

How Local Are Labor Markets?: Evidence from a Spatial Job Search Model

Alan Manning

London School of Economics and CEP (LSE)

Barbara Petrongolo

Queen Mary University and CEP (LSE)

May 2011

Abstract

We use data at the Census ward level to investigate the extent to which labor markets are ‘local’. We present some non-structural estimates in which the probability of filling a vacancy is influenced by unemployment and vacancies in the surrounding areas, but we argue that these estimates cannot adequately estimate the true cost of distance. We then present a simple model of job-search across space that allows us to estimate a matching process with a very large number of segments. We find that the cost of distance is relatively high. That is, the utility of being offered a job decays at exponential rate around 0.25 with distance (in km) to the job. Also, workers are discouraged from applying to jobs where they expect a large number of applications. Finally, we find evidence of constant returns to scale in matching markets. The estimated model seems to replicate fairly accurately actual commuting patterns across Census wards, although it tends to underpredict the proportion of individuals who live and work in the same ward. Finally, we find that, despite the fact that labor markets are relatively ‘local’, local stimulus to labor demand is fairly ineffective in raising the local unemployment outflow, because labor markets overlap, and the associated ripple effects in applications largely dilute the effect of local stimulus across space.

1 Introduction

How local are labour markets? A number of important questions in labour economics turn on the answer. For example, the spatial mismatch hypothesis (Kain, 1968) suggested that the unemployment rate of blacks in the inner city was so high because many jobs had moved to the suburbs and these jobs were no longer in the local labour market of those living in the city. And policies that aim to improve labour market outcomes in disadvantaged areas (e.g. the empowerment zones studied by Busso, Gregory and Kline, 2010) need to know about the size of local labour markets to decide about the appropriate nature of the intervention. If labour markets are very local then an effective intervention will have to be targeted on the disadvantaged areas themselves even if those areas are not conducive to generating employment. But if labour markets are not as local then there is less need for the targeted intervention and a targeted intervention may simply attract workers from other more advantaged areas. In recent years there seems to have been a resurgence of interest in such ‘place-based’ policies (e.g. see Glaeser and Gottlieb, 2008; Moretti, 2010, for overviews). There is also a sizeable literature on the incidence of local shocks to labour demand (see Notowidigdo, 2010, for a recent contribution). Such research needs a clear definition of a ‘place’.

Most academic research on the topic and government statistical agencies often divide their jurisdictions into non-overlapping areas e.g. cities that are then simply assumed to be single labour markets. Examples would be the BEA’s 179 Economic Areas or the 722 Commuting zones in the US or the 320 Travel to Work Areas (TTWAs) in the UK. While these efforts are understandable and useful they do have their problems. For examples, UK TTWAs are constructed so that (as far as possible) at least 75% of the population resident in a TTWA actually work in the area, and at least 75% of those who work in the area also live in the area. Because people commute from large distances to central London to work, this means that the whole of the Greater London area is classified as a single labour market. But those who live in the northern suburbs of London do not really think of the far southern suburbs as part of ‘their’ local labour market. And the non-overlapping nature of local labour markets causes inevitable discontinuities around the boundaries. Someone living just inside the London TTWA will be classified as living in an enormous labour market while someone living just across the border in the Luton TTWA will be classified as being in a modestly-sized local labour market. In reality, these people are in essentially the same labour market. A proper analysis needs to recognize the continuous nature of geographic space. The problem is how to model geographical space in a continuous way while preserving tractability - one of the contributions of this paper is to show how one might approach this problem.

The bottom line is that we do not have an enormous amount of existing evidence about

how local are labour markets. The most common approach is to simply assume that a particular area is a single labour market. Within those areas, spatial inequalities would be interpreted as the outcome of residential sorting but not to have labour market consequences. Let us briefly consider how one might approach the question of seeking to estimate how local are labour markets.

Commuting data might be expected to contain useful information about the size of local labour markets as they tell us about how far workers seem prepared to travel to jobs. But the cross-section of commuting patterns represents the outcome of a lot of decisions (e.g. residential location) that muddy the waters. To give a specific example, consider the academic job market. From commuting patterns one would observe that most faculty live reasonably close to their place of work and perhaps conclude that the labour market for academics was relatively local. But, of course, it is better described, albeit with some hyperbole, as global. What information would allow us to detect that? The argument of our paper is that one could detect that by looking at the address of the job market candidate when they applied for a job and looking at the distances they are considering. In the academic market a job market candidate in a specific current residential location is prepared to take a job over a very large geographical area but will then change residential location to be close to whatever job they obtain. In this situation it makes sense to think of the individual being in a very large local labour market as they will consider a very wide range of jobs but they will end up with a low commute.

Our research design is intended to try to avoid these potential problems and get to the heart of the question of how local are labour markets. By a local labour market we mean the set of jobs that an unemployed worker, currently in a particular location, will apply for. It may be that, if the application is successful, the individual chooses to change residential location but that does not concern us. The data we use is high-frequency data (monthly) on unemployment and vacancies (stocks and flows) in small neighborhoods in the UK (8850 in total). The high-frequency nature of the data means that it is reasonable to think that the location of the unemployed represents the location when currently applying for jobs. It also means that one can reasonably think of whether there is a vacancy in a particular area in a particular month as essentially random. The large number of neighborhoods means it is appropriate to model space as more or less continuous.

The ideal data set would contain information on the location of jobs applied for by individual workers. We do not have such information but we present a model in which using data on the filling of vacancies can be used to infer the distance over which workers look for work. The intuition for our approach is the following. Consider a vacancy in area A. It is plausible to think that the ease of filling this vacancy depends on the number of unemployed workers for whom the vacancy is in their local labour market (and the number of other vacancies, something our framework accounts for but complicates the intuitive discussion

here). If the ease of filling a vacancy in A is influenced by the number of the unemployed in area B but not in area C then it is a reasonable conclusion that area A is in the local labour market for people who live in area B but not for those who live in area C. Of course, reality is not as simple as the ‘in’ or ‘out’ of the local labour market we have used to give the intuition – in what follows we model a cost of distance so the size of the local labour market is measured by the cost of distance. In the academic job market example, we would expect the ease of filling a position at LSE is affected by the supply of PhDs in the US if the market is indeed global, but also influenced by the vacancies in the US and, possibly, throughout the world.

The plan of the paper is as follows. In the next section we describe the data we use and present some estimates of matching functions allowing for geographic spillovers that are similar to existing models in the literature. However, we argue that such equations are limited in what they can tell us about the size of local labour markets. We then present a more structural model of job search that provides an explicit micro-foundation for how to model the linkages between a very large number of areas in a way that preserves tractability.

2 Data and preliminary evidence

We use data on unemployment and vacancies, disaggregated at the Census ward level (CAS 2003 classification). These data are available on NOMIS (nomisweb.co.uk) and run monthly since April 2004. There are 10,072 wards in Great Britain, of which 7,969 in England, 881 in Wales, and 1222 in Scotland, with an average population of 5,670. Although unemployment and vacancy data are available for Scottish wards, commuting data, which we will also use below, are not, and thus we restrict our sample to the 8,850 wards in England and Wales.

Our data cover registered unemployment (the “claimant count”) and job vacancies advertised at Jobcentres. The UK Jobcentre Plus system is a network of government funded employment agencies, where each town or city typically has at least one Jobcentre. A Jobcentre’s services are free of charge to all users, both to job seekers and to firms advertising vacancies. To be entitled to receive welfare payments, unemployed benefit claimants are required to register at a Jobcentre, and “sign-on” every two weeks.

Employers wishing to advertise job vacancies can submit a form with detailed job specifications to a centralized service called Employer Direct. The job vacancy is then assigned to the employer’s local Jobcentre, and will have a dedicated recruitment adviser, who can assist the employer with the recruitment process. Regardless of the Jobcentre in charge, the Census ward for each vacancy is defined using the full postcode of the job location. Each job vacancy is advertised in three ways: (a) on the centralized employment website www.direct.gov.uk; (b) through the Jobcentre Plus phone service for job applicants; and (c) on the Jobcentre Plus network, which can be accessed at Jobcentre offices around the

country. Jobseekers can sample job openings via one or more of these methods, using various search criteria (sector, occupation, working hours, etc.). In particular, they can select jobs located within a certain radius (1 mile, 2 miles, 5 miles etc.) from a given postcode.

While the monthly series run from April 2004 to date, we restrict our sample period to April 2004-April 2006, because from May 2006 Jobcentre Plus introduced changes to its vacancy handling procedures, and the vacancy series for the post-2006 period are not fully comparable to those for the earlier period.¹

The data we use on the unemployed and vacancies cover a very large number of job-seekers and vacancies and at a much more disaggregated level than available through any other source. In more aggregate form these data have been used in other studies of the UK matching process (see, among others, Coles and Smith, 1998, Burgess and Profit, 2001, and Coles and Petrongolo, 2008). But, it is important to realize they do not represent the universe of job-seekers or vacancies, and it is important to get some idea of how much of the matching process is being captured by these data.

On the worker side, not all job-seekers are claimant-count unemployed, as job-seekers may also be employed, or unemployed but not claiming benefits; and not all the claimant unemployed may be job-seekers (though they are meant to be according to the rules for benefit entitlement). To get some ideas of the numbers involved, we turn to the UK Labour Force Survey (LFS), which asks a direct question about job search both of those who are currently in and out of employment. In the Spring of 2008 (to give one example) the LFS suggests there were about 3.6 million job-seekers in the UK, and total employment was about 29.5 million. Almost exactly half of the job-seekers were not currently employed, and at that time the official figures for the claimant count was about 850,000. In the LFS, approximately 20% of the claimant unemployed do not report looking for work in the past 4 weeks, suggesting that the claimant unemployed represent about 20% of total job-seekers in the economy.

It may be argued that the claimants are among the most intensive job-seekers (see, among others, Flinn and Heckman, 1983, Jones and Riddell, 1999), and thus we weight job-seeker figures in the LFS by the number of reported search methods used. During the 2002-2008 period, the unweighted share of claimants in total job-seekers was 17.6%, while the weighted share was 23.7%. As one would expect, the share of claimants in job-seekers also varies markedly with levels of education, being 15% among college graduates, 21.8% among those with ‘A levels’ (high school graduates), 24.9% among GCSEs (who left school at 16), and

¹Prior to May 2006, vacancies notified to Jobcentre Plus were followed up with the employer until they were filled, and the number of vacancies filled at Jobcentres was used as one of the main indicators of their performance. From May 2006, the Jobcentre Plus performance evaluation is no longer based on vacancies being filled, thus vacancies notified to Jobcentre Plus are not followed-up, and have an ex-ante closure date agreed with the employer, upon which they are automatically withdrawn. This systematically understates the stock of unfilled vacancies from may 2006 onwards.

35.2% among those with no qualifications. This means our study is best interpreted as being about lower-skill labour markets that probably tend to be more ‘local’. So it has to be acknowledged that the claimant count represents only a fraction of job-seekers in the economy, and this will cause a bias if such fraction varies systematically across areas, something on which unfortunately we have no information.

For our purpose it is also important to know the fraction of job-seekers who is looking at vacancies recorded in our data, i.e. vacancies advertised at Jobcentres. Using reported information on the job-search methods used, during 2002-2007, 92% of claimants use Jobcentres, and 45.2% of them quote Jobcentres as their most important job-search method. These proportions fall to 44.4% and 18.3% for the non-claimant unemployed, and to 19.1% and 5.9% respectively for the employed. So, Jobcentres are widely-used by the job-seekers in our sample. In this regard, it is important to realize that the UK Public Employment Service is much more widely used than the US equivalent. Manning (2003, Table 10.5) shows that only 22% of the US unemployed report using the PES compared to 75% of the UK unemployed, and OECD (2000, Table 4.2) shows that the market share of the PES in the US in vacancy coverage and total hires is substantially lower than in the UK. So, unlike the US, UK job centres do play an important role in matching job-seekers and vacancies.

On the job vacancy side there is fairly limited external evidence that we can use to assess the representativeness of our Jobcentre data. Since 2001 the Vacancy Survey of the Office for National Statistics provides comprehensive estimates of the number of job vacancies in the UK, obtained from a sample of about 6,000 employers every month. Employers were asked how many job vacancies there were in their business, for which they were actively seeking recruits from outside the business. These vacancy data cover all sectors of the economy except agriculture, forestry and fishing, but are not disaggregated at the occupation or area level, so we can only make aggregate comparisons between ONS and Jobcentre vacancy series.

On average, since April 2004, the Jobcentre vacancy series in the UK is about two thirds the ONS series, but there seem to be reasons to believe that such proportion may be overstated (Machin, 2001). In particular, in May 2002, an extra question has been added to the ONS Vacancy Survey, on the number of vacancies notified at Jobcentres, and based on this information the ratio of total vacancies advertised at Jobcentres was 44%. While one should allow for sampling variation (this information is only available for May 2002, and for only 420 respondents), this 44% proportion seems to be markedly lower than the two thirds recorded for the post 2004 period. According to Machin (2001), the major reason for this discrepancy is that Jobcentre vacancies obtained from the computerized system may include vacancies which are “awaiting follow-up”, but which have already been filled by employers, or which have been “suspended” by the Jobcentres as it appears that sufficient potential recruits have already been referred. Our vacancy series obtained from Jobcentres (“live

unfilled vacancies”) excludes suspended vacancies, but “may still include some vacancies which have already been filled or are otherwise no longer open to recruits, due to natural lags in procedures for following up vacancies with employers”², thus one can still imagine that two-thirds is indeed an upper bound for the fraction of job openings that are effectively available to job-seekers at Jobcentres. As no occupation information is available for the ONS vacancy series, it is not possible to determine how the skill distribution of our vacancy data compares to that of the whole economy, but of course it is realistic to imagine that Jobcentre vacancies tend to over-represent less-skilled jobs.

From this discussion it should be clear that we capture an important section of job search in the UK, especially for low-skilled workers, but it is also clear that we cannot provide a fully comprehensive picture of the job search process. This introduces a bias if the portion of the job search process covered by our data varies systematically across areas, something on which unfortunately we have no information. As a check against the possibility of gross biases we also investigate how well our model explains the commuting flows across wards using census data that covers everyone in employment, no matter how they searched for jobs.

In the data presented below and in all estimated specifications, we obtain the vacancy and unemployment outflows as differences between the corresponding inflows and the monthly variations in the stocks. For the unemployed, the outflow series predicted by the stock-flow accounting identity was virtually identical to that reported, while for vacancies the correlation was 0.81. Such discrepancy may arise because the reported outflow does not include cases of “suspended” vacancies, or cases of vacancies “awaiting follow-up”, but these may well be cases of positions being filled without keeping the Jobcentre informed. Due to measurement error, for about 0.5% of observations the vacancy outflow implied by the stock-flow accounting identity is negative, and thus we drop the corresponding observations.

Table 1 presents some simple descriptive statistics on unemployment and vacancies stocks and flows from May 2004-April 2006.³ English wards have on average 104 unemployed and 47 vacancies. Taken across the whole period, both unemployment and vacancy inflows and outflows seem very similar but with vacancy outflow slightly lower than the inflow and the unemployment inflow slightly above the outflow. There is also very wide spatial variation in unemployment and vacancies. [insert maps] Because of the high-frequency nature of our data there is also very considerable variation in unemployment and vacancies within wards from one month to the next. One can think of it as essentially random whether a firm has a vacancy this month or next month. We intend to experiment with instruments for local demand shocks in the future.

We start our investigation of the data by estimating a conventional log-linear matching

²<https://www.nomisweb.co.uk/articles/showArticle.asp?title=warning: limitations of data&article=ref/vacs/warning-unfilled.htm>

³We lose the initial month in the sample period because we need lags of the vacancy and unemployment stock measures, i.e. values measured at the end of the previous month.

function in unemployment and vacancies, possibly augmented by local spillovers:

$$\begin{aligned} \log\left(\frac{M_b}{V_b}\right) &= \alpha_0 + \alpha_1 \log(U_b + \beta_1 U_{5b} + \beta_2 U_{10b} + \beta_3 U_{20b} + \beta_4 U_{35b}) \\ &\quad + \alpha_2 \log(V_b + \gamma_1 V_{5b} + \gamma_2 V_{10b} + \gamma_3 V_{20b} + \gamma_4 V_{35b}) + \varepsilon_b \end{aligned} \quad (1)$$

where U_b is the number of unemployed in ward b , U_{5b} is the number of unemployed in wards within 5km of b (excluding b itself), U_{10b} is the number of unemployed in wards between 5km and 10km of ward b etc., and similarly for vacancies. M_b is the vacancy outflow from ward b , so the dependent variable is the outflow rate. The basic idea behind this specification is that the probability of filling a vacancy in b depends on local unemployment and on unemployment in the surrounding areas, but that more distant unemployed workers are less effective in filling a vacancy in b i.e. we would expect $\beta_i < 1$. Similarly, more vacancies in area b and neighboring wards might be expected to reduce the outflow rate in b , but more distant vacancies have a smaller effect, i.e. we expect $\gamma_i < 1$.

Next define the total number of unemployed and vacancies within 10km of b to be:

$$\tilde{U}_{10b} = U_b + U_{5b} + U_{10b}; \quad \tilde{V}_{10b} = V_b + V_{5b} + V_{10b};$$

and approximate (1) by

$$\begin{aligned} \log\left(\frac{M_b}{V_b}\right) &\approx \alpha_0 + \alpha_1 \log \tilde{U}_{10b} + \alpha_1 \left(\frac{1 - \beta_2}{\beta_2} \frac{U_b}{\tilde{U}_{10b}} + \frac{\beta_1 - \beta_2}{\beta_2} \frac{U_{5b}}{\tilde{U}_{10b}} + \frac{\beta_3 - \beta_2}{\beta_2} \frac{U_{20b}}{\tilde{U}_{10b}} + \frac{\beta_4 - \beta_2}{\beta_2} \frac{U_{35b}}{\tilde{U}_{10b}} \right) \\ &\quad + \alpha_2 \log \tilde{V}_{10b} + \alpha_2 \left(\frac{1 - \gamma_2}{\gamma_2} \frac{V_b}{\tilde{V}_{10b}} + \frac{\gamma_1 - \gamma_2}{\gamma_2} \frac{V_{5b}}{\tilde{V}_{10b}} + \frac{\gamma_3 - \gamma_2}{\gamma_2} \frac{V_{20b}}{\tilde{V}_{10b}} + \frac{\gamma_4 - \gamma_2}{\gamma_2} \frac{V_{35b}}{\tilde{V}_{10b}} \right). \end{aligned} \quad (2)$$

This specification has the advantage that one can simply read off the returns to scale by a comparison of the coefficients on $\log \tilde{U}_{10b}$ and $\log \tilde{V}_{10b}$, while the coefficients on the share variables tell us about the relative effectiveness of unemployment and vacancies at different distances. The decision to ‘normalize’ on unemployment and vacancies within 10km is essentially arbitrary but it is important to choose a normalization for which β_2 and γ_2 are not zero and for which the ‘share’ variables are not too large. In experimentation, 10km seemed about right to us. On average, about 5% of unemployment and vacancies within 10km are in the local ward, one-third are within 5km. Moving beyond the 10km ring, there are about 4.5 times the number of unemployed and vacancies between 10 and 20 km as within 10km and 16 times as many within 20km.

Estimates of (2) are reported in Table 2. In the first column we simply pool all months and wards without time or ward effects. The estimates are in line with what we would expect. More unemployed raise the probability of filling a vacancy while more vacancies reduce it. The coefficients on the unemployment and vacancy variables suggest something very close to constant returns – the implied returns to scale parameter being 0.988. This is

significantly different from one but that is largely a result of the large number of observations. It is not just the level of unemployment and vacancies within 10km that affect the outflow rate but also their geographical mix. As one might expect, the more the unemployed are close to the ward, the higher the probability of filling it. From the coefficients on the share of unemployment in the local ward and within 5km one can derive an estimate of β_2 of 0.21 and β_1 of 0.46, i.e. unemployed workers outside the ward but with 5km have 46% of the effectiveness of generating matches as those within the ward and the unemployed in the 5-10km ring have an effectiveness of 21%. Unemployed in the 20k and 35k rings have tiny effects on the vacancy outflow, though are significantly different from zero. For vacancies, the more local the vacancies the lower the outflow rate as one would expect as such vacancies are closer substitutes. Vacancies outside the ward but with 5km have 27% of the effectiveness of those within the ward, and vacancies in the 5-10km ring have an effectiveness of 19%. Vacancies in the 20k and 35k rings have very small effects on the vacancy outflow rate.

The second column introduces time dummies: the main consequence of this is a very slight attenuation of all the coefficients but the qualitative conclusions remain similar. The third column instruments the vacancy variables using the one-month lags of the vacancy variable. Our reason for showing this specification is that the dependent variable is obtained by dividing the recorded outflows by the local stock of vacancies and this local stock also appears in the construction of some of the right-hand side variables. This means that a division bias problem might occur if there are measurement problems with the current vacancy stock. The estimates in the second and third column are very similar, with the possible exception of vacancies within 10km that are in the local ward. It is exactly this variable where the local vacancy stock has the most influence so this is perhaps some indication that there are modest issues of division bias in the estimates.

The fourth column introduces ward fixed effects. Comparing the estimates in the second and fourth columns one notes that the coefficient on the share of vacancies in the local ward becomes much more negative. This is perhaps what one would expect if again there are division bias issues, as it is now only the within-ward variation in vacancies that is being exploited, and that probably has more transitory components. One also notes that the coefficient on the total unemployment in the 10km ring rises but the coefficients on the unemployment mix variables become insignificantly different from zero. The fifth and sixth columns present IV estimates with ward fixed effects, the fifth column instrumenting both the unemployment and vacancy variables, while the sixth column instruments only the vacancy variables.

The log-linear matching functions estimated in Table 2 are standard in the literature but have the disadvantage that the dependent variable is not defined when the outflow rate is zero. Although this is not an issue in existing empirical studies of the matching function because of the level of aggregation in those studies, it becomes a potential issue when using

data on very small areas, and indeed in our sample 6.2% of observations have zero outflows. There are a number of approaches one might take to dealing with this. Here, we take the approach that is most similar to the log-linear function (2), i.e. to estimate in levels instead of logs. In a later section we will present a model in which the functional form in levels can be thought of as a legitimate specification of the expected outflow rate given unemployment and vacancies. The functional form used in this section has the disadvantage that the ‘predicted’ value need not lie in the unit interval, but has the advantage that one can compare estimates with the log-linear matching functions.

The first column of Table 3 presents estimates of a log-linear matching function but excluding unemployment and vacancies more than 10km distant, as Table 2 has suggested that the impact of these distant unemployment and vacancies was negligible. The second column then estimates the level version of this equation by non-linear least squares, excluding observations with zero vacancy outflow, thus on the same sample as in the first column. The estimates are qualitatively similar but one does notice a considerable reduction in the size of the coefficients on all ratio variables. Finally, the third column presents the levels model but includes the ‘zeroes’, i.e. the estimation method is the same as the second column, but with a larger sample size. The estimates obtained are very close to those reported in the second column. Finally the fourth and fifth columns report results for the log-linear and linear models estimated for one month only (February 2005), that figures in the structural estimates below.

The results of Table 2 and 3 are consistent with a simple matching model with spatial spillovers, and the results are qualitatively similar whether one includes or excludes time or ward fixed effects. However, these estimated equations do have their limitations for making inferences about the size of local labour markets. First, they do not allow us to estimate where those who are filling the vacancies actually live, whether they are predominantly local or more distant. Data that provided information on where the successful job applicant lived could answer that question. But the estimated equations are also not very informative about the reasons for the spill-overs – at best, they represent a description of the data. When it is shown that an increase in the number of unemployed 10km away raises the probability of filling a vacancy in area b , is this because those unemployed workers apply for vacancies in b or because they apply for vacancies local to them which then become harder to obtain causing workers 5km away from b to shift their job search efforts towards vacancies in area b ? To answer this question we need a more structural model of job search and that is what the next section provides.

3 The Structural Model

The key ingredient of our methodology consists in relating job matches in a given area to the number of applications received by job vacancies in that area. The novel element with respect to most of the matching literature is to model applications per job in each area based on optimizing job search behavior across space, rather than use local unemployment as a proxy for applications. Our approach is to use the expected number of applicants for a job in an area as a measure both of how easy it will be for an employer to in that area to fill a vacancy and a measure of how much competition for a job in that area there is for a worker who is considering applying for a vacancy there. We next outline a model of the process by which unemployed workers determine the number of job applications they make and their distribution over space, and we will then relate applications to job matches.

3.1 The application process

At any moment in time there are U_a unemployed workers and V_a vacancies in each area a of the economy. Denote by (U, V) the vector of unemployed workers and vacancies across areas.

Suppose that individuals are deciding how many of the existing vacancies to apply for. Because of the time lag in the process of filling jobs, they cannot apply sequentially to vacancies, thus applications must be simultaneous. Assume that an application to vacancy i has a probability p_i of being successful, and generates utility u_i in that case. Assume further that the probability of more than one application being successful is infinitesimal so that expected return for a worker can be written as:

$$\sum_i D_i p_i u_i,$$

where D_i is a binary variable taking the value 1 if an individual applies to the job and zero if they do not. Individuals have a cost function for N applications of the form:

$$C(N) = \frac{c}{1 + \eta} N^{1+\eta},$$

so that the net expected utility can be written as:

$$\sum_i D_i p_i u_i - \frac{c}{1 + \eta} \left(\sum_i D_i \right)^{1+\eta}.$$

The optimal application rule is to apply for a vacancy if the expected utility from doing so is higher or equal to the marginal cost $C'(N)$. This happens if

$$p_i u_i \geq c \left(\sum_i D_i \right)^\eta = c N^\eta. \tag{3}$$

This result says that the attractiveness of vacancies is determined by the expected utility they offer, and job-seekers apply first for jobs offering the highest expected utility and continue to do so until expected utility is below the marginal cost of an extra job application. Another implication of this is that whether the individual applies for a particular vacancy or not depends only on the expected utility it offers and the marginal cost of an application. Other vacancies only affect this decision through the effect on the marginal benefit of an application. While extremely convenient, it is important to note that the assumption that the probability of more than two applications being successful is infinitesimal plays a critically important role here – if this assumption is not met then one cannot rank vacancies by their expected utility and the decision-problem is far more difficult.

In what follows we will assume that the probability of filling a vacancy and of success in a particular application depends on the expected number of applications to that job, that we denote by A . Denote the probability of being the successful applicant by $p(A)$, with $p'(A) < 0$. The extent to which the probability of success is related to the number of applicants is an important parameter in the model - we will refer to it as the congestion parameter.

The parameter η (or, more accurately, a transformation of it) will turn out to be very important in determining the returns to scale in the model so it is worth taking some time to understand its role.

If $\eta = 0$ there is a constant marginal cost of an application and an unemployed worker will apply for a vacancy if the expected utility is above this marginal cost. In this situation a doubling in the number of vacancies will lead to a doubling in the number of applications each unemployed worker makes. The average number of applicants per vacancy will remain unchanged so it is plausible to think that the probability of filling each vacancy will remain unchanged. The total number of matches will then also double. In this situation there are constant returns to scale to vacancies alone. If one doubles both vacancies and the number of unemployed workers then the number of applications will rise four-fold as the applications per worker will double and the number of workers double. This implies increasing returns to scale.

At the other extreme, consider $\eta = \infty$. This should really be thought of as the case where each unemployed worker has a fixed number of applications to make and will apply to those vacancies that offer the highest expected utility. In this case a doubling of vacancies and unemployment will lead to a doubling of applications as applications per worker are unchanged and the number of the unemployed has risen. Hence applications per vacancy are unaltered, the probability of filling a vacancy is unaltered and the total number of matches will double. This corresponds to the case of constant returns to scale.

Our set-up makes it harder to rationalize the possibility of decreasing returns to scale for which we would have to introduce some extra form of congestion in the model. But,

the estimates presented so far are very close to constant returns to scale for the economy as a whole. However, this is perfectly consistent with decreasing returns to vacancies and unemployment in individual areas - typically doubling vacancies and unemployment in a particular area will result in a lower probability of filling jobs in that area.

To put more structure on the problem, assume that the utility from a job in area b for someone from area a is given by:

$$u_{ab} = f_{ab}\varepsilon,$$

where f_{ab} represents the intrinsic attractiveness of a job in area b for someone in area a and ε is an idiosyncratic component which we assume has a Pareto distribution with exponent k . A natural specification for f_{ab} is a function declining with the distance between a and b so that jobs in more distant areas are less attractive. Otherwise one could use information on commuting time or commuting costs between a and b .

Hence, using (3), an individual in a applies to a vacancy in b if

$$p(A_b)f_{ab}\varepsilon \geq cN_a^\eta,$$

where N_a denotes the total number of applications made by each worker in a . Given the assumption that ε has a pareto distribution, this happens with probability:

$$\Pr(p(A_b)f_{ab}\varepsilon \geq cN_a^\eta) = \left(\frac{p(A_b)f_{ab}}{cN_a^\eta}\right)^k. \quad (4)$$

Although there is some uncertainty about whether an individual applies to a particular vacancy (because of the idiosyncratic component to utility), let us assume we can apply a law of large numbers so that the total number of applications can be treated as non-stochastic. We will thus have that the number of applications sent from unemployed workers in a to job vacancies in b are given by:

$$N_{ab} = V_b \left(\frac{p(A_b)f_{ab}}{cN_a^\eta}\right)^k.$$

Adding up the N_{ab} 's across all possible destination areas b gives the total number of applications sent by unemployed workers in a :

$$N_a = \sum_b V_b \left(\frac{p(A_b)f_{ab}}{cN_a^\eta}\right)^k,$$

which can be solved for the number of applications N_a :

$$N_a = \left[c^{-k} \sum_b V_b (p(A_b)f_{ab})^k \right]^\gamma, \quad (5)$$

where $\gamma = (1 + \eta k)$. $\eta = 0$, the case of a constant marginal cost of an application, corresponds to $\gamma = 1$, while $\eta = \infty$, the case of a fixed number of applications, corresponds to $\gamma = 0$.

Using (4) and (5), one can compute the total number of applications made by the unemployed in a to vacancies in b , N_{ab} , as

$$N_{ab} = c^{-k\gamma} V_b (p(A_b) f_{ab})^k \left[\sum_{b'} V_{b'} (p(A_{b'}) f_{ab'})^k \right]^{\gamma-1}. \quad (6)$$

The number of applications received by vacancies in b is equal to all applications that unemployed workers decide to send to area b from all areas a . Thus the ratio of applications per vacancy in b , denoted by A_b , is given by

$$\begin{aligned} A_b &= \frac{\sum_a N_{ab} U_a}{V_b} \\ &= c^{-k\gamma} \sum_a U_a (p(A_b) f_{ab})^k \left[\sum_{b'} V_{b'} (p(A_{b'}) f_{ab'})^k \right]^{\gamma-1}. \end{aligned} \quad (7)$$

To estimate expression (7), we make two further functional form assumptions, and namely $p(A_b) = A_b^{-\tilde{\beta}}$, where $\tilde{\beta} > 0$ denotes the effect of job competition on applications to jobs in a given area,⁴ and $f_{ab} = \exp(-\tilde{\delta} d_{ab})$, where d_{ab} may be proxied by distance between a and b , and $\tilde{\delta}$ measures the exponential rate of decay of the attractiveness of a given job with distance to that job.⁵ Under these assumptions we can solve for A_b :

$$A_b = \left\{ c^{-k\gamma} \sum_a U_a \exp(-\delta d_{ab}) \left[\sum_{b'} V_{b'} A_{b'}^{-\beta} \exp(-\delta d_{ab'}) \right]^{\gamma-1} \right\}^{1/(1+\beta)}, \quad (8)$$

where $\beta = k\tilde{\beta}$ and $\delta \equiv k\tilde{\delta}$.

Equation (8) is the key relationship delivered by our spatial search model, and captures all the inter-dependencies between areas. In particular, the number of applicants to jobs

⁴The functional form $p(A_b) = A_b^{-\tilde{\beta}}$ can be derived from a more structural urn-ball model describing how vacancies are filled. Let's denote by π the probability that any particular candidate is acceptable for a given job vacancy. The probability that a firm does not fill the vacancy is $(1 - \pi)^{A_b}$; the probability that the firm does fill the vacancy is $1 - (1 - \pi)^{A_b}$; and the probability that any particular applicant is selected is $[1 - (1 - \pi)^{A_b}] / A_b$. We approximate this expression as $A_b^{-\tilde{\beta}}$. Both expressions are decreasing and convex in A_b , so we would be fitting similar functional forms to our data, but using the $A_b^{-\tilde{\beta}}$ approximation instead of the exact formula makes the model a great deal more tractable.

⁵While we are only taking into account horizontal heterogeneity between workers and jobs, represented by distance, this model could be generalized to allow for some form of vertical heterogeneity between jobs at different locations. For example, workers at all locations positively value the wage attached to a job offer, and thus, other things equal, receive higher utility from applying to jobs in high-wage areas than to jobs in low-wage areas. In this case one could have $f_{ab} = \exp(-\delta d_{ab}) + \alpha w_b$, where w_b denotes destination-specific characteristics, and α is the associated effect on utility. While local wages are the most natural proxy for w_b , at the moment we do not have access to earnings information at the ward level, and thus simply model the attractiveness of a location as a function of the distance to that location.

in b is likely to be influenced (even if only very slightly) by unemployment and vacancies in all other areas, because they are ultimately linked through a series of overlapping labor markets. This expression might be thought impossibly difficult to solve as, if we have 8,850 wards, it has 8,850 equations in 8,850 unknowns. But, under reasonable conditions, and namely $|\beta| < 1$, (8) is a contraction mapping, in which case it can be solved iteratively and economically to obtain A_b .

A useful result that can be obtained from (8) is that A_b is homogeneous of degree $\gamma/(1 + \beta\gamma)$ in U and V , and this relates to the returns to scale in the matching process. In particular, when $\gamma = 0$, the matching process displays constant returns. $\gamma/(1 + \beta\gamma) > 0$ would imply increasing returns to scale. This can be seen more clearly in the special case in which areas are isolated, so that $f_{ab} = f > 0$ for $a = b$, and $f_{ab} = 0$ for $a \neq b$. In this case it can be shown that the number of applications per vacancy in an area can be written as a function of the U-V ratio in the area and the overall level of vacancies (though one could also re-write it as a function of the total level of unemployment):

$$\ln A_b = \frac{1}{1 + \beta\gamma} \ln \left(\frac{U_b}{V_b} \right) + \frac{\gamma}{1 + \beta\gamma} [\ln(V_b) + \ln(f)]. \quad (9)$$

As the vacancy outflow rate in an area depends on the number of applications per job in that area, expression (9) implies a relationship between the vacancy outflow rate, the local U-V ratio, and the level of vacancies, which is very similar to the log-linear matching function usually estimated in the literature (see Petrongolo and Pissarides, 2001). When $\gamma = 0$ the number of applications per job only responds to the ratio of unemployment to vacancies, and is not affected by the size of the labor market, represented by V_b , implying constant returns.

To summarize, one can think of our model as having three key parameters:

- δ , the cost of distance;
- γ , the returns to scale of the matching function;
- β , the congestion parameter, that measures how much workers are deterred from applying to jobs where they expect a large number of applications.

One might also notice that (8) also contains a ‘constant’, c . But, one can normalize this to one without loss of generality as the number of applications per vacancy is not something that is actually observed, just a theoretical construct. This also means that it makes no sense to actually discuss the computed number of applications per vacancy as a guide to whether the model is ‘plausible’ or not. What is observed is the actual number of matches that we posit to be related to the number of applications per vacancy - we next turn to that relationship.

3.2 From applications to job matches

We use the vacancy outflow in an area as a proxy of job matches, and we express the vacancy outflow rate as a function of the number of applications per vacancy, i.e.

$$E\left(\frac{M_b}{V_b}\right) = \Psi(A_b), \quad \Psi'(\cdot) > 0. \quad (10)$$

Various functional forms have been used in the literature for estimating Ψ , based on possible microfoundations of the matching function and empirical tractability. The simplest way to justify a matching function like (10) is to think of an urn-ball problem,⁶ in which firms play the role of urns and applications the role of balls. Because of a coordination failure, a random placing of the balls in the urns implies that some urns will end up with more than one ball and some with none. Thus an uncoordinated application process will lead to overcrowding in some jobs and no applications in others.

Conditional on receiving an application, a vacancy may still remain unfilled if one allows for worker heterogeneity and thus the possibility that the applicant may not be suitable for the job. We denote by $p(A_b) = A_b^{-\beta}$ the probability that a given job applicant is selected for the job. Thus the probability that a given vacancy is not filled by any applicant is $(1 - A_b^{-\beta})^{A_b}$, and the vacancy outflow rate is $M_b/V_b = 1 - (1 - A_b^{-\beta})^{A_b}$. For small enough $A_b^{-\beta}$, $(1 - A_b^{-\beta})^{A_b} \simeq \exp(-A_b^{1-\beta})$, and thus we estimate

$$\frac{M_b}{V_b} = 1 - \exp\left[-\exp(\alpha)A_b^{1-\beta}\right] + e_b \quad (11)$$

where we have added a non-negative multiplicative constant $\exp(\alpha)$ and an error term e_b . The term $\exp(\alpha)A_b^{1-\beta}$ represents the continuous-time hazard at which vacancies are filled.

Alternatively, a simple log-linear specification can be estimated, i.e.

$$\frac{M_b}{V_b} = \exp(\alpha)A_b^{1-\beta} + e_b. \quad (12)$$

The nice feature of the urn-ball specification is that it ensures a the vacancy outflow rate between 0 and 1, while this is of course not imposed by the log linear specification. However, in the log linear specification has the advantage to yield a constant elasticity of the vacancy outflow with respect to the number of job-seekers and vacancies, and this property allows us to more easily assess the returns to scale in matching. As we will note below, the results are virtually identical with the two specifications. Whether estimating (11) or (12), A_b is implicitly defined by (8), and thus δ , β , γ are further parameters to be estimated. In practice we estimate (11) and (12) by maximum likelihood, and at every iteration of the maximization solve the contraction mapping in (8).⁷

⁶See Butters (1977) and Pissarides (1979) for early microfoundations of the matching function based on an urn-ball model.

⁷Again, to avoid dropping observations with zero outflows, both (11) and (12) are estimated in levels instead of logs.

In both of these specifications one can see that the normalization of the number of applications discussed above is, indeed, without loss of generality. If one changed the normalization one would simply change the parameters relating the number of applications to the vacancy outflow and the overall fit of the model would be the same.

One might wonder about the relationship between our model of the job search process that is based on vacancies receiving a number of applications and then, possibly, choosing one of the applicants, and the more common modelling strategy in which there is an arrival rate of job applicants and the first acceptable one is chosen. However, one could reinterpret the number of applications in our modelling strategy as a decision about the rate at which to apply for jobs and there is then the distribution of these applications over vacancies in different areas. That would also lead to a specification that related the outflow rate to the number of ‘applicants’ but the number of applicants should be re-interpreted as the rate at which job applicants apply to the firm.

Our overall approach has some similarities to the way in which economists in Industrial Organization have modelled markets. One think of a ‘product’ as being a job in a particular area. Compared to most applications we have a very large number of ‘products’ but we also have some a priori information on which of these products are the closest substitutes based on geographical distance. Consumers are also differentiated - in our application, this is by space, the same differentiation of the products though there is nothing inevitable about this. One can think of our information on unemployment and vacancies as being information on the level of demand by different types of consumers and the level of supply of different products. Our variable ‘applications per vacancy’ functions rather like a price in the sense that greater applications discourage consumers from purchasing a product of a particular type and encourage them to take their demand to other products. Our outcome variable - the number of matches - can be thought of as representing the market outcome in a quantity space. The equation we estimate is essentially a reduced-form equation for the quantity traded as a function of the demand and supply fundamentals. One hopes to retrieve the estimates of the demand functions because of the assumption that the supply of vacancies is exogenously fixed. Given this discussion, one might wonder why we do not include an explicit price - the wage - in the model. By raising the offered wage, employers would be expected to be able to attract more workers. That would be an interesting extension of the model but we currently have no data on wages at the ward level, so we abstract from them.

4 Results

4.1 Main estimates

Our first set of results is based on an urn-ball specification of the matching function, as shown in equation (11). For reasons of computing capacity, we cannot estimate our regression equations on the whole sample period. We thus estimate our main specification separately for each month from May 2004–April 2006, and we then run a number of specification and robustness checks on a representative month.

In Figure 1 we plot point estimates of the parameters of interest over the 24 months in our sample, and in Table 4 we report their time averages, together with their standard deviations, minimum and maximum values. The congestion parameter β is positive, implying that the probability of being selected for a given job opening falls with the number of applicants, with an average elasticity of about 0.75. As a corollary, job-seekers respond to strong job competition in a ward by reducing applications to that ward. The elasticity of the vacancy filling hazard with respect to applications is given by $1 - \beta$ (see equation (11)), thus vacancy duration falls by 25% when the number of applications doubles. The average estimate for γ is negative, implying decreasing returns in matching, although the low point estimate suggests a scenario very close to constant returns. The utility of a job in b for a worker located in a is modeled based on the geographic distance between a and b , and thus δ represents the exponential rate of decay of a job’s attractiveness with distance in km to that job. An average δ of 0.3 is consistent with relatively fast decay of job utility with distance. For example, the attractiveness of a job falls by about 4.5 times whenever one moves the job 5 km further afield from a job-seeker’s location. Overall, both β and δ appear to be precisely estimated over the sample period, but there is slightly more variation in γ . If one is willing to make the hypothesis that the month-to-month variation in the relevant variables is largely driven by independent, random shocks, then the average parameter estimates and associated standard deviations can be used for bootstrap inference. Thus one can conclude that while both β and δ are highly statistically significant, γ is not statistically different from zero.

In Table 5 we report estimates of alternative specifications of the job application model for February 2005. The simple criteria used for picking a reference month is that it should not be December or a summer month, and that the parameter estimates for this month should be quite close to the sample averages. Column 1 in Table 5 estimates the basic specification of an urn-ball matching function, with the attractiveness of jobs represented by distance to the job-seeker’s location, corresponding to the specification of Table 3. The parameter estimates reported in column 1 are all very close to their sample averages reported in Table 3. The associated standard errors are corrected for some (arbitrary) structure of spatially correlated shocks.⁸ Both β and δ are highly statistically significant, while γ is not

⁸In particular, we assume that spatial correlation across wards decays at rate δ with ward distance or

statistically different from zero. To determine the returns to scale in the matching function, recall that A_b is homogeneous of degree $\gamma/(1+\beta\gamma)$ in U and V . Thus the returns to scale can be obtained multiplying by $\gamma/(1+\beta\gamma)$ the elasticity of matches with respect to applications. Such elasticity is equal to $(1 - M_b/V_b)/(M_b/V_b)(1 - \beta) \exp(\alpha)A_b^{1-\beta}$, and can be computed using estimates of α and β , and predicted values for M_b/V_b and A_b . The sample average of this expression equals -0.011, implying a returns to scale estimate of 0.989, which is very close to constant returns.

Column 2 tries to assess whether job applications are a sufficient statistics for describing local job matches. In other words we test whether local unemployment still retains some explanatory power on local job matches, once one controls for applications per job as predicted by the model. For this purpose we estimate the following urn-ball matching function:

$$\frac{M_b}{V_b} = 1 - \exp \left[- \exp(\alpha) A_b^{1-\beta} \left(\frac{U_b}{V_b} \right)^{\alpha_1} \right] + e_b,$$

where A_b is obtained from the contraction mapping (8) and the local unemployment to vacancy ratio is included as an extra regressor in the matching equation. Column 2 shows that the main parameter estimates β , γ and δ stay virtually unchanged from the specification of column 1, and that the local unemployment to vacancy ratio has a small, though statistically significant, impact on the matching rate. Although the coefficient on the unemployment to vacancy ratio is much lower than the coefficient on applications, given by $1 - \beta = 0.263$, this finding would point at a failure of our job application model, namely there are some local effects in matching that a simple job application model across space fails to capture.

Similarly as we noted for the log linear matching functions estimated in Section 2, there may be a problem of division bias here if the vacancy stock is measured with some error, as it appears at the denominator of both the dependent variable and of one of the right-hand side variables. The simplest way to address the division bias problem in this context consists in including the vacancy stock among right-hand side variables, with exponent α_2 . This would reveal whether the positive estimated impact of U_b/V_b in column 2 stems from its numerator or denominator. Column 3 shows that the impact of the U_b/V_b ratio on the vacancy outflow rate is somewhat reduced, and becomes non significantly different from zero when one controls for the total vacancy stock.

In columns 4 we estimate a similar specification to that of column 1, having expressed the job matching rate as a log-linear function of applications per job, as in equations (12). While estimates are very similar to those obtained on an urn-ball matching function, the

commuting cost. In particular, the estimated variance-covariance matrix of the parameters is given by $\hat{V} = \hat{\sigma}^2(\hat{X}'\hat{X})^{-1}(\hat{X}'\hat{\Omega}\hat{X})(\hat{X}'\hat{X})^{-1}$, where $\hat{\sigma}^2$ is the sum of squared residuals divided by the number of observations, \hat{X} is the matrix of partial derivatives of the regression function with respect to right-hand side variables, and the spatial correlation matrix $\hat{\Omega}$ is proxied by $\exp(-\hat{\delta}D)$, where D is given by the distance or commuting cost matrix, according to specifications, and $\hat{\delta}$ is the associated parameter estimate.

log-linear specification has the advantage of delivering a constant elasticity of the matching rate with respect to applications, equal to $1 - \beta$. As A_b is homogeneous of degree $\gamma/(1 + \beta\gamma)$ in U and V , this would in turn deliver an elasticity of the matching rate with respect to U and V equal to $(1 - \beta)\gamma/(1 + \beta\gamma)$. Using estimates from column 4, this is equal to -0.008 . A wald test on this statistics gives a χ^2 value of 0.193, which falls below the 5% critical value of 3.84, thus the hypothesis of constant returns to scale cannot be rejected.

In column 5 we estimate an urn-ball matching function, having expressed the utility of jobs at different locations as a function of commuting costs.⁹ The results are fairly similar to those based on geographic distance, with the job congestion estimate at 0.7, and again close to constant returns to scale. What differs from column 1 is of course the estimate of the δ parameter, being based on a different distance metrics. To give an idea of magnitudes, the attractiveness of a given job is reduced by a half for each extra £1 added to the one-way commuting cost (at 2001 prices).

In the specification of column 6 we consider that target labor markets may differ not only in terms of geographic distance (or commuting costs) from an applicant's location, but they may also differ in terms of the skill composition of available jobs, as measured by occupations. We thus construct an index of mismatch between the skill composition of each origin labor market and that of each destination labor market, based on the occupational composition of claimants and job vacancies.¹⁰ In particular, we extract data on claimants and job vacancies by CAS ward and 1-digit occupation, and construct the following mismatch index between origin area a and destination area b :

$$m_{ab} = \sum_{k=1}^8 \left| \frac{U_{ka}}{U_a} - \frac{V_{kb}}{V_b} \right|, \quad (13)$$

where the occupation categories considered are: (1) managers and professionals; (2) associate professionals and technical occupations; (3) administrative and secretarial occupations; (4) skilled trades occupations; (5) personal service occupations; (6) sales and customer service occupations; (7) process, plant and machine operatives; (8) elementary occupations. We then express the utility of a job in area b for an unemployed in area a as a function of both the geographic distance and the occupational mismatch index in (13):

$$f_{ab} = \exp(-\delta d_{ab} - \mu m_{ab}),$$

where μ is an extra parameter to be identified. The estimate for μ reported in column 6 is positive and significant, implying that the unemployed would be discouraged to send applications in areas where the array of jobs available would not match their sought occupation.

⁹These data were obtained by Daniel Graham at Imperial College and have their origins in transport planning.

¹⁰For the unemployed the occupation refers to the type of job sought.

Quantitatively though, the impact of occupational mismatch on job applications is relatively modest. The mismatch index obtained has an average of 1.067, with a standard deviation of 0.042. Thus a one standard deviation increase in the mismatch index would imply a fall in the utility of applying to a given area by 3.2%, while a one standard deviation increase in distance would imply a fall in such utility by 98.7%.

To compare the relative merits of the job application model with the simple matching function in vacancies and unemployment only, the specification of column 7 only includes U_b/V_b as a regressor. The coefficient on U_b/V_b is positive and significant, although the adjusted R^2 is much lower than that obtained when estimating the job application model of column 1. Thus the job application model seems to perform better at explaining the variation in job matching rates than the simple matching function in unemployment and vacancies only. Again we tried to address any division bias issue by also including the vacancy stock as a regressor (results not reported), but this hardly affected our conclusions.

We finally assess whether there are systematic differences in the job application and matching process of skilled and unskilled workers, by estimating our main specification for low and medium-high occupations separately. The low-occupation group includes categories (6)-(8) above, and the medium-high occupation group includes categories (1)-(5). This partition allows us to have about half of the vacancies and half of the unemployed (on average) in each group. While available unemployment and vacancy data would allow for a much finer segmentation of skills and occupations, the finer the classification used the higher the number of wards with no vacancies in the occupation considered, and these would have to be dropped from our analysis. This may become a serious issue in a spatial job application model, because dropping several wards from the sample would leave adjacent wards with no effective neighbour labour market where unemployed could apply for jobs, thus distorting our estimate for the cost of distance. Just to give an example, the proportion of wards with no vacancies at all is only 1.1%, but the proportion of wards with no vacancies in each 1-digit occupations ranges between 10% (for elementary occupations) and 38% (for associate professionals and technical occupations). We then compromise and simply divide the market into low-skill occupations and medium-high skill occupations, losing 2.7% of wards for the former and 5.4% for the latter.

The results are reported in Table 6, using geographic distance and commuting costs in turn to represent the utility of jobs at different locations. The results are remarkably similar across occupational levels, and in particular the market for low-skill occupations does not seem to be any more local than the market for medium and high-skill occupations. If anything, the γ parameter is higher in absolute value for medium-high skill occupations than for the less-skilled, but in neither case is it significantly different from zero. These estimates by occupation should be probably accompanied by an important qualification, namely that medium- and high-skill jobs and workers may be more negatively selected into

the Jobcentres dataset than low-skilled jobs and workers, and different selection patterns across skill groups may affect the comparability of estimates.

4.2 Commuting flows

Our estimated model of job applications across space has predictions for commuting patterns among wards in our sample. In particular, the share of applications to ward b that come from ward a is given by the number of applications that the unemployed in a send to jobs in b , divided by the total number of applications received by jobs in b , i.e.

$$\frac{U_a N_{ab}}{A_b V_b}. \quad (14)$$

As firms are assumed to select job-seekers randomly within the pool of job applicants, the ratio in (14) also denotes the proportion of total matches in ward b that involve job-seekers from ward a . Thus the number of vacancies in ward b that are filled by job-seekers in ward a is given by

$$\frac{U_a N_{ab}}{A_b V_b} M_b. \quad (15)$$

Finally one can obtain the distribution of commutes predicted by the model as the share of workers who live in ward a and work in ward b , for all possible pairs (a, b) . Given (15), this is equal to

$$\frac{N_{ab} M_b / A_b V_b}{\sum_{b'} N_{ab'} M_{b'} / A_{b'} V_{b'}}. \quad (16)$$

We draw data on commuting patterns from the Special Workforce Statistics of the 2001 Census. This provides a count of all those who live in one ward and work in another, albeit with some noise deliberately induced to preserve anonymity in cells with small numbers. Census data encompass commutes of all workers, while our model is meant to explain the commutes associated to newly-filled Jobcentre vacancies. These two concepts of commuting may not coincide if, for example, workers who move from one job to another tend, on average, to shorten their commute, or if job-seekers filling Jobcentre vacancies have different commuting patterns from job-seekers who find their jobs via other methods.

However, we do have some indirect evidence that this potential concern is not a major one. The LFS contains data on commuting times for those in new jobs and continuing jobs and, for those in new jobs, how that job was obtained. Table 6 presents the data on the average length of commute for these groups. One notices very little variation in the average commute between the group of workers whom we model – those who have recently got a job through a Jobcentre – and the whole economy. As the characteristics of workers in different categories may differ, and they may be related to commuting times, we also regressed commuting times on the method used to find the current job, controlling for age, gender, region and year (results not reported), and we found no significant difference between

commuting times of those who found jobs via Jobcentres and those who are not on new jobs. So, we feel justified in comparing the commutes predicted by our model with the data for all workers.

Using estimates from a job applications model with an urn-ball matching function (column 1 in Table 5), we estimate that the correlation between actual and predicted commuting flows is 0.71. Interestingly, this rises to 0.83 if one excludes “locals” from the sample, i.e. individuals who live and work in the same ward. Thus our model provides a fairly good representation of commuting patterns, but it reproduces the behavior of those who live and work in the same ward less accurately than that of commuters. In particular, the model predicts that about 10% of individuals live and work in the same area, while in reality this proportion is about 24%. Thus our model overestimates the number of commuters and it underestimates the number of locals. This is consistent with the finding that the local unemployment to vacancy ratio still plays some role in explaining variations in matching rates, having controlled for applications per job as predicted by the model.

4.3 The propagation of local shocks

A key policy question is whether one can alleviate unemployment in a depressed area using local stimulus to labor demand, or whether local stimulus is ineffective because it becomes diluted across space through a chain reaction of local spillovers. To answer this question we introduce a labor demand shock in a given ward, and we use model predictions to simulate the effect of this shock locally and its decay with distance from the target ward.

As a descriptive example, we consider an increase in the number of job openings in Stratford and New Town ward in East London, which is the main venue of the 2012 Olympic Games. In February 2005, Stratford and New Town had ratio of claimant unemployment to resident population ratio of 6%, which was nearly three times higher than the average ratio for England and Wales. We pick this example because it would show the impact of very strong, local labor demand stimulus (with thousands new job openings announced in connection to the building of the Olympic Village) in one of the highest-unemployment areas of the country.

Specifically, we simulate the impact of a doubling in the number of vacancies in Stratford and New Town Ward, from 464 to 928, under the assumption of constant returns to scale, i.e. imposing $\gamma = 0$. This assumption was not rejected in our estimates above. In the case of constant returns the total number of applications made by unemployed workers at all locations is independent of the size of the economy, and thus it remains unaffected by the shock considered (see equation (5)). Values used for β and δ are those obtained in column 1 of Table 5. With these estimates, the model predicts a total increase in the vacancy outflow, and thus in the unemployment outflow, of 212.

What is more interesting than the global effect is its spatial distribution around the target ward. The results of this exercise are reported in the upper panel of Table 7, showing the predicted percentage change in applications per job (obtained from equation (8)), in the vacancy outflow (obtained from equation (11)), and the unemployment outflow (obtained from equation (15)), within alternative distance cutoffs from Stratford. As total applications in the economy stay constant, applications per job on average fall. In Stratford, where vacancies double, applications per job fall by about 2.2%. Around Stratford, applications per job also fall because Stratford attracts job applications from surrounding areas. This spillover effect decays with distance, and the percentage change in applications per job is below 1% beyond 10 km from Stratford, and virtually zero beyond 35 km. The number of vacancies filled in Stratford rises by 98.9%, with a very slight decline in the probability of filling any one vacancy, given that the number of vacancies has doubled. There is a very tiny decrease in the number of vacancies filled in surrounding areas as well, because of increased competition for applicants in and around Stratford, but again this effect is virtually negligible.

But when we look at the change in the unemployment outflow, we find no evidence at all of any sharp local effect, with the unemployment outflow in Stratford only rising by 0.4%. If anything, the unemployment outflow within 20 km rises slightly more than in Stratford,¹¹ and beyond this cutoff distance the change in the unemployment outflow becomes negligible. The bottom line is that, while labor markets are quite ‘local’, in the sense that the attractiveness of job offers strongly declines with distance, local labor markets do overlap; thus the ripple effect generated by local shocks implies that their propagation is fairly wide. One should therefore conclude that even strong local stimulus has a limited bite on the local outflow rate from unemployment, because a series of spatial spillovers would greatly dilute any local shock across space. Specifically, unemployed workers living relatively close to Stratford divert some of their job search effort from their local wards towards Stratford. This reduces local job competition and attracts applications from elsewhere, and so on. This mechanism explains the spatial propagation of local shocks in the presence of relatively high costs of distance.

5 Conclusions

In this paper we have used high-frequency, geographically very disaggregated data on unemployment and vacancies to investigate the extent to which labour markets are ‘local’. We have presented some non-structural estimates in which the probability of filling a vacancy

¹¹This non-monotonicity comes from the fact that the $\left[\sum_{b'} V_{b'} A_{b'}^{-\beta} \exp(-\delta d_{ab'})\right]^{\gamma-1}$ term, capturing the extent of job competition, falls more in Stratford than elsewhere. This term determines the number of applications sent from each area a to each area b , and thus the unemployment outflow in each area a , according to equation (15).

is influenced by unemployment and vacancies in the surrounding area, though more distant unemployment and vacancies have a diminishing impact. However, we argued that such estimates cannot adequately estimate the true cost of distance. We presented a simple job search model that allows, in a tractable way, to estimate a market process with a very large number of market segments. Our estimates of that model suggest that the cost of distance is relatively high. That is, the utility of being offered a job decays at exponential rate around 0.25 with distance (in km) to the job. Also, workers are discouraged from applying to jobs where they expect a large number of applications. Finally, constant returns in matching markets are not rejected, and in particular the total number of job applications made in this economy does not respond to the absolute size of the vacancy. The estimated model seems to replicate fairly accurately actual commuting patterns across Census wards, although it tends to underpredict the proportion of individuals who live and work in the same ward. Finally, we find that, despite the fact that labor markets are relatively ‘local’, local stimulus to labor demand is fairly ineffective in raising the local unemployment outflow, because labor markets overlaps and the associated ripple effects in applications largely dilute the effect of local stimulus across space.

Does any of this matter? If residential mobility is costless, one could argue that it does not. Workers who are unfortunate enough to find themselves in an area with few vacancies would simply up and move to somewhere with better opportunities. But we have evidence that residential mobility is limited especially for low-skill workers. This seems true in the US but is likely to be even more so in the UK where residential mobility rates are lower and those who in subsidized public housing are especially immobile.

References

- [1] Burgess, S. M. and S. Profit. (1998). “Externalities in the Matching of Workers and Firms in Britain,” *Labour Economics* 8:3, pp. 313-333.
- [2] Busso, Matias, Jesse Gregory and Patrick Kline (2010), “Assessing the Incidence and Efficiency of a Prominent Place Based Policy”, UC Berkeley, mimeo.
- [3] Butters, G. R. (1977). “Equilibrium Distributions of Sales and Advertising Prices,” *Review of Economic Studies*, 44:3, pp. 465-91.
- [4] Coles, M. G. and B. Petrongolo (2008). “A Test Between Stock-flow Matching and the Random Matching Function Approach,” *International Economic Review* 49, pp. 1113-1141.
- [5] Coles, M. G. and E. Smith. (1996). “Cross-Section Estimation of the Matching Function: Evidence from England and Wales,” *Economica*, 63:252, pp. 589-98.
- [6] Flinn, C. and J. Heckman. (1983). “Are Unemployment and Out of the Labor Force Behaviorally Distinct Labor Force States?,” *Journal of Labor Economics* 1(1), pp. 28-42.
- [7] Glaeser, Edward L. and Joshua D. Gottlieb (2008) “The Economics of Place-Making Policies”, *Brookings Papers on Economic Activity*, Spring 2008, 155-253.
- [8] Kain, John F. (1968) “Housing Segregation, Negro Employment, and Metropolitan Decentralization”, *Quarterly Journal of Economics* 82, 175-197.
- [9] Machin, A. (2001). “The New Vacancy Survey: A New Series of National Statistics”, National Statistics Features, ONS. Available at http://www.statistics.gov.uk/articles/labour_market_trends/VacancySurvey_LMT0703.pdf.
- [10] Manning, A. (2003). *Monopsony in Motion. Imperfect Competition in Labor Markets*. Princeton, Princeton University Press.
- [11] Moretti, Enrico (2010) “Local Labor Markets”, UC Berkeley, mimeo.
- [12] Notowidigdo, Matthew J. (2010) “The Incidence of Local Labor Demand Shocks”, MIT, mimeo.
- [13] OECD (2000). *The Public Employment Service in the United States*.
- [14] Petrongolo, B. and C. A. Pissarides. “Looking into the Black Box: A Survey of the Matching Function,”. *Journal of Economic Literature* 39, 390-431, 2001.

- [15] Pissarides, C. A. (1979). "Job Matchings with State Employment Agencies and Random Search," *Economic Journal*, 89:356, pp. 818-33.
- [16] Jones, S. R. G., and C. Riddell. (1999). "The Measurement of Unemployment: An Empirical Approach," *Econometrica* 67, pp. 147-161.




Figure 1: Main parameter estimates over the sample period May 2004-April 2006.

Table 1
Descriptive statistics on local labor markets.

Variable	Mean	St. dev.	No. Obs.
Unemployment stock	105.7	147.4	210755
Unemployment inflow	20.4	24.6	210755
Unemployment outflow	19.7	23.8	210755
Vacancy stock	91.0	227.4	210755
Vacancy inflow	28.1	72.1	210755
Vacancy outflow	28.8	73.2	210755

Notes. Source: NOMIS. Sample: CAS 2003 Wards in England and Wales, May 2004-April 2006.

Table 2
Log-linear matching functions

	(1)	(2)	(3)	(4)	(5)	(6)
$\text{Log } \hat{V}_{10b}$	-0.222*** (0.00561)	-0.207*** (0.00598)	-0.221*** (0.00354)	-0.169*** (0.00947)	-0.135*** (0.0164)	-0.126*** (0.0161)
$\text{Log } \hat{U}_{10b}$	0.210*** (0.00490)	0.197*** (0.00524)	0.206*** (0.00300)	0.0727*** (0.0189)	0.0634** (0.0292)	0.113*** (0.0208)
V_b / \hat{V}_{10b}	-0.974*** (0.0486)	-0.962*** (0.0485)	-0.879*** (0.0245)	-1.491*** (0.0480)	-0.995*** (0.0740)	-1.006*** (0.0739)
V_{5b} / \hat{V}_{10b}	-0.0891*** (0.0238)	-0.0842*** (0.0238)	-0.0741*** (0.0127)	-0.0317 (0.0345)	0.0266 (0.0472)	0.0290 (0.0472)
V_{20b} / \hat{V}_{10b}	-0.000502 (0.000919)	-0.000303 (0.000840)	-0.00359*** (0.000418)	-0.000135 (0.000264)	-0.00249** (0.00126)	-0.00237* (0.00123)
V_{35b} / \hat{V}_{10b}	-3.03e-06 (9.01e-05)	2.17e-05 (8.37e-05)	-3.35e-06 (9.22e-05)	-1.42e-07 (5.82e-05)	-1.51e-05 (0.000312)	4.33e-05 (0.000309)
U_b / \hat{U}_{10b}	0.779*** (0.0576)	0.781*** (0.0570)	0.688*** (0.0310)	-0.205** (0.0874)	0.340** (0.156)	0.229** (0.0911)
U_{5b} / \hat{U}_{10b}	0.252*** (0.0272)	0.244*** (0.0271)	0.225*** (0.0144)	0.0715 (0.0909)	0.0453 (0.143)	0.0712 (0.0941)
U_{20b} / \hat{U}_{10b}	-0.00147* (0.000783)	-0.00175** (0.000762)	0.000499 (0.000444)	-0.000686 (0.000885)	0.000266 (0.00276)	0.000632 (0.00139)
U_{35b} / \hat{U}_{10b}	0.000282*** (0.000109)	0.000237** (0.000109)	0.000178** (7.48e-05)	0.000217 (0.000136)	-2.10e-05 (0.000366)	0.000219 (0.000217)
	197579	197579	188651	197579	188594	188651
R-squared	0.049	0.073	0.073	0.046		
Time Effects	No	Yes	Yes	Yes	Yes	Yes
Ward Effects	No	No	No	Yes	Yes	Yes
Instruments	No	No	Vacancy Variables	No	Vacancy/ Unemploym.	Vacancy Variables

Notes. The Table provides estimates for equation (2) in the text. Sample: England and Wales, May 2004-April 2006.

Table 3
Log linear and linear matching functions.

	(1)	(2)	(3)	(4)	(5)
$\text{Log } \hat{V}_{10b}$	-0.206 ^{***} (-0.00535)	-0.213 ^{***} (0.0076)	-0.202 ^{***} (0.0078)	-0.199 ^{***} (-0.0129)	-0.155 ^{***} (0.0132)
$\text{Log } \hat{U}_{10b}$	0.197 ^{***} (-0.00484)	0.202 ^{***} (0.0079)	0.201 ^{***} (0.0080)	0.196 ^{***} (-0.0114)	0.169 ^{***} (0.0115)
V_b / \hat{V}_{10b}	-0.962 ^{***} (-0.0484)	-0.420 ^{**} (0.1043)	-0.349 ^{**} (0.1056)	-0.853 ^{***} (-0.1140)	-0.521 ^{***} (0.1155)
V_{5b} / \hat{V}_{10b}	-0.0846 ^{***} (-0.0238)	0.018 (0.0387)	0.0036 (0.0394)	-0.0153 (-0.0575)	-0.062 (0.0574)
U_b / \hat{U}_{10b}	0.781 ^{***} (-0.0575)	0.307 ^{**} (0.0934)	0.397 ^{**} (0.0943)	0.642 ^{***} (-0.1450)	0.311 ^{***} (0.1405)
U_{5b} / \hat{U}_{10b}	0.252 ^{***} (-0.0269)	0.092 ^{**} (0.0448)	0.086 [*] (0.0456)	0.192 ^{***} (-0.0648)	0.121 ^{***} (0.0640)
Observations	197579	197579	208717	8282	8708
R-squared	0.0724		0.5449	0.0434	0.0347
Funct. Form	Log-Linear	Linear	Linear	Log-Linear	Linear
Time Effects	Yes	Yes	Yes	Only Feb 2005	Only Feb 2005
Sample	Non-zero Outflow	Non-zero Outflow	All	Non-zero Outflow	All

Notes. Columns (1) and (4) provide estimates for equation (2) in the text. Estimation method is OLS. Columns (2), (3) and (5) provide estimates for the exponential of equation (2) in the text. Estimation method is nonlinear least squares. Sample: England and Wales, May 2004-April 2006 in columns (1)-(3) and February 2005 in column (5).

Table 4
Estimates of a job application model with an urn-ball matching function.
Sample averages over May 2004-April 2006

	Mean	St. Dev.	Min	Max	No. months
β	0.745	0.035	0.652	0.793	24
γ	-0.068	0.075	-0.175	0.064	24
δ	0.300	0.064	0.209	0.470	24
Constant (α)	-0.766	0.192	-1.055	-0.419	24

Notes. The Table reports mean estimates of the parameters β , γ , δ , and α across the 24 months from May 2004-April 2006, together with standard deviations, minimum and maximum values.

Table 5
Estimates of a job application model
Alternative specifications for February 2005

Dependent variable: Vacancy outflow rate							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
β	0.747 ^{***} (0.064)	0.793 ^{***} (0.076)	0.770 ^{***} (0.087)	0.799 ^{***} (0.050)	0.756 ^{***} (0.042)	0.745 ^{***} (0.044)	
γ	-0.043 (0.087)	-0.067 (0.121)	-0.009 [*] (0.131)	-0.038 (0.086)	-0.053 (0.084)	-0.048 ^{**} (0.057)	
Δ	0.259 ^{***} (0.058)	0.226 ^{***} (0.070)	0.241 ^{***} (0.069)	0.263 ^{***} (0.059)	0.766 ^{***} (0.117)	0.260 ^{***} (0.038)	
μ						0.783 (0.182)	
Constant (α)	-0.771 ^{***} (0.141)	-0.770 ^{**} (0.167)	-0.777 ^{**} (0.151)	-1.004 ^{**} (0.110)	-0.762 ^{***} (0.123)	-0.758 ^{***} (0.083)	-0.899 ^{***} (0.026)
U_b/V_b		0.044 ^{***} (0.008)	0.026 (0.025)				0.087 ^{***} (0.013)
V_b			-0.020 (0.025)				
Observations	8709	8709	8709	8709	8709	8709	8709
Adjusted R ²	0.0395	0.0302	0.0305	0.0391	0.0374	0.0410	0.0109
Distance concept	Geographic	Geographic	Geographic	Geographic	Cost	Geographic	-
Matching Function specification	Urn-ball	Urn-ball	Urn-ball	Log-linear	Urn-ball	Urn-ball	Urn-ball

Table 6
Estimates of a job application model with an urn-ball matching function
Low- and medium-high occupations

Dependent variable: Vacancy outflow rate				
	Occupations			
	Low	Medium-high	Low	Medium-high
	(1)	(2)	(3)	(4)
β	0.738 ^{***} (0.063)	0.841 ^{***} (0.071)	0.740 ^{***} (0.045)	0.858 ^{***} (0.235)
γ	-0.081 (0.088)	-0.153 (0.151)	-0.104 (0.077)	-0.185 (0.125)
δ	0.266 ^{***} (0.058)	0.278 ^{***} (0.095)	0.746 ^{***} (0.123)	0.855 ^{***} (0.042)
Constant (α)	-0.627 ^{***} (0.143)	-0.689 ^{***} (0.147)	-0.597 ^{***} (0.123)	-0.689 ^{***} (0.104)
Observations	8546	8336	8546	8336
Adjusted R ²	0.0437	0.0140	0.0429	0.0138
Distance concept	Geographic	Geographic	Cost	Cost
Matching function specification	Urn-ball	Urn-ball	Urn-ball	Urn-ball

Notes. Low occupations include: sales and customer service occupations; process, plant and machine operatives; elementary occupations. Medium-high occupations include: managers and professionals; associate professionals and technical occupations; administrative and secretarial occupations; skilled trades occupations; personal service occupations.

Table 6.
Average commuting times in the UK.

	Mean	Std. Dev.	No. Obs.
Not on new job	24.5	22.2	612787
On new job, found via:			
Reply to advert	24.5	21.6	16059
Job centre	24.5	20.2	4491
Careers office	30.2	26.1	453
Jobclub	25.6	25.6	61
Private agency	34.6	26.4	4859
Personal contact	23.2	23.0	15523
Direct application	22.4	21.7	9646
Some other method	27.7	26.7	5618
Total	24.5	22.3	669497

Notes. Figures report one-way daily commuting times (in minutes). New jobs are defined by tenure up to three months. Source: Labour Force Survey, 1993-2007.

Table 7: The propagation of local shocks

Distance	Percentage change		
	Applications per job	Vacancy outflow	Unemployment outflow
0 km	-2.16	98.90	0.40
(0,5] km	-1.88	-0.48	0.66
(5,10] km	-1.33	-0.34	0.66
(10,20] km	-0.76	-0.19	0.45
(20,35] km	-0.27	-0.07	0.18
(35,50] km	-0.04	-0.01	0.04
50+ km	-0.00	-0.00	0.00

Notes: The Table shows the simulated effect of a doubling in the number of vacancies in Stratford and New Town Ward, using the estimates from column 1 of Table 5 (having set $\gamma=0$).