BANK OF ENGLAND

Staff Working Paper No. 667

Matching efficiency and labour market heterogeneity in the United Kingdom

Carlo Pizzinelli and Bradley Speigner

August 2017

BANK OF ENGLAND

Staff Working Paper No. 667

Matching efficiency and labour market heterogeneity in the United Kingdom

Carlo Pizzinelli[(1)] and Bradley Speigner[(2)]

Abstract

This paper investigates how compositional changes in the UK labour market affect the matching process between vacancies and job seekers. We augment a state space representation of the aggregate matching function with a measure of job seekers' 'search intensity' that is recovered from micro-data on individual unemployment-to-employment transitions, in line with recent developments in the literature. The baseline results show that matching efficiency declined by around 15% between 1995 and 2010 but subsequently recovered by about 5 percentage points in the last six years. Compositional changes in the labour force that improved aggregate search intensity prior to the 2008 recession will tend to obscure the decline in aggregate matching efficiency unless controlled for properly. Considering broader definitions of job seekers that include marginally-attached workers and on-the-job searchers exacerbates the registered decline in matching efficiency. Changes in 'recruiting intensity' and the share of vacancies posted by different industries provide a potential explanation for some, but not all, of the initial fall in matching efficiency that preceded the 2007–08 recession. Finally, we quantitatively analyse how labour force heterogeneity and changes in matching efficiency have affected the shape and location of the UK Beveridge Curve.

Key words: Unemployment, labour heterogeneity, matching function, Beveridge Curve.

JEL classification: E24, E32, J64, J82.

(1) University of Oxford. Email: carlo.pizzinelli@economics.ox.ac.uk
(2) Bank of England. Email: bradley.speigner@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. We thank Will Abel for his essential support and advice with data preparation and seminar participants at the Bank of England for useful comments.

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/research/Pages/workingpapers/default.aspx

Publications and Design Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 email publications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

Introduction

The rate at which job seekers are matched with job vacancies is a key labour market indicator which contains useful information for understanding unemployment dynamics both over the business cycle and the longer term. The job finding rate is determined by numerous factors, including demographic trends, cyclical shifts in labour force composition, as well as other structural changes in the matching process between workers and firms. In this work, we undertake a quantitative analysis of the job matching process in the UK using micro-data on individual labour market status transitions. Specifically, we estimate a matching function—one of the key building blocks of the Diamond-Mortensen-Pissarides (DMP) model—that is generalised to incorporate heterogeneity in individual characteristics, which determine search ability as well as time-varying matching efficiency at the aggregate level.

Our sample covers the period from 1994 to 2016 and is therefore largely characterised by the recovery from the early 90s recession and the business cycle associated with the Great Recession. Although the rise in UK unemployment following the Great Recession was relatively limited in comparison to some other advanced economies, the UK unemployment rate did remain persistently elevated for several years after the initial shock. **Figure 1** (left panel) shows that the unemployment rate rose from around 5 percent to over 8 percent during the Great Recession as the vacancy rate fell, only beginning to decline in 2013 and fully recovering by the end of 2016.[1] The rise in unemployment was driven by changes in both job creation and destruction. There was a sharp spike in job destruction at the onset of the recession and a protracted decline in the job finding probability (**Figure 1**, right panel).[2]

Figure 1: Unemployment, vacancies, and transition rates in the UK

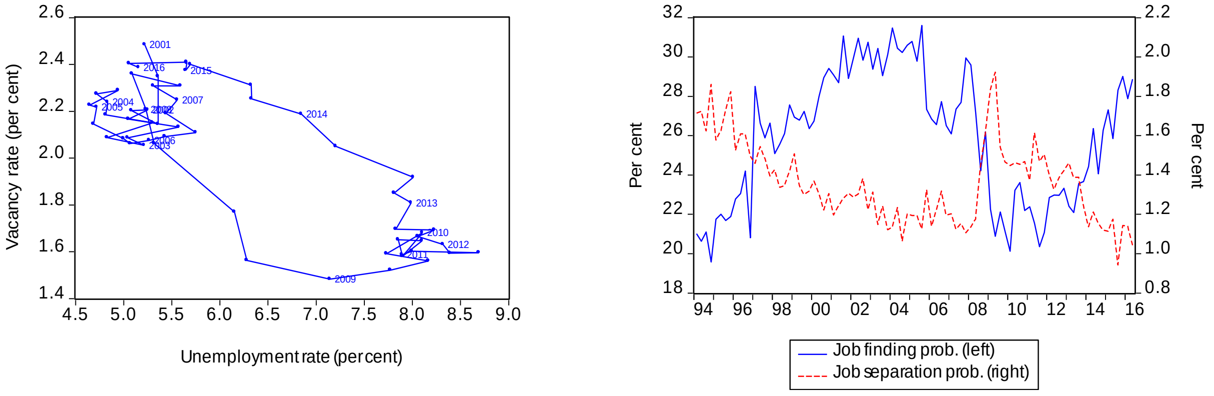

[1] The vacancy rate is defined as the ratio of the stock of vacancies to the sum of vacancies and employment.
[2] The job destruction probability is defined here as the probability of transitioning from employment to unemployment. The job finding probability is the probability of transitioning from unemployment to employment.

The primary focus of this work is on explaining the dynamics of the job finding rate in the UK both over the more recent period covering the recession as well as over the longer run covered by our available data sample. In the most basic DMP model, purely cyclical fluctuations in unemployment are generated by fluctuations in the level of labour demand, moving the labour market along a stable Beveridge curve. In this paper, we are concerned with measuring how changes in the individual characteristics of workforce participants as well as structural aggregate conditions affect the performance of the labour market in terms of matching job seekers to vacant jobs, thereby shifting the Beveridge curve.

Figure 2: Compositional changes in the pool of unemployed workers in the UK

Compositional changes in the characteristics of the searcher pool can arise from cyclical shocks disproportionately affecting workers of certain types and from longer term underlying demographic trends. **Figure 2** shows the varying composition of the pool of unemployed workers along four characteristics: age, education, sex and duration of unemployment. While the latter follows a cyclical pattern, the other three variables present clear secular trends. As employment prospects covary strongly with observable demographic characteristics, the aggregate propensity to find a job, which we call "search intensity", depends on the composition of the pool of job seekers.

From a technical standpoint, the effect that aggregate search intensity has on the labour market is similar to an efficiency improvement in the aggregate matching technology. We adopt a two-step estimation strategy based on Barnichon and Figura (2015) with some extensions. In the first step, we augment the aggregate matching function of the canonical DMP model with a term representing the average search "intensity", or "effectiveness", of job seekers that is aggregated up from individual observations at the micro-level, thereby introducing heterogeneity in the searcher pool. It is worth remarking at the outset that search "intensity"—equivalently, "ability" or "effectiveness"—is interpreted broadly. It is meant to parameterise exogenous differences in unemployment exit probabilities that are attributable to differences in observable characteristics: we do not attempt to distinguish between effort and opportunity.[3]

It then remains to estimate the path of aggregate matching efficiency, which is commonly defined as the scale parameter in the aggregate matching function. In order to do this, the second stage of the estimation process casts the model in state space form, treating aggregate matching efficiency as an unobserved time-varying state variable. Maximum Likelihood Estimation (MLE) of the state space model, conditional on the estimated path of search intensity from the first step, generates an estimate of the path of structural matching efficiency, which helps to explain changes in the job finding probability that are unaccounted for by fluctuations in labour demand and aggregate search intensity. The idea is that, by controlling for compositional changes in the observable individual characteristics of the searcher pool, an attempt is made to filter out a "purified" path of matching efficiency that is not distorted by changes in search effort.

Our main findings are summarised as follows. First, matching efficiency commenced a downward trajectory in the late 1990's, which reached its trough in 2009. Although the downward trend has no longer been apparent since then, there has only been a partial recovery in the level. Second, we do not find a lasting deterioration in matching efficiency that is associated with the 2008 recession. There was also a pronounced rise in search

[3] We recognise that using the word "intensity" may be susceptible to misinterpretation, and sometimes "employability" may be a better descriptor. The estimation process remains agnostic about the causes of such differences in employment prospects.

intensity in the first part of the sample, during the 90's, peaking in 2005 and later falling during the Great Recession. This pro-cyclicality is accounted for by changes in mean unemployment duration. Over the longer term, education is the main factor underpinning a secular upward trend in search effectiveness, but age also has some effect. Gender is found to have very little effect on employment transitions.

In the second part of the paper we consider extensions to the baseline specification in two principal directions. First, rather than focusing narrowly on unemployment, we broaden the definition of job seekers to include individuals who are "marginally attached" to the labour force and employed workers that are searching for another job. Our main conclusions are robust to this extension. Additionally, in recognition of the fact that the baseline model accounts for heterogeneity solely on the labour supply side, we also make an attempt to model heterogeneity in the *demand* for labour by accounting for variation in the "recruiting intensity" of different industries. Due to the lack of firm-level micro data, this approach is not as comprehensive as our modelling of the supply side. However, we do uncover some tentative evidence that sectoral shifts in vacancies may help to explain the fall in matching efficiency in the period 1995-2002 that is detected by the baseline model.

Our final exercise is to map our results on changing labour force composition and matching efficiency into an analysis of the Beveridge Curve. Improvements in matching efficiency and search intensity both cause inward movements of the curve and we quantify how the Beveridge curve may have shifted over time as a result. Furthermore, we also demonstrate that labour force heterogeneity implies that the composition of the pool of job seekers changes *along* the curve as the level of labour demand varies in the steady state. This channel causes the Beveridge Curve to pivot, rather than to shift, relative to the canonical model with homogenous job seekers.

The rest of the paper is structured as follows. Section 2 relates our work to numerous other studies on labour heterogeneity and matching models and outlines which strands of the literature we bring together and contribute to. Section 3 describes the estimation procedure. Section 4 presents the baseline results and Section 5 discusses extensions to the baseline. In Section 6 we report the implications of our matching function analysis for the location of the Beveridge Curve. Section 7 concludes.

2. Relation to the literature

The idea that fluctuations in the aggregate job finding rate are driven by compositional changes in the pool of job seekers dates back to at least Darby et al. (1985, 1986), who note that "cyclical unemployment is concentrated in groups with low normal exit probabilities". This was challenged by Shimer (2012), whose reassessment of the "ins and outs" of unemployment led him to conclude that job seeker heterogeneity was not a major factor driving fluctuations in the overall unemployment exit rate, which in turn was the largest determinant of unemployment fluctuations.[4] Several subsequent studies seem to go against the conclusion that shifts in the observable composition of the unemployment pool could not explain the fall in the job finding probability of US workers during the Great Recession.[5]

There is a large empirical literature focusing on matching functions—see Petrongolo and Pissarides (2001) for a useful overview of early contributions. What emerged early on from this literature is the somewhat counterintuitive finding that the rate of job matching, conditional on the level of labour market tightness, appears to be falling rather than increasing over time. Compared to the initial attempts at matching function estimation, what distinguishes the recent works is a significantly more sophisticated approach to explicitly incorporating labour force heterogeneity using detailed micro-data sets. Such "generalised matching functions" incorporate heterogeneity at the individual level and demonstrate that shifts in the average characteristics of the workforce can account for much of the residual in standard matching rate regressions (Barnichon and Figura, 2015).[6] Hall and Schulhofer-Wohl (2015) build on this work by constructing an efficiency index for a definition of job seekers that is wider than just the unemployed. By doing so, they reveal the presence of pre-crisis trends that have continued to put pressure on the Beveridge curve to shift outwards even after the 2008 recession. They do not find evidence suggesting a fall in matching efficiency related to the 2008 recession itself, which is in line with what we conclude about the UK labour market.

[4] The main conclusion of his paper was that the "outs" from unemployment were found to be quantitatively more important than the "ins" in terms of accounting for cyclical unemployment dynamics.

[5] Among the contributions, key lines of analysis include general demographic characteristics (Barnichon and Figura, 2015; Kroft et al., 2016; Bachmann and Sinning, 2016), long-term unemployment and duration-dependence of employment probability (Krueger et al., 2014; Kroft et al., 2016), unobserved heterogeneity (Ahn and Hamilton, 2016; Morchio, 2016), the labour force participation margin (Elsby et al. 2015), job-to-job transitions (Sedláček, 2014), firm size (Gavazza et al., 2016), variable search effort (Hornstein et al., 2016), firm hiring standards (Sedláček, 2014), and sectoral segmentation (Sahin et al., 2014).

[6] Barnichon and Figura (2015) also quantify an additional source of variation in matching efficiency that results from dispersion in labour market tightness across sub-sectors of the labour market.

Our two-step empirical approach fuses insights from the literature on generalised matching function estimation with other work which applies time-varying parameter methods to labour market flows data and the estimation of aggregate matching efficiency. It has been recognised for a while that matching efficiency is subject to time variation. Earlier work tended to model this instability through the use of deterministic time trends (Petrongolo and Pissarides, 2001). More recently, Sedláček (2016) applied a more flexible latent variable technique to the estimation of matching efficiency, while also specifying a generalised matching function that included a broad pool of job seekers. Hornstein and Kudlyak (2016) also make use of a Kalman filter to infer aggregate matching efficiency as the unobserved state in a similar setup. Their model endogenises variable search effort in addition to accounting for exogenous heterogeneity of the searcher pool. While most of the analyses have focused on the US[7], several pieces of work also investigate the dynamics of UK labour flows (Smith, 2011; Gomes, 2012; Sutton 2013) and the resulting compositional changes in the pool of job seekers (Elsby et al., 2011; Singleton 2017).

Sedláček (2016) found that restricting the definition of job seeker to just the unemployed biases downwards the estimated contribution of fluctuations in matching efficiency to unemployment dynamics. The intuition is that the measure of non-unemployed job seekers moves pro-cyclically, so recessions are periods during which there is less congestion from on-the-job or passive searchers in the overall matching process. We therefore check our baseline results for robustness to a different definition of the pool of job seekers, including the “marginally attached” and those searching while employed. The main findings from this extension corroborate the conclusions of Sedlacek (2016) that the estimated aggregate matching parameters can be very sensitive to the specific choices about the considered pool of searchers.

Finally, we also make a first attempt—to our knowledge—at considering the role of shifts in the composition of labour demand across industries, an issue which has not yet been investigated in the UK. The majority of the work elsewhere in the literature has tended to focus on labour supply heterogeneity as opposed to heterogeneity in labour demand, probably due to data availability. Notable exceptions include Davis et al. (2013) and Gavazza et al. (2016) who look in detail at the macro implications of firm-side recruiting behaviour.

[7] For studies on German data, see Kohlbrecher et al. (2016) and Klinger and Weber (2016).

3. Generalising the matching function

Following a longstanding theoretical literature, the DMP model of matching frictions has become the canonical approach to modelling the aggregate job finding probability in the economy. It has now become standard to assume the existence of a matching technology that relates the flow of new job matches to the stocks of vacancies and unemployment. Given the considerable supporting evidence in the applied literature, we make the conventional assumption of a Cobb-Douglas function with constant returns to scale. Defining m_t as the flow of newly formed job matches (i.e. worker-job pairs) within period *t*, the basic matching function is given by

$$m_t = \mu V_t^{1-\eta} U_t^{\eta} \tag{1}$$

where V_t denotes the supply of available job vacancies, U_t is unemployment—or the pool of job seekers more generally—and η is a positive fraction. Matching efficiency is defined as the scale parameter μ in equation (1).

In its basic form, equation (1) is an aggregate relation which abstracts from heterogeneity in the input variables V_t and U_t. However, data on the variation of job finding propensities across observable individual characteristics means that this assumption is restrictive in practice. We therefore generalise equation (1) to allow for time variation in μ and heterogeneity in the "search intensity" of different groups of job seekers.[8] Denoting the search intensity of worker type *j* as s_j, assumed to be time-invariant, the matching function in generalised form is

$$m_t = \mu_t V_t^{1-\eta} \ (s_t U_t)^{\eta} \ , \tag{2}$$

where $s_t = \sum_j \frac{U_{jt}}{U_t} s_j$. Aggregate search intensity fluctuates over time due to changes in the unemployment shares of worker types, $\frac{U_{jt}}{U_t}$. For now, attention is restricted to the unemployment pool and we will consider expanding the searcher pool in Section 5 below.

[8] There is the additional question of the underlying industrial structure of the labour market, and whether a single aggregate matching technology which pools workers and jobs from all industries is realistic. Barnichon and Figura (2015), for example, assumed that the labour market is segmented, with individuals searching only for jobs within their occupation of previous employment and geographic location, with each segment characterised by a separate matching function. We do not make this assumption because we lack data on vacancies by occupation. Despite having data on vacancies by industry, we opted for not including industry of previous employment among the first-stage controls because the data show that a surprisingly large number of individuals who lose their jobs end up finding work in a different industry.

Under the assumption of random matching—meaning that each job seeker of a given type has the same probability of being matched to a vacant job—the job finding rate of the type-*j* unemployed worker is

$$f_{jt} = \frac{s_{jt}}{s_t}\frac{m_t}{U_t} \tag{3}$$

Data on individual employment transitions are combined with a parameterisation of the function s_{jt} in order to estimate the dependence of job finding rates on individual traits using the relation in (3), following Barnichon and Figura (2015). The results are used to generate a time series for aggregate search intensity, s_t. Once this has been carried out, dividing both sides of (2) by unemployment and taking logs yields a regression equation for the aggregate job finding rate;

$$\ln f_t = \ln\mu_t + \eta\ln s_t + (1-\eta)\ln\theta_t \tag{4}$$

where $\theta_t = V_t/U_t$ is labour market tightness. We now describe in more detail how we take equations (3) and (4) to the data with a two-stage estimation approach.

First stage: Micro-estimation

Individual search intensity, s_j, is parameterised using micro-data on individual labour market transitions from the Labour Force Survey. The search function s_j is simply assumed to be an exponential function of observable characteristics;

$$s_j = \exp(\beta X_j)$$

where X is a vector of worker characteristics. From our data, we observe whether each job seeker transitioned into employment in a given time period. The log-likelihood function is therefore set up as

$$l(\beta) = \sum_t \sum_j \sum_{i\,\in\{1,N_j\}} \{y_{it}\ln(F_{jt}) + (1-y_{it})\ln(1-F_{jt})\} \tag{5}$$

where y_{it} takes a value of 1 if the individual finds a job in period *t* and 0 otherwise and F_{jt} is the discrete time-adjusted job finding probability. Given that the data are only observed at discrete intervals, the continuous time job finding rate (which is assumed to be constant within each quarter) is converted to a discrete-time quarterly probability. Formally, $F_{jt} = 1 - e^{-\frac{s_{jt}}{s_t} f_{jt}}$. The parameter vector β is then estimated by maximising the likelihood function.

Given the structure of the matching process, what matters for individual transitions is *relative* search effectiveness. That is, search effectiveness is only identified up to a normalising constant in our model. We therefore normalise aggregate search intensity to 1 over the time period used for estimation, which is four periods of quarterly data for the year 1994. As mentioned previously, the assumption is then that the estimated impact of individual characteristics on job success do not change over the rest of the sample (i.e. β is fixed over time). Changes in aggregate search intensity over time therefore occur only through changes in the shares of job seekers across the different categories. We discuss in detail the robustness of this assumption below. Once s_j has been estimated for each worker type, the time-varying aggregate s_t can be computed by multiplying each s_j by the respective share of worker type *j* for all time periods.

Second stage: Macro-estimation

The second stage of the process is to estimate the matching elasticity parameter, η, and time-varying efficiency, μ_t, conditional on s_t. In order to do so, we cast the model in state space form, treating μ_t as an unobserved time-varying state variable. The state equation for matching efficiency is assumed to be a random walk,

$$\ln\mu_t = \ln\mu_{t-1} + \epsilon_t, \quad \epsilon_t \sim N(0, \sigma_\epsilon) \tag{6}$$

where ϵ_t represents innovations to matching efficiency.

In practice, we also treat the true stock of vacancies as an additional unobserved state variable. This is due to the fact that prior to the introduction of the current national survey in 2001, vacancies data were only available from job centres. The issue that this presents is that job centre vacancies typically suffer from incomplete coverage, since it is not generally required for firms to post their job openings at job centres. Improvements in coverage over time due to modernisation can then induce a false perceived upward trend in the vacancies stock, which would bias the measured matching efficiency path. Our technical assumption is that the job centre data on vacancies is only an imperfect signal of the true underlying stock of vacancies for the sample period to which this applies (i.e., prior to 2001). The numerator of tightness, V_t, is therefore also an unobserved state, and is assumed to follow a random walk with an independently distributed error term.

Data

We restrict attention to the sample period from 1994q1 to 2016q2. This is the period for which it is possible to measure labour force transitions at the quarterly frequency using the two-quarter longitudinal UK Labour Force Survey (LFS). For each quarter, individual observations contain information regarding demographic characteristics as well as labour force status for the previous and current quarters. The linked data allow us to observe transitions from unemployment to employment as well as continued spells of unemployment. The job finding probability F_t is defined as the fraction of workers transitioning into employment from unemployment in a given quarter. We apply the recommended survey weights to make the sample representative of the UK population. The LFS is also used to construct the aggregate unemployment level, U_t, needed to compute labour market tightness $\theta_{t.}$.

The vacancies data that we use come from two sources. For the more recent part of our sample, covering 2001q3 to 2016q2, a national survey of job vacancies is available from the ONS Vacancy Survey. For the initial part of the sample period, data on the stock of vacancies is obtained from vacancies at job centres.

4. Baseline Results

First stage (micro) estimation

Table 1 presents the coefficients from the first stage MLE. The "reference group" is comprised of male individuals with no GCSE qualification aged 16 to 25 who are short-term unemployed. Individuals in this category are 1.5 times (i.e., exp(0.401)) more likely to find a job than the average searcher in 1994. Educational attainment is positively associated with the chances of finding a job. Meanwhile, being female, older age groups, and duration of unemployment are all negatively correlated with job finding probability. Figure A1 in the appendix shows the implied distribution of relative search intensity across the population based on our regression results for the year 1994, illustrating large variation across groups.

Table 1: **Maximum-likelihood coefficients from the first (micro) stage estimation**

	Coef.	**Std. Err.**	**P-val.**
Reference group	0.401	0.051	0.000
Female	-0.047	0.036	0.184
Age 26-35	-0.100	0.044	0.022
Age 36-45	-0.020	0.049	0.681
Age 46-55	-0.230	0.055	0.000
Age 56-65	-0.645	0.083	0.000
Other qual.	0.242	0.055	0.000
GCSE qual.	0.331	0.053	0.000
A-level qual.	0.331	0.052	0.000
Higher educ.	0.601	0.058	0.000
3-6 months	-0.264	0.047	0.000
6 month-1 year	-0.570	0.049	0.000
1-2 years	-0.915	0.055	0.000
2+ years	-1.435	0.061	0.000

Number of obs. = 16,961

Wald chi2(13) = 1131.60

Prob. > chi2 = 0.0000

Log pseudolikelihood = -4885.2914

Once the vector β has been estimated, we obtain a time series of aggregate search intensity using the appropriate unemployment weights, which is plotted in **Figure 3** together with the job finding rate. Comparing the two series, aggregate search intensity is found to follow cyclical fluctuations similar to those of the job finding rate but also has an upward trend that is absent from the latter. **Figure 4** decomposes aggregate search intensity into the observable characteristics that we measure. Fluctuations in the long-term share of unemployment account for the swings in search effectiveness at the business cycle frequency. On the other hand, the secular trend is accounted for mainly by education and, to a much lesser extent, age. Over time, the unemployment pool has become better educated, yielding an overall improvement in search effectiveness. Gender has not had a measurable impact on search intensity.

Figure 3: Estimated aggregate search intensity and the job finding rate

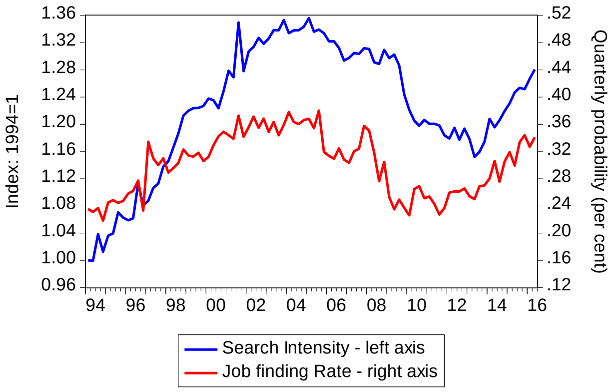

Figure 4: Decomposing changes in search intensity (4-quarter moving average)[9]

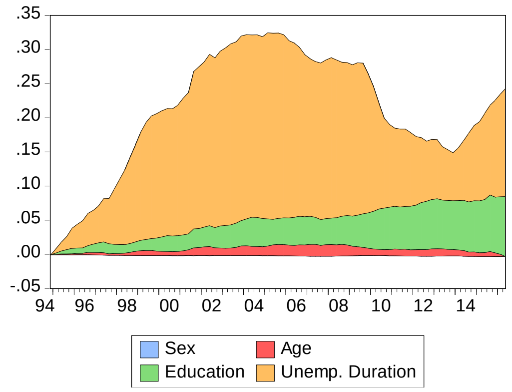

Second stage (macro) estimation

The search intensity index which was computed in the first stage of the estimation process is now treated as an observable variable in the aggregate matching regression (4). Standard Kalman filtering techniques are then applied to estimate the path of matching efficiency. For comparison, results are also reported for a version of the state space model excluding search intensity in **Table 2**. The elasticity of the job finding rate with respect to tightness is close to 0.3 in both cases, which is consistent with other results in the empirical literature, but the point estimate is about 20% higher when search intensity is omitted from the model. This suggests that specifications which fail to control for fluctuations in search intensity will have an upwardly biased matching elasticity estimate, deriving from the pro-cyclicality of search effectiveness caused by fluctuations in average unemployment duration as described previously.

Figure 5 plots the path of matching efficiency with and without controlling for variations in the search intensity of the unemployed. Failing to control for time variation in the aggregate composition of job seekers significantly affects the resulting path of matching efficiency. When search intensity is not controlled for, it appears as though matching efficiency was stable, or even slightly rising, in the earlier part of the sample, before beginning a decline

[9] To compute the decomposition, we slightly changed the estimation of the likelihood function, using the linear functional form $s_{ij} = \beta X_i$ rather than the exponential one. Given the fact that all X's are dummy variables, this change allows for aggregate search intensity to be computed simply as $s_t = \sum_j s_j U_{jt} \ / U_t$, where each j represents a category of sex, age, education, and unemployment duration. The reason why we do not use this approach in the main model is that the linear formulation does not guarantee that the estimated s_j's are positive, although they all turn out to be for the estimation run our sample of data.

prior to the Great Recession which has stalled but not unwound since. Controlling for heterogeneity, the deterioration in matching efficiency commences earlier in the sample, during the mid-90's. The overall drop is larger, falling fairly consistently from the beginning of the sample until the 2008 recession. We fail to find a sustained negative impact of the 2008 recession on matching efficiency, as the decline started well before. The estimated decline in matching efficiency predating the recession is large; the labour market was estimated to be about 10% less efficient at matching workers with job openings in 2008 than it was in 1995. However, there has been a partial recovery since then, implying that matching efficiency is about 5 percent lower at the end of the sample compared to the beginning.

Table 2: Matching elasticity estimates

Matching elasticity	**(1)** **With search intensity**	**(2)** **Without search intensity**
1- η	0.277	0.333
	(0.048)	(0.042)
Log likelihood	234.77	234.82
Observations	90	90
Search Intensity	Yes	No

Figure 5: The estimated path of matching efficiency (1995q1=1)

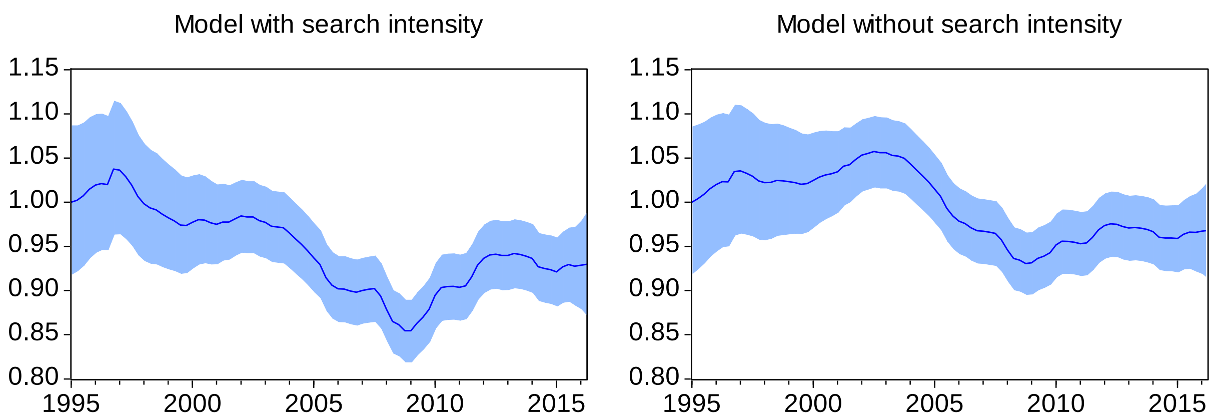

Sensitivity analysis

Before discussing extensions to the model, we test the robustness of the baseline results with a few small variations in the first-stage estimation. In the first case, we exclude Government Training Schemes (GTS) from the unemployment-to-employment transitions. GTS inflows are reflective of government policies rather than labour demand. Therefore, flows into GTS's are arguably not subject to the same frictional matching process as regular job openings, and omitting such flows from the definition of the job finding probability may have implications for our results, with our baseline treatment of the data possibly overstating the true vacancy yield particularly during the recession. Although GTS's constitute only a small portion of all transitions into employment, the share of newly employed workers in GTS's increased significantly over the 2008 recession, particularly for the long-term unemployed. The first column of Table A1 shows that the coefficients of the first stage are not particularly affected by this restriction on the type of exits from unemployment.

The second robustness check addresses the issue of true duration of a joblessness spell. The LFS includes a question on the time since an individual last had a job. For a non-trivial portion of the sample, the duration of unemployment is shorter than the total time out of a job, implying at least one transitional period out of the labour force. Being out of a job, whether inactive or seeking, is detrimental for human capital and hence employment prospects. Only considering unemployment duration can therefore neglect long periods of non-employment for some individuals.

We assess the results' sensitivity to this issue with two different specifications. In the first one, we generate a binary variable for having spent time in the inactive state. In the second one, we replace unemployment duration with joblessness duration, regardless of whether the individual was searching or not for its entirety. Columns 2 and 3 of **Table A1** show that the coefficients of the first-stage MLE for these two specifications are very similar to those of the baseline model. In column 2, the dummy variable for having had a spell of inactivity is negative and statistically significant. This result may either indicate an adverse impact of having spent time outside the labour force or simply account for the extra length of the joblessness spell for this group of workers. Interestingly, in this specification the coefficients of the duration categories are similar in value to the baseline model. However, when using the effective time out of a job (column 3), the magnitude of the coefficients is attenuated, implying a less adverse effect of duration. The possible explanation for the attenuation is that a portion of workers reporting a short duration of unemployment in fact had spent a long time out of the labour force, which compresses the true disparities in job finding prospects.

Moving on to the second stage, **Table A2** reports the estimated elasticity parameter of the matching function for the three robustness exercises both with and without search intensity. In all cases, the estimates are in line with the baseline results from **Table 2**: the elasticity with respect to vacancies estimated with search intensity is lower than in the canonical model. Finally, **Figure A2** shows the smooth forecast of the path of matching efficiency. Consistent with the baseline result, for all three checks matching efficiency in the intensity-augmented models is lower than in the canonical model starting from around the late 90's. Interestingly, when excluding government training schemes, the canonical model yields an upward path for matching efficiency in the period 1995-2003. The following downward trajectory, however, is in line with the other specifications.

The choice of the reference year for the first stage is also a potentially important factor in the results. The key assumption for the first stage is that the β's, estimated on a baseline year only, are constant over the years, and hence the contribution of different individual characteristics to search intensity does not vary with time. To assess the robustness of this assumption, we repeat the MLE on each year separately and plot the estimated coefficients for each year (with 95% confidence intervals) in **Figure A3**. The plot shows that most coefficients exhibit only minor variations over the years. Only three coefficients show substantial changes or clear trends: age 56-65, unemployed for longer than 2 years, and the reference group. In itself, variation in the values is not necessarily a problem as it may simply result from changes in the composition of job seekers. For instance, as the unemployed pool becomes more educated, the advantage of high-education workers in finding a job relative to the "average" worker falls and hence the respective coefficient would fall too. However, changes may also be driven by true shifts in relative search intensities.

The latter case would result in changes in the path of the aggregate s_t over time compared to our baseline results. To assess the impact of different coefficients on the aggregate analysis, **Figure A4** plots aggregate search intensity, normalized to 1 in 1994q1, using the β's from different years: 1994, 2000, 2005, 2010, and 2015. It is clear from the graphs that qualitatively the results do not change based on the year used for the first stage. In all cases, search intensity presents an upward path until the early 2000's, followed by a dip and partial recovery after the Great Recession. Quantitatively, the main difference arises from the magnitude of the fluctuations. For 2005, intensity peaks at 1.4, while for 2010 and 2015 the maximum value is slightly below 1.3. The 1994 series is somewhat in the middle of these extremes, meaning that it can be interpreted as being a more balanced candidate to serve as reference year.

5. Extensions

So far, we have associated "job seekers" with the ILO definition of unemployment. In this section, we first present an extension that involves an expansion to the standard definition of unemployment which incorporates "marginally attached" (MA) individuals and workers undertaking on-the-job search (OJS). Our second extension proposes a preliminary attempt to account for heterogeneity on the labour demand side, which remains relatively unexplored in the wider literature.

5.a Expanding the set of job searchers

In the baseline model, we only consider transitions from unemployment to employment. However, this flow accounts for less than half of all newly employed workers. Therefore, we now expand the definition of job seekers to include job-to-job moves and employment inflows from inactivity. Technically, we re-interpret U_t in the matching function (2) as a measure of all job seekers, not just the unemployed, but also on-the-job searchers (OJS) and "marginally attached" (MA) inactive workers, each with an associated level of search intensity. What distinguishes these groups of seekers from the unemployed is the intensity of job search in the matching function, and possibly the degree to which observable characteristics affect their job finding probability.[10]

Following Gomes (2012), OJS individuals are defined as those employed workers who state that they are looking for another job. For these workers, a job-to-job transition is identified in the LFS micro-data as an OJS worker who is employed in both quarters but whose employment tenure in the second quarter is below three months. This measure possibly overstates the total number of job-to-job flows because it does not account for possible spells of involuntary unemployment between the two jobs. The marginally attached category is

[10] For this reason, we also allow for returns to education, age, and sex to differ for each group.

defined as those individuals not in the labour force who are not actively searching for work but would be willing and able to start a job in the next two weeks.

Some descriptive analysis of the main differences between unemployed searchers and MA and OJS individuals, provided by **Figure 6,** offers several insights. While OJS activity is highly procyclical, the level of MA searchers remains fairly stable over the years.[11] As a result, the composition of the entire pool of job seekers varies over time, with the unemployed and OJS workers accounting for the main changes. The share of OJS workers rose from 30 percent in 1994 to more than 40 percent by 2007. It then dropped sharply during the Great Recession and eventually recovered beyond its 2007 point by 2016. The flow rate into employment varies dramatically across the three groups, both in its level and in the magnitude of its fluctuations. The job finding rate of the unemployed is the highest and the most pro-cyclical. OJS workers have lower chances of finding new employment, MA searchers lower still. These two transition rates are also less volatile than the UE transition probability. Furthermore, while the latter has almost returned to its pre-2007 peak, the job OJS finding rate remains below it.

The bottom right panel of **Figure 6** shows the estimated path of aggregate search intensity for each of the different job seeker definitions (the estimated coefficients are reported in **Table A3**). When including MA and OJS types, aggregate search intensity (black line) only rises moderately in the first decade of the sample and by 2016 it is almost back to the original level. Compared to the baseline series, s_t has a lower pre-2007 peak and a short-lived spike due to the sharp fall in OJS in 2008.

Table 3 presents the second-stage state space estimation results for the extended sample with and without search intensity. Unlike the baseline case, the addition of search intensity leads to an estimate of the matching elasticity with respect to vacancies that is higher (by about 0.05). This suggests that not controlling for fluctuations in the number of OJS job seekers can bias downwards the elasticity estimate since their share in the aggregate searcher pool is procyclical and they have a low likelihood of matching.

[11] A caveat: the relative acyclicality of the raw figure of MA workers may mask cyclical fluctuations in the transitions from inactivity to unemployment. For instance, if MA individuals were to be more likely to actively search for work (and hence be classified as unemployed) during recoveries, the constant MA level would imply a commensurate increase in the inflows into MA status from other labour force statuses. Singleton (2017) provides a comprehensive analysis of flows in and out of inactivity.

Figure 6: Comparison of unemployed, MA, and OJS: level, share, job finding probability and estimated search intensity.

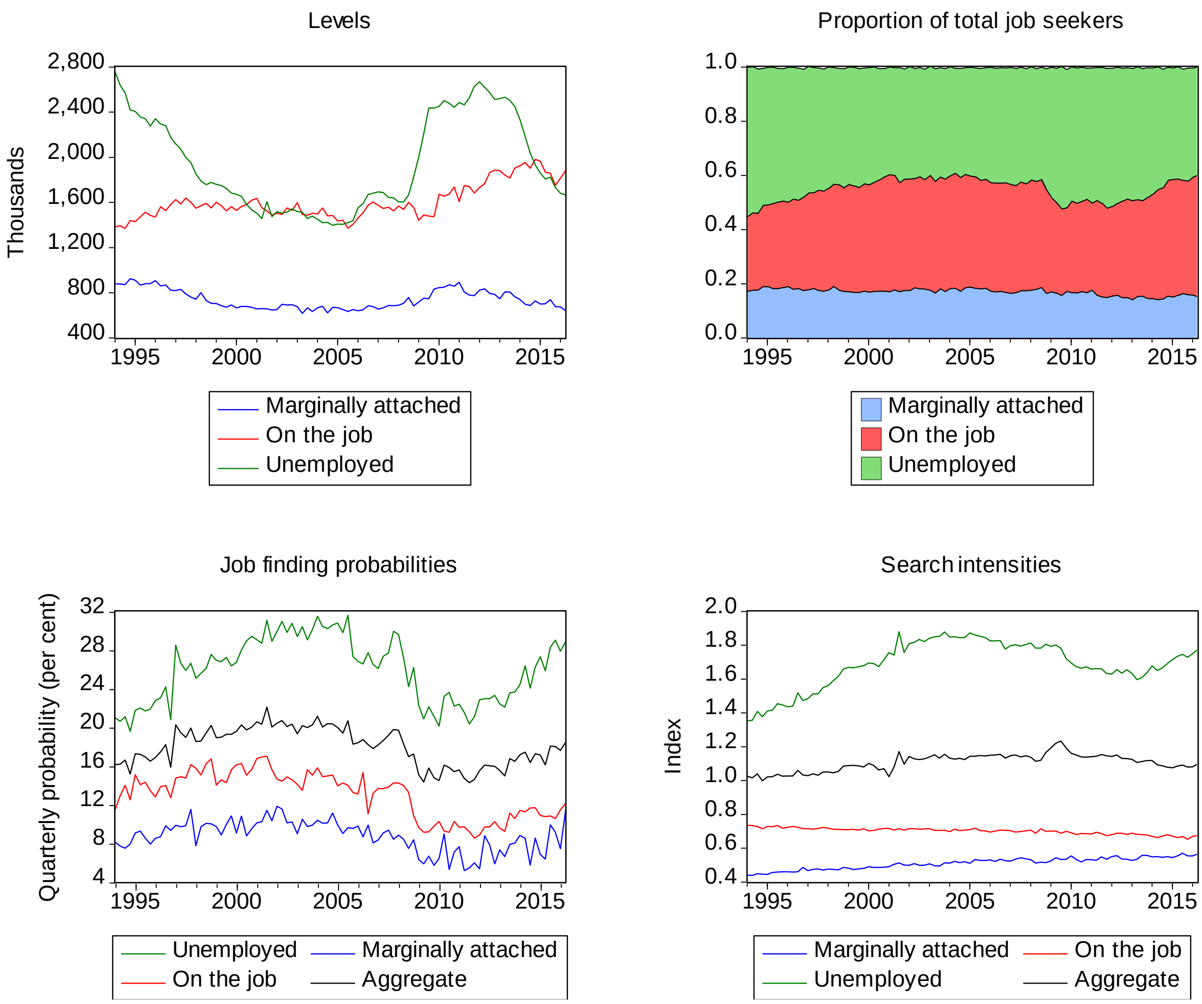

Table 3: Second-stage estimation results for the model with the expanded set of job seekers.

Matching elasticity	**(1)**	**(2)**
1- η	0.341	0.292
	(0.050)	(0.048)
Log likelihood	246.3781	246.3871
Observations	90	90
Search Intensity	Yes	No
Sample	U+MA+OJS	U+MA+OJS

The resulting Kalman smoothed path of matching efficiency for the full U+MA+OJS sample is shown in **Figure 7**. Comparing the results of this model with the baseline one from **Figure 5** (reproduced as the red lines in **Figure 7**), two differences are visible. First, the path of matching efficiency prior to the Great Recession has a sharper downward slope, reaching a deeper trough. In the intensity-augmented specification, the minimum search intensity is 76 percent of the initial value. This trough is 10 percentage points lower than in the baseline sample that is restricted to unemployed searchers. Second, the path of the model without search intensity remains almost flat until 2005 and does not present the slight increase visible in **Figure 5** for the same period.

Figure 7: Matching efficiency assuming an expanded job seeker definition (1995=1).

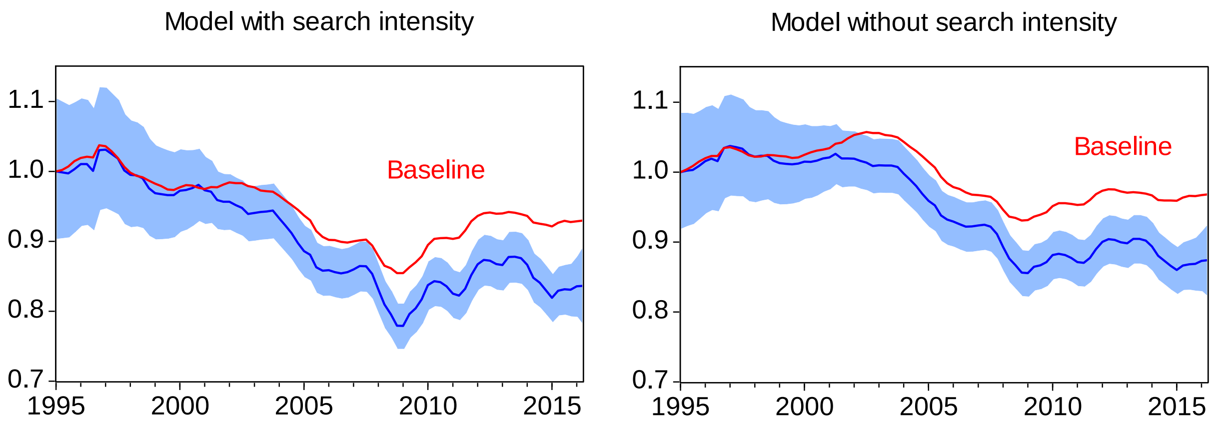

The explanation for the second observation originates from the differences between unemployed workers and the OJS pool. As outlined above, the dynamics of the job finding rate are affected by the compositional changes occurring within the expanded pool of job seekers. Because of the increase in the share of OJS workers in the period 1995-2005, the aggregate job finding rate does not rise as steeply as for when only the unemployed are included in the searcher pool.

The results of this section indicate that the halt in the decline of matching efficiency in 2008 appears to be robust to different methods of controlling for variable search. However, the resulting dynamics from adding MA and, more importantly, OJS workers show that focusing narrowly on unemployment is not representative of job search behaviour in the broader labour market.

So far, the analysis has only controlled for variable search intensity along the supply side of the labour market. In particular, we have assumed that all vacancies are supplied by firms which recruit with homogenous effectiveness. In the next section, we describe how we attempt to relax this assumption and how it influences our results.

5.b Recruiting Intensity

While our model controls for observable characteristics which influence the search effectiveness of job seekers, this only captures half of the story in a two-sided labour market. Recruiting intensity of firms can play an equally important role. Recent research for the US (Gavazza et al., 2016) has demonstrated that firms tend to adjust recruiting effort as the number of job seekers per vacancy changes, leading to fluctuations in vacancy yields as recruitment effort drops during recessions.

One direction in which we can take a tentative step is by exploiting cross-sectional variation in recruiting effectiveness across industries. To the extent that vacancy yields differ across sectors, then either business cycle fluctuations, trend shifts in industrial composition, or changes in the recruiting intensity within each industry may all affect measured aggregate matching efficiency.

Figure 8 shows the shares of vacancies and of flows from unemployment to employment for each of ten industries.[12] Two observations can be made. First, for most industries there is a wedge between the vacancy share and the hiring share, indicating differences in recruiting intensity across sectors. Second, some industries display large changes over time in the way the two shares track each other, particularly around the Great Recession. For instance, Financial Services experienced a permanent drop in the share of matches but not in the share of vacancies, indicating a potential fall in recruiting intensity. The opposite occurred for Construction and the Public Sector.

[12] Vacancies by industry are produced using the Jobs Centre Data from NOMIS until 2001 and the ONS series VACS02 from 2001 onwards. To compute recruiting intensity over the whole period we chain-link the NOMIS series to the ONS ones. The industry classification, which changed is harmonized to create a consistent set of industries throughout the sample. The LFS contains a question on current industry of employment, which is used to measure flows from unemployment into different industries using the same classification as for vacancies.

Figure 8: Match and vacancy shares by industry[13]

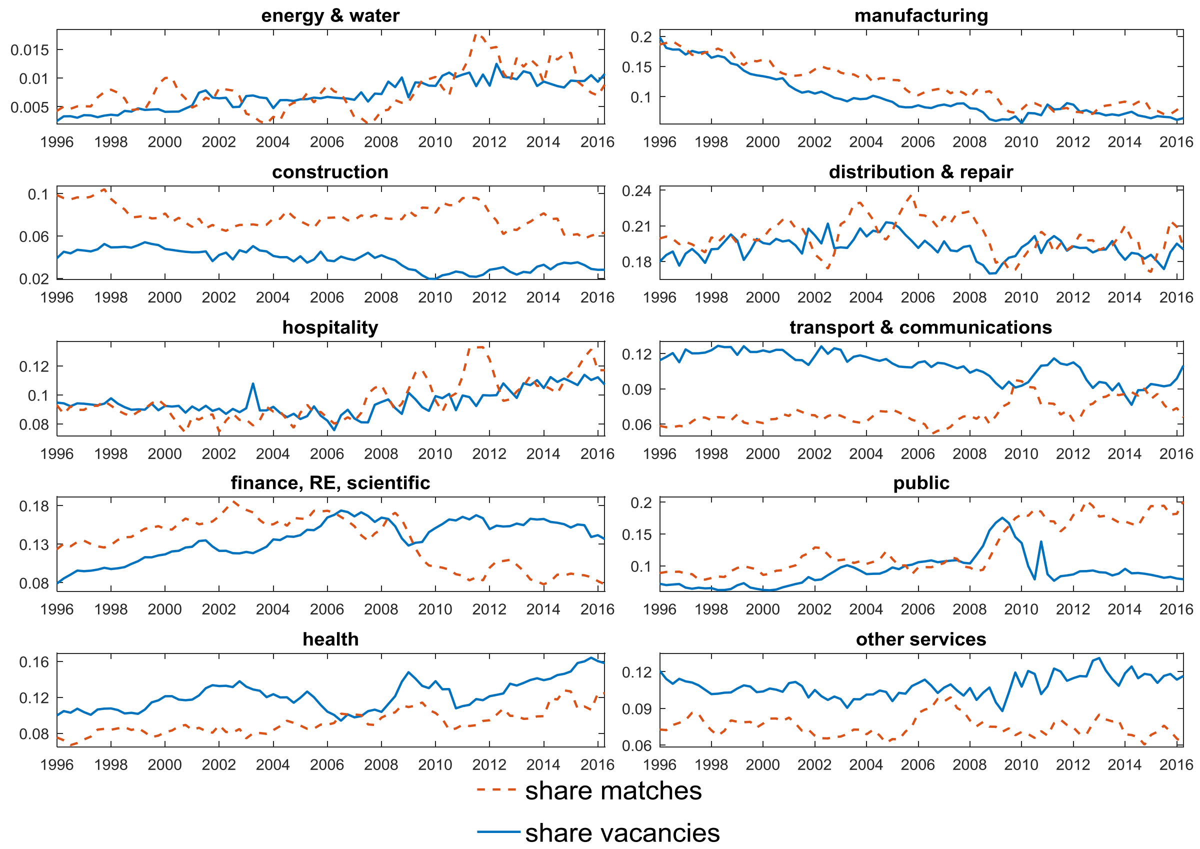

Under the assumption that job seekers search in all industries (i.e. search is not segmented), the aggregate matching function can be expanded to include recruiting intensity through a vacancy-augmenting term as follows:

$$m_t = \mu_t (r_t V_t)^{1-\eta} (s_t U_t)^{\eta} , \tag{9}$$

where $r_t = \sum_i \frac{V_{it}}{V_t} r_{it}$. The term r_{it} represents the industry-i specific recruiting intensity, which—unlike for search intensity—we assume to be time-varying. The first-stage estimation for search intensity allows for a wide set of covariates in age, education, duration, and sex. For recruiting intensity, industry is the only dimension of heterogeneity. In practice, we choose to allow for the r_{it}'s to vary over time in relation to a reference industry. Consequently, the discrete-time probability for an unemployed individual of type j to find a job in industry i is

[13] Both shares are reported as 4-quarter moving averages.

$$F_{ijt} = \frac{r_{jt}V_{jt}}{r_t V_t} F_{jt} = \frac{r_{jt}V_{jt}}{r_t V_t}(1 - e^{-f_{jt}}) \tag{10}$$

Hence the probability of entering a job in industry *i* conditional on finding a job is independent of the worker's individual characteristics, which implies that the relative recruiting intensities of different industries can be estimated separately from the search intensity of individuals. In each period *t,* the set of relative recruiting intensities r_{it}'s can be estimated as the set of coefficients that satisfy the equations

$$\frac{r_{it}V_{it}}{r_{\hat{\imath}t}V_{\hat{\imath}t}} = \frac{m_{it}}{m_{\hat{\imath}t}} \tag{11}$$

where one industry $\hat{\imath}$ is taken as the reference industry in all periods, such that $r_{\hat{\imath}t} = 1 \; \forall \; t$.

The r_{it}'s recovered through this method represent recruiting efforts *relative* to the reference industry. However, under the assumption that the intensity in the reference industry has not changed over time, the aggregate r_t can also be interpreted as absolute changes in recruiting intensity. Based on the graphic evidence of **Figure 8**, manufacturing is the industry that seems closest to satisfying this assumption. The shares of vacancies and of hiring have a correlation of 0.85, signalling a stable relationship between the two. Taking manufacturing as the reference group, whose recruiting effort is constant and normalized to 1, changes in all the other r_{it}'s and in the vacancies V_{it} determine the progression of aggregate recruiting intensity.

Figure 9 reports the estimated changes in *r*'s and the derived aggregate intensity. Aggregate recruiting intensity was on a slow downward trend for the period 1995-2008, initiating a recovery in the aftermath of the Great Recession. Clearly, the relative ability of firms to hire in different sectors seems to vary over time, such that compositional shifts cannot be considered the only drivers of recruiting intensity.

Including our measure of recruiting intensity in the state-space estimation of the matching function is straightforward, as the signal equation becomes:

$$\ln f_t = \ln \mu_t + \eta \ln s_t + (1-\eta) \ln r_t + (1-\eta) \ln \theta_t \tag{12}$$

Table 4 presents the estimates of the parameters adding recruiting intensity to the aggregate matching function. We carry out the MLE on the sample comprised only of the unemployed. As shown in the left panel of **Figure 10**, over the period 1995-2002, matching efficiency remains fairly constant, before turning downwards until the outbreak of the Great Recession.

The inclusion of recruitment intensity tempers, but does not eliminate, the pre-crisis decline in matching efficiency compared to the baseline model. This is because over the period 1995-2002 the rise in search intensity and the fall in recruiting intensity effectively offset each other, leaving matching efficiency relatively stable. While this extension can therefore explain some of the decline in matching efficiency at the beginning of the sample, the key result remains that we observe deteriorating matching efficiency before the crisis, at least to some extent, which has only partially unwound since.

Figure 9: Recruiting intensity by industry

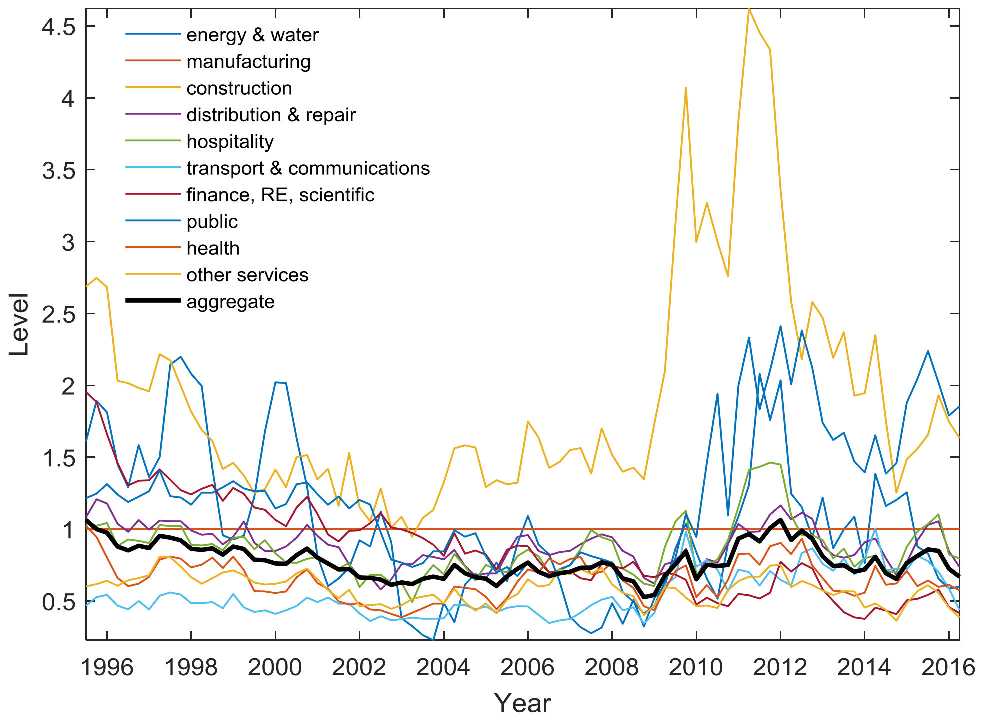

Table 4: Second-stage estimation results with recruiting intensity

Matching elasticity	**(1)**	**(2)**
1- η	0.230	0.248
	(0.041)	(0.041)
Log likelihood	232.540	228.992
Observations	84	84
Search Intensity	Yes	No
Recruiting Intensity	Yes	Yes
Sample	U	U

Figure 10: Matching efficiency controlling for recruiting intensity (1995=1)

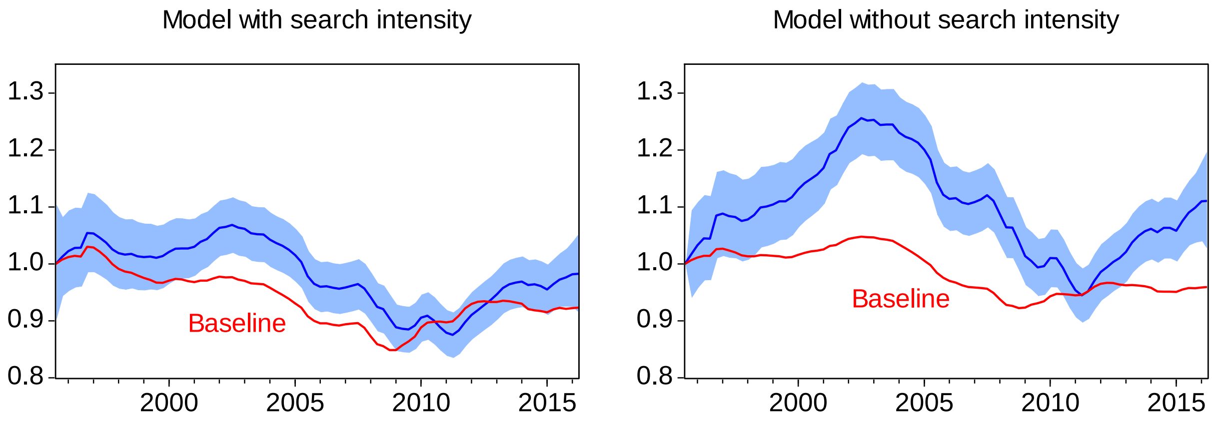

6. Beveridge Curve

In this section, we map changes in matching efficiency and search intensity into shifts in the location of the Beveridge curve. In the standard DMP model with homogeneous labour, the implicit steady-state relationship between vacancies and unemployment is given by

$$\bar{U} = \frac{\delta}{\delta + F\left(\frac{\bar{V}}{\bar{U}}\right)}, \tag{13}$$

where δ is the exogenous separation rate, $F\left(\frac{\bar{V}}{\bar{U}}\right)$ is the job finding rate, and the bar over unemployment and vacancies implies steady-state values. In our model with heterogeneous job searchers, the steady state composition of unemployment is itself a function of labour market tightness. The inclusion of worker heterogeneity entails a much larger set of steady-state conditions. For a given value of vacancies, the aggregate level of unemployment is computed by solving a set of equations similar to (13), one for each worker type. Furthermore, while most worker characteristics are assumed to be fixed, duration dependence of employment prospects implies that transitions from short-term to long-term unemployment (more specifically, across the different duration categories available in the LFS) must also satisfy steady state restrictions.[14] To measure the impact of labour force heterogeneity, we compute the Beveridge Curve with the set of conditions outlined above using the matching efficiency value for 2005q1 and the labour force composition from the LFS for the same quarter. For comparison, we then compute the homogenous Beveridge Curve using equation

[14] For simplicity, 10-year age groups are assumed to be fixed in the exercise.

(13), calibrated to have the same value of V when U is equal to 5.5 percent. As shown in **Figure 11**, heterogeneity *pivots* the Beveridge curve relative to a model with homogenous unemployment, effectively flattening the curve. When vacancies are high (low), unemployment is lower (higher) than in homogeneous case. Heterogeneity therefore increases the implicit elasticity of steady-state unemployment with respect to vacancies.

Figure 11: The Beveridge Curve with heterogeneity

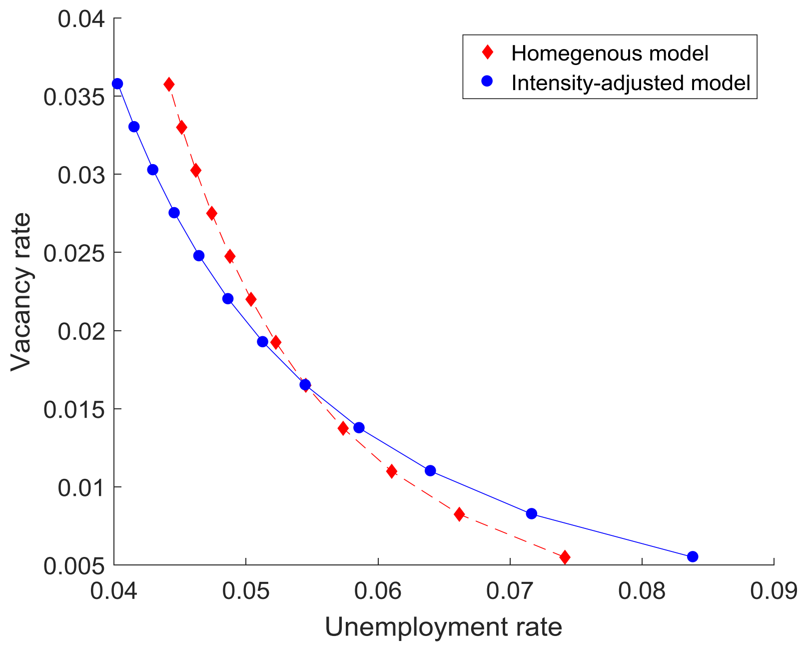

Figure 12 shows how the composition of steady state unemployment changes with the level of labour demand, illustrating that variation in the duration distribution of unemployment accounts for the pivot in **Figure 11**. The intuition is as follows. Higher vacancies raise job finding prospects for all groups, but in the steady state this shift implies that fewer workers reach long-term unemployment and so the long-term share declines, raising the average level of search intensity.

This framework can shed light on possible shifts of the Beveridge Curve over time. Improvements in matching efficiency or in the search intensity of the labour force shift the Beveridge curve inwards, so that a given level of job openings will be associated with a lower steady state unemployment rate. Our results from the previous section indicate that over the initial part of the sample, before the 2008 recession, the improvement in job seeker quality was offset by a decline in general matching efficiency. These counterbalancing forces tended to keep the Beveridge curve relatively stable.

Figure 12: Changing composition along the Beveridge curve

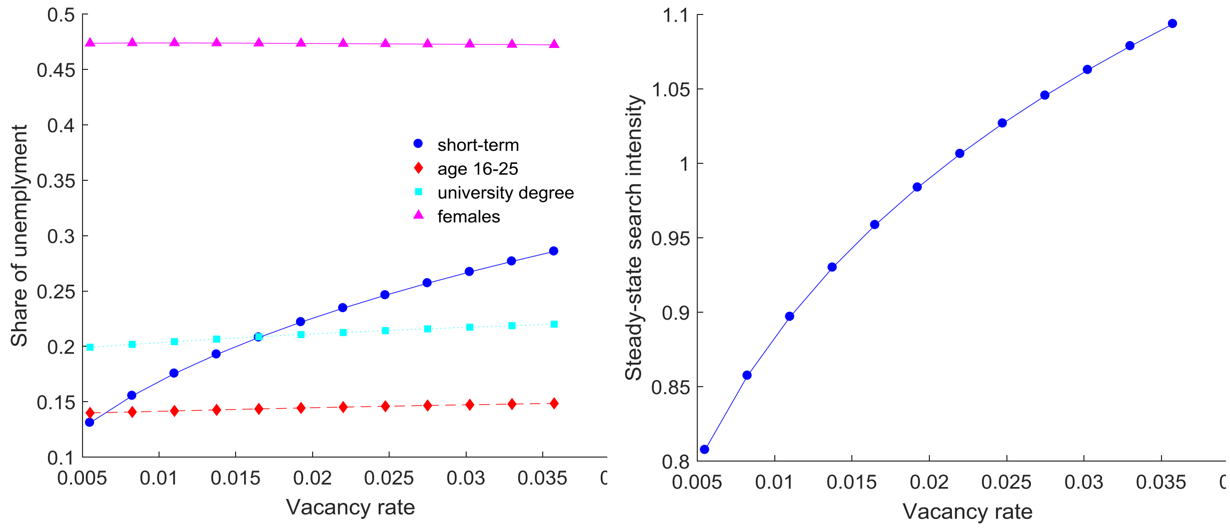

In more recent years, matching efficiency stopped declining but the search intensity of the labour force continued to rise. The implication is that the Beveridge curve may have shifted inwards due to the ongoing improvement in search quality (**Figure 13**).

Figure 13: Shifts in the Beveridge curve

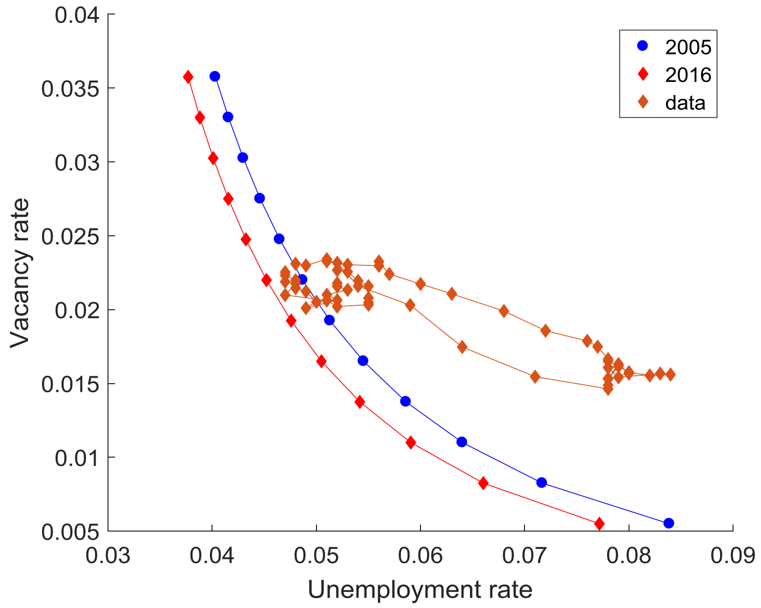

7. Conclusion

In this work, we investigate the role of labour force heterogeneity and matching efficiency in determining unemployment-to-employment flows in the UK. Following the approach of Barnichon and Figura (2015), we expand a canonical aggregate matching function with a term representing the average search intensity of the unemployment pool. Search intensity is estimated with micro-data on individual transitions into employment using UK Labour Force Survey data. In the second stage of our two step procedure, the aggregate matching function is estimated as a state space model in which efficiency is an unobservable time-varying process.

We find that search intensity has progressively increased since the mid-90s along with the rise in educational attainment. Additionally, the duration dependence of unemployment exit rates means that aggregate search intensity also contains a quantitatively important cyclical component associated with the share of short-term unemployment. The estimated path of matching efficiency shows a downward trend that begins prior to the 2008 recession, partially recovering since then.

We also consider two extensions to the baseline model. First, we expand the sample to on-the-job-search workers and marginally attached individuals, two groups of workers with very different characteristics from the unemployed. The results on the expanded pool show an even larger drop in matching efficiency in the first part of the sample. The key takeaway from this exercise is that the outcome of the analysis depends on the boundaries of the definition of "job seeker". The second extension models heterogeneity in recruiting capacity across industries with a similar approach. We find that a fall in aggregate recruiting intensity may provide a partial explanation for the decrease in matching efficiency up to 2002 found in the baseline model.

In the final section, we quantify how our estimates of search intensity and matching efficiency may have shifted the Beveridge curve over time. Labour heterogeneity is shown to increase the elasticity of steady state unemployment with respect to vacancies due to endogenous changes in the steady state duration distribution of unemployment along the curve, which acts to flatten the Beveridge curve relative to a canonical model of homogenous job searchers. In the last decade, improvements in the composition of the labour force and a partial recovery in matching efficiency are likely to have shifted the Beveridge Curve inward.

References

Ahn, H. J., & Hamilton, J. D. (2016). Heterogeneity and Unemployment Dynamics. *NBER Working Paper Series, No. 22451*.

Bachmann, R., & Sinning, M. (2016). Decomposing the Ins and Outs of Cyclical Unemployment. *Oxford Bulletin of Economics and Statistics*, *78*(6), 853–876.

Barnichon, R., & Figura, A. (2015). Labor Market Heterogeneities and the Aggregate Matching Function. *American Economic Journal: Macroeconomics*, *7*(4), 222–249.

Darby, M. R., Haltiwanger, J., & Plant, M. (1985). Unemployment Rate Dynamics and Persistent Unemployment under Rational Expectations. *American Economic Review*, *75*(4), 614–637.

Darby, M., Haltiwanger, J., & Plant, M. (1986). The Ins and Outs of Unemployment: The Ins Win. *NBER Working Paper, No. 1997.*

Davis, S. J., & Haltiwanger, J. (2014). Labor Market Fluidity and Economic Performance. *NBER Working Paper Series, No. 20479*.

Elsby, M. W., Hobijn, B., & Şahin, A. (2015). On the importance of the participation margin for labor market fluctuations. *Journal of Monetary Economics*, *72*, 64-82.

Elsby, M. W. L., Smith, J. C., & Wadsworth, J. (2011). The role of worker flows in the dynamics and distribution of UK unemployment. *Oxford Review of Economic Policy*, *27*(2), 338–363.

Gavazza, A., Mongey, S., & Violante, G. L. (2016). Aggregate Recruiting Intensity. NBER Working Paper No. 22677 .

Gomes, P. (2012). Labour market flows: Facts from the United Kingdom. *Labour Economics*, *19*(2), 165-175.

Hall, R. E., & Schulhofer-Wohl, S. (2015). Measuring Matching Efficiency with Heterogeneous Jobseekers, 1–30. Mimeo.

Hornstein, A., Kudlyak, M., (2016). Estimating Matching Efficiency with Variable Search Effort. Federal Reserve Bank of San Francisco Working Paper Series, 2016-24.

Klinger, S., & Weber, E. (2016). Decomposing Beveridge Curve Dynamics By Correlated Unobserved Components. *Oxford Bulletin of Economics and Statistics*, *78*(6), 877–894.

Kohlbrecher, B., Merkl, C., & Nordmeier, D. (2016). Revisiting the matching function. *Journal of Economic Dynamics and Control*, *69*, 350–374.

Kroft, K., Lange, F., Notowidigdo, M. J., & Katz, L. F. (2016). Long-Term Unemployment and the Great Recession: The Role of Composition, Duration Dependence, and Nonparticipation. *Journal of Labor Economics*, *34*(S1), S7–S54.

Krueger, A., Cramer, J., & Cho, D. (2014). Are the long-term unemployed on the margins of the labor market? *Brookings Papers on Economic Activity*, 229–300.

Morchio, I. (2016). Work Histories and Lifetime Unemployment. Mimeo.

Petrongolo, B., & Pissarides, C. (2001). Looking into the black box: A survey of the matching function. *Journal of Economic Literature*, *39*(2), 390–431. Retrieved from

Şahin, A., Song, J., Topa, G., & Violante, G. L. (2014). Mismatch unemployment. *The American Economic Review*, *104*(11), 3529-3564.

Sedláček, P. (2014). Match efficiency and firms' hiring standards. *Journal of Monetary Economics*, *62*, 123–133.

Sedláček, P. (2016). The aggregate matching function and job search from employment and out of the labor force. *Review of Economic Dynamics*, *21*, 16–28.

Shimer, R. (2012). Reassessing the ins and outs of unemployment. *Review of Economic Dynamics*, *15*(2), 127–148.

Singleton, C. (2017). Long-term unemployment and the Great Recession: Evidence from UK stocks and flows. Mimeo.

Smith, J. C. (2011). The Ins and Outs of UK Unemployment. *Economic Journal*, *121*, 402–444.

Sutton, A. (2013). On the determinants of UK unemployment and the Great Recession: analysing the gross flows data. *Applied Economics*, *45*(25), 3599–3616.

Appendix

Figure A1: distribution of estimated search intensity among job seekers in 1994.

The histogram displays the distribution of search intensities (s_j) in the 1994q1 sample obtained from the first-stage micro-level estimation.

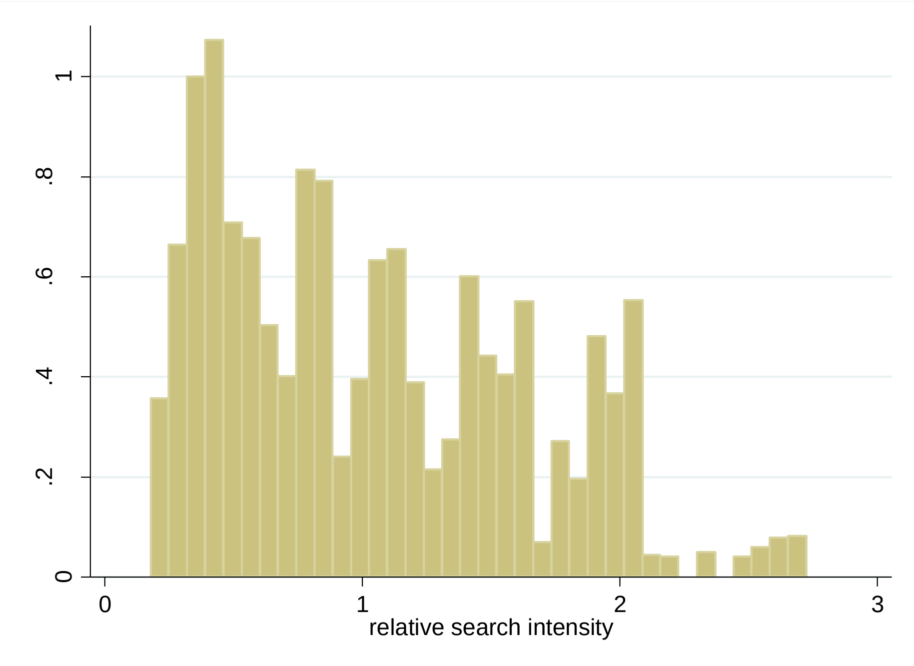

Figure A2: smooth forecast path of matching efficiency for three robustness checks.

a) Excluding government training schemes from unemployment-to-employment transitions

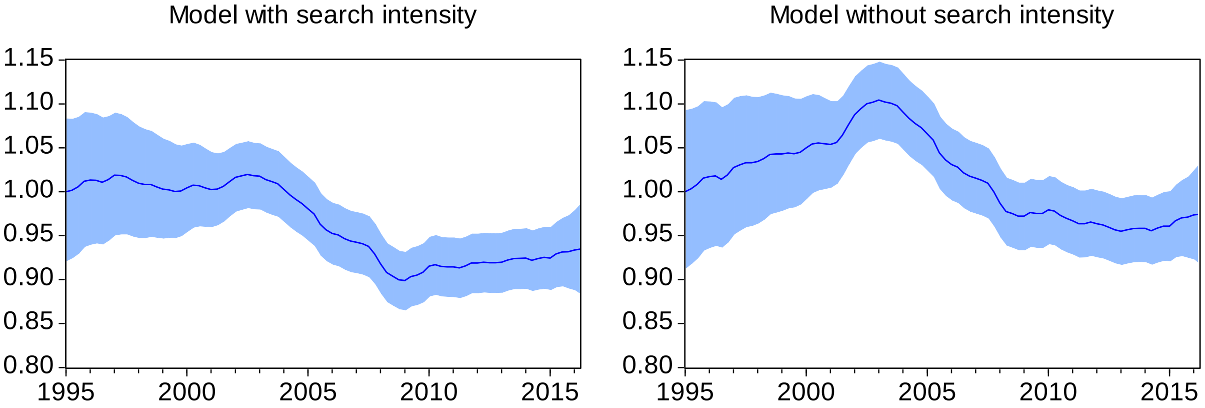

b) Adding dummy variable for spell of inactivity

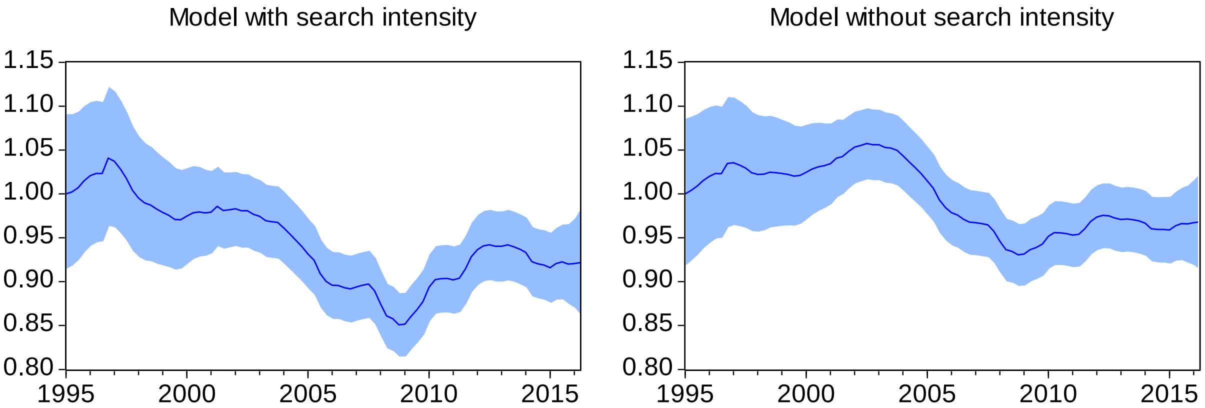

c) Using effective duration of joblessness instead of duration of unemployment

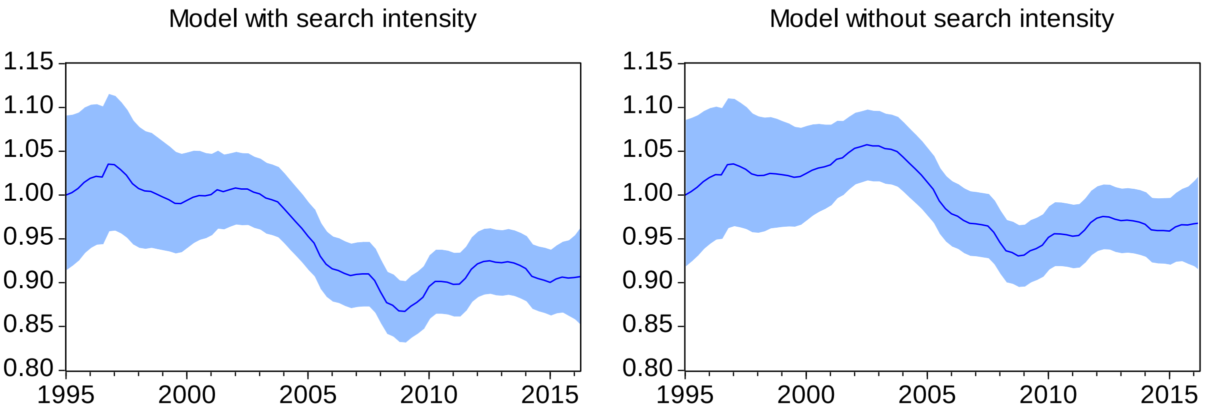

Figure A3: Plot of the coefficients from the first-stage MLE in different years.

Each plot reports the estimated value for the coefficient of the respective variable in the first-stage MLE carried out over different years. The 95 percent confidence intervals are reported in grey.

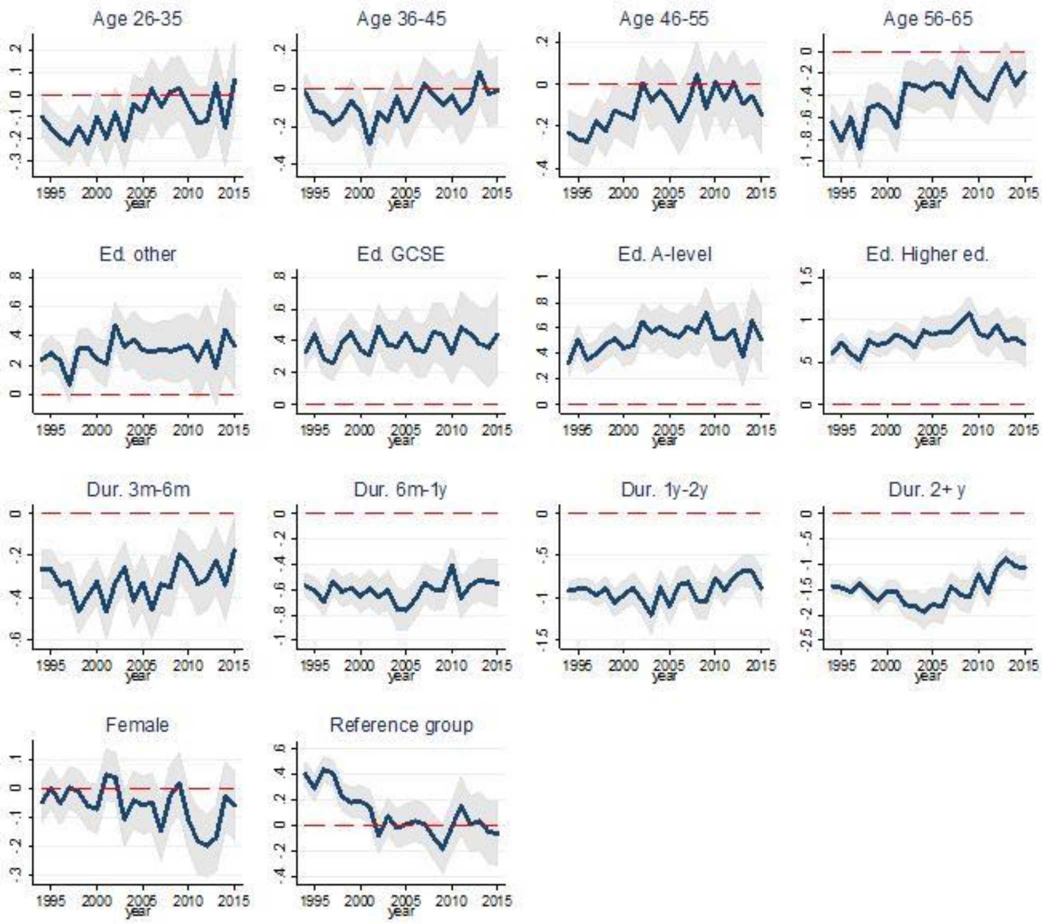

Figure A4: Aggregate search intensity with parameters estimated in different years.

Each line represents the path of aggregate search intensity *st* computed using the coefficients estimated on different baseline years. All series are normalized to be equal to one in 1994q1 and are reported in 4-quarter moving averages.

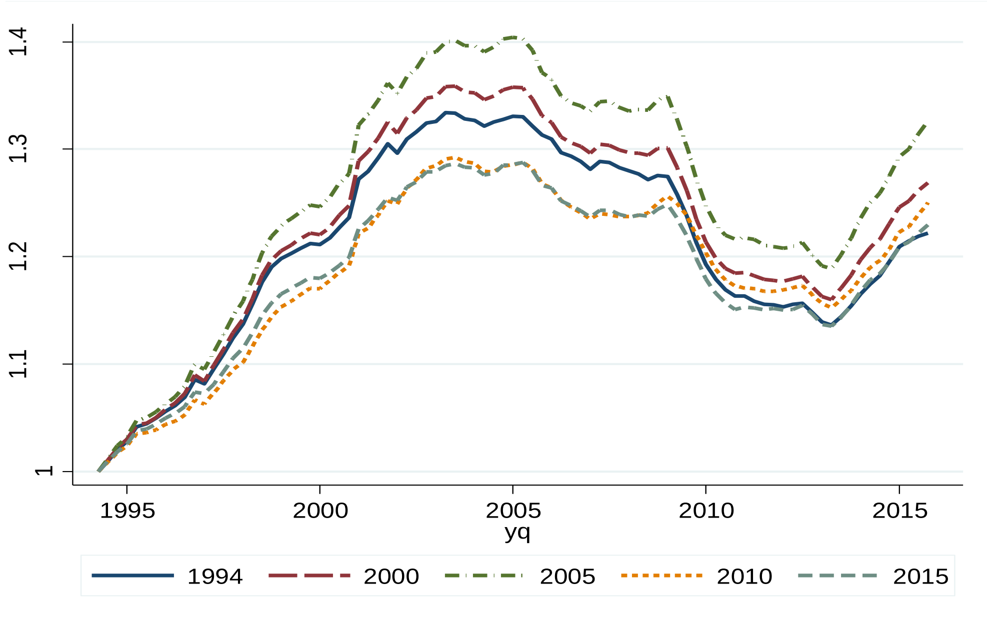

Table A1: Sensitivity analysis of first-stage MLE estimation.

The reference group is comprised of male individuals with no GCSE qualification aged 16 to 25 who are short-term unemployed. The first column excludes Government Training Schemes from the unemployment exit flows. The second column includes a dummy variable for having spent time out of the labour force prior to unemployment. The third column uses effective duration of the joblessness spell rather than effective duration of unemployment. The reported *p*-values are for a two-tailed test.

	(1) No GTS UE flow		(2) Previously Inactive Control		(3) Duration since last job	
	Coef.	P-Val.	Coef.	P-Val.	Coef.	P-Val.
Reference group	0.245	0.00	0.465	0.00	0.257	0.00
Female	0.005	0.90	0.033	0.363	-0.010	0.77
Age 26-35	-0.054	0.25	-0.037	0.405	-0.122	0.01
Age 36-45	0.014	0.78	0.038	0.44	-0.078	0.11
Age 46-55	-0.183	0.00	-0.190	0.001	-0.287	0.00
Age 56-65	-0.545	0.00	-0.611	0.00	-0.697	0.00
Other qual.	0.268	0.00	0.235	0.00	0.246	0.00
GCSE qual.	0.370	0.00	0.324	0.00	0.371	0.00
A-level qual.	0.415	0.00	0.321	0.00	0.380	0.00
Higher educ.	0.650	0.00	0.608	0.00	0.664	0.00
3-6 months	-0.306	0.00	-0.272	0.00	-0.015	0.77
6 month-1 year	-0.663	0.00	-0.589	0.00	-0.257	0.00
1-2 years	-1.025	0.00	-0.962	0.00	-0.681	0.00
2+ years	-1.724	0.00	-1.493	0.00	-1.192	0.00
Previously Inactive			-0.384	0.00		
Sample size	19,961		19,961		19,961	
Log-pseudolikelihood	-4,928		-4,858		-4,928	

Table A2: Second-stage state space estimation results for the robustness checks.

	No GST		Inactivity Spell		Effective Duration	
Matching elasticity	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
1- η	0.296	0.360	0.294	0.332	0.314	0.332
	(0.04)	(0.040)	(0.05)	(0.04)	(0.04)	(0.04)
Log likelihood	235.245	230.249	232.916	234.820	234.820	232.917
Observations	90	90	90	90	90	90
Search Intensity	Yes	No	Yes	No	Yes	No
Job seeker definition	U	U	U	U	U	U

Table A3: coefficients from the first-stage Maximum Likelihood Estimation in the expanded sample.

The expanded pool of job seekers includes marginally attached (MA) individuals and on-the-job-search (OJS) workers. The reference group is comprised of male individuals with no GCSE qualification aged 16 to 25 who are short-term unemployed. All coefficients on sex, education, and age are estimated separately for each type of job seeker. The reported p-values are for a two-tailed test.

	Unemployed		Marginally		On-the-job Search	
	Coef.	P-val.	Coef.	P-val.	Coef.	P-val.
Reference group	0.673	0.000				
MA or OJS dummy			-1.298	0.000	-0.409	0.000
Female	-0.047	0.188	0.081	0.475	0.127	0.057
Age 26-35	-0.097	0.027	-0.316	0.013	-0.407	0.000
Age 36-45	-0.016	0.740	-0.570	0.000	-0.520	0.000
Age 46-55	-0.227	0.000	-0.559	0.000	-0.525	0.000
Age 56-65	-0.640	0.000	-0.692	0.001	-0.521	0.055
Other qual.	0.244	0.000	-0.087	0.564	-0.456	0.000
GCSE qual.	0.333	0.000	0.186	0.162	-0.483	0.000
A-level qual.	0.331	0.000	0.205	0.161	-0.505	0.000
Higher educ.	0.599	0.000	0.344	0.037	-0.914	0.000
3-6 months	-0.259	0.000				
6 month-1 year	-0.568	0.000				
1-2 years	-0.912	0.000				
2+ years	-1.433	0.000				

Number of obs. = 33,065

Wald chi2(33) = 2182.9

Prob. > chi2 = 0.0000

Log pseudolikelihood = -7910.47