

**What is stopping you?
The falling employment-to-employment mobility in the UK**

See-Yu Chan

April 2024

No: 1496

Warwick Economics Research Papers

ISSN 2059-4283 (online)

ISSN 0083-7350 (print)

What is stopping you? The falling employment-to-employment mobility in the UK.

See-Yu Chan

University of Warwick

April 24, 2024

Abstract

What contributed to the decline in employment-to-employment (EE) transition rate in the UK in recent decades? This paper empirically examines potential channels that caused the sluggish EE mobility from 2000-2019. First, I break down the observed fall in EE mobility relative to unemployed-to-employment (UE) transition into changes in relative search intensity and worker's acceptance rate. I find the vast majority of the persistent decline after 2010 was due to fall in job acceptance. Second, I estimate a dynamic job ladder model using UK survey data to examine the relative importance of changes in employment and job offer distribution in reducing job acceptance. Results reveal that the falling job acceptance in the 2000s was attributed to workers moving up the job ladder; while acceptance remained low in the 2010s as a result of deterioration in offer qualities. Counterfactual exercise shows that if the "attractiveness" of poaching offers did not deteriorate after 2010, the EE mobility would have returned to levels in early 2000s. Finally, I test the contribution of composition changes to the fall in EE rate by implementing a between-within decomposition using a structural framework, which accounts for both worker heterogeneity and sectoral compositions. Results rule out demographic changes or structural transformation as main drivers of the fall in EE rates.

Keywords: employment-to-employment transition; job ladder; offer acceptance rate; job offer distribution.

JEL Codes: E20, E24, J62, J63

1 Introduction

The employment-to-employment (EE) transition, a process in which employees switch to a new job without go through a spell of unemployment, is a key component of labour market dynamics in the UK. In terms of total counts, there are at least 50 percent more employed workers switching jobs than people getting hired out of unemployment in a given quarter, according to the UK Labour Force Survey (UKLFS) from 1997 to 2019. However, the UK EE transition rate exhibited a persistent fall in level since 2000, as shown in figure 1. Specifically, the quarterly rate declined from an average level of 3.2 percent in late 1990s to less than 2.5 percent in late 2010s. This represents an decline of 24 percent in the average level.

While EE transitions take up the largest share of reallocation of employment, the EE rate is economically important because of the following reasons. First, frequent poaching and hiring of workers across firms are essential for productivity-enhancing reallocation of human resources (Davis and Haltiwanger, 1992). A slowdown of labour reallocation across firms would potentially lead to slower firm growth, either in firm sizes or productivity. Second, from individual worker’s perspective, EE transitions provide them with an opportunity to renegotiate their contract and subsequently generate wage growth (Postel-Vinay and Robin, 2002). A slowdown in such process could result in stagnant wages. From a macroeconomic point of view, due to the strong associations between EE rate and wage growth – while UE rate has comparatively less explanatory power for wage growth – understanding what drives the EE rate can provide macroeconomic implications on inflation and monetary policy (Moscarini and Postel-Vinay, 2017). Finally, persistent fall in EE transition is not a unique phenomenon to the UK only. Similar falls are also observed in other developed countries, such as the United States and Australia (Bagga, 2022; Deutscher, 2019).

This paper aims to shed lights on how changes in the labour market can lead to this persistent decline in UK EE transition rates from 2000 to 2019. I first investigate the changes in search behaviour of employed workers. To study this, we need a benchmark for comparison and the unemployed workers evolve as a suitable candidate. While the EE rate exhibited a persistent decline in past decades, the unemployment-to-employment (UE) transition rate were fluctuating within a range without any persistent changes in levels, as shown in the lower panel of figure 1. Hence, understanding the decline in the EE-UE ratio, or the relative EE rate, should be equivalent to understanding the decline in EE rate itself. Intuitively, employed workers can switch jobs at a lower frequency relative to the unemployed along two dimensions. First, employed workers might be reducing their search efforts over time relative to the unemployed. For instance, they can submit less job applications or attending less job fairs. Alternatively, employed workers might become increasingly picky on their job offers




Figure 1: UK Labour Market Dynamics

Note: Series computed using the UK Labour Force Survey. A EE transition is counted when an worker was employed at the beginning of the quarter as well as the at the end, but the current employment spell is less than three months by the end of the quarter. EE rate displayed in the upper panel is the total number of recorded EE transition over the average number of employment during the quarter. An UE transition is recorded when a worker who was unemployed to start the quarter became employed by the end of that quarter in the UKLFS. UE rate displayed in the lower panel is computed as the UE transition counts over the number of initial unemployed workers. Shaded area indicates recession periods in the UK.

over this period. In other words, they could be sending the same number of job applications, but either the job offers received were not attractive enough or the offer they were looking for did not materialise. One can examine the relative contribution of “search” and “acceptance” by decomposing the observed EE-UE ratio into these two components.

Since worker’s decision to search on the job is endogenous to their expected likelihood to accept a job offer, one needs to control for worker’s search decision before inferring their acceptance rate from the EE-UE ratio, or vice versa. To do this, I utilize the additional information in UKLFS to identify active on-the-job searchers in the sample. By focusing on workers who explicitly state that they are searching on the job, their decision to search — the extensive margin — is predetermined. In addition, by further controlling the worker’s search efforts with information from the UK time-use survey, I compute changes in the job acceptance rate from variations in the ratio of transition rates between active searchers and the unemployed. Subsequently, the contribution of relative search efficiency can be backed out from the observed EE-UE ratio. Results show that worker’s job acceptance probability




Figure 2: EE-UE ratio

Note: Series computed using the UK Labour Force Survey. The EE-UE ratio is computed by dividing the quarterly EE transition rate with the quarterly UE transition rate displayed in the figure 1.

exhibits a persistent fall from around 55 percent in late 1990s to only about 45 percent in late 2010s. Quantitatively, this explains roughly 80 percent of the difference in levels of relative EE rate between 2000 and 2019. If one looks at the two components' relative importance over time, reduction in search and acceptance each contributes half of the decline in the relative EE rates in 2000s, while continuous fall in job acceptance causes the EE rate to be remained low after 2010.

Why are employees less likely to accept a job offer when they receive one? There can be two possible reasons. It could be the case that workers were increasingly well-matched to jobs at the top of the job ladder. In other words, it could be an enhancement in workers' allocation across jobs. Otherwise, acceptance could fall if workers were receiving less "attractive" offers over time while they stay at the same position on the job ladder. Specifically, this would imply the sampling probability of offers from higher-ranking jobs reduced in past decades. To evaluate the relative importance of these two channels, I quantitatively estimate a dynamic job ladder model of the UK labour market as in Moscarini and Postel-Vinay (2016a). This framework allows me to infer the job offer distribution over time from the law of motion of the employment (or wage) distribution. To generate a common job ladder for the whole economy, I specify the job ladder using residualised wages — residuals obtained through

running a Mincer-type regression on worker’s observed characteristics, industries, regions and occupations. Results show that the fall in job acceptance should be examine in two separated 10-year periods. For the first 10 years before 2010, the job acceptance declined as the employment distribution was more concentrated at the top of job ladder. However, from 2010 to 2019, acceptance remains low due to deterioration of the job offer distribution. As a robustness check, I estimate job offer sampling distribution using residualised real wages with different specifications in the regression. The results are robust and it indicates that the “offer” channel has a major role in the persistent decline in job acceptance after 2010.

There can be other reasons that cause this phenomenon. One of the main alternative explanations is composition changes. For example, an ageing labour force may reduce EE rates as older workers are, on average, less likely to change jobs than younger workers. It could also be a result of structural changes in the UK economy towards professional services, where average recruitment process is longer than, say, hiring a catering worker in a restaurant. Also, migration of job seekers towards geographical regions in which job openings are relatively scarce would reduce EE mobility. In other words, EE rate can decline without any changes in individual’s search behaviour, but only due to shifts in compositions of the economy.

I conduct some quantitative exercises to evaluate empirical importance of composition shifts. I investigate this using a structural framework inspired by Barnichon and Figura (2015). This framework allows for a between-within decomposition of the residualised EE rate — EE series after controlling for fluctuations associated with aggregate number of vacancies — that jointly account for composition changes in multiple dimensions, including various individual characteristics, migrations and sectoral structural changes. Specifically, I examine the contribution of four main components on the decline in EE rate: 1) worker composition effect, which accounts for changes in average characteristics of jobseekers; 2) job composition effect, which measures shifts in relative importance of industrial sectors in the economy; 3) dispersion effects, which measures how efficiently vacancies and search effort are allocated across sectors; and 4) within-sector variations. I apply this framework to data in the UK Labour Force Survey and Vacancy Survey using maximum likelihood estimation. Results suggest that worker and job composition effects, as well as dispersion effects, play only minimal roles in reducing EE transition rates. In other words, most of the fall occurs within sectors, i.e. changes in individual search behaviour. This exercise provides us with important negative results that rule out mechanisms associated with composition changes. For instance, we now know that ageing labour force or structural change towards professional services are not the main reasons for falling EE rates.

The empirical results highlighted one main reason for the persistent decline in the EE rate:

a deterioration in the sampling probability of high-quality poaching offers after 2010. But why did these better job offers disappear? While EE mobility was persistently low in 2010s, the UK economy was also suffering from a period of substantially low productivity growth after the Great Recession. Goodridge et al. (2018) finds that such decline in productivity was concentrated among the most productive firms. Theoretically speaking, as firms grew slower in their productivity, the poaching firms could be constrained in making offers “attractive” enough to lure workers away from other firms. Hence, with the UK leading firms growing at a much slower pace in terms of their productivity, average offer “quality” deteriorated and EE rate declined.

Related Literature. This paper contributes by examining empirical importance of different mechanisms to the persistent drop in the job mobility. Mercan (2017), Bagga (2022) and Baksy et al. (2024) study the fall in EE mobility in the US. This paper first presents new empirical facts on worker’s job switching behaviour in the UK. Earlier studies utilises time-use surveys to documents facts associated with on-the-job search (Aguiar et al., 2013; Mukoyama et al., 2018). Recently, Faberman et al. (2022) develops a comprehensive survey to collect detailed information on on-the-job search in the US. With this new information, researchers find that the elasticity of search effort is highly sensitive to current wages of the workers. In this paper, I introduce a less data-demanding method to disentangle the relative importance of search efficiency and job acceptance in driving the EE-UE ratio, under reasonable assumptions.

There is a large literature which examines the interaction between EE transition and macroeconomic outcomes using job ladder models. Burdett and Mortensen (1998) and Postel-Vinay and Robin (2002), among others, document the importance of EE transitions to wage dispersions. In addition, various studies provide evidence on the close positive associations between EE transitions and wage growth (e.g. Topel and Ward, 1992; Postel-Vinay and Robin, 2002; Karahan et al., 2017; Moscarini and Postel-Vinay, 2017; Faberman and Justiniano, 2015). In a recent paper, Moscarini and Postel-Vinay (2023) suggests that the EE-UE ratio and job acceptance are more relevant measures than the unemployment rate in predicting inflationary pressure. This paper builds on the literature of job ladder model (Burdett and Mortensen, 1998; Moscarini and Postel-Vinay, 2016b), and provides new empirical evidence on how job acceptance can be affected along two channels of employment and job offer qualities.

Besides, a large literature delivers facts on how labour market dynamic changes over time. Fallick and Fleischman (2004) uses the US survey data to estimate monthly EE transition rate in recent decades. Fujita et al. (2020) builds on this method and adjusts for missing

answers in the survey. They find that the EE transition rate in the US did not exhibit a downward trend after the Great Recession once researchers adjust for proxy respondents. In the UK context, Gomes (2012), Smith (2011) and Postel-Vinay et al. (2019) present features of labour market dynamics using various datasets. Barnichon and Figura (2015) and Hall and Schulhofer-Wohl (2018) introduce structural frameworks to study how worker heterogeneity and market segmentation can affect labour market dynamics, with the main focus on transition out of unemployment. This paper builds on these frameworks to study how worker characteristics and job compositions can affect EE transitions.

Regarding the root cause of declining EE mobility, Bagga (2022) and Baksy et al. (2024) study the declining EE mobility in the US and they suggest that it was mainly due to the decline in market competition — measured as the fall in firms per workers. Although it is suggested that the deterioration of job quality after 2010 was related to the slowdown of productivity growth, what caused the declined growth rate is beyond the scope of this paper. One potential reason could be because leading firms had gained enough market power and thus reduced their innovation effort to shield competitions in their sectors. Hence, it is not suggested that the potential channel through increasing labor market power is ruled out.

The rest of this paper proceeds as below. In section 2, I analyse the relative contribution of search and job acceptance in reducing the EE-UE ratio. Section 3 studies what caused changes in job acceptance in the UK. Section 4 provides a discussion on the other potential mechanisms, including vacancy creation and composition changes, in affect the EE rate. Section 5 concludes.

2 Search versus Acceptance

Figure 1 shows that a persistent decline in the transition rate only occurred among the employed workers but not with the unemployed. This enables us to investigate the search behaviour of employees with the unemployed workers as a benchmark. Why do we see a persistent decline in the probability of job switching but not in the job finding rate of the unemployed? There are two potential reasons behind this observation. First, it could be the case that employed workers were reducing their search efforts relative to the unemployed. Second, it could be because employed workers were becoming more picky in accepting a poaching offer relative to the unemployed.

To characterise these concepts, I write down a simple random search model with on-the-job search. Suppose all workers sample offers from a common and exogenous offer distribution F which ranks from zero to one. Without loss of generality, I normalise the search effort of unemployed workers to one, and denote s_t the relative search intensity of employed workers.

Employed and unemployed job seekers sample offers at rate λ_t^e and λ_t^u respectively, given their search efforts.

To simplify notation, I assume unemployed workers accept all job offers with probability one. Employed workers only move to another job that is ranked higher than their current job on the job ladder. Hence the average job acceptance rate of employed workers can be expressed as $\mathcal{A}_t = \int_0^1 (1 - F_t(x)) dN_t(x)$, where $N_t(x)$ is the share of employed workers hired in firms up to rank x .

With this setting, the EE-UE ratio (or the relative EE rate) at time t can be expressed as

$$\frac{ee_t}{ue_t} = \frac{s_t \lambda_t^e \int_0^1 (1 - F_t(x)) dN_t(x)}{\lambda_t^u} = \underbrace{\frac{s_t \lambda_t^e}{\lambda_t^u}}_{\text{Relative search efficiency}} \underbrace{\int_0^1 (1 - F_t(x)) dN_t(x)}_{\text{Acceptance Rate}} = \phi_t \mathcal{A}_t \quad (1)$$

To further simplify notations, denote $\phi_t = s_t \lambda_t^e / \lambda_t^u$ the relative search efficiency. Equation 1 provides an analytical expression in which the EE-UE ratio can fall. First, it can decline through the relative search efficiency ϕ_t . This happens either by a reduction in worker's effort to search on-the-job, or an increased duration to receive an offer given the search intensity. Second, there can be a drop in job acceptance rate \mathcal{A}_t . To examine the relative contribution to the decline in the EE-UE ratio, one needs to obtain an empirical counterpart for either ϕ_t or \mathcal{A}_t from data.

The EE and UE transition rates can be obtained from the UK Labour Force Survey. The UKLFS is a quarterly survey which provides a representative sample of the whole UK working population. The survey follows individuals in a household for five consecutive quarters while one fifth of them are replaced every quarter. In the main analysis, the sample is restricted to contain only working population with age between 16 and 65. An EE transition is recorded when an individual is employed in both consecutive quarters, but her current employment spell is shorter than three months. Quarterly EE rate is then the fraction of individuals who recorded an EE transition in the average employed population in the quarter. As a limitation to the quarterly LFS, EE transitions also include those labour market transitions where an employed individual has a short unemployment spell within the quarter. Similarly, the UE transition is recorded when someone was unemployed to begin the quarter and declared employed by the end of that quarter. The UE rate is computed as the ratio of UE transitions to the number of unemployed workers at the beginning of the quarter.

Examining the EE-UE ratio instead of the raw EE rate allows one to control for aggregate shocks that commonly affect both employed and unemployed job seekers. For instance, a reduction in total number of vacancies in the market would reduce both EE rate and UE rate

separately, but it would not cause a fall in the EE-UE ratio in a random search framework. In addition, if both employed and unemployed workers look for jobs with a common search technology, an deterioration in the search technology will reduce each transition rate but not the EE-UE ratio.

Although I assume in equation 1 that unemployed workers accept all offers drawn from offer distributions $F(\cdot)$, this is only to simplify the mathematical expression of \mathcal{A}_t . Equation 6 can also include an additional integral in the denominator which represents the average sampling probability for an offer to be above the reservation wage of unemployed workers. In that case, \mathcal{A}_t would be interpreted as the relative job acceptance rate of the employed to that of the unemployed upon receiving an offer. One can also think of this as the offer endogenous adjusts such that job at rank zero always delivers the reservation wage to unemployed workers, as in Burdett and Mortensen (1998). Again, having this assumption is mainly to simplify interpretation and is not critical to any of the main results.

2.1 Estimation of job acceptance rate

Direct measures of acceptance rate requires information on number of offers that workers receives before accepting one. Unfortunately, this information is not readily available for the 20-year period in the UK. As an alternative, the UKLFS explicitly asks employed workers whether they are currently searching for a new job to replace their current one. This provides potentially two ways to control for ϕ_t and subsequently compute \mathcal{A}_t from the observed EE-UE ratio.

A straightforward way to utilize this piece of information is to compute the proportion of workers who declared that they are searching on-the-job among all employed workers. This measures the extensive margin of on-the-job search effort. One can then assume that the employed and unemployed share the same matching technology so they receive job offers at the same rate, i.e. $\lambda_t^e = \lambda_t^u$, given their search efforts. This allows one to compute ϕ_t and \mathcal{A}_t as

$$\phi_t = \frac{s_t \lambda_t^e}{\lambda_t^u} = \frac{OJS_{t-1}}{E_{t-1}} \Rightarrow \mathcal{A}_t = \frac{ee_t}{ue_t} \frac{OJS_{t-1}}{E_{t-1}} \quad (2)$$

correspondingly.

However, there are several issues with this approach other than the intensive margin of job search is ignored. First, the decision to search on-the-job is endogenous to the expected probability of offer acceptance. Employed workers would decide not to search if they expected to only receive offers that they would reject. Since the search effort is positively associated with job acceptance, contribution of \mathcal{A}_t would be underestimated if ϕ_t is measured using equation 2. Second, this method suffers from time-aggregation bias in the survey data.

Table 1: Average daily hours on job-seeking.

	2000	2014
In employment	1.21 (.241)	2.06 (.483)
Unemployed	1.02 (.182)	2.28 (.351)
Difference	-.034 (.232)	0.22 (.422)

Note: Standard error in parenthesis. Values state the mean number of daily hours on job seeking activities in UK Time-Use Survey 2000-01 and 2014-15.

If one looks at the composition of realised EE transitions, only about half of all switches were coming from workers who declared that they were searching in the previous quarter. Hence, this approach would severely underestimate the level of search intensity and thus overestimate the acceptance rate. Subsequently, the relative contribution of \mathcal{A}_t to the changes in EE-UE ratio would be again underestimated.

2.1.1 Step 1: Job acceptance for active seekers

A more reliable approach to utilize this information is to estimate average job acceptance rate in two steps. I first estimate job acceptance of active seekers \mathcal{A}_t^s by comparing their transition rates with the UE rate directly. After that, I estimate acceptance of non-active employees \mathcal{A}_t^{ns} and obtain the average acceptance of all employed workers as a weighted average.

I first focus on estimating \mathcal{A}_t^s because active seekers' decision to search is predetermined. This mitigated the endogeneity issue of the search decision. To control for the intensive margin of search, I compare the average daily hours that an active searcher spends on job seeking activities relative to the unemployed using the UK Time-Use Survey in 2000-01 and 2014-15, as shown in table 1. The statistics shows that employed and unemployed workers spent very similar duration in a day on job searching, given that they had spent time doing it on the day. The difference in search duration in a given survey year is statistically indifferent. In other words, the intensive margin of job search for active on-the-job seeker is similar to that of unemployed workers. Essentially, this suggests that the relative search intensity of the active seekers s_t^s is close to unity in the UK.

With search efforts accounted for, if job contact rate of active on-the-job seekers λ_t^s is homothetic to the offer rate of unemployed λ_t^u , the proportional changes in ratio between




Figure 3: Job acceptance rate of active seekers

Note: Series is estimated using the ratio in transition rates between active on-the-job seekers and unemployed workers using the UKLFS. An active seeker is defined as an employee who declares to be searching for a job to replace current position. Shaded area indicates recession periods in the UK.

their job transition rate will infer the movement in \mathcal{A}_t^s across time:

$$\frac{ee_t^s}{ue_t} = \frac{\lambda_t^s}{\lambda_t^u} \mathcal{A}_t^s \Rightarrow \Delta \frac{ee_t^s}{ue_t} = \Delta \mathcal{A}_t^s \text{ if } \lambda_t^s = \xi \lambda_t^u \text{ for some constant } \xi > 0 \quad (3)$$

where ee_t^s is the EE rate for active job searchers. Intuitively, if active on-the-job searchers received a constant fraction ξ of offers relative to the unemployed over this twenty-year period, the changes in their relative transition rates would be solely a result of changes in \mathcal{A}_t^s .

If one takes a step further and assumes active seekers and unemployed workers search for jobs with the same technology, i.e. $\xi = 1$ and thus $\phi_t^s = 1$, then the ratio of ee_t^s to ue_t will give the levels of \mathcal{A}_t^s . This assumption has some empirical support from Faberman et al. (2022), in which they find that ratio of offer arrival rate for active job seeker relative to the unemployed, conditional on their search effort, is around 1.3¹. The estimated \mathcal{A}_t^s is shown in figure 3. The series demonstrates a clear secular fall since late 1990s. Specifically, active seekers would accept around 55 to 60 percent of new offers when she received one in late 1990s. Yet, this probability fell to around 45 percent by late 2010s.

¹In table IV of Faberman et al. (2022), the mean offer received by an active searcher to replace current job is 0.258 (standard error: 0.024) and the mean offer received by an unemployed worker in a given month is 0.666 (0.286). Meanwhile in table II, average number of application sent for active seeker is 3.06 (0.29); while unemployed submit 10.39 (1.37) number of job applications.




Figure 4: Comparing wages of active seekers vs. non-active employees.

Note: The left panel shows the time series of mean real wages (at 2020 value) of active seekers vs. non-active employees. The right panel presents the proportional difference in real wages between active seekers and non-active employees. Shaded area indicates recession periods in the UK.

2.1.2 Step 2: Average job acceptance of all employed workers

After obtaining a time series of \mathcal{A}_t^s , the next step is to compute \mathcal{A}_t as the weighted average of job acceptance for active seekers \mathcal{A}_t^s and the non-active employees \mathcal{A}_t^{ns} :

$$\mathcal{A}_t = \frac{OJS_t}{E_t} \mathcal{A}_t^s + \frac{NS_t}{E_t} \mathcal{A}_t^{ns} \quad (4)$$

where OJS_t and NS_t is the number of active seekers and non-active searchers in period t , respectively. The next step now is to retrieve \mathcal{A}_t^{ns} .

Since we don't have any information on the search behaviour of non-active employees apart from their transition rate, we need to make further assumptions to retrieve their job acceptance \mathcal{A}_t^{ns} . Specifically, since non-active to active transition ratio can be expressed as

$$\frac{EE_t^{ns}}{EE_t^s} = \frac{\phi_t^{ns} \mathcal{A}_t^{ns}}{\phi_t^s \mathcal{A}_t^s},$$

we need two assumptions for ϕ_t^{ns} and \mathcal{A}_t^{ns} correspondingly.

Since \mathcal{A}_t^s is retrieved in step 1, I first check if the wages of active seekers evolves in a different path as the non-active employees. As shown in figure 4, average wages of active seekers are always lying below the level of the non-active employees. In addition, the difference in wages for active and non-active employees are quite stable over this 20-year period. Since non-active workers earns more, they should be located at a higher position in the job ladder and be more picky than active seekers. Hence, \mathcal{A}_t^{ns} should always be less than \mathcal{A}_t^s at

any given period t . In addition, since the earning gap between active seekers and non-active is stable over time, one can estimate the unobserved \mathcal{A}_t^{ns} to be scaling from \mathcal{A}_t^s with a constant coefficient $\kappa < 1$. With this assumption, the ratio of EE transition rate of non-active employees and active seekers is specified as

$$\frac{EE_t^{ns}}{EE_t^s} = \frac{\phi_t^{ns} \kappa \mathcal{A}_t^s}{\phi_t^s \mathcal{A}_t^s} = \kappa \frac{\phi_t^{ns}}{\phi_t^s}.$$

In addition to imposing structure on \mathcal{A}_t^{ns} , we also need to discipline the unobserved ratio ϕ_t^{ns}/ϕ_t^s before estimating the scaling parameter κ . Intuitively, this ratio can be interpreted as relative fraction of non-active employees received an offer relative to active seekers. Denote ν the average of ϕ_t^{ns} over time, i.e. $\nu = T^{-1} \sum_t \phi_t^{ns}$. With reference to the result in Faberman et al. (2022), I set this ratio ν equals 0.275². This value implies for every active seeker who received an offer, there are 0.275 of non-active employees receiving an job offer. Given ν , the value of κ can be obtained with the average ratio of EE transition rate between active and non-active employees

$$T^{-1} \sum_t \left(\frac{EE_t^{ns}}{EE_t^s} \right) = T^{-1} \sum_t \frac{\phi_t^{ns}}{\phi_t^s} \kappa = \nu \kappa \quad (5)$$

Subsequently, acceptance rate of non-active employees $\mathcal{A}_t^{ns} = \kappa \mathcal{A}_t^s$, with \mathcal{A}_t^s estimated in step 1.


2.2 Relative importance of the acceptance rate

Using the method illustrated in step 2, the estimated value of κ is 0.527. In other words, non-active seekers are about 52 percent less likely to accept an offer relative to an active seeker. Average acceptance rate of all employed workers is computed using equation 4 and shown in figure 5a. Results show that \mathcal{A}_t declined from 32 percent in late 90s to around 24 percent in 2019. This represents a 25-percent fall in levels since late 90s.

In terms of the relative contributions of search and acceptance, this is displayed in figure 6. One can clearly separate the result into two periods, before and after 2010, for interpretation. Before 2010, search and acceptance were almost equally important to the decline in EE-UE ratio; while acceptance accounted for most of the decline in EE-UE ratio after 2010.

The explicit role of average acceptance rate in the decline in relative EE rate is shown in figure 5b, which I compute a counterfactual EE-UE ratio if acceptance probability were held

²In table IV of Faberman et al. (2022), there are 0.051 (standard error: 0.004) share of non-active employees receiving job offers in a month; for active seekers, a proportion of 0.173 (0.013) receives a job offer to replace their current job.



(a) Average acceptance rate with $\kappa = 0.527$. (b) Counterfactual EE-UE ratio without variations in job acceptance.

Figure 5: Estimated job acceptance rate and its relative importance.

Note: The left panel presents the average offer acceptance rate for all employees computed using equation 3. The solid line in the right panel shows the EE-UE ratio observed in the UKLFS; the dash line shows the counterfactual EE-UE ratio computed from equation 1 in which the value of job acceptance is kept at the 1998 level. Shaded area indicates recession periods in the UK.

constant since 1998. Hence, variations in the counterfactual series were purely driven by changes in ϕ_t . Despite some declines during 2000s, the counterfactual series would recover in 2010s. Essentially, there would have been no persistent decline in the EE-UE ratio as observed in the data if job acceptance was not falling. Quantitatively, as the aggregate EE-UE ratio dropped by about 30 percent since 1998, fall in workers' job acceptance accounted for about 80 percent of the decline in average level.

To check the sensitivity of this result, I estimate the contribution of average acceptance rate by categorizing workers into different groups by gender, age and education qualifications. The results are robust as not only workers in all groups experience secular decline in the EE-UE ratio, but the acceptance rate is also the main driver in all of them. Results by age groups and qualifications are shown in figure 7³. This is consistent with the negative results in section 4.2 that changes in worker composition cannot explain the fall in EE rate.

The fact that the persistent fall in relative EE rate is predominantly due to a declining offer acceptance rate instead of relative job offer rate is an important result. It rules out a potential mechanism that there is an increase in employee's search costs over time which leads to a decline in the EE rate. Otherwise, this would have reflected in a decline in ϕ_t .

The discussion so far highlighted the role of job acceptance in keeping the EE rate low after 2010. This is based on the assumption that relative job offer rate ϕ_t^s of active

³Age group is separated at age 38 because it is the median age of all EE transitions.




Figure 6: Relative contribution of search vs. acceptance.

Note: Blue (gray) area shows proportion of the decline in the EE-UE ratio since 1998 that is driven by the fall in relative search efficiency (average job acceptance).

seekers didn't deteriorate over time. While Faberman et al. (2022) provides some evidence that the average offer rate is statistically indifferent between active seekers and unemployed workers in the US, there is no direct supporting evidence for the case in the UK. Hence, one cannot completely rule out the possibility that the fall in EE mobility was due to reduced job contact rate to the employed. For instance, there can be a deterioration in matching technology that discriminate against employed workers. While this possibility still remains, this do not undermine the result that the reduction in job acceptance is the main driving force of reducing EE flows because of the following reasons. First, let's consider an extreme case. If job acceptance was kept constant and equalled to one, the fall in EE-UE ratio would be solely driven by the reduction in job arrivals to employed workers. In other words, on-the-job searchers would accept every offer that came their way and this is essentially the setting of directed search in Menzio and Shi (2011). Subsequently, based on the assumption that workers only transit to higher ranking jobs, this implies employed workers would only apply for jobs that were better than their current jobs. Yet, "a reduction in job arrivals of these better job offers" in the directed search setting is equivalent as saying "employed workers are less likely to receive a job offer that they would accept" in the random search setting. Second, the fact that UE rate does not exhibit notable increases in this period implies that a reduction in the relative offer rate would have to be a result of reduced job arrival rate against employed workers. In reality, it is hard to imagine a change to the matching technology that would not affect unemployed workers but only the employed. Even if such




Figure 7: Contribution of Acceptance Rate by Worker Characteristics.

Note: The left panel shows the observed EE-UE ratios and the counterfactual series in which variations in job acceptance are shutdown, by age below and above 38 years old. 38 years old is used as a cut off because it is median age of job switchers recorded in UKLFS. The right panel shows the observed EE-UE ratios and the counterfactual series without variations in acceptance, by workers' qualification attainments. Shaded area indicates recession period in the UK.

changes in matching technology existed, this effect would have to reduce employed workers' offer arrival rate massively to eliminate the role of job acceptance.

3 Changes in Job Offer Distribution

Now we know that EE rate was persistently low in 2010s because employed workers were being more picky against poaching offers. Yet, why were workers not accepting those offers? Extracting the expression of \mathcal{A}_t in equation 1 and restate here:

$$\mathcal{A}_t = \int_0^1 (1 - F_t(x)) dN_t(x) \quad (6)$$

This expression of \mathcal{A}_t allows us to further decompose acceptance into two channels. First, more workers turned down offers because more of them were positioned around the top of the job ladder over time. Denote this the “employment” channel. Analytically, this would imply an increase of employment distribution $N_t(x)$ of higher ranking jobs x . Second, it could also be a deterioration in the sampling probability of better offers disregarding how workers were allocated along the job ladder. I call this the “job offer” channel. This works through a reduction in the probability of sampling better offers from the job offer distribution $F_t(x)$ in equation 6, for a given rank x on the job ladder.

It is important to highlight that these two channels carry drastically different economic implications. If employment channel dominated, it would be indeed an enhancement of worker's labour market outcomes as more workers were matched with better jobs. Otherwise, if offer channel were the main contributor, it would be a worrying sign of weakening labour market. To quantitatively examine the relative importance of these two channels, I estimate a dynamic job ladder model as in Moscarini and Postel-Vinay (2016a) with the UK data.

3.1 Accounting framework of the dynamic job ladder model

While the employment distribution is observed in the data, the job offer distribution is not directly observed. Moscarini and Postel-Vinay (2016a) provides an accounting framework to infer the job offer distribution using observed transitions in employment distribution using a dynamic job ladder model. It starts with assuming workers agree on a common rank of jobs from zero to one and all job seekers engage in random search. Unemployed workers receive job offer at Poisson rate λ_{t+1}^u . Employed workers have relative search efficiency ϕ_t with respect to the unemployed. Hence, offer arrival rate to employed workers at time t is $\phi_t \lambda_t^u$. Same as in Moscarini and Postel-Vinay (2016a), I allow an employed worker receives a reallocation shock at rate ρ_t and transfers to a random job on the job ladder, independent of her initial job rank x . Reallocation shock can be seen as incidents where a family reallocates to another location due to personal reasons. This also accounts for any transition to jobs that were paid less than the original job, including receiving a wage cut. Adopting the same notations for employment and offer distribution as in equation 6, the net change in the total employment up to rank x across two periods is

$$\begin{aligned}
 N_{t+1}(x) - N_t(x) = & - [\delta_{t+1} + \rho_{t+1} + \phi_{t+1} \lambda_{t+1}^u (1 - F_{t+1}(x))] N_t(x) \\
 & + \{ \rho_{t+1} N_t(1) + \lambda_{t+1}^u [1 - N_t(1)] \} F_{t+1}(x)
 \end{aligned} \tag{7}$$

where δ_{t+1} is the exogenous separation rate in the labour market.

Note that since $x = 1$ indicates the top rank, $N(1)$ measures the total employment. Hence, $U_t = 1 - N(1)$ is the unemployment rate. The first line of equation 7 accounts for outflows from employment in firms ranked up to x . These outflows include separations to unemployment, reallocations and workers moving up to higher ranking firms. The second line of the equation accounts for inflows into firms up to rank x . They are coming from reallocations, and transitions out of unemployment and into firms with rank x or below. Under some assumptions, there exists a rank-preserving equilibrium of this job-posting model, in which initially top firms always offer better contracts to workers regardless of the realised

aggregate states⁴. This implies the existence of an ergodic employment distribution along the job ladder. Hence, the corresponding equilibrium job offer distribution $F_{t+1}(\cdot)$ can be derived using equation 7.

A main data limitation in estimating the job offer distribution is that employment shares on the whole continuous rank of firms are not completely observed. Instead of estimating the entire distribution over a continuous support, I discretize the support and estimate densities at K cut-off quantile points X_k plus zero, where $k = 0, \dots, K$, with $X_0 = 0$ and $X_K = 1$. Denote $\bar{F}_t(X_k) = 1 - F_t(X_k)$ the acceptance distribution at some cut-off rank X_k at time t . This can be expressed as

$$\bar{F}_{t+1}(X_k) = \frac{[N_{t+1}(1) - N_{t+1}(X_k)] - (1 - \rho_{t+1} - \delta_{t+1}) [N_t(1) - N_t(X_k)]}{\rho_{t+1}N_t(1) + \phi_{t+1}\lambda_{t+1}^u N_t(X_k) + \lambda_{t+1}^u U_t}. \quad (8)$$

The numerator of equation 8 measures employment density between $X_K = 1$ and X_k , net of stayers who were already employed in this range in last period; the denominator accounts for the total number of workers at rank X_k or below (including all reallocating and unemployed workers) who received a new job offer. Hence, the acceptance distribution $\bar{F}_{t+1}(X_k)$ has a direct economic interpretation: the share of workers who are just hired by firms ranked above X_k out of all workers who wished to move to rank X_k or above and received a new job offer.

3.2 Estimation of the dynamic job ladder model

Job ladder are specified using real wages from the UKLFS. I discretise the global range of real wages in UKLFS from 1997-2019 into a 500-point grid. Under this specification, employment distribution is essentially the observed wage distribution and the wage density at each grid point can be directly observed in the survey data. I then estimate and track the changes in the cross-sectional offer distribution over time. Figure 8 shows the evolution of real wage distribution in the UK from 1997-2019. The general upward trend of the distribution is consistent with the fact that real wages are growing as a result of economic growth, while wage growth was notably slower after the Great Recession.

Other elements in equation 8 can either be directly taken from the UKLFS, or estimates from previous section. Specifically, unemployment rate U_t , job offer rate λ_t^u of the unemployed and the separation rate δ_t can be estimated from the UKLFS. In addition, re-allocation rate ρ_t is specified as the EE transition rates of non-active employees, normalised with their acceptance rate \mathcal{A}_t^{ns} in that period. This ensures the arrival rate of reallocation

⁴See Moscarini and Postel-Vinay (2013) for a in-depth theoretical discussion on the rank-preserving equilibrium.




Figure 8: UK wage distribution from 1997-2019.

Notes: Character P in the legend refers the corresponding percentile of the wage distribution. For instance, $P50$ corresponds to the median of the wage distribution. Shaded area indicates the Great Recession.

shock does not take acceptance into accounts. For relative search efficiency ϕ_t , I take the series estimated in the section 2 and thus the job offer rate to employed workers is $\phi_t \lambda_t^u$.

The reliability of the estimated offer sampling distribution depends on a correct specification of the job ladder. Hence, the assumption that all workers agree on a single job ladder is a strong one. For instance, Individuals can sort into different markets and for each specific market there is a corresponding job ladder. To mitigate this, instead of specifying a single job ladder using real wages, I compute residualised wages using a Mincer regression. The goal is to remove wage difference across different job ladders by controlling for both worker and sectoral characteristics. Essentially, this allows me to collapse numerous job ladders with different observed characteristics into a single one through residualisation. In the regression model, I control for various individual characteristics, including gender, age, education, and the indication of whether they are actively searching on the job. The indicator of active on-the-job search acts as a proxy to whether the workers think they are mismatched. Beside, I also control for 20 geographical regions, 20 industry sectors and 9 occupation categories. The offer distribution estimated using the residualised wages is the preferred job ladder specification for this paper.

3.3 Results on job offer distribution




Figure 9: Wage Offer Distribution estimated from UKLFS

Notes: Character P in the legend refers the corresponding percentile of the wage offer distribution. For instance, $P50$ corresponds to the median of the wage offer distribution. Shaded area indicates the Great Recession.




Figure 10: Changes in acceptance rate along offer and wage distribution

Notes: Areas underneath each curve to the x-axis is the average acceptance rate in a given year.

Results using wages. I first specify the job ladder using real wages and estimate the wage offer distribution using equation 8. The result is displayed in figure 9. A surprising observation is that the wage offer distribution does not follow similar growth path as the wage distribution in figure 8. In fact, the wage offer distribution appeared to be stable during 2000s and arguably regressed after 2010.

To have another look at how acceptance rate changes along the whole distribution, I explicitly present \mathcal{A}_t of a given year estimated using the dynamic job ladder model in figure 10. The acceptance probability $1 - F_t(\cdot)$ is put on the y-axis; and wage distribution $N_{t-1}(\cdot)$ is on the x-axis. The area underneath each curve represents \mathcal{A}_t of that year. A downward shift of the curve from 2000 to 2009 shows the decline in \mathcal{A}_t . Meanwhile, \mathcal{A}_t was kept at a fairly stable level after 2010 and this is consistent with movement of \mathcal{A}_t estimated in figure 5a. For instance, the result shows that a worker at the 40th percentile of the wage distribution was accepting around 50 percent of incoming job offers in 2000. This acceptance rate fell to around 30 percent in 2009 and was remained to be around 25-30 percent by 2019. This representation in figure 10 also allows us to examine the relative contribution of “offer” and “employment” channel. As shown in figure 8, since the wage distribution was shifting upwards from 2000-2009 while the offer distribution was stable, it suggests the employment channel contributed to the decline in \mathcal{A}_t from 2000 to 2009. While both the wage and offer distribution were relative stable from 2010-2019, \mathcal{A}_t was then kept at a similar level as in 2009.

Results using residualised wages. As discussed above, the job ladder specified using the residualised wages is the preferred setup since it collapses segregations in the labour market on a single ranking. Job offer distribution estimated using residualised wages from the Mincer regression is shown in figure 11. While the entire offer distribution was stable during the 2000s, there was a sharp downward shift in the offer distribution after 2010. As shown in figure 12, the residualised wages exhibited a growing trend from 2000-2009. After that, there was a tumble in the residualised wage distribution followed by a period of stagnant growth in the 2010s.

Given how the residualised wage and offer distribution changed over time, one can see the decline in job acceptance from 2000-2010 was due to upward movement in the wage distribution of workers. Hence, the employment channel dominated the decline in job acceptance in the first 10 years since 2000. Yet, there was a shape downward shift in the wage distribution after 2010. By itself, this would have generated a positive employment effect which would increase workers’ acceptance rate. However, such effect was offset almost entirely by the downward shift in the offer distribution and the negative effect of the offer channel was the




Figure 11: Residualised wage offer distribution



Figure 12: Residualised wage distribution

Notes: Character P in the legend refers the corresponding percentile of the distribution. For instance, $P50$ corresponds to the median of the wage distribution. Shaded area indicates the Great Recession.

reason for job acceptance to remain persistently low from 2010-2019.

3.4 Reconciling results on job acceptance

So far I introduce two different approaches in estimating job acceptance rate \mathcal{A}_t . The first approach is through comparing transition rates between employed and unemployed, illustrated in section 2. The second approach in this section involves estimating offer distribution using changes in wage distributions. This second approach allows us to further decompose job acceptance along the two dimension of “offer” and “employment” channels.

To reconcile results from these two approaches, I compute \mathcal{A}_t using equation 6 and compare it with the series using transition rates from equation 4 in section 2. As shown in figure 13, \mathcal{A}_t estimated from two different approaches are able to produce similar movement in levels over the two decades. \mathcal{A}_t computed from the wage distribution method gives greater fluctuations because of the noisiness in wage observations in the UKLFS.

Subsequently, I construct the EE-UE ratio with equation 1 and compare this model-constructed series with the one observed in figure 14a. Again, while the constructed series generates more quarter-to-quarter fluctuations, it is able to regenerate a similar path in levels throughout the two decades. If one shut down the “offer” channel in job acceptance, as shown in figure 14b, the counterfactual EE-UE ratio would exhibit a similar fall as in the observed series between 2000 and 2010. This is because workers were cutting back on search efforts and becoming more picky as they moved up the job ladder. However, without changes in the offer distribution, the EE-UE ratio would have recovered most of its fall after 2010. This, again, indicates that the deterioration in job offers is the main driver for the persistently low EE mobility in the past 10 years.

Yet, how did these better offers just disappear after 2010? The answer might lie in the sluggish productivity growth in the UK after the Great Recession, as shown in figure 15a. From 1980s to 2008, productivity growth in the UK was between 1 to 2 percent annually. In the 10-year period after 2010, the productivity growth rate sharply declined and was mostly kept below 1 percent. According to a result in Schneider (2018) in figure 15b, much of the reduction in the productivity growth after 2010 was concentrated among the most productive British firms. Hence, the deterioration of offer quality could be due to the slowdown in the creation of new “best” jobs at the top of the job ladder.





Figure 13: Reconciling job acceptance from two approaches.

Notes: Solid line represents 4-quarter moving average of \mathcal{A}_t estimated using ratios of job transition rate. Dash line is the 8-quarter moving average of \mathcal{A}_t estimated using changes in wage offer distribution. Shaded area indicates the Great Recession.



(a) Reconciling the two series.

(b) Counterfactual series without offer channel.

Figure 14: Comparison between observed and model-constructed EE-UE ratio.

Notes: Left panel reconciles the EE-UE ratio constructed from the model with the observed series in the data. Solid line represents 4-quarter moving average of the observed EE-UE ratio. Dash line is the 8-quarter moving average of ratio constructed using \mathcal{A}_t estimated using wage distribution. Right panel shows the counterfactual EE-UE ratio when offer channel in acceptance is shut down. Shaded area indicates the Great Recession.

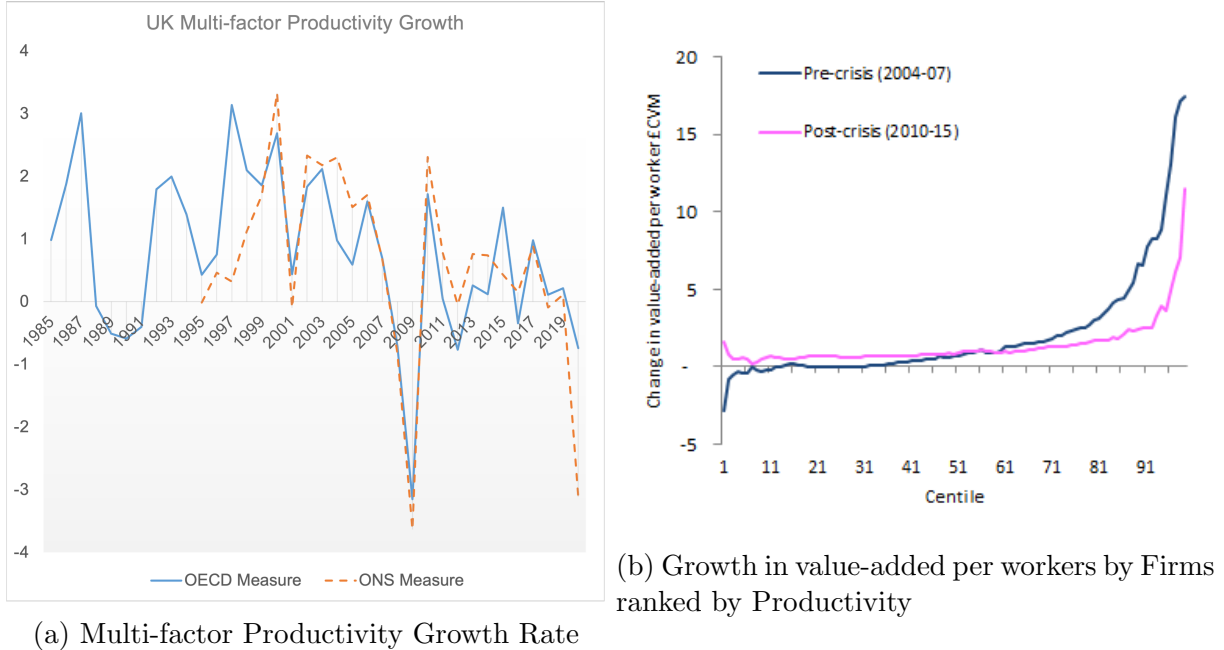


Figure 15: UK Technological Growth and Firm Dynamics

Note: The left panel shows the annual multi-factor productivity growth rate of the UK. The solid line represents the measure given by the OECD, and the dash line is the same measure from the Office of National Statistics (ONS). The right panel shows the productivity growth by firms ranked by their productivity level. This is a result borrowed from Schneider (2018). The navy line displays firms' growth profile before the financial crisis in 2008; the pink line presents the growth of firms after the crisis from 2010 to 2015.

4 Additional Facts

There are other potential explanations to the declining EE transition rate in past decades. First, I have a look at how much of this decline in mobility was simply due to the lack of job openings over the business cycle. Besides, composition changes in the economy can also play major roles in driving down the EE rate. For instance, EE mobility can be affected by demographic changes in the labour force and structural transformation of the economy. I examine the contribution of composition changes using a between-within decomposition framework.

4.1 Declined EE rate is not due to poor market condition

How much of these decline in job mobility is due to cyclical fluctuations in vacancies? I examine this by specifying a standard matching function and fits it to the UK data. Specifically, by adopting a Cobb-Douglas functional form for the matching function, total EE transitions M^e within a given period t is

$$M_t^e = \mu_t (s_t^e E_t)^\sigma V_t^{1-\sigma} = \tilde{\mu}_t E_t^\sigma V_t^{1-\sigma} \quad (9)$$

where μ_t measures the efficiency of the matching technology; s_t is the search efficiency of employed workers; and σ is the match elasticity with respect to number of employed workers. Since aggregate search efficiency of the employed s_t^e is not directly observed, I re-specify $\tilde{\mu}_t = \mu_t(s_t^e)^\sigma$ as the residualised EE rates. Subsequently, the aggregate EE transition rate is

$$ee_t = \frac{M_t^e}{E_t} = \tilde{\mu}_t \left(\frac{V_t}{E_t} \right)^{1-\sigma} = \tilde{\mu}_t \theta_t^{e1-\sigma}$$

where $\theta_t^e = \frac{V_t}{E_t}$ is the vacancy rate, a cyclical measure of labour market conditions. Hence, given a value of matching elasticity σ , one can obtain a log series of residualised EE rates $\tilde{\mu}_t$ from this expression

$$\ln \tilde{\mu}_t = \ln ee_t - (1 - \sigma) \ln \theta_t^e \quad (10)$$




Figure 16: UK vacancy rate and labour market tightness

Note: Quarterly series computed using UKLFS and Vacancy Survey. Vacancy rate defined as unfilled vacancy over total number of employed workers in the quarter. Labour market tightness is defined the ratio between unfilled vacancy and unemployed population.

Variations in $\tilde{\mu}_t$ can either be a shift in workers' search behaviour s_t^e or a change in the matching technology μ_t . The former component can be further decomposed into 1) search efficiency of employed workers and 2) their corresponding acceptance rate when they receive an offer, which we discussed in section 2. The latter component is seen as a change in the complexity of recruitment process. For instance, more rounds of interviews before delivering



Figure 17: Residualised EE transition rates

Note: The left panel shows the residualised log EE series at different value of σ . The red solid line displays the series with σ estimated using the variations during the Great Recession; the green dash line shows the σ estimated using OLS regression of the log EE rate on the log vacancy in the UK. The northeast panel replicates the same residualised EE series with the value of σ from OLS. The southeast panel provides a residualised UE series obtained from the OLS regression of the quarterly log UE rate on log vacancy. Shaded area indicates the recession period.

an job offer would prolong the matching process and reduce the matching efficiency.

Equation 10 shows that $\tilde{\mu}_t$ can be computed using observed EE transition rates and the vacancy rate. I take the quarterly EE rates and employment series from the UKLFS. Vacancy information is obtained from the UK Vacancy Survey, which provides monthly numbers of 3-month moving average of unfilled vacancies from 2002 to 2019. The Vacancy Survey began in 2001 by conducting interviews with around 6,000 enterprises in the Great Britain each month for their vacancy numbers. Figure 16 displays two measures of labour market condition. In the upper panel, Although UKLFS provides workers information in 1990s, to match the availability of vacancy data, estimates of the residualised series only cover period from 2002Q1 to 2019Q4.

Finally, one needs the elasticity parameter σ to compute the series of $\tilde{\mu}_t$ from observables. A straightforward method of estimating σ is to apply a linear regression of log EE rates on the log vacancy rate θ_t^e and back out $\tilde{\mu}_t$ as residuals. Another way is to consider that changes in the EE rates during the global financial crisis from 2008Q3 to 2009Q2 was a pure result of the deterioration in market conditions. Specifically, $1 - \sigma \simeq \frac{1}{4} \sum_{t=2008Q3}^{2009Q2} \frac{\Delta \ln ee_t}{\Delta \ln \theta_t^e}$. Besides, I also include two other reference values of σ for sensitivity purpose.

The left panel of figure 17 shows the residualised EE series in the UK at different value of σ . It exhibits a persistent decline over the two decades and this is robust across different σ values. This shows that the sluggish performance in EE mobility since 2000 is not a cyclical phenomenon driven by lack of job openings. Meanwhile, if one controls for the effect of vacancies on UE rate similarly, the residualised UE transition rate doesn't demonstrate any

secular trend despite some fluctuations, as shown in the right panel of figure 17.

4.2 Composition changes in labour market cannot explain the fall in EE rates

Since the persistent fall in EE transition rate is a long-term phenomenon, a natural question to ask is whether this is associated with composition shifts in the UK economy. A potential explanation is changing demography. For instance, having an ageing labour force, in which older workers are generally less likely to switch jobs, would imply a fall in aggregate EE rate, even if general job search behaviour at different ages didn't change over time. Another potential cause is structural transformation towards professional services in the UK. The growing importance of professional services could prolong the average recruitment process. For instance, since hiring a financial manager takes more rounds of interviews than recruiting a construction worker, average EE rate falls. To jointly examine these potential channels, I introduce a theoretical framework for a between-within decomposition of the residualised EE rate that accounts for multiple dimensions of composition changes.

4.2.1 Theoretical Framework of Decomposition

As discussed when I introduce the residualised series $\tilde{\mu}_t$ in section 4.1, composition shifts in average worker characteristics and sector shares can contribute to changes in aggregate EE rate. Barnichon and Figura (2015) shows that fundamental changes in smaller sub-markets can contribute to fluctuations of $\tilde{\mu}_t$ and the authors illustrate it with the UE transitions.

Building on the framework as in Barnichon and Figura (2015), I specify total EE transition M_t^e to be the sum of matches from J non-overlapping local markets. In addition, I allow each individual job seeker i in each local market j to differ in search intensity s_{ijt} based on their individual characteristics. Hence, aggregate EE transition M_t^e is the sum of transitions in all local markets:

$$M_t^e = \sum_j \mu_{jt} (s_{jt} E_{jt})^\sigma V_{jt}^{1-\sigma} \quad (11)$$

where s_{jt} is the weighted average search intensity across individuals i within sector j , i.e. $s_{jt} = \sum_i \frac{E_{ijt}}{E_{jt}} s_{ijt}$, where E_{ijt} is individual weight in the survey and $E_{jt} = \sum_{i \in j} E_{ijt}$; and μ_{jt} is local-market-specific matching efficiency. One can manipulate equation 11 to take the form of equation 9, and obtain the following expression

$$M_t^e = E_t^\sigma V_t^{1-\sigma} \left(\sum_{j=1}^J \frac{E_{jt}}{E_t} \mu_{jt} (s_{jt})^\sigma \left(\frac{\theta_{jt}^e}{\theta_t^e} \right)^{1-\sigma} \right)$$

where $E_t \equiv \sum_j E_{jt}$; $V_t \equiv \sum_j V_{jt}$; and $\theta_{jt}^e \equiv \frac{V_{jt}}{E_{jt}}$. This implies that the residualised EE rate $\tilde{\mu}_t$ can also be expressed as

$$\tilde{\mu}_t = \sum_{j=1}^J \frac{E_{jt}}{E_t} \mu_{jt} (s_{ijt})^\sigma \left(\frac{\theta_{jt}^e}{\theta_t^e} \right)^{1-\sigma} \quad (12)$$

To conduct a between-within decomposition in the residualised EE rate, I apply 2nd order approximation to equation 12 with s_{ijt} around 1, μ_{jt} around μ_0 , and θ_{jt}^e around θ_t^e . Here, I normalise the average search intensity s_{ijt} to unity without loss of generality. In Barnichon and Figura (2015), it is the average search intensity of the unemployed which normalised to one. While average search intensity of employed workers is usually considered less than the unemployed in the literature, the objective here is to study the relative changes in search intensities among employed workers across time. Hence, changing the anchor point of the normalization doesn't change the qualitative results. As a robustness check, I extended this framework by incorporating both employed and unemployed job seekers in the appendix. In that extended framework, only average search intensity of unemployed workers is normalised to one as in Barnichon and Figura (2015). Results remain robust in that setting.

The resulting expansion indicates that $\tilde{\mu}_t$ can be decomposed into the following components⁵:

$$\tilde{\mu}_t \approx \mu_0 \left\{ 1 + \psi_t^i + \psi_t^j - \frac{\sigma(1-\sigma)}{2} \left(\text{Var} \left(\frac{\theta_{jt}^e}{\theta_t^e} \right) + \text{Var} (s_{ijt}) \right) \right\} + \text{Covariance Terms} \quad (13)$$

with

$$\begin{cases} \psi_t^i = \sigma \sum_{i,j} \frac{E_{ijt}}{E_t} (s_{ijt} - 1), & \psi_t^j = \sum_j \frac{E_{jt}}{E_t} \left(\frac{\mu_{jt}}{\mu_0} - 1 \right) \\ \text{Var} \left(\frac{\theta_{jt}^e}{\theta_t^e} \right) = \sum_j \frac{E_{jt}}{E_t} \left(\frac{\theta_{jt}^e}{\theta_t^e} - 1 \right)^2, & \text{Var} (s_{ijt}) = \sum_{i,j} \frac{E_{ijt}}{E_t} \frac{E_{ijt}}{E_{jt}} (s_{ijt} - 1)^2 \end{cases}$$

Note that μ_0 is the weighted average of local-level matching efficiency across time. Specifically,

$$\mu_0 = T^{-1} \sum_{t=1}^T \sum_{j=1}^J \frac{E_{jt}}{E_t} \mu_{jt} = T^{-1} \sum_{t=1}^T \sum_{j=1}^J \frac{E_{jt}}{E_t} \mu_j$$

where the second equality follow from the assumption that μ_{jt} is stationary around its mean for each local market j .

Each component in equation 13 carries economic implications as below⁶:

⁵The first-order terms in the expansion with respect to θ_{jt}^e is dropped since it averages to zero.

⁶Covariance terms consist of three pair-wise associations between average search intensity, market tightness and matching efficiency within local markets. Since their overall contributions is negligible, their implication is abstracted from here.

1. ψ_t^i measures the worker effect;
2. ψ_t^j measures the job effect;
3. $\text{Var}\left(\frac{\theta_{jt}^e}{\theta_t^e}\right)$ measures the job dispersion effect across local markets; and
4. $\text{Var}(s_{ijt})$ measures the worker dispersion effect in their search intensities.

The implications of first-order effects ψ_t^i and ψ_t^j measures the direct effect of composition changes on the residualised EE series, which shall be discussed extensively below. The second-order terms also carry essential economic meanings. First, as illustrated in Herz and Van Rens (2020), job dispersion $\text{Var}\left(\frac{\theta_{jt}^e}{\theta_t^e}\right)$ measures the level of misallocation between vacancies and job seekers across submarkets. Intuitively, greater dispersion in the vacancy-to-job seeker ratio hampers worker reallocation rates since some firms and workers are not taking advantage of better opportunities in forming a match in other sectors. Second, worker dispersion $\text{Var}(s_{ijt})$ measures the degree of worker heterogeneity in their search intensities across local markets. Conditional on the sector-level matching efficiency, greater dispersion in search intensity would reduce job-to-job mobility because workers in some sectors are exerting more efforts on job search than they efficiently needed in other sectors. In technical terms, due to the concavity of the matching function, greater variations in market conditions and workers' search behaviour across local markets would reduce the number of overall matches. Meanwhile, conditional on the level of dispersion, the overall labour market would also be more efficient if workers exert greater search efforts in a local market with tighter market conditions or higher local matching efficiency. Hence, positive covariance relationships across sector-level matching efficiency μ_{jt} , search intensity s_{jt} and market condition θ_{jt}^e would enhance transition rates. However, since these covariance effects are quantitatively small, they have limited contributions to changes in EE rates.

Worker Effects. The worker effect ψ_t^i reflects how composition changes in individual's characteristics on their search behaviour relative to the average search intensity. To explicitly estimate this with observed worker characteristics, I specify individual search intensity to take the form of

$$s_{ijt} = \exp(\beta X_{ijt} + \varepsilon_{ijt})$$

where $X_{ijt} = [1, x_{ijt}^1, \dots, x_{ijt}^K]$ is a vector of observed worker characteristics and ε_{ijt} is the unobserved error term, which by construction has mean zero. Subsequently, a linearisation

of ψ_t^i in equation 13 gives the following expression

$$\psi_t^i \approx \sigma_{EE} \sum_{i,j} \frac{E_{ijt}}{E_t} \sum_{k=0}^K \beta_k (x_{ijt}^k - \bar{x}^k) + \sigma_{EE} \sum_{i,j} \frac{E_{ijt}}{E_t} \varepsilon_{ijt}$$

where $x^k = T^{-1} \sum_t \sum_{i,j} \frac{E_{ijt}}{E_t} x_j^k$ is the weighted average of observed worker characteristics k across time. By assuming the expected individual error in each period is identically and independently drawn from a standard normal distribution in the cross section, the second term in equation 4.2.1 should converge to zero as observations grow ⁷.

In the estimation, I account for Individual characteristics including gender, age groups, education qualifications, and job tenure at current employer in last quarter in X_{ijt} . The incorporation of job tenure at current employer is key as it allows us to control for the workers’ “growing attachment” with their current employer. This can either come from the worker’s seniority as they climb up the ladder internally; or some unobserved attachment with the company. These include some amenities or social attachments that workers gain with tenure.

Job Effect. Without inserting any restrictive structure, job effect ψ_t^i can be further break down into two components:

$$\psi_t^j = \underbrace{\sum_j \frac{E_{j t}}{E_t} \left(\frac{\mu_j}{\mu_0} - 1 \right)}_{\text{Pure Composition Shifts}} + \underbrace{\sum_j \frac{E_{j t}}{E_t} \left(\frac{\mu_{j t}}{\mu_0} - \frac{\mu_j}{\mu_0} \right)}_{\text{Within-sector Variations}} \quad (14)$$

where μ_j is the mean value of matching efficiency μ_{jt} in local sector j across time. The first term of equation 14 captures the pure composition shifts of jobs across local markets at time t . For instance, increasing shares of jobs that require a shorter recruiting process in general (i.e. higher μ_j) would result in more transitions over time. While the second term is the weighted average deviations of μ_{jt} from their corresponding means μ_j at a given time t . Without making any assumption on the evolution of local matching efficiency μ_{jt} , one can estimate μ_j with local-sector fixed effects. Within-sector variation of matching efficiency can be back out from the residuals at sector levels.

Unlike error terms in estimating individual search intensity, within-sector variation do

⁷The specifications of s_{ijt} also affects the estimation of second-order components in equation 13. For instance, the unobserved error term of individual search intensity would generate interaction terms with local within deviations and local market tightness. Yet, if we believe individual are small relative to the size of local markets, these covariance terms between unobserved individual components and local-sector-level variables should be negligible.

not vanish in the cross-section without making very strong assumptions that local matching efficiency is time-invariant. Since within-sector variation would also capture impacts of sectoral shocks, assuming local matching efficiencies to be time-invariant would risk omitting an important component of the decomposition.

4.2.2 Estimation of the decomposition framework

I now present the estimation process of the decomposition framework illustrated in section 4.2.1. First, I specify a local market as an industry in a region of the UK. In total, there are 44 local sectors with four big industry group and eleven local regions in the UK. These four industry classes are “professional”, “sales”, “production” and “services”; while the eleven geographical areas include all nine official regions of England, together with Scotland and Wales⁸. These regions are defined broadly due to two reasons. First, they are classified to fulfil the requirement of being an “island”, where between-sector transition is limited. With the current classification of local sectors, about 90 percent of all EE transitions are within sectors. Second, sector definition needs to accommodate data availability. An adequate number of recorded transitions in each period is needed to compute the EE series for every sub-markets. Finer definition of sectors would demand for dramatically bigger data to implement the framework.

Parameters to be estimated in this framework include 1) a vector β which measures the semi-elasticity of individual search intensity s_{jit} of each worker type i in each sector j at time t with respect to their characteristics X_{jit} ; 2) average sector-specific matching efficiency μ_j for each local sector j across time; and 3) matching elasticity parameter σ .

Structurally estimating these parameters requires a measure of individual EE transition probability. Yet, individual’s EE rates is not directly observed from the data. To overcome this limitation, I assume individual workers to be small in their corresponding sector and a person’s own search intensity have negligible effect on the sector-level EE rates. Given this assumption, individual EE rate is the sector-level EE rate adjusted by the ratio between individual search intensity and their average sectoral search intensity. Specifically, individual EE transition rate has this expression $ee_{ijt} = \frac{s_{ijt}}{s_{jt}} ee_{jt}$. As a general practice in the literature, I also assume the instantaneous EE rate to be constant throughout each time period. Hence,

⁸The composition of four main industry sectors: 1) *Professional* includes “Information and Communication”, “Financial and Insurance Activities”, “Real Estate Activities”, “Professional, Scientific and Technical Activities”, “Administrative and Support Service Activities” and “Education”; 2) *Sales* includes “Wholesale and Retail Trade; Repair of Motor Vehicles and Motorcycles”, “Transportation and Storage” and “Accommodation and Food Service Activities”; 3) *Production* includes “Utilities”, “Manufacturing” and “Construction”; 4) *Services* includes “Human Health and Social Work Activities”, “Arts, Entertainment and Recreation” and “Other Service Activities”

average EE transition probability within period is given by

$$EE_{ijt} = 1 - \exp\left(-\frac{s_{ijt}}{s_{jt}} ee_{jt}\right) = 1 - \exp\left(-\frac{s_{ijt}}{s_{jt}} \mu_{jt} \theta_{jt}^{e, 1-\sigma}\right). \quad (15)$$

Individual characteristics X_{ijt} include duration of job at period $t - 1$, gender, qualification level, and age groups. These variables are available for each worker in the UKLFS.

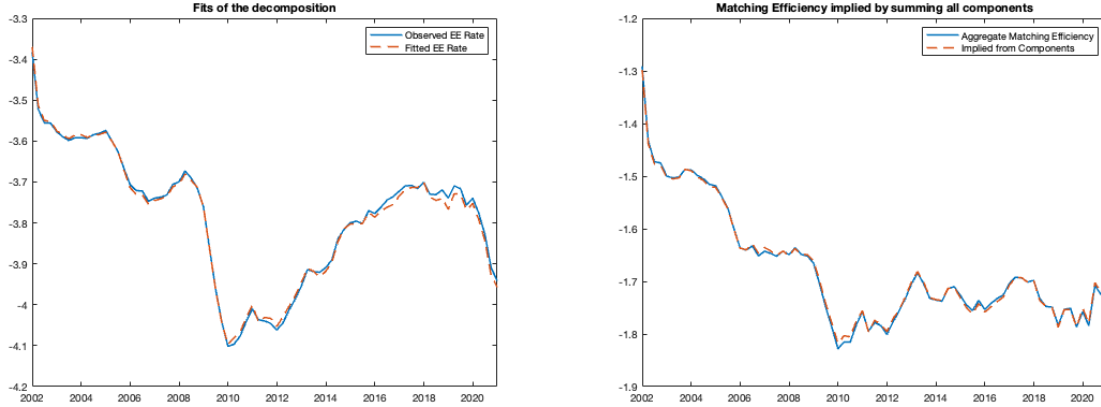
Cross-sectional difference in worker characteristics and their corresponding search outcomes shall help us identify β . The second equality of equation 15 also indicates that parameters μ_{jt} and σ can be estimated altogether using maximum likelihood. However, since the dynamics of μ_{jt} is unknown, I first estimate the average sector-specific matching efficiency μ_j using a sector-level fixed effect. I then recover the deviation of μ_{jt} from its mean value μ_j as a residual in the log linearised expression of sector-level EE rate given the estimated parameters. Specifically, I exploit the log specification of average local EE transition rate:

$$\ln ee_{jt} = \ln \mu_j + \sigma \ln s_{jt} + (1 - \sigma) \ln \theta_{jt}^e + \ln(1 + \Delta \mu_{jt})$$

where $\Delta \mu_{jt}$ is the percentage deviation from the sector mean μ_j . Using this expression and the sets of estimated parameters from the maximum likelihood estimator, I can back out the series of within-sector deviations $\Delta \mu_{jt}$. Subsequently, sector-level match efficiency μ_{jt} can be computed as $\mu_{jt} = \mu_j(1 + \Delta \mu_{jt})$.

The local-level market tightness θ_{jt}^e is defined as the ratio between numbers of vacancy in the local industry and employed workers in local sector j at time t . Although the UK Vacancy Survey provides the numbers of job openings at industry-level, regional breakdown of vacancy counts is not available. To allocate vacancies to different regions, I turn to the UK Jobcentre Plus vacancy data from 2004Q2 - 2012Q4 for regional information on job openings. The Jobcentre Plus was an employment agency ran by the Department of Work and Pensions of the UK government. It reports monthly numbers of active unfilled job openings that were posted on this government platform at local-industry level. I compute average shares of vacancy going into a region for each industry using this data from 2004Q2 - 2012Q4. Vacancy counts of each industry in the UK Vacancy Survey is then allocated according to this average regional share obtained from Jobcentre Plus data. The reason that Jobcentre Plus dataset is not a preferred source of vacancy data since it only contains a fraction of all job openings in the economy and is not as representative as the UK Vacancy Survey. In addition, Jobcentre Plus data does not cover period after 2012.

The set of parameters $\{\beta, \mu_j, \sigma\}$ can then be calculated with maximum likelihood estimator using the observed individual transition outcome. Specifically, the log likelihood function



(a) Implied EE Rate

(b) Residualised series implied by components

Figure 18: Fits of the Decomposition Framework

Note: The left panel compares the fit of constructed EE rate with the observed EE rate. The right panel compares the fit of the residualised series with σ equals the estimated value from the decomposition framework.

is

$$\ln \mathcal{L} = \sum_{j,i,t} Y_{ijt} \ln(EE_{ijt}) + (1 - Y_{ijt}) \ln(1 - EE_{ijt}) \quad (16)$$

where $Y_{ijt} = 1$ if a EE transition occur for individual i .


4.2.3 Result on composition changes

Implementing the decomposition framework illustrated above delivers insights on the relative importance of each components. In short, composition shifts in worker and sector shares play little role in explaining the fall in EE rates. Most declines in the residualised EE series occurs within each local labour market. Since the main purpose of the exercise is to study what drives or doesn't drive the reduction in EE rates, discussions on the estimated coefficients of each decomposition component are left in the appendix. Here, the discussion is focused on each component's impact on the EE rates.

Figure 18a illustrates how the decomposition framework is able to recover the aggregate EE rate. Overall, the framework provides a close fit to the observed aggregate EE rates, as it only generates around 1% deviation from the observed series. The remaining deviation exists because of the assumptions made in allocations of vacancies to different regions using the Jobcentre Plus data. Since vacancies of an industry are allocated to a region by the average vacancy share within the industry of that region in the Jobcentre Plus data, some cross-sectional variations in the local market tightness are not accounted for. How this limitation affects our estimated EE series depends on the covariance between the regional

vacancy share of an industry and the vacancy dynamics of that industry at a given time. If the vacancy opening in a period is more concentrated in a region with already high average vacancy share for an industry, this positive association will enlarge the dispersion of market conditions. If this was the case, since a greater dispersion would hamper the EE rate, our estimation would overestimate $\tilde{\mu}_t$. This appears to be the plausible reason that the fit is a relatively off during the Great Recession, when vacancy allocation is generally found to be more dispersed (Herz and Van Rens, 2020).

Figure 18b presents the residualised EE rates computed using equation 12 and the same series implied by summing up all components from equation 13. The main takeaway from figure 18b is that this approximation method does not leave out any essential component of $\tilde{\mu}_t$. Hence, this enables us to conduct counterfactual exercises to test their corresponding contribution of each component to falling EE rates in equation 13.



(a) Counterfactual EE rate

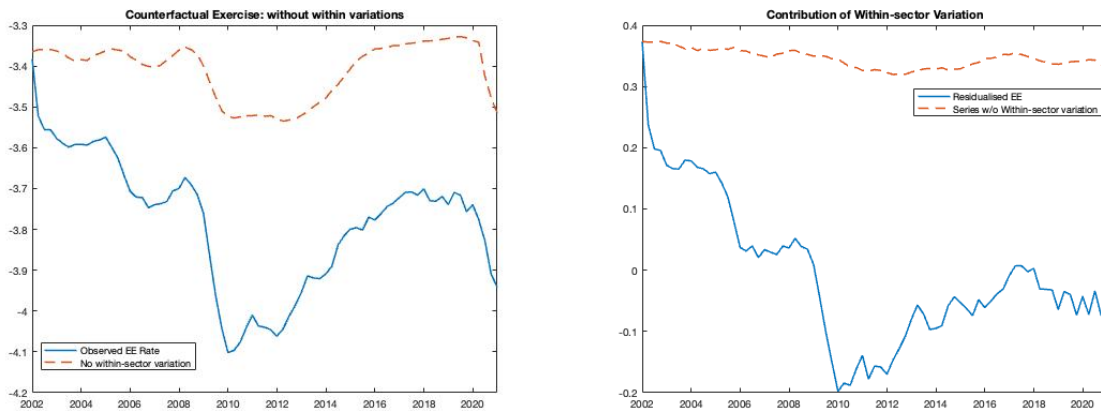
(b) Counterfactual residualised EE rate

Figure 19: Contribution of Composition Shifts

Note: The left panel compares the counterfactual EE rate with the observed EE rate (solid line). The dashed line indicates the counterfactual series without variations associated with worker (red) and job (yellow) composition shifts. The right panel compares the residualised series (solid line) with counterfactual series (dash line) which shuts down the variation caused by composition shifts.

First, I test the importance of the job and worker composition effects. Surprisingly, I find that shifts in neither worker characteristics nor sectoral shares in past decades explain much of the decline in EE rates. This is illustrated in figure 19, where variations associated with worker and job composition shifts are shut down respectively. Muting the variations in both components has only minimal effect in enhancing the EE rates. Hence, this allows us to dismiss the idea of composition shifts, including ageing population, the influx of college graduates, or structural transformation, to be the main cause of reduced job reallocation rate.

Besides, second-order components in equation 13 also have little to no contribution to the fall in EE rates. This is because these components either exhibit no apparent trending movement despite their fluctuation during the two decades; or the variation magnitude is essentially too small to generate actual effects on the EE series. As a result, it shows that job dispersion and heterogeneity in worker search efforts are not important contributors to the reduction in EE flows.



(a) Counterfactual EE rate

(b) Counterfactual residualised EE rate

Figure 20: Contribution of Within-sector Variations

Note: The left panel compares the counterfactual EE rate with the observed EE rate (solid line). The dashed line indicates the counterfactual series (dash line) without within-sector variations. The right panel compares the residualised series (solid line) with counterfactual series (dash line) which shuts down within-sector variations.

As composition variations have limited explanatory power to the persistent decline in EE rates, within-sector variations evolve to be the main force that drives down the EE rates in the UK. Indeed, by shutting down within-sector variations, as shown in figure 20, we would not see a downward trending EE transition in past decades. This is consistent with the result in section 3 that the persistent fall in EE mobility is due to reduction in job acceptance. As both phenomena that workers climbing up the job ladder from 2000 to 2010 and the subsequent deterioration of poaching offer quality after 2010 occur within sectors.

In other words, although the between-within decomposition exercise does not deliver direct evidence on what drives within-sector changes, it delivers valuable negative results that rule out some potential mechanisms to low EE mobility in the UK. For instance, ageing population and structural changes towards professional service sectors are not the main drivers behind the phenomenon.

5 Conclusion

This paper tries to shed lights on the root cause of the falling EE transition rate in the UK since 2000 with several empirical exercises. These analyses deliver a couple facts on the sluggish EE rates. First, using the job finding behaviour of unemployed worker as a benchmark, I find the persistent decline in the EE mobility is mainly due to the reduced likelihood of employee accepting a poaching offer when they receive one. Second, I further examine potential channels that may lead to decline in job acceptance rate using a dynamic job ladder model. By fitting it to the UK data, I am able to further decompose of job acceptance along the changes in employment and offer distributions. Result shows that the former channel was responsible for the decline before 2010, while the later kept the acceptance low in the 2010s. In other words, the fall in job acceptance from 2000 to before the Great Recession was attributed to workers becoming more picky as they moved up the job ladder. Meanwhile, acceptance was kept low after 2010 due to deterioration in job offer distribution. Subsequently, I show that the EE-UE ratio would have recovered to average level in early 2000s if not because of the declined quality in poaching offers after the Great Recession. Finally, I also show that fluctuations in vacancies and composition changes in the UK economy were not main drivers of the persistent fall in EE rate. Conducting a between-within decomposition with a structural framework rules out composition changes in the labour market as the main driver. This is true for both composition shifts in worker characteristics and sectoral structural transformation.

References

- Aguiar, Mark, Erik Hurst, and Loukas Karabarbounis (2013), “The life-cycle profile of time spent on job search.” *American Economic Review*, 103, 111–116.
- Bagga, Sadhika (2022), “Firm market power, worker mobility, and wages in the us labor market.”
- Baksy, Aniket, Daniele Caratelli, and Niklas Engbom (2024), “The long-term decline of the us job ladder.”
- Barnichon, Regis and Andrew Figura (2015), “Labor market heterogeneity and the aggregate matching function.” *American Economic Journal: Macroeconomics*, 7, 222–49.
- Burdett, Kenneth and Dale T Mortensen (1998), “Wage differentials, employer size, and unemployment.” *International Economic Review*, 257–273.

- Davis, Steven J and John Haltiwanger (1992), “Gross job creation, gross job destruction, and employment reallocation.” *The Quarterly Journal of Economics*, 107, 819–863.
- Deutscher, Nathan (2019), *Job-to-job transitions and the wages of Australian workers*. 2019-07, Treasury Working Paper.
- Faberman, R Jason and Alejandro Justiniano (2015), “Job switching and wage growth.” *Chicago Fed Letter*, 1.
- Faberman, R Jason, Andreas I Mueller, Ayşegül Şahin, and Giorgio Topa (2022), “Job search behavior among the employed and non-employed.” *Econometrica*, 90, 1743–1779.
- Fallick, Bruce and Charles A Fleischman (2004), “Employer-to-employer flows in the us labor market: The complete picture of gross worker flows.”
- Fujita, Shigeru, Giuseppe Moscarini, and Fabien Postel-Vinay (2020), “Measuring employer-to-employer reallocation.” Technical report, National Bureau of Economic Research.
- Gomes, Pedro (2012), “Labour market flows: Facts from the united kingdom.” *Labour Economics*, 19, 165–175.
- Goodridge, Peter, Jonathan Haskel, and Gavin Wallis (2018), “Accounting for the uk productivity puzzle: a decomposition and predictions.” *Economica*, 85, 581–605.
- Hall, Robert E and Sam Schulhofer-Wohl (2018), “Measuring job-finding rates and matching efficiency with heterogeneous job-seekers.” *American Economic Journal: Macroeconomics*, 10, 1–32.
- Herz, Benedikt and Thijs Van Rens (2020), “Accounting for mismatch unemployment.” *Journal of the European Economic Association*, 18, 1619–1654.
- Karahan, Fatih, Ryan Michaels, Benjamin Pugsley, Ayşegül Şahin, and Rachel Schuh (2017), “Do job-to-job transitions drive wage fluctuations over the business cycle?” *American Economic Review*, 107, 353–57.
- Menzio, Guido and Shouyong Shi (2011), “Efficient search on the job and the business cycle.” *Journal of Political Economy*, 119, 468–510.
- Mercan, A Yusuf (2017), “Fewer but better: The decline in job mobility and the information channel.” Technical report, Working paper, University of Melbourne.
- Moscarini, Giuseppe and Fabien Postel-Vinay (2013), “Stochastic search equilibrium.” *Review of Economic Studies*, 80, 1545–1581.

- Moscarini, Giuseppe and Fabien Postel-Vinay (2016a), “Did the job ladder fail after the great recession?” *Journal of Labor Economics*, 34, S55–S93.
- Moscarini, Giuseppe and Fabien Postel-Vinay (2016b), “Wage posting and business cycles.” *American Economic Review*, 106, 208–13.
- Moscarini, Giuseppe and Fabien Postel-Vinay (2017), “The relative power of employment-to-employment reallocation and unemployment exits in predicting wage growth.” *American Economic Review*, 107, 364–68.
- Moscarini, Giuseppe and Fabien Postel-Vinay (2023), “The job ladder: Inflation vs. reallocation.” Technical report, National Bureau of Economic Research.
- Mukoyama, Toshihiko, Christina Patterson, and Ayşegül Şahin (2018), “Job search behavior over the business cycle.” *American Economic Journal: Macroeconomics*, 10, 190–215.
- Postel-Vinay, Fabien and Jean-Marc Robin (2002), “Equilibrium wage dispersion with worker and employer heterogeneity.” *Econometrica*, 70, 2295–2350.
- Postel-Vinay, Fabien, Alireza Sepahsalari, et al. (2019), “Labour mobility and earnings in the uk, 1992-2016.” Technical report, Understanding Society at the Institute for Social and Economic Research.
- Schneider, Patrick (2018), “Decomposing differences in productivity distributions.” *Bank of England Working Paper*.
- Smith, Jennifer C (2011), “The ins and outs of uk unemployment.” *The Economic Journal*, 121, 402–444.
- Topel, Robert H and Michael P Ward (1992), “Job mobility and the careers of young men.” *The Quarterly Journal of Economics*, 107, 439–479.