnings changes in response to shocks (Petrosky-Nadeau and Zhang, 2021) — precisely the variation we study in the tails of the earnings growth distribution. Because the discrete occupation-mobility choice makes value functions intersect, it can create non-concave regions of the state space and long (or infinite) convergence times; we therefore smooth this choice with a steep logit.

We then draw 64 histories of 10,000 agents each over 13 years, parallelised to match a multiple of the cores on our cluster. Each history draws its own aggregate-shock path, with 20% of periods in recession on average and the cyclical auto-correlation preserved; simulated values lie on the shock grids rather than being interpolated. Workers are initialised at the ergodic distribution of x_h and occupations. We convert the simulated histories to the SIPP's four-month frequency and average across histories to compute the moments $\mathbf{M}^S(\cdot)$.

Search across occupations and the Gumbel-type shocks

In Section 4.1 we parameterized the probability of obtaining a z from a given occupation as $\alpha^i(s_{\tilde{o}}) = \alpha_0 e^{\alpha_{\tilde{o}}\alpha_1^i} s_{\tilde{o}}^{1-\alpha_1^i}$, where occupation $\tilde{o} \in O^-$ denotes the search direction, i the worker's labor force status, $i = U, E$, and $s_{\tilde{o}}$ denotes search intensity. Workers have to chose a $s_{\tilde{o}}$ for each $\tilde{o} \in O^-$ to maximise the probability of receiving a z given that $\sum_{\tilde{o}\in O^-} s_{\tilde{o}} = 1$. The first order condition for such maximisation is given by $\alpha'(s^*_{\tilde{o}})\Phi^i(\tilde{\Omega}_1) = \mu$, where μ denotes the multiplier on the constraint $\sum_{\tilde{o}\in O^-} s_{\tilde{o}} = 1$, $\tilde{\Omega}_1 = \{\tilde{z}, x_1, \tilde{o}, A, \mathcal{P}_O\}$, and $\Phi^i(\tilde{\Omega}_1)$ denotes the net conditional return to searching in direction $\tilde{o}$: the expected value of drawing a new z in occupation $\tilde{o}$ and searching there, net of the value of not reallocating. These per-direction returns are the building blocks of the reallocation values R^U and R^E of Section 3.2. Substituting the functional form for $\alpha(.)$ into the first-order condition, solving for $s^*_{\tilde{o}}$, and imposing $\sum_{\tilde{o}\in O^-} s_{\tilde{o}} = 1$ (which cancels the common factor $\big((1-\alpha_1^i)\alpha_0/\mu\big)^{1/\alpha_1^i}$), one obtains

$$s^*_{\tilde{o}} = \frac{e^{\alpha_{\tilde{o}} + \frac{1}{\alpha_1^i}\log\left(\Phi^i(\tilde{\Omega}_1)\right)}}{\sum_{\tilde{o}\in O^-} e^{\alpha_{\tilde{o}} + \frac{1}{\alpha_1^i}\log\left(\Phi^i(\tilde{\Omega}_1)\right)}}.$$

If the directional terms $\alpha_{\tilde{o}}$ are all equal, this reduces to the convenient form

$$s^*_{\tilde{o}} = \frac{e^{\frac{1}{\alpha_1^i}\log\left(\Phi^i(\tilde{\Omega}_1)\right)}}{\sum_{\tilde{o}\in O^-} e^{\frac{1}{\alpha_1^i}\log\left(\Phi^i(\tilde{\Omega}_1)\right)}}, \tag{20}$$

Equation (20) is convenient because it allows more flexibility in matching net mobility across occupations. It does so by breaking the symmetry that the random-utility model typically used to model occupational mobility imposes. It nonetheless has a direct random-utility counterpart. Suppose each occupation $o^j \in O^-$ carries a Gumbel-distributed shock ϖ^j with dispersion α_1 that enters *multiplicatively*, so that searching in direction o^j yields expected payoff $\Phi(\Omega_1^j)e^{\varpi^j}$ (we leave the index $i = U, E$ implicit).

The population share choosing o^j, $v_{o^j} = \Pr[\Phi(\Omega_1^j)e^{\varpi^j} > \Phi(\Omega_1^k)e^{\varpi^k}\ \forall k]$, then takes the standard closed form obtained by taking the monotone log-transformation and integrating over the shocks,

$$v_{o^j} = \frac{e^{\frac{1}{\alpha_1}\log\Phi(\Omega_1^j)}}{\sum_{k\in O^-} e^{\frac{1}{\alpha_1}\log\Phi(\Omega_1^k)}},$$

which is exactly (20); the directional terms $\alpha_{\tilde{o}}$ are recovered by scaling each return $\Phi(\Omega_1^k)$ by $e^{\alpha_{\tilde{o}}}$.

Finally, this technology nests the canonical on-the-job-search model with endogenous search intensity (Burdett, 1978), in which a worker chooses a scalar effort s to raise the arrival probability $\lambda(s) \leq 1$ of a single wage draw, with λ continuous, weakly increasing and weakly concave, subject to a convex search cost. Here the worker instead allocates one unit of search intensity across the O^- occupations, which share the offer distribution F, and receives a z with probability $\sum_{\tilde{o}\in O^-} \alpha(s_{\tilde{o}}; o) \leq 1$, where α has the same curvature properties in s and $\sum_{\tilde{o}\in O^-} s_{\tilde{o}} = 1$.

Bargaining protocol

In Section 4.1, equation (7), we propose an extension to the earnings equation used in the theoretical model in order to incorporate any additional persistence not accounted for by the productivity process. This equation approximates a bargaining setup combining Gottfries (2025) and Nagypal (2008). As in Gottfries (2025), renegotiation opportunities arrive with per-period probability $1-\eta > 0$ (independent of the job-offer arrival rates $\lambda_i(A)$ and $\lambda_i^c(A)$); if renegotiation does not occur, the worker is paid the previous period's earnings. When it does, and as in Nagypal (2008), the worker and firm set w by Nash bargaining during the production stage,

$$w_t = \underset{w}{\operatorname{argmax}} \Big[W_A(.) - W_D(.)\Big]^{\Lambda} \Big[J_A(.) - J_D(.)\Big]^{1-\Lambda},$$

W_i and J_i denote the expected values of agreement A and disagreement D for the worker and firm, respectively, and Λ is the exogenous worker bargaining power. If the worker and firm agree, production takes place and the worker receives w_t, while the firm receives $(A_t + \tilde{p}_o + p_{o,t} + x_h + z_t + \epsilon_t) - w_t$ at time t.

Unlike in DMP models, the disagreement payoffs are not the values of unemployment and a vacancy. Instead the match continues: during a disagreement the worker receives b and the firm zero, the worker keeps searching on the job, and all the shocks and stages of Section 3 unfold as usual; a disagreement at t also triggers a forced renegotiation at $t + 1$, regardless of the $1 - \eta$ process. This preserves Nagypal's (2008) threat-point logic — continuation values from $t + 1$ onwards coincide under agreement and disagreement up to a wedge of order η — and yields (7) as a Nash outcome. The worker's agreement and disagreement values are

$$\begin{aligned} W_i(\epsilon, x, z, o, \Omega) \;\; &= \;\; \hat{w}_i + \beta\mathbb{E}\Big[\delta_z(A')R^U(x', \underline{z}, o, \Omega') + \delta_\epsilon(A')W^U(x', z', o, \Omega') + (1-\delta_z(A')-\delta_\epsilon(A')) \\ &\times \;\; \max\Big\{W^U(x', z', o, \Omega'), \max\Big\{R_i^E(\epsilon', x', z', o, \Omega'), \hat{W}_i(\epsilon', x', z', o, \Omega')\Big\}\Big\}\Big], \end{aligned}$$

where $\hat{W}_i$ is the value of staying and searching within the same occupation, defined exactly as $\hat{W}^E$ in Section 3.2 but with the continuation value N_i in place of W^E in the no-meeting and offer-rejection branches; $\hat{w}_A = w_t$ and $\hat{w}_D = b$. Here $N_i(\epsilon, x, z, o, \Omega)$ is the worker's value of continuing with the firm in the same occupation, measured at the production stage: under agreement it carries w_t over with probability η and is renegotiated afresh with probability $1 - \eta$, whereas under disagreement the forced-renegotiation assumption makes the next period a renegotiation with probability one. We assume renegotiation also occurs whenever the worker changes occupation within the same employer.

The firm's agreement and disagreement values J_A and J_D are defined symmetrically, with per-period flow profit $\pi_A = (A_t + \tilde{p}_o + p_{o,t} + x_h + z_t + \epsilon_t) - w_t$ under agreement and $\pi_D = 0$ under disagreement, and the same continuation structure (the match is either terminated, with the firm obtaining the vacancy value $V(o, \Omega')$; retained with the worker switching occupations within the firm; retained in the same firm-occupation match; or lost to another firm).

The continuous-renegotiation case ($\eta = 0$) When $\eta = 0$, every period is a renegotiation event, so the worker's continuation value $N_i(.)$ at the production stage equals the expected value of bargaining afresh in the next period. This continuation value is identical under agreement and disagreement,

$N_A(.) = N_D(.)$, so the worker's and firm's contributions to the Nash product collapse to the per-period flow differences $w_t - b$ and $(A_t + \tilde{p}_o + p_{o,t} + x_h + z_t + \epsilon_t) - w_t$. The Nash outcome is then

$$w_t = \Lambda(A_t + \tilde{p}_o + p_{o,t} + x_h + z_t + \epsilon_t) + (1 - \Lambda)b,$$

which is exactly the per-period Nash earnings analyzed by Nagypal (2008). This is also the earnings formula consistent with the value functions in Section 3.1 of the main text, where past earnings do not enter the worker's or firm's state space precisely because under continuous renegotiation the earnings is reset every period.

The infrequent-renegotiation case ($\eta > 0$) When $\eta > 0$, the continuation values $N_A(.)$ and $N_D(.)$ generally differ. Under disagreement at t, the forced-renegotiation assumption ensures that the earnings at $t + 1$ equals the fresh Nash value $w_{\text{Nash}}(\Omega_{t+1})$. Under agreement at t, with probability $1 - \eta$ a scheduled renegotiation event at $t + 1$ likewise sets a fresh Nash earnings, but with probability η the earnings carry over to w_t. Denote by $V_R(\Omega_{t+1})$ the worker's continuation value at $t + 1$ under renegotiation and by $V_C(w_t, \Omega_{t+1})$ the worker's continuation value under carryover; analogously J_R and J_C for the firm. The worker's wedge in N between the agreement and disagreement branches is then $\eta\,[V_C(w_t, \Omega_{t+1}) - V_R(\Omega_{t+1})]$, and the firm's wedge is $\eta\,[J_C(w_t, \Omega_{t+1}) - J_R(\Omega_{t+1})]$. Substituting into the Nash first-order condition yields the bargained earnings

$$w_t = \Lambda y_t + (1-\Lambda)b + \beta\eta\left[\Lambda\big(J_C(w_t, \Omega_{t+1}) - J_R(\Omega_{t+1})\big) - (1-\Lambda)\big(V_C(w_t, \Omega_{t+1}) - V_R(\Omega_{t+1})\big)\right], \quad (21)$$

where $y_t \equiv A_t + \tilde{p}_o + p_{o,t} + x_h + z_t + \epsilon_t$ and the correction in square brackets is the deviation of the Nash earnings from equation (7) of Section 4.1.

Magnitude of the correction Under risk neutrality and the geometric series implied by the η-stickiness, the worker's and firm's wedges are linear in the difference between the carried-over wage and the fresh Nash wage. Hence the worker's wedge satisfies

$$V_C(w, \Omega_{t+1}) - V_R(\Omega_{t+1}) \;\approx\; \frac{w - w_{\text{Nash}}(\Omega_{t+1})}{1 - \beta\eta}, \quad J_C(w, \Omega_{t+1}) - J_R(\Omega_{t+1}) \;\approx\; -\frac{w - w_{\text{Nash}}(\Omega_{t+1})}{1 - \beta\eta}. \quad (22)$$

Taking expectations over Ω_{t+1} conditional on Ω_t, the magnitude of the correction depends on the productivity process driving y_t. When productivity is constant, $w_{\text{Nash}}(\Omega_{t+1}) = w_t$, the wedge vanishes, and (7) holds exactly. When y_t follows an AR(1) with persistence ρ and long-run mean $\bar{y}$, then $\mathbb{E}[w_{\text{Nash}}(\Omega_{t+1}) \mid \Omega_t] = \rho w_t + (1 - \rho)\bar{w}$ with $\bar{w} \equiv \Lambda\bar{y} + (1 - \Lambda)b$ (using the conjecture $w_t = \Lambda y_t + (1 - \Lambda)b$); substituting into (22) and (21) yields

$$w_t \;=\; \Lambda y_t + (1 - \Lambda)b \;-\; \beta\eta(1 - \rho)\left(w_t - \bar{w}\right) \;+\; O(\eta^2).$$

The deviation from $w_t = \Lambda y_t + (1 - \Lambda)b$ is thus of order $\eta(1 - \rho)$. With our estimated $\eta \approx 0.09$ and the highly persistent AR(1) productivity processes A_t, $p_{o,t}$, and z_t of Section 4.1, the approximation error in (7) is quantitatively small in our setting.

Bargaining and labor market tightness A potential limitation of using the expected value of waiting as the disagreement payoff in the bargaining protocol above is that the resulting earnings equations do not allow the state of the labor market to determine earnings dynamics, as would happen in the

DMP framework. To investigate the extent to which the earnings process in Section 4.1 is affected by this omission, we extend it by directly incorporating the role of labor market tightness θ, so that bargained earnings take the form

$$w_{o,t} = \begin{cases} w_{o,t-1} & \text{with probability } \eta, \\ \Lambda(A_t + \tilde{p}_o + p_{o,t} + x_h + z_t + \epsilon_t) + (1-\Lambda)b + k\Lambda\theta_t & \text{with probability } 1-\eta, \end{cases}$$

where k is the per-period vacancy-posting cost. We do not derive this equation from a bargaining protocol; it is a reduced-form device for the effect of tightness on cyclical earnings. Since our model takes meeting rates as exogenous and abstracts from firm entry, θ is not pinned down in equilibrium, so we recover it from an approach motivated by the DMP free-entry condition.

We use the model's expected values of employment W^E and unemployment W^U to back out the firm's value of a filled job, J, through the DMP-Nash sharing rule $(1-\Lambda)(W^E - W^U) = \Lambda J$. Aggregating across p_o, z, ϵ, x_h gives an economy-wide $J(A)$, so the free-entry condition $k/q(\theta) = J(A)$ with a Cobb–Douglas matching function $q(\theta) = m_o\theta^{\chi-1}$ yields $\theta(A_t) = (J(A_t)m_o/k)^{1/(1-\chi)}$, consistent with free entry and capturing the estimated meeting process through W^E and W^U.

Calibration To implement this approach we set k to 20% of average productivity (see Silva and Toledo, 2009) and $\Lambda = \chi = 0.5$, as is typical in DMP-style models. The matching efficiency parameter m_o is calibrated by minimising the distance between the average value of labor market tightness (computed from the ratio of vacancies to unemployment) in the data and the model-implied average, obtained from $\theta(A_t) = (J(A_t)m_o/k)^{1/(1-\chi)}$. We build a monthly V/U series for 1990–2013, the same period as in our SIPP sample, using the Barnichon (2010) composite Help-Wanted Index spliced to JOLTS (anchored at December 2000), over civilian unemployment. This procedure implies a value of $m_o = 0.0117$ given that the average values of labor market tightness in the data and model are 0.533. The value of θ in expansion periods is 0.602 in the model and 0.612 in the data; while in recession periods is 0.294 in the model and 0.259 in the data. Hence, the cyclicality of θ, measured as the ratio between its expansion to recession values is 2.05 in the model and 2.36 in the data.

Results Using the benchmark productivity processes and value of b, we find that adding labor market tightness leaves the earnings process quantitatively close to the benchmark of Section 4.1. Figure 16 plots, for both wage equations and the data, the cross-sectional earnings growth quantile function and the U-shape describing its cyclical difference: the two wage equations deliver very similar cross-sectional quantile functions and very similar cyclical changes.

The key feature underlying these results is that variation in z-productivities is by far the dominant driver of cyclical earnings dynamics, so additional cyclical variation in θ does not play a quantitatively important role in shaping the earnings growth distribution.

The relationship between occupational mobility and earnings growth

In Online Appendix C, Figure 7 we showed the probability of an occupational change and its relation with a given value of earnings growth quantiles. We find that in the data the tails of the annual earnings growth distribution are associated with a high probability of an occupation change. Small earnings changes centered around zero, however, are associated with a much smaller probability of

Figure 16: Earnings growth quantile function under alternative wage equations

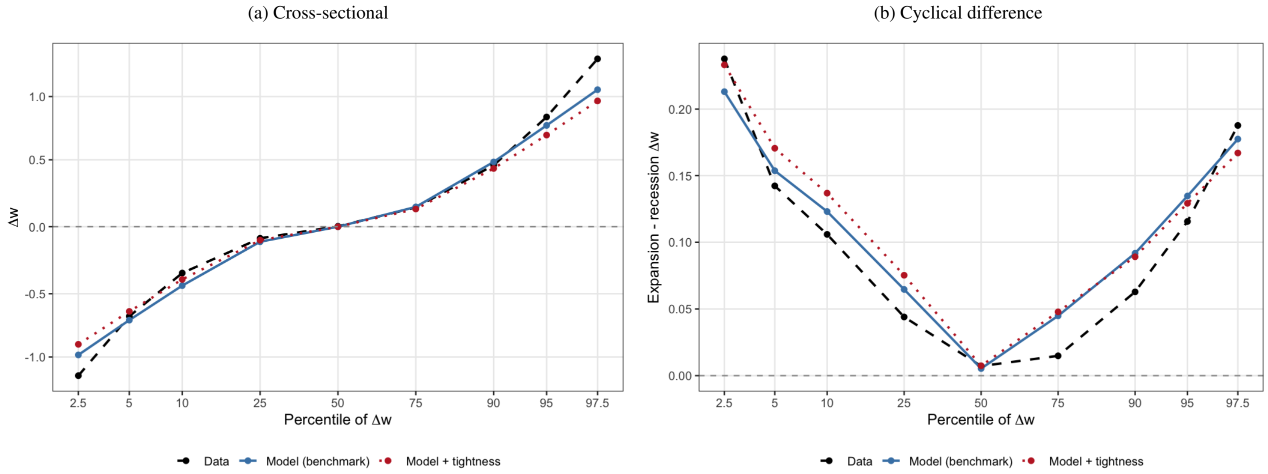

an occupation change. The importance of occupation mobility in the tails of the earnings growth distribution can be also observed in the model.

Figure 17: Earnings growth and occupation mobility - model vs data

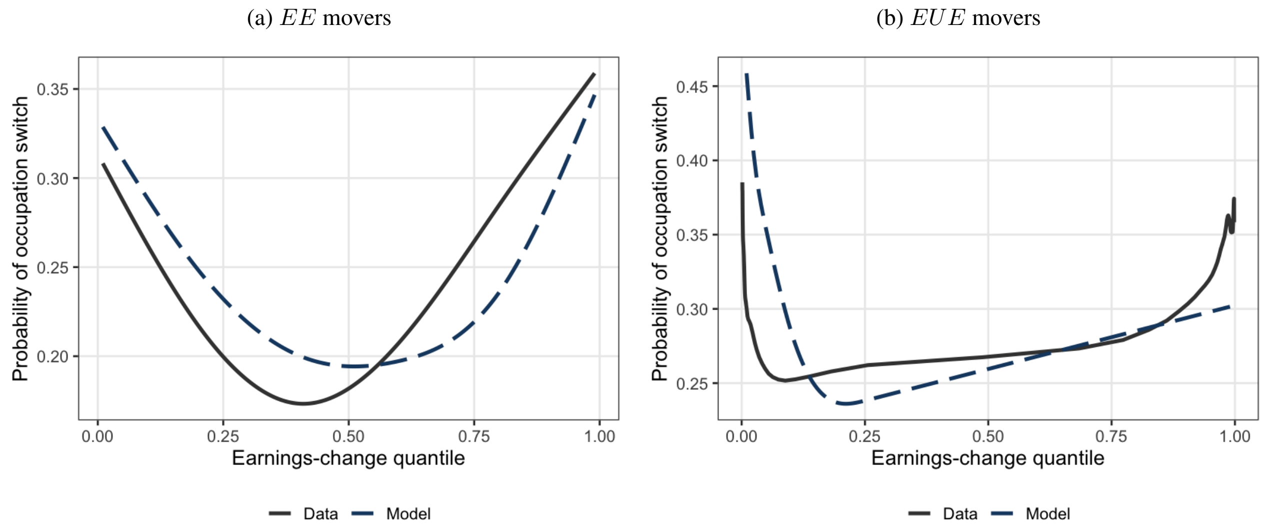

Figure 17 shows this relationship for EE and EUE movers. These relationships are untargeted and overall show very similar patterns to the data: occupational mobility contributes disproportionately to the largest changes in earnings. Among EE movers, the model generates a clear U-shape, but exhibits a higher occupational change probability below the first quartile relative to the data. Among EUE movers the model generates the same steep drop at the far end of the left tail and then a steady (nearly linear) increase as in the data. It misses, however, the steep (but short) uptick at the far end of the right tail.

References

[1] Autor, D. and D. Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market". *American Economic Review*, 103(5): 1553-1597.

[2] Barnichon, R. 2010. "Building a Composite Help-Wanted Index", *Economics Letters*, 109(3): 175-178.

[3] Busch, C., D. Domeij, F. Guvenen, and R. Madera. 2022. “Skewed Idiosyncratic Income Risk over the Business Cycle: Sources and Insurance”, *American Economic Journal: Macroeconomics*, 14(2): 207-242.

[4] Burdett, K. 1978. “A Theory of Employee Search and Quits”, *American Economic Review*, 68(1): 212-220.

[5] Burdett, K. and D. T. Mortensen. 1998. “Wage Differentials, Employer Size, and Unemployment”, *International Economic Review*, 39(2): 257-273.

[6] Carrillo-Tudela, C. and L. Visschers. 2023. “Unemployment and Endogenous Reallocations Over the Business Cycle”. *Econometrica*, 91(3): 1119-1153.

[7] Dorn, D. 2009. “Essays on Inequality, Spatial Interaction, and the Demand for Skills”. Dissertation, University of St. Gallen.

[8] Fujita, S. and G. Moscarini. 2017. “Recall and Unemployment”. *American Economic Review*, 107(12): 3875-3916.

[9] Halvorsen, E., H. A. Holter, S. Ozkan, and K. Storesletten. 2024. “Dissecting Idiosyncratic Earnings Risk,” *Journal of the European Economic Association*, 22(2): 617-668.

[10] Harris, J. (1992). Algebraic Geometry: A First Course, Graduate Texts in Mathematics 133, Springer.

[11] Hartshorne, R. (1977). Algebraic Geometry, Graduate Texts in Mathematics 52, Springer

[12] Hudomiet, P. 2015. “Career Interruptions and Measurement Error in Annual Earnings”. Mimeo, Department of Economics, University of Michigan.

[13] Hubmer, J. 2018. “The Job Ladder and Its Implications for Earnings Risk,” *Review of Economic Dynamics*, 29, 172-194.

[14] Gottfries, A. 2025. “Bargaining with Renegotiation in Models with On-the-Job Search”, *Review of Economic Dynamics*, forthcoming.

[15] Gottschalk, P. 2005. “Downward Nominal-Wage Flexibility: Real or Measurement Error?”.*The Review of Economics and Statistics*, 87(3): 556-568.

[16] Griffiths, P., and Harris, J. (1994). Principles of algebraic geometry. John Wiley & Sons.

[17] Groes, F., P. Kircher, and I. Manovskii. 2015. “The U-Shapes of Occupational Mobility,” *Review of Economic Studies*, 82(2): 659-692.

[18] 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.

[19] Guvenen, F., S. Ozkan, and J. Song. 2014. “The Nature of Countercyclical Income Risk,” *Journal of Political Economy*, 122(3): 621-660.

[20] Kambourov, G. and I. Manovskii. 2008. “Rising Occupational and Industry Mobility in the United States: 1968-97”. *International Economic Review*, 49(1): 41-79.

[21] Kambourov, G. and I. Manovskii. 2009. “Occupational Specificity of Human Capital”. *International Economic Review*, 50(1): 63-115.

[22] Kurmann, A., and E. McEntarfer. 2019. “Downward Wage Rigidity in the United States: New Evidence From Administrative Data,” Census Working Paper No. CES-19-07.

[23] Moscarini, G. and K. Thomsson. 2007. “Occupational and Job Mobility in the US”. *Scandinavian Journal of Economics*, 109(4): 807-836.

[24] Nagypál, É. 2008. “Worker Reallocation over the Business Cycle: The Importance of Employer-to-Employer Transitions”. Mimeo, Department of Economics, Northwestern University.

[25] Okamoto M. (1973). Distinctness of the eigenvalues of a quadratic form in a multivariate sample. The Annals of Statistics, 1(4), 763–765. https://doi.org/10.1214/aos/1176342472

[26] Ozkan, S., J. Song and F. Karahan. 2023. “Anatomy of Lifetime Earnings Inequality: Heterogeneity in Job-Ladder Risk versus Human Capital”, *Journal of Political Economy Macroeconomics*, 1(3): 506-550.

[27] Petrosky-Nadeau, N. and L. Zhang. 2021. “Unemployment Crises”. *Journal of Monetary Economics*, 117: 335-353.

[28] Rothenberg, T. J. 1971. “Identification in Parametric Models”, *Econometrica*, 39(3): 577-591.

[29] Shimer, R. 2005. “The Cyclical Behavior of Equilibrium Unemployment and Vacancies”. *American Economic Review*, 95(1): 25-49.

[30] Silva, J. I. and M. Toledo. 2009. “Labor Turnover Costs and the Cyclical Behavior of Vacancies and Unemployment”, *Macroeconomic Dynamics*, 13(S1): 76-96.