DISCUSSION PAPER SERIES

No. 8686

HOW LOCAL ARE LABOR MARKETS? EVIDENCE FROM A SPATIAL JOB SEARCH MODEL

Alan Manning and Barbara Petrongolo

LABOUR ECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: **www.cepr.org/pubs/dps/DP8686.asp**

ISSN 0265-8003

HOW LOCAL ARE LABOR MARKETS? EVIDENCE FROM A SPATIAL JOB SEARCH MODEL

Alan Manning, London School of Economics
Barbara Petrongolo, Queen Mary, University of London; CEP, LSE and CEPR

Discussion Paper No. 8686
December 2011

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 8686

December 2011

ABSTRACT

How local are labor markets? Evidence from a spatial job search model*

This paper uses data on very small UK geographies to investigate the effective size of local labour markets. Our approach treats geographic space as continuous, as opposed to a collection of non-overlapping administrative units, thus avoiding problems of mismeasurement of local labour markets encountered in previous work. We develop a theory of job search across space that allows us to estimate a matching process with a very large number of areas. Estimates of this model show that the cost of distance is relatively high - the utility of being offered a job decays at exponential rate around 0.3 with distance (in km) to the job - so that labour markets are indeed quite `local'. Also, workers are discouraged from applying to jobs in areas where they expect relatively strong competition from other jobseekers. The estimated model replicates fairly accurately actual commuting patterns across neighbourhoods, 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 development policies are fairly ineffective in raising the local unemployment outflow, because labor markets overlap, and the associated ripple effects in applications largely dilute the impact of local stimulus across space.

JEL Classification: J61, J63, J64 and R12
Keywords: job search, local labour markets, location-based policies and ripple effects

Alan Manning
Professor of Economics
Economics Department
London School of Economics
Houghton Street
LONDON WC2A 2AE

Email: a.manning@lse.ac.uk

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=108595

Barbara Petrongolo
Professor of Economics
School of Economics and Finance
Queen Mary University
and Centre for Economic
Performance
London School of Economics
Houghton Street
London WC2A 2AE

Email: b.petrongolo@lse.ac.uk

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=137529

* We wish to thank Stéphane Bonhomme and seminar participants at the London School of Economics, Paris School of Economics, University of Toulouse, the University of Essex, City University, the NBER Summer Institute Labor Studies 2010, the NBER Summer Institute Macro Perspectives 2011, the CEPR ESSLE Conference 2010, the Bank of Portugal Conference on Wages, Job Turnover and Education 2011, and the Cergy-Pontoise Conference on Labor Markets 2011 for helpful comments. We also thank Timothée Carayol for excellent research assistance.

Submitted 25 November 2011

1 Introduction

How local are labor markets? A number of important questions in labor economics turn on the answer. In recent years there has been a resurgence of interest in the consequences of localization of economic activity for workers' welfare (see Moretti, 2011, for a recent overview) and in policies aimed to improve labor market outcomes in disadvantaged areas (Glaeser and Gottlieb, 2008). In the US, federal, state and local governments combined spend nearly $50 billion per year on local development policies. Notable examples are the Moving to Opportunities program (Katz, Kling and Liebman, 2001) and the Empowerment Zones program (Busso, Gregory and Kline, 2011). These policies need to know about the size of local labor markets to decide about the appropriate nature of the intervention. If labor markets are very local then an effective intervention will have to be targeted to the disadvantaged areas themselves even if those areas are not conducive to generating employment. But if labor 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. There is also a sizeable related literature on the incidence of local shocks to labor demand and their impact on labor mobility and labor market equilibrium (see among others Blanchard and Katz, 1992, Bound and Holzer, 2000, and Notowidigdo, 2011). Such research needs a clear definition of a 'place'.[1]

Most academic research on the topic and government statistical agencies divides geographical space into non-overlapping areas, which are then assumed to be single labor markets. Examples would be the 367 Metropolitan Statistical Areas, the BEA's 179 Economic Areas and the 720 Commuting zones for the US, or the 320 Travel to Work Areas (TTWAs) for the UK. Within labor markets, spatial inequalities would be interpreted as the outcome of residential sorting. While these efforts are understandable and useful they do have their problems. For example, 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 labor 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 labor market. And the non-overlapping nature of local labor markets constructed in this way causes inevitable discontinuities around the boundaries. Someone living just inside the London TTWA will be classified as living in an enormous labor market while someone living just across the border

[1] Another related issue is the spatial mismatch hypothesis (Kain, 1968), suggesting 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 labor market of those living in the city (see also Hellerstein, Neumark and McInerney, 2008, and Boustan and Margo, 2009, for recent studies).

in the Luton TTWA will be classified as being in a modestly-sized local labor market. In reality, these people are in essentially the same labor market.

The fundamental cause of these problems is a failure to recognize the continuous nature of geographic space, that the labor market for one individual at one location will overlap with that for another individual in a different but not too distant location. Consequently, there is no way to segment an economy into non-overlapping areas without mismeasuring local labor markets. But, how can one model geographical space in a more realistic way while preserving tractability? One of the contributions of this paper is to show how one might approach this problem.

The approach of this paper is to treat geographic space as continuous, as opposed to a collection of non-overlapping administrative units. This avoid problems of mismeasurement of local labor markets described above, and turns out to have crucial implications for the impact of local policies. We model the size of local labor markets by the cost of distance, whether measured as geographic distance, commuting time, or commuting cost. If the cost of distance is high then workers will be more reluctant to consider jobs at more distant locations than if the cost of distance is small. This approach means that the boundary of a local labor market is fuzzy but that is the right way to think about a jobseeker's decision problem.

Let us briefly consider how one might approach the question of estimating how local are labor markets from individual behavior. Commuting data might be expected to contain useful information about the size of local labor 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 bunch 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 labor 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 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 labor market as they will consider a very wide range of jobs. But they will end up with a low commute, so commuting data would not reveal the true extent of the local labor market.

Our research design is intended to try to avoid these potential problems and get closer to the heart of the question of how local are labor markets. By a local labor 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 would not concern us. We use monthly data on unemployment and vacancies in small neighborhoods in England and Wales (8,850 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. The large number of neighborhoods means that our data provide a much closer approximation to the continuous nature of geographic space than all existing studies.

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 propose a model in which 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 labor 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 labor market for people who live in area B but not for those who live in area C.

The job search model that we propose provides an explicit micro-foundation for how to model the linkages between a very large number of areas in a way that preserves tractability. Unemployed workers decide whether to apply to job vacancies at different locations based on the cost of distance to target jobs and on the likelihood that applications to such jobs be successful, in turn depending on how many other jobseekers across the economy would find these jobs attractive. Inter-dependencies across areas arise because the number of applicants to jobs in a given area is likely to be influenced (even if only very slightly) by unemployment and vacancies in all other areas, insofar they are ultimately linked through a series of overlapping markets. Key parameters of this framework are the rate of decay of the utility obtained from a job with the distance from that job, and the way in which job competition in a given area discourages jobseekers from applying to jobs in that area. The resulting vector of job applications in each area is the central ingredient of the process by which local vacancies are filled.

We estimate our model using data on unemployment and vacancies disaggregated at the Census ward level. We use very small spatial units such as wards as building blocks for overlapping local labor markets, and let job matching patterns at the local level guide us as to the effective size of local labor markets. Estimates of this model suggest that the cost of distance is relatively high. Specifically, the attractiveness of a job falls by about 4.5 times if one pulls the job 5 km further afield from a jobseeker's location. Also, workers are discouraged from applying to jobs in areas where they expect relatively strong competition from other jobseekers. An interesting side result is that constant returns in search markets are not rejected, implying that larger-scale markets would

not systematically offer more efficient matching of workers to jobs. The estimated model predicts commuting patterns across UK wards that replicate fairly accurately actual commuting patterns obtained from the 2001 Census, 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', location-based policies in the form of local stimulus to labor demand or improved transport links are rather ineffective in raising the local job finding rate, because labor markets overlaps and the associated ripple effects in applications largely dilute the effect of local policies across space.

The plan of the paper is as follows. In the next Section we describe the data we use and in Section 3 we 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 labor markets. In Section 4 we then present our structural model of job search across space, which is then estimated in Section 5. Section 6 uses structural estimates to illustrate the simulated impact of location-based policies on the spatial distribution of the unemployment outflow. Section 7 finally concludes.

2 Data and descriptive statistics

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 1,222 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[2].

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 neighborhood within a city has at least one Jobcentre. A Jobcentre's services are free of charge to all users, both to jobseekers 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

[2]Because the border between England and Scotland is very sparsely populated, the commuting flows across the border are small so it is not a problem to regard Scotland as distinct from England.

defined using the full postcode of the job location. Each job vacancy is advertised in three ways: on the centralized employment website www.direct.gov.uk; through the Jobcentre Plus phone service for job applicants; and 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, distance from a given postcode etc.). The detailed geographic information on both claimant unemployment and job vacancies recorded at Jobcentres makes them a unique data source for studying job matching patterns at the very local level.

While the monthly series run from April 2004 onwards, 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 since May 2006 are not fully comparable to those for the earlier period.[3]

The data we use on the unemployed and vacancies cover a very large number of jobseekers 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, one should realize they do not represent the universe of jobseekers 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 jobseekers are claimant-count unemployed, as jobseekers may also be employed, or unemployed but not claiming benefits; and not all the claimant unemployed may be jobseekers (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 labor 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 2005 (to give one example) the LFS suggests there were about 3.1 million jobseekers in the UK, and total employment was about 28.1 million. Almost exactly half of the jobseekers were not currently employed, and at that time the official figures for the claimant count was about 875,000. 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 nearly a quarter of total jobseekers in the economy.

It may be argued that the claimants are among the most intensive jobseekers (see, among others, Flinn and Heckman, 1983, Jones and Riddell, 1999), and thus we weight jobseeker figures in the

[3]The vacancy data after May 2006 are less suitable for our purposes. Prior 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.

LFS by the number of reported search methods used. During the 2002-2007 period,[4] the unweighted share of claimants in total jobseekers was 17.6%, while the weighted share was 23.7%. As one would expect, the share of claimants in jobseekers 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 35.2% among those with no qualifications. This means our study is best interpreted as being about lower-skill labor markets that probably tend to be more 'local'.

For our purpose it is also important to know the fraction of jobseekers who are looking at the 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 report 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 jobseekers in our sample. In this regard, it is also 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 jobseekers 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 are asked how many job vacancies there are in their business, for which they are 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 are reasons to believe that such proportion may be overstated (Machin, 2001). In particular, in May 2002, an extra question was added to the ONS Vacancy Survey, on whether vacancies reported had also been 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%

[4] We need to expand the sample period here with respect to that used in the main analysis in order to improve precision of the statistics reported.

proportion is markedly lower that the two thirds recorded for the post-2004 period. According to Machin (2001), the main 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",[5] thus one can still imagine that two-thirds is indeed an upper bound for the fraction of job openings that are effectively available to jobseekers at Jobcentres. As no occupation breakdown 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 it is very likely that Jobcentre vacancies over-represent less-skilled jobs.

From this discussion it should be clear that we capture an important section of both supply and demand of the job search process in the UK, especially for low-skilled workers and jobs, but it is also clear that we cannot provide a fully comprehensive picture. This would introduce 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, a period of historically low and stable unemployment.[6] English wards have on average 106 unemployed and 91 vacancies. Taken across the whole period, both unemployment and vacancy inflows and outflows seem very similar but with vacancy outflow slightly

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

[6] We 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.

lower than the inflow and the unemployment inflow slightly above the outflow. There is also very wide spatial variation in unemployment and vacancies, which may be best grasped on a map. Figure 1 is a map of England and Wales, that shows spatial variation in the unemployment-to-vacancy ratio for a representative month in our sample, February 2005. Different wards are shaded according to the quartile of the corresponding U-V ratios, with darker shades corresponding to higher quartiles. The striking feature that emerges from this map is that there is no simple pattern - rather we observe a patchwork of very different labor market outcomes across quite small areas, e.g. many high-unemployment wards are adjacent to low-unemployment wards so that one cannot detect one large region in which, say, all high-unemployment wards are clustered together. Figure 2 shows the same picture for London - one can observe the central business areas where the U-V ratio is low and residential areas (especially in inner London) where the U-V ratios are high. But, again, there is a patchwork quality to the picture - wards with unemployment/vacancy ratios in the top quartile coexist next to areas with unemployment/vacancy ratios in the bottom quartile.

3 Non-structural estimates

We start our investigation of the data by estimating a conventional log-linear matching function where the dependent variable is the vacancy outflow rate and the regressors are unemployment and vacancies, possibly augmented by local spillovers:

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

where M_b is the vacancy outflow from ward b, 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. The dependent variable is thus the vacancy 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 vacancy outflow rate in b, but more distant vacancies have a smaller effect, i.e. we expect $\gamma_i < 1$. Specifications similar to (1) have been estimated by Burda and Profit (1996) for Czech districts, and Burgess and Profit (2001) and Patacchini and Zenou (2007) for UK TTWAs.

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

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

and approximate (1) by

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

This specification has the advantage to be linear in parameters, so that we can estimate various specifications using instrumental variables and/or ward fixed-effects. Moreover, one can simply read off the returns to scale in the matching function by a comparison of the coefficients on $\log \widetilde{U}_{10b}$ and $\log \widetilde{V}_{10b}$, while the coefficients on the share variables $U_b/\widetilde{U}_{10b}$,..., $U_{35b}/\widetilde{U}_{10b}$, and $V_b/\widetilde{V}_{10b}$,..., $V_{35b}/\widetilde{V}_{10b}$ 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 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 0.22 and β_1 of 0.48, i.e. unemployed workers outside the ward but with 5km have 48% of the effectiveness of generating matches as those within the ward and the unemployed in the 5-10km ring have an effectiveness of 22%. Unemployed in the 20k and 35k rings have tiny effects on the vacancy outflow, though these are statistically different from zero. For vacancies, the more local the vacancies the lower the outflow rate - as one would expect - as such jobs are closer substitutes to local ones. Vacancies outside the ward but within 5km have 28% of the effectiveness of those within the ward, and vacancies in the 5-10km ring have an effectiveness of 21%. 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 their one-month lags. 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 issue 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 lagged unemployment as an instruments for the unemployment variables and the results are virtually unchanged from those of column 3.

The fifth column introduces ward fixed effects. Comparing the estimates in the second and fifth columns one notes that the coefficient on the share of vacancies in the local ward becomes much more negative and that the coefficients on the unemployment variables are attenuated and even perversely signed. The change in the coefficient of the local vacancy share is 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. The sixth column re-estimates using instrumental variables and the coefficient on the local vacancy share is reduced as one would expect. The attenuation of the coefficients on the unemployment variables in specifications with ward fixed effects is most likely caused by the fact that unemployment rates within wards are much more stable than vacancy rates, i.e. some neighborhoods have persistently high unemployment rates, some persistently low unemployment rates. This implies that most of the useful variation in investigating spatial matching is cross-sectional and that is what we are going to exploit in our structural estimates. To allay fears that we are simply picking up fixed ward characteristics that are correlated with unemployment rates we do experiment with including other ward-level controls.

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 their higher level of aggregation, 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 will estimate vacancy outflow rate equations like (2) in levels instead of logs. In the next 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 be between zero and one 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 in 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 will feature in some of the structural estimates below.

The results of Table 2 and 3 are consistent with a simple matching model with spatial spillovers. However, these estimated equations do have their limitations for making inferences about the size of local labor 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 spillovers – 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.

4 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 is to base the empirical specification on a model of optimizing job search behavior across space that makes predictions about the number of applications from unemployed workers in one area to vacancies in every other area. Our approach is then 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 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 to 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.

4.1 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 are 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 returns from search 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:

$$C(N) = \frac{c}{1+\eta} N^{1+\eta}, \tag{3}$$

so that the net expected utility from job search 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

$$p_i u_i \geq c \left(\sum_i D_i \right)^{\eta} = c N^{\eta}. \tag{4}$$

This result says that the attractiveness of vacancies is determined by the expected utility they offer, and jobseekers 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.[7]

[7] While extremely convenient, it is important to note that the assumption that the probability of more than one application being successful is infinitesimal plays an important role here – if this assumption is not met then one cannot rank vacancies by their expected utility and the decision-problem does not lead to such a simple rule. To see this more formally, suppose that we can order jobs in terms of utility, with job 1 offering a higher utility than job 2 and so on. Furthermore, assume that the jobs that offer a higher utility have a lower probability of success so that $p_1 < p_2 < ...$ (i.e. that dominated jobs are not applied for). In this case a worker will only accept job i if

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 key in determining the returns to scale in the model. The issue of constant versus increasing returns to scale is a recurrent question in the matching literature, as increasing returns to scale lead to the possibility of multiple equilibria (Diamond, 1982) and are a potential explanation for agglomeration, as large-scale markets would offer a more efficient matching process that can, in principle, off-set higher land and labor costs of locating in agglomerations. So it is worth taking some time to understand the role of the η parameter in this framework and its link to returns to scale.

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,

no job applications to lower jobs have been successful. The expected utility from applying to a set of jobs is thus $\sum_i D_i p_i u_i \Pi_{j=1}^{i-1}(1-p_j)^{D_j}$, which leads to a decision rule in which the marginal benefit of applying to vacancy i can be written as $p_i\left(u_i - E_i\right)\left(1 - Q_i\right)$, where Q_i is the probability of getting a better job than i and E_i is the expected utility from jobs worse than i, conditional on a better job not being obtained. The effect of other applications on the decision to apply to vacancy i is no longer only contained in the effect on marginal costs. But the difference between this specifciation of marginal benefit and the one we use will be small if Q_i and E_i are small.

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 to have 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.

Hence, using (4), 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

$$\Pr\left(p(A_b)f_{ab}\varepsilon \geq cN_a^{\eta}\right) = \left(\frac{p(A_b)f_{ab}}{cN_a^{\eta}}\right)^k. \tag{5}$$

Although there is some uncertainty about whether an individual applies to a particular vacancy (because of the idiosyncratic component to utility), let us assume that 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:

$$N_{ab} = V_b\left(\frac{p(A_b)f_{ab}}{cN_a^{\eta}}\right)^k. \tag{6}$$

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 :

$$N_a = \left[c^{-k}\sum_b V_b\left(p(A_b)f_{ab}\right)^k\right]^{\gamma}, \tag{7}$$

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

Using (6) and (7), one can compute the total number of applications made by the unemployed in a to vacancies in b, N_{ab}, as

$$N_{ab} = c^{-k\gamma} V_b \left(p(A_b) f_{ab}\right)^k \left[\sum_{b'} V_{b'} \left(p(A_{b'}) f_{ab'}\right)^k\right]^{\gamma-1}. \tag{8}$$

The intuition behind expression (8) is that the number of applications sent from area a to area b depends on job opportunities in area b (V_b), how attainable they are ($p(A_b)$), and how far they are located from a (f_{ab}). The term in square brackets can be interpreted as a weighted average of vacancies everywhere in the country, where weights are given by a combination of their attainability and distance to a. This term captures the 'effective' size of the whole economy, and would simply work as a normalization in the case of constant returns ($\gamma = 0$).

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

$$\begin{aligned} A_b &= \frac{\sum_a N_{ab} U_a}{V_b} \\ &= c^{-k\gamma} \sum_a U_a \left(p(A_b) f_{ab}\right)^k \left[\sum_{b'} V_{b'} \left(p(A_{b'}) f_{ab'}\right)^k\right]^{\gamma-1}. \end{aligned} \tag{9}$$

Equation (9) simply tells that the number of applications per job in area b depends on the distribution of the unemployed across all possible origin areas a (U_a), how far they are located from b (f_{ab}), and how attainable they perceive job vacancies in b to be ($p(A_b)$).

To take equation (9) to the data, we make two further functional form assumptions, and namely $p(A_b) = A_b^{-\widetilde{\beta}}$, where $\widetilde{\beta} > 0$ denotes the effect of job competition on applications to jobs in a given area,[8] and $f_{ab} = \exp(-\widetilde{\delta} d_{ab})$, where d_{ab} may be proxied by distance between a and b, and $\widetilde{\delta}$ measures the exponential rate of decay of the attractiveness of a given job with distance to that job.[9] Under

[8] The functional form $p(A_b) = A_b^{-\widetilde{\beta}}$ 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 vacancy is filled 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 $A_b^{-\widetilde{\beta}}$. Both expressions are decreasing and convex in A_b, so we would be fitting similar functional forms to our data, but using the $A_b^{-\widetilde{\beta}}$ approximation instead of the exact formula makes the model a great deal more tractable.

[9] While 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 $f_{ab} = \exp(-\widetilde{\delta} d_{ab}) + \alpha w_b$, 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.

these assumptions we can solve for A_b :

$$A_b = \left\{ c^{-k\gamma} \sum_a U_a \exp(-\delta d_{ab}) \left[\sum_{b'} V_{b'} A_{b'}^{-\beta} \exp(-\delta d_{ab'}) \right]^{\gamma-1} \right\}^{1/(1+\beta)}, \tag{10}$$

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

Equation (10) is the key relationship delivered by our spatial job search model, and captures all the inter-dependencies between areas. In particular, the number of applicants to jobs 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, it can be shown that (10) is a contraction mapping, in which case it can be solved iteratively and economically to obtain A_b.

To see this, let's take logs of (10) to write it as:

$$\ln A_b = \frac{1}{1+\beta} \left\{ -k\gamma \ln c + \ln \sum_a U_a \exp(-\delta d_{ab}) + (\gamma - 1) \ln \left[\sum_{b'} V_{b'} e^{-\beta \ln A_{b'}} \exp(-\delta d_{ab'}) \right] \right\}. \tag{11}$$

Think of this as a mapping from one set of log applications across areas to a new set - denote this mapping by $T(\ln A)$. To apply Blackwell's sufficient conditions consider $T(\ln A + z)$. Simple algebraic manipulation shows that we have that:

$$T(\ln A + z) = T(\ln A) + \frac{\beta(1-\gamma)}{1+\beta} z < T(\ln A) + z. \tag{12}$$

It is worth discussing why we can only identify $\beta = k\widetilde{\beta}$ and $\delta \equiv k\widetilde{\delta}$, and not all the underlying structural parameters. The reason is that an increase in the size of the idiosyncratic component of the utility from a job – measured by k – is observationally equivalent to a change in the cost of distance or the effect of congestion on the probability of obtaining a job. If the idiosyncratic component is very small, then, for the same number of applicants, a worker is very likely to apply to a closer job rather than one that is more distant. In other words, this scenario has observable consequences equivalent to one with a higher cost of distance but a more important idiosyncratic component. Similarly for the effect of the number of applicants, holding distance to a job constant.

A useful result that can be obtained from (10) 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 to scale, while $\gamma/(1+\beta\gamma) > 0$ would imply increasing returns. This can be seen more clearly in the special case in which areas are isolated, such 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):

$$\ln A_b = -\frac{k\gamma \ln c}{1+\beta} + \frac{1}{1+\beta\gamma} \ln\left(\frac{U_b}{V_b}\right) + \frac{\gamma}{1+\beta\gamma} \left[\ln\left(V_b\right) + \ln(f)\right]. \tag{13}$$

As the vacancy outflow rate in an area depends on the number of applications per job in that area, expression (13) 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, a measure of the cost of distance;
- β, a measure of the congestion parameter, that measures how much workers are deterred from applying to jobs in areas where they expect a large number of applications;
- γ, the returns to scale of the matching function.

Finally, one might also notice that (10) 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 in our data, 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.

4.2 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 expected number of applications per vacancy, i.e.

$$E\left(\frac{M_b}{V_b}\right) = \Psi(A_b), \qquad \Psi'(\cdot) > 0. \tag{14}$$

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 (14) is to think of an urn-ball problem,[10] in which firms play the role of urns and applications the role of balls. Because of a coordination failure, a random placing of

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

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. The probability that a given job applicant is selected for a job is $A_b^{-\beta}$. 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

$$\frac{M_b}{V_b} = 1 - \exp\left[-\exp(\alpha)A_b^{1-\beta}\right] + e_b, \tag{15}$$

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.

$$\frac{M_b}{V_b} = \exp(\alpha)A_b^{1-\beta} + e_b. \tag{16}$$

The nice feature of the urn-ball specification is that it ensures a vacancy outflow rate between 0 and 1, while this is of course not imposed by the log linear specification. However, the log linear specification has the advantage that it yields a constant elasticity of the vacancy outflow with respect to the number of jobseekers 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 (15) or (16), A_b is implicitly defined by (10), and thus δ, β, γ are further parameters to be estimated. In practice we estimate (15) and (16) by maximum likelihood, and at every iteration of the maximization solve the contraction mapping in (10).[11]

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 may 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 (e.g. Pissarides, 2000). 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

[11] Again, to avoid dropping observations with zero outflows, both (15) and (16) are estimated in levels instead of logs.

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 can think of a 'product' as being a job in a particular area. Compared to most applications in Industrial Organization we have a very large number of 'products' but we also have a priori information on which of these products are the closest substitutes - those closer in space - which allows us to reduce the dimensionality of product heterogeneity. Consumers are also differentiated - in our application, this is by space, the same differentiation as 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 more 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 are forced to abstract from them.

5 Results

5.1 Main estimates

Our first set of results is based on an urn-ball specification of the matching function, as shown in equation (15). For reasons of computing capacity, we cannot estimate our regression equation on the whole sample period, and we thus estimate it separately for each month from May 2004-April 2006. This, however, does have the advantage that we can think of each month's estimate as a draw from the data (not necessarily independent, and we can look for serial correlation in the estimates), so giving us an idea of the standard error of our estimates from the different months, which we can then compare with that produced by our structural estimation method. We thus estimate our main specification separately for each month, and then run a number of specifications and robustness checks on a representative month.

In Table 4 we report time averages of the parameters of interest, together with their standard

deviations, minimum and maximum values. 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 jobseeker's location. 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, jobseekers 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 (15)), 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. Overall, both δ and β 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 δ and β are highly statistically significant, γ is not statistically different from zero.

In Figure 1 we plot point estimates of these parameters over the 24 months in our sample. The series fluctuate somewhat over the sample period, but show no definite trend, and we could detect no significant serial correlation of either first or second order in δ, β or γ.

In Table 5 we report estimates of alternative specifications of the job application model for February 2005, and Table A1 in the Appendix reports the corresponding estimates for the whole sample period, obtained again as averages of monthly estimates for May 2004-April 2006. The simple criteria used for picking a reference month in Table 5 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 to make the estimates of Table 5 well representative for the whole sample period. Here we will not comment the average estimates of Table A1 separately, because indeed they are very close to those reported in Table 5, both in terms of parameter estimates and their standard errors.

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 jobseeker's location, corresponding to the specification of Table 3. The associated standard errors are corrected for some (arbitrary) structure of spatially correlated shocks.[12] Both β and δ are highly statistically significant, while γ is not

[12] In particular, we assume that spatial correlation across wards decays at rate δ with ward distance or commuting cost. Our estimated variance-covariance matrix of the parameters is given by $\widehat{V} = \widehat{\sigma}^2(\widehat{X}'\widehat{X})^{-1}(\widehat{X}'\widehat{\Omega}\widehat{X})(\widehat{X}'\widehat{X})^{-1}$,

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 statistic 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:

$$\frac{M_b}{V_b} = 1 - \exp\left[-\exp(\alpha)A_b^{1-\beta}\left(\frac{U_b}{V_b}\right)^{\alpha_1}\right] + e_b, \tag{17}$$

where A_b is obtained from the contraction mapping (10) 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 δ, β and γ 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 $1-\beta = 0.207$, 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 3, 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. A simple way to address the division bias problem in this context (analogous to a 'control function' approach) 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 insignificantly different from zero, when one controls for the total vacancy stock.

In column 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 (16). While estimates are very similar to those obtained on an urn-ball matching function, the log-linear specification has the advantage of delivering a constant elasticity of the matching rate with respect to applications,

where $\widehat{\sigma}^2$ is the sum of squared residuals divided by the number of observations, $\widehat{X}$ is the matrix of partial derivatives of the regression function with respect to right-hand side variables, and the spatial correlation matrix $\widehat{\Omega}$ is proxied by $\exp(-\widehat{\delta}D)$, where D is given by the distance matrix, and $\widehat{\delta}$ is the associated parameter estimate.

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.

5.2 Robustness tests

In our main estimates we have modelled distance between wards using physical distance. In this section we explore different alternative ways of modelling distance. In column 5 of Table 5 we estimate an urn-ball matching function, having modelled the utility of jobs at different locations as a function of commuting times, expressed in one-way commuting minutes.[13] The results are fairly similar to those based on geographic distance, with the job congestion estimate at 0.75, 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 factor of 2.7 should one increase by 5 minutes the one-way commuting time to a job. In column 6, distance is measured by one-way commuting costs, and the corresponding δ estimate implies that the attractiveness of a job is halved for each extra £1 added to the one-way commuting cost (at 2001 prices).

All three measures of distance - physical distance, commuting time and commuting costs - yield a very high decay of the probability of applying to a given job with distance. While these estimates are obtained on a relatively unskilled sample, for which the labor market may be more local than for the universe of jobseekers, the estimates of Bonhomme and Jolivet (2009) are suggestive of very high distance costs on a sample that is representative of the overall population. Using information on job satisfaction from the ECHPS, they find that workers in Europe are typically willing to forgo large fractions of their salaries to become satisfied with their commuting distances/costs, ranging from 40% in France to 14% in Austria. Unfortunately, they do not report estimates for the UK.

We finally 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.[14] In particular, we extract data on claimants and job vacancies by CAS ward and 1-digit occupation, and construct the

[13] The data on commuting costs were obtained from Daniel Graham at Imperial College and have their origins in transport planning.

[14] For the unemployed the occupation refers to the type of job sought.

following index of occupation dissimilarity between origin area a and destination area b :

$$m_{ab} = \sum_{k=1}^{8} \left| \frac{U_{ka}}{U_a} - \frac{V_{kb}}{V_b} \right|, \tag{18}$$

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 (18):

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

where μ is an extra parameter to be identified. The results are reported in column 7, where the estimate for μ 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. 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 conclude, we compare the relative merits of the job application model with the conventional matching function in vacancies and unemployment in the specification of column 8, which 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 substantially 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.

5.3 Worker interactions and commuting flows

One idea that has received a fair amount of attention in recent years is that networks are important in labor markets and that a good source (though not the only source) of contacts is one's neighbors. In this case one would expect to see clusters of flows from one area to another, which are the outcome of these networks. The empirical studies by Topa (2001), Bayer, Ross and Topa (2008) and Hellerstein, McInerney and Neumark (2011) all contain evidence on the importance of residence-based networks. Topa (2001) presents evidence of positive correlation between unemployment in a neighborhood and others that are physically close to it. While social interactions are one possible

explanation for this, so is our distance based model of the labor market and so might be residential segregation. Bayer et al. (2008) show that workers who live on the same block are significantly more likely to work in the same block. However the probability that pairs who reside in the same block also work in the same block is 0.48%, i.e. 1 in 200 workers, perhaps not as high as one might expect if networks are very important. Hellerstein et al. (2011) take this further as they have matched employer-employee data that enables them to show that workers who work in the same firm are significantly more likely to live in the same census tract than those who work in the same census tract but in different firms. However, the baseline probabilities are, again, quite small.

The way in which we investigate the social network hypothesis is the following. From the Special Workforce Statistics of the 2001 Census we have data on commuting flows between every ward (albeit with some noise deliberately introduced to preserve anonymity in cells with small numbers). We think of these commuting flows as linkages between wards that existed prior to the period of our data. So, if networks are important we might expect to see that, controlling for distance, unemployed workers are more likely to apply for jobs in wards towards which there is a large commuting flow from their own location. So we include the commuting flow in addition to our distance measure, i.e.

$$f_{ab} = \exp(-\delta * d_{ab} + \zeta * comm_{ab}),$$

where $comm_{ab}$ denotes the number of individuals resident in a who commute to b. The commuting flow $comm_{ab}$ is, of course, endogenous, but the likely bias is in an upward direction. For example, if two wards are linked by a superfast bus service then this will lead to large commuting flows, given the (mis)measured distance.

The results for February 2005 are reported in Table 6, where each column uses a different proxy for the cost of distance. In all specifications, commuting flows have a negative rather than positive impact on the attractiveness of jobs at various locations. When distance is measured in kilometers, the coefficient on commuting flows is high, negative and significant, and when distance is measured in either commuting time of commuting cost, the coefficient on commuters falls in both absolute size and significance, but it remains firmly negative. The corresponding average estimates over the sample period are reported in Table A2 in the Appendix, and the results are very close to those reported in Table 6, although it has to be noted that the coefficients on the commuters variable display quite a bit more variation over the sample period than other coefficients in the regression. Overall, we find little evidence here for residence-based networks being quantitatively important, which is fairly consistent with the evidence presented in the other studies discussed above.

5.4 Predicted 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.

$$\frac{U_a N_{ab}}{A_b V_b}. \tag{19}$$

As firms are assumed to select jobseekers randomly within the pool of job applicants, the ratio in (19) also denotes the proportion of total matches in ward b that involve jobseekers from ward a. Thus the number of vacancies in ward b that are filled by jobseekers in ward a is given by

$$\frac{U_a N_{ab}}{A_b V_b} M_b. \tag{20}$$

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 (20), this is equal to

$$\frac{N_{ab} M_b / A_b V_b}{\sum_{b'} N_{ab'} M_{b'} / A_{b'} V_{b'}}. \tag{21}$$

We can compare these predictions with the Census data on commuting used in the previous section. 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 jobseekers filling Jobcentre vacancies have different commuting patterns from jobseekers who find their jobs via other methods.

However, we do have some indirect evidence that this potential concern is not a major one. The UK LFS contains data on commuting times for those in new jobs and those in continuing jobs and, for those in new jobs, on how that job was obtained. Table 7 presents evidence 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 overall employed population. As the characteristics of workers in different categories may differ, and they may be related to commuting times, we also compared differences in commuting times controlling for the method used to find the current job, 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 – as implied by (21) – is 0.71. This is the pairwise correlation between two matrices of commuting

flows (actual and predicted), which include several zeros, thus one may worry that a relatively high correlation between the two could be driven by the vast proportion of cells in either matrix with zero commuters. But when we restrict to cells with nonzero commuters we still obtain a correlation of 0.69.

Interestingly, the correlation between actual and predicted commuting rises to 0.83 if one excludes 'locals' from the sample, i.e. individuals who live and work in the same ward (and this stays unchanged if we further exclude cells with zero commuters). 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. This can be seen more clearly comparing the distributions of actual and predicted commuting flows, as shows in Table 8. Both distributions are hump-shaped, with a peak in the (0,5] km range, but the model tends to underpredict locals and as a consequence to overpredict short-distance 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.

6 Evaluating Location-Based Policies

6.1 Local Labor Demand Stimulus

A key policy question for addressing spatial inequalities 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 an 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 of 6%, which was nearly three times higher than the average ratio for England and Wales. We pick this example because it combines very large increases in numbers of vacancies as a result of Olympic-related projects with a relatively depressed local labor market.

Specifically, we simulate the impact of a doubling in the number of vacancies in Stratford and New Town Ward in a given month, from 464 to 928, under the assumption of constant returns to scale, i.e. imposing $\gamma = 0$, an assumption that 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 (7)). Values used for δ and β 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 Table 9, showing the predicted percentage change in applications per job (obtained from equation (10)), in the vacancy outflow (obtained from equation (15)), and the unemployment outflow (obtained from equation (20)), 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,[15] and beyond this cutoff distance the change in the unemployment outflow becomes negligible. A similar picture can be grasped graphically from the map in Figure 4, in which wards around Stratford are shaded according to the average percentage change in the unemployment outflow.

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 job competition in their local wards and attracts applications from elsewhere, and so on. This mechanism explains the spatial

[15] This 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, according to (8), and thus the unemployment outflow in each area a, according to (20).

propagation of local shocks in the presence of relatively high costs of distance.

How does this prediction about the impact of a local labor demand shock compare with what has actually happened in Stratford in the run-up to the 2012 Olympic Games? Much of the increase in labor demand is yet to take place (e.g. the jobs associated with running the Olympics itself), while some has built up steadily over time (e.g. in construction). But there is one instance of a sharp increase in labor demand that has taken place: on 13 September 2011, a new shopping centre, Westfield Stratford City, opened next to the London Olympic Park. This is one of the largest urban shopping centres in Europe, and has been expected to contribute significantly to local development, "with the creation of up to 10,000 permanent jobs, including 2,000 jobs going to people in the local area."[16] We can then look in the current unemployment and vacancy data (again extracted from NOMIS) for any evidence of the corresponding labor demand shock on vacancy creation and the unemployment outflow around Stratford.

Figure 5 presents a time series for new vacancies advertised in Stratford, in wards within 3km of Stratford, and in all of London, all normalized to their January 2009 values.[17] One can clearly notice in Panel A the sharp rise in vacancies in Stratford in the summer of 2011 that is associated with the opening of the shopping centre - vacancy inflows are running at over three times the usual level. The other areas show no such trend. What about the unemployment outflow? - Panel B shows this. Even though the vertical axis in Panel B is on a different scale to that of Panel A, one can see little or no evidence of an increase in outflows in Stratford or surrounding areas as a result of the spike in vacancies. Not enough time has elapsed since the opening of the shopping centre to do a proper statistical analysis of its effects on the surrounding labor market, but the early indications are exactly in line with the predictions of our model and are certainly consistent with negligible local effects of targeted labor demand stimulus, as shown in Table 9.

6.2 Reduction in Transportation Costs

We next assess the importance of transportation costs by simulating the impact of a sizeable reduction in the cost of distance between Stratford and central London - and in order to focus on distance costs alone we leave labor demand in Stratford unchanged. The idea is to evaluate the impact of an improved transport link between a high-unemployment area and the city centre, with relatively higher supply of jobs. Specifically, we simulate the effect of a faster connection between Stratford and Kings Cross Station in central London, by halving the distance between the two corresponding wards. Stratford and Kings Cross are located 8.4 km apart, and we build

[16] http://en.wikipedia.org/wiki/Westfield_Stratford_City

[17] While information on the stock of vacancies in the NOMIS is not comparable before and after May 2006 (see footnote 3), the procedure for registering the inflow of newly advertised vacancies (reported in Figure 5) remains the same.

a new distance matrix in which the distance between the two wards is set at 4.2 km, and we allow the distance between any two other wards a and b to be affected if either the distance $a-$Kings Cross$-$Stratford$-b$ or $a-$Kings Cross$-$Stratford$-b$ is shorter than the original distance ab. Essentially this is equivalent to introducing a fast, non-stop service between Stratford and Kings Cross, allowing individuals residing and searching near either location to re-optimize their travel schedule accordingly. As a consequence, we would expect some jobseekers in Stratford (the high-unemployment area) to choose to search for jobs in central London (the high-vacancy area), thus raising the unemployment outflow in Stratford and at the same time raise applications per job in central London, then lowering the unemployment outflow around Kings Cross.

We show in Table 10 the impact of this improved transport link on applications per job, the vacancy outflow and the unemployment outflow, at various distances from King's Cross and Stratford. As expected, applications per job rise both in King's Cross and its close vicinity (see column 1, rows 1 and 2), as central London is now attracting more jobseekers from Stratford and surrounding areas. As a consequence, the vacancy outflow increases (column 2) and the local jobseekers are less likely to find jobs (column 3) as they face stronger job competition from new applicants attracted by the faster transport link. Moving down to row 3, we find that applications per job also rise in Stratford. Despite the fact that some jobseekers are quitting Stratford to search for jobs in central London, the more efficient transport link now attracts some jobseekers from other surrounding areas. But the unemployment outflow still increases as the opportunity to find jobs away from Stratford dominates the effect of increased job competition from elsewhere. Row 4 shows that within 3 km from Stratford (and excluding Stratford itself), applications per job and the vacancy outflow fall, and the unemployment outflow increases, all driven by the possibility of finding new jobs elsewhere. It should be noted however that, quantitatively, the impact on the unemployment outflow is always very small, whether positive or negative, and becomes negligible beyond 3 km from either transport node (rows 5 and 6).

Spillovers on the unemployment outflow around King's Cross and Stratford are illustrated in more detail on a map in Figure 6, where darker and lighter shades correspond to an increase and a decrease, respectively, in the unemployment outflow. The map once again shows that making it easier for workers in a suburban, high-unemployment area to reach the urban centre raises the unemployment outflow in the suburbs and depresses the unemployment outflow in the centre, with declining intensity as one moves away from either transport node. Overall, as a consequence of ripple effects, the impact on the unemployment outflow is very modest (within a range -0.66% to +0.55%), but it propagates quite widely around the target.

7 Conclusions

In this paper we have used high-frequency, geographically very disaggregated data on unemployment and vacancies to investigate the extent to which labor markets are 'local'. These data allow us to approximate the continuous nature of geographic space, and build overlapping local labor markets based on optimizing job search behavior.

We have first presented some non-structural estimates in which the probability of filling a vacancy 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 convey evidence on the true cost of distance.

We then proposed a model of job search across space that allows, in a tractable way, estimation of 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 0.3 with distance (in km) to the job, and similar qualitative conclusions are obtained when we measure distance using commuting time or commuting costs. Also, workers are significantly discouraged from applying to jobs for which 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 pool. Commuting flows predicted by the estimated model replicate fairly accurately actual commuting patterns across Census wards, although our model tends to underpredict the proportion of individuals who live and work in the same ward.

We finally used the estimated model to simulate the impact of local development policies like local stimulus to labor demand or improved transportation links. Despite the fact that labor markets are relatively 'local', location-based policies turn out to be rather ineffective in raising the local unemployment outflow, because labor markets overlaps and the associated ripple effects in applications largely dilute the effect of local shocks across space. We argued that early indications about the impact of the opening of a new shopping centre close to the site of the 2012 London Olympics are consistent with this result.

References

[1] Blanchard, Olivier and Larry Katz (1992), "Regional Evolutions", *Brookings Papers on Economic Activity* 1: 1-75.

[2] Bound, John and Harry Holzer (2000), "Demand Shifts, Population Adjustments and Labor Market Outcomes during the 1980s", *Journal of Labor Economics* 18: 20-54.

[3] Boustan, Leah P. and Robert Margo (2009), "Race, Segregation, and Postal Employment: New Evidence on Spatial Mismatch", *Journal of Urban Economics* 65: 1-10.

[4] Burda, Michael and Stefan Profit (1996), "Matching across Space: Evidence on Mobility in the Czech Republic", *Labour Economics* 3: 255-378.

[5] Burgess, Simon and Stefan Profit (2001), "Externalities in the Matching of Workers and Firms in Britain", *Labour Economics* 8: 313-333.

[6] Busso, Matias, Jesse Gregory and Patrick Kline (2011), "Assessing the Incidence and Efficiency of a Prominent Place Based Policy", UC Berkeley, mimeo.

[7] Butters, Gerard R. (1977). "Equilibrium Distributions of Sales and Advertising Prices," *Review of Economic Studies,* 44:3, pp. 465-91.

[8] Coles, Melvyn G. and Barbara Petrongolo (2008). "A Test Between Stock-flow Matching and the Random Matching Function Approach," *International Economic Review* 49, pp. 1113-1141.

[9] Coles, Melvyn G. and Eric Smith. (1996). "Cross-Section Estimation of the Matching Function: Evidence from England and Wales," *Economica*, 63:252, pp. 589-98.

[10] Diamond, Peter (1982). "Aggregate Demand Management in Search Equilibrium". *Journal of Political Economy* 90: 881-894.

[11] Flinn, Christopher and James Heckman. (1983). "Are Unemployment and Out of the Labor Force Behaviorally Distinct Labor Force States?," *Journal of Labor Economics* 1(1), pp. 28-42.

[12] Glaeser, Edward L. and Joshua D. Gottlieb (2008). "The Economics of Place-Making Policies", *Brookings Papers on Economic Activity,* Spring 2008, 155-253.

[13] Hellerstein, Judith K., David Neumark and Melissa McInerney (2008). "Spatial Mismatch or Racial Mismatch?", *Journal of Urban Economics* 64: 464-479.

[14] Kain, John F. (1968) "Housing Segregation, Negro Employment, and Metropolitan Decentralization", *Quarterly Journal of Economics* 82, 175-197.

[15] Katz, Larry, Jeffrey Kling and Jeffrey Liebman (2001), "Moving to Opportunity in Boston: Early Results of a Randomized Mobility Experiment", *Quarterly Journal of Economics* 116: 607-654.

[16] Machin, Andrew (2001). "The New Vacancy Survey: A New Series of National Statistics", National Statistics Features, ONS. Available at http://www.statistics.gov.uk/articles/labor_market_trends/VacancySurvey_LMT0703.pdf.

[17] Manning, Alan (2003). *Monopsony in Motion. Imperfect Competition in Labor Markets.* Princeton, Princeton University Press.

[18] Moretti, Enrico (2011) "Local Labor Markets", in O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*, vol. 4B Amsterdam: Elsevier-NorthHolland, pp. 1237-1314.

[19] Notowidigdo, Matthew J. (2011) "The Incidence of Local Labor Demand Shocks", Chicago Booth, mimeo.

[20] OECD (2000). *The Public Employment Service in the United States.*

[21] Patacchini, Eleonora and Yves Zenou (2007). "Spatial Dependence in Local Unemployment Rates", *Journal of Economic Geography* 7: 169-191.

[22] Petrongolo, Barbara and Christopher A. Pissarides. "Looking into the Black Box: A Survey of the Matching Function,". *Journal of Economic Literature* 39, 390-431, 2001.

[23] Pissarides, Christopher A. (1979). "Job Matchings with State Employment Agencies and Random Search," *Economic Journal*, 89:356, pp. 818-33.

[24] Pissarides, Christopher A. (2000). *Equilibrium Unemployment Theory.* Cambridge (MA): The MIT Press.

[25] Jones, Stephen R. G., and Craig W. Riddel. (1999). "The Measurement of Unemployment: An Empirical Approach," *Econometrica* 67, pp. 147-161.

Figure 1
Unemployment to Vacancy ratios in England and Wales
Shades correspond to quartiles.

Figure 2
Unemployment/Vacancy ratios in Greater London
Shades correspond to quartiles

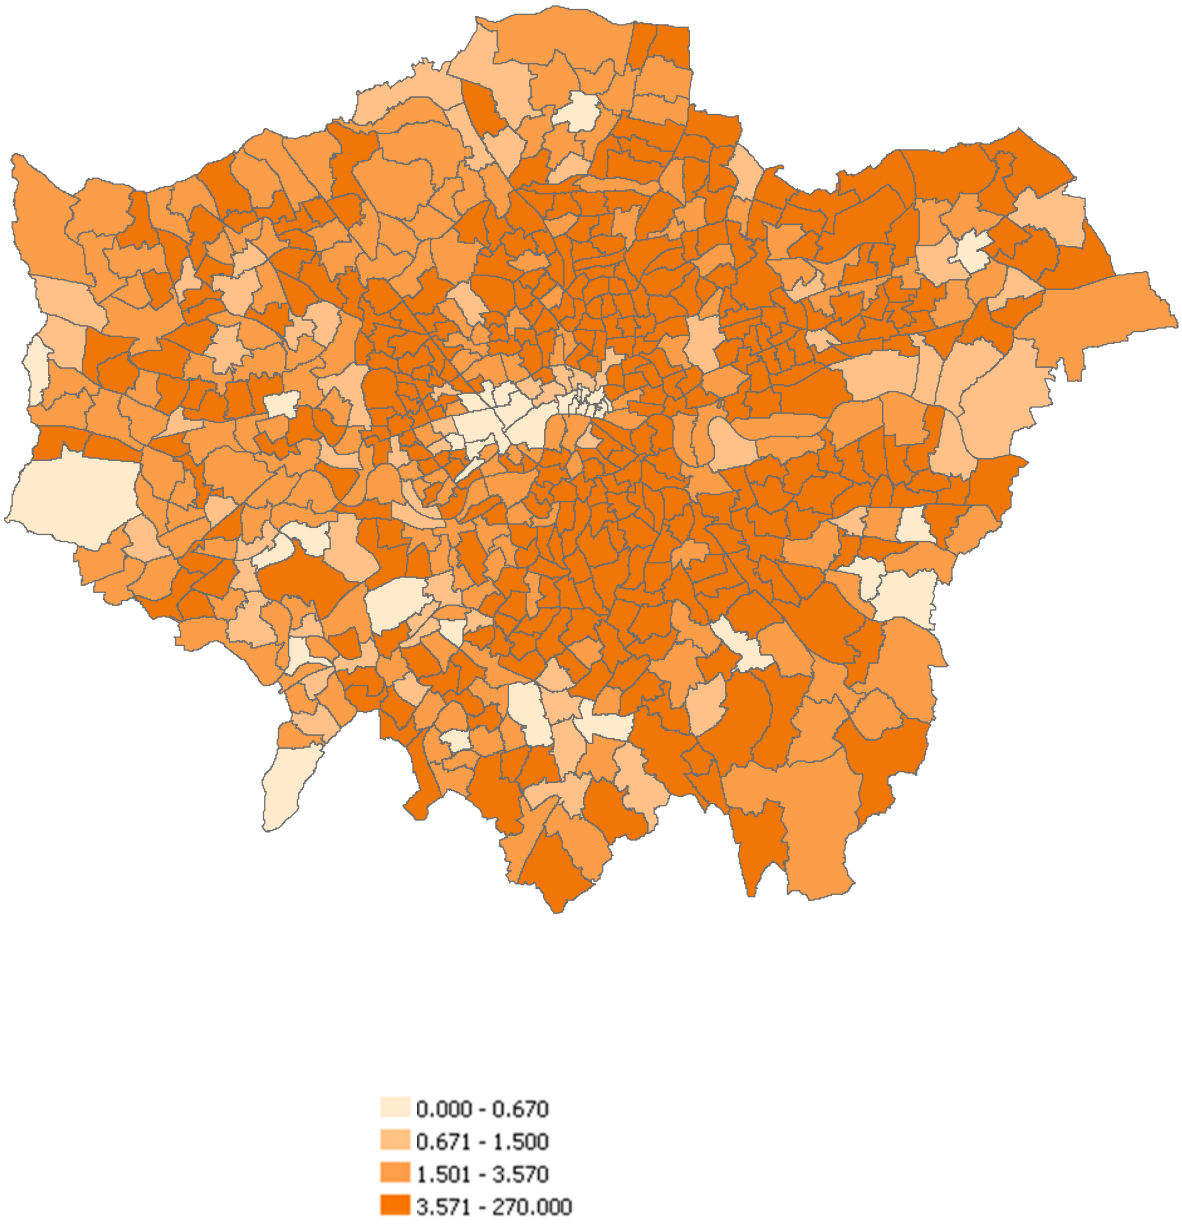

Figure 3
Main parameter estimates over the sample period May 2004-April 2006

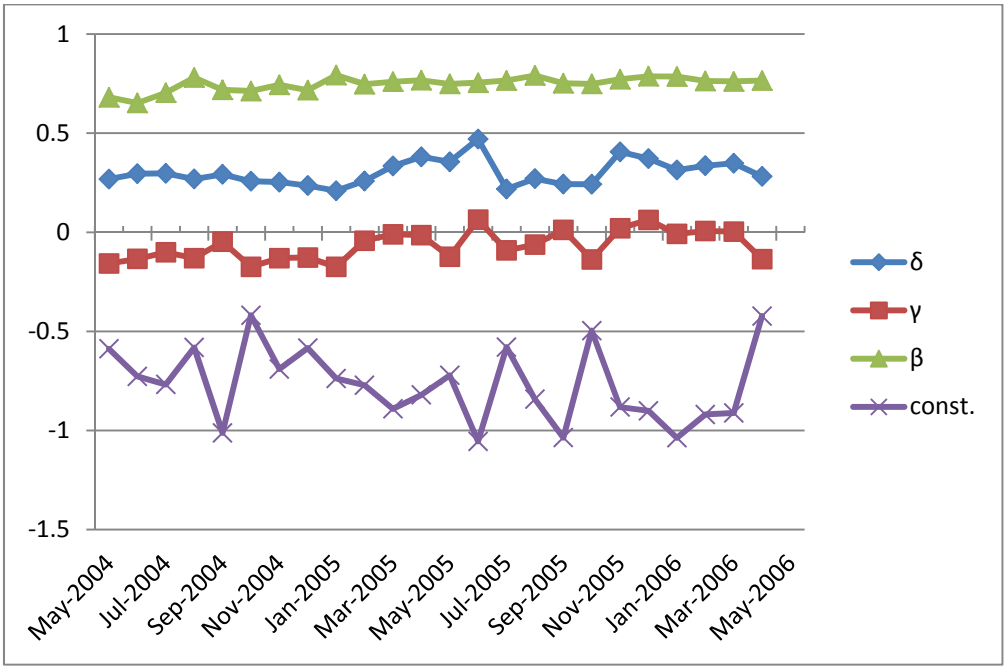

Figure 4
Effect of a doubling in the number of vacancies in Stratford on the unemployment outflow, percentage change

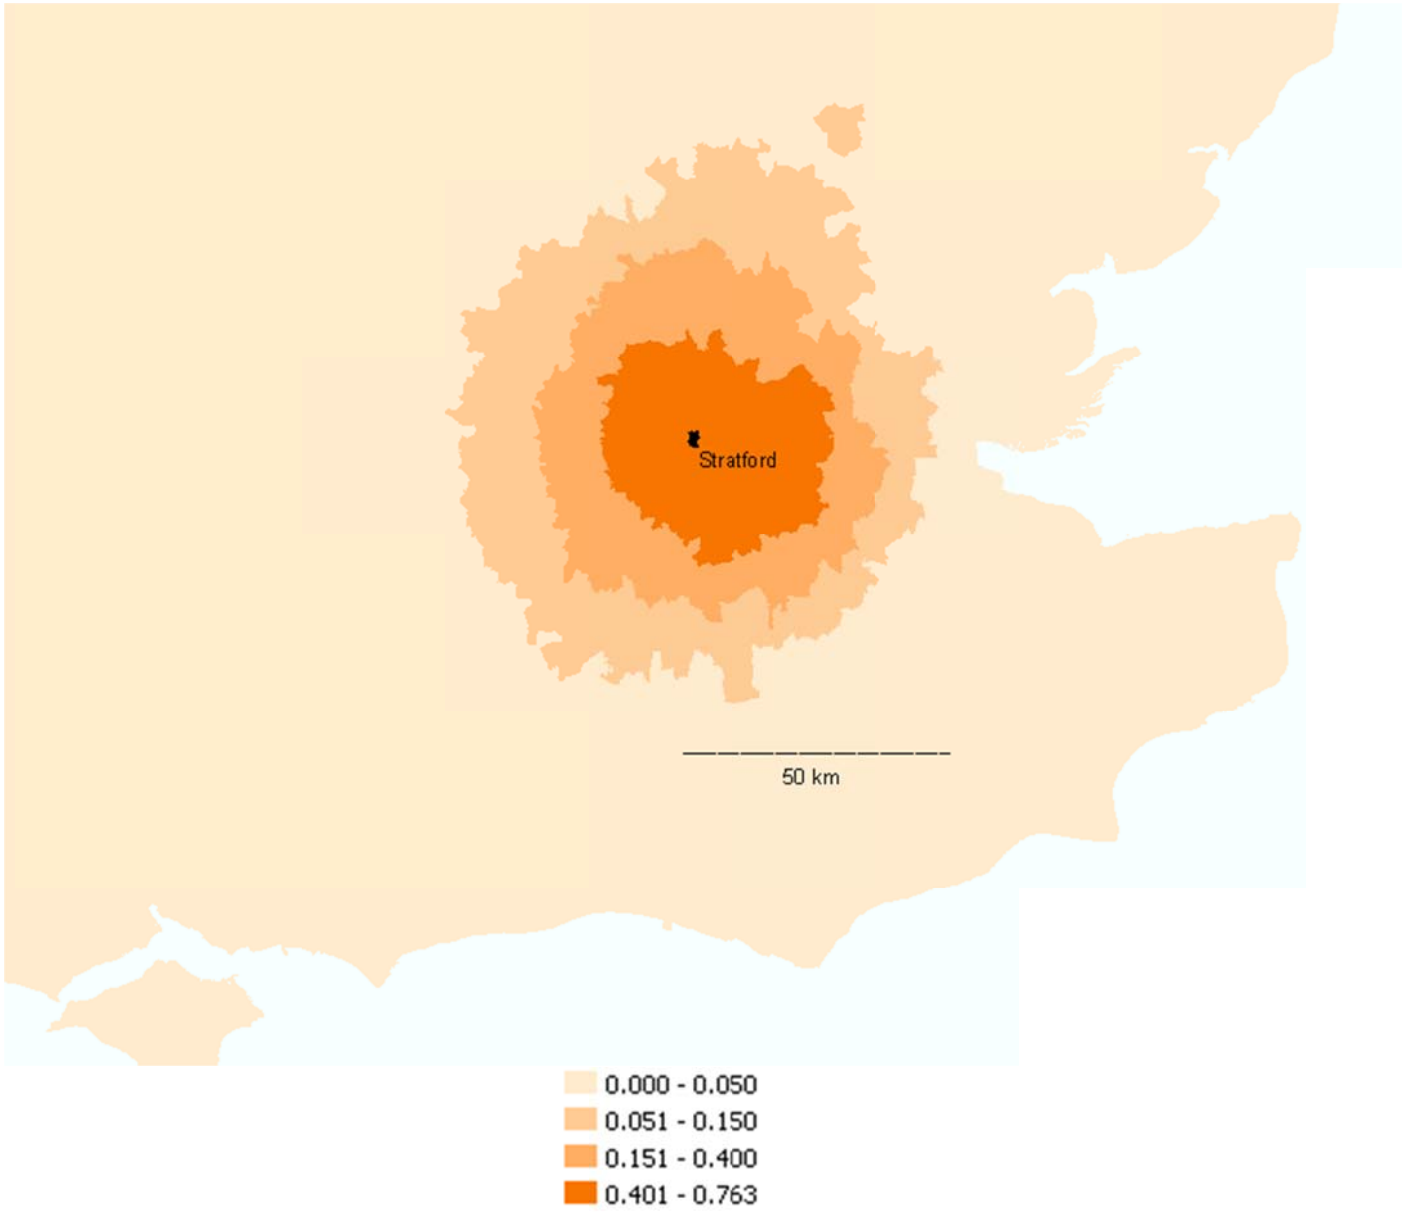

Figure 5
Recent changes in the vacancy inflow (Panel A) and the unemployment outflow (Panel B) in and around Stratford

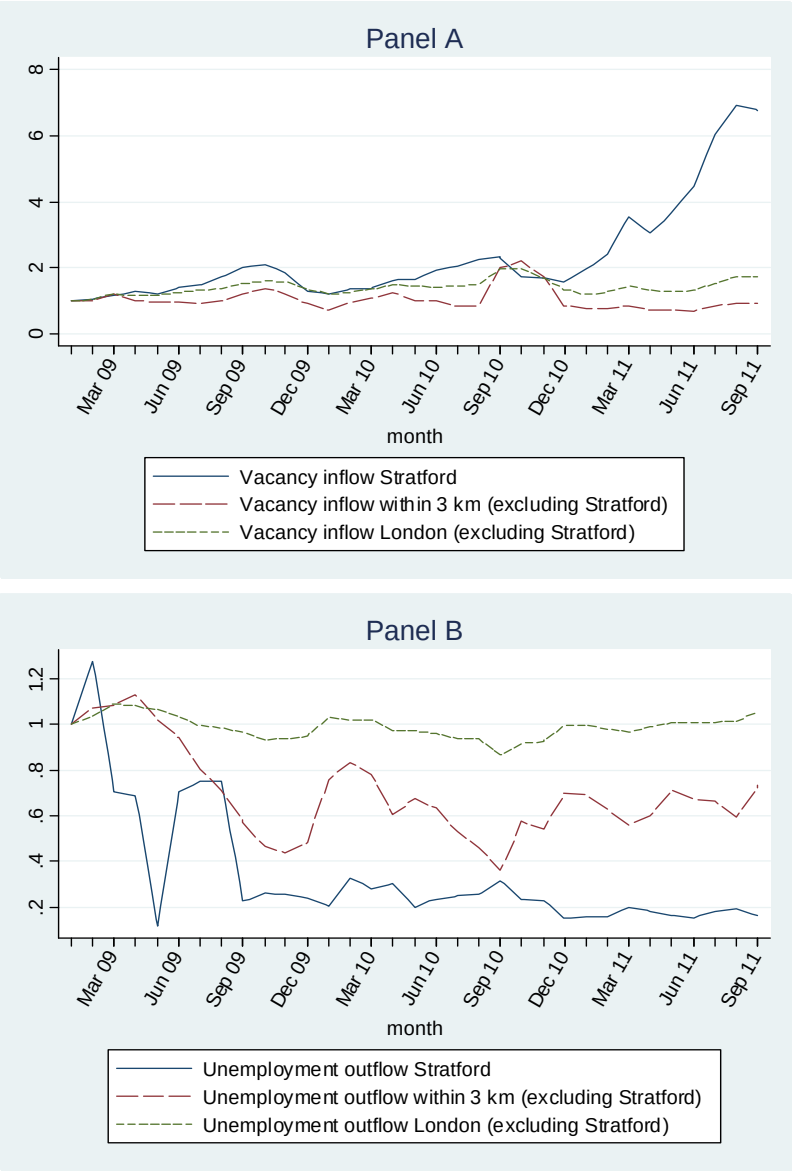

Notes: All series are smoothed using moving averages with a 3-month window and equal monthly weights, and normalized to their January 2009 values.

Figure 6
Effect of halving the cost of distance between King's Cross and Stratford on the unemployment outflow, percentage change.

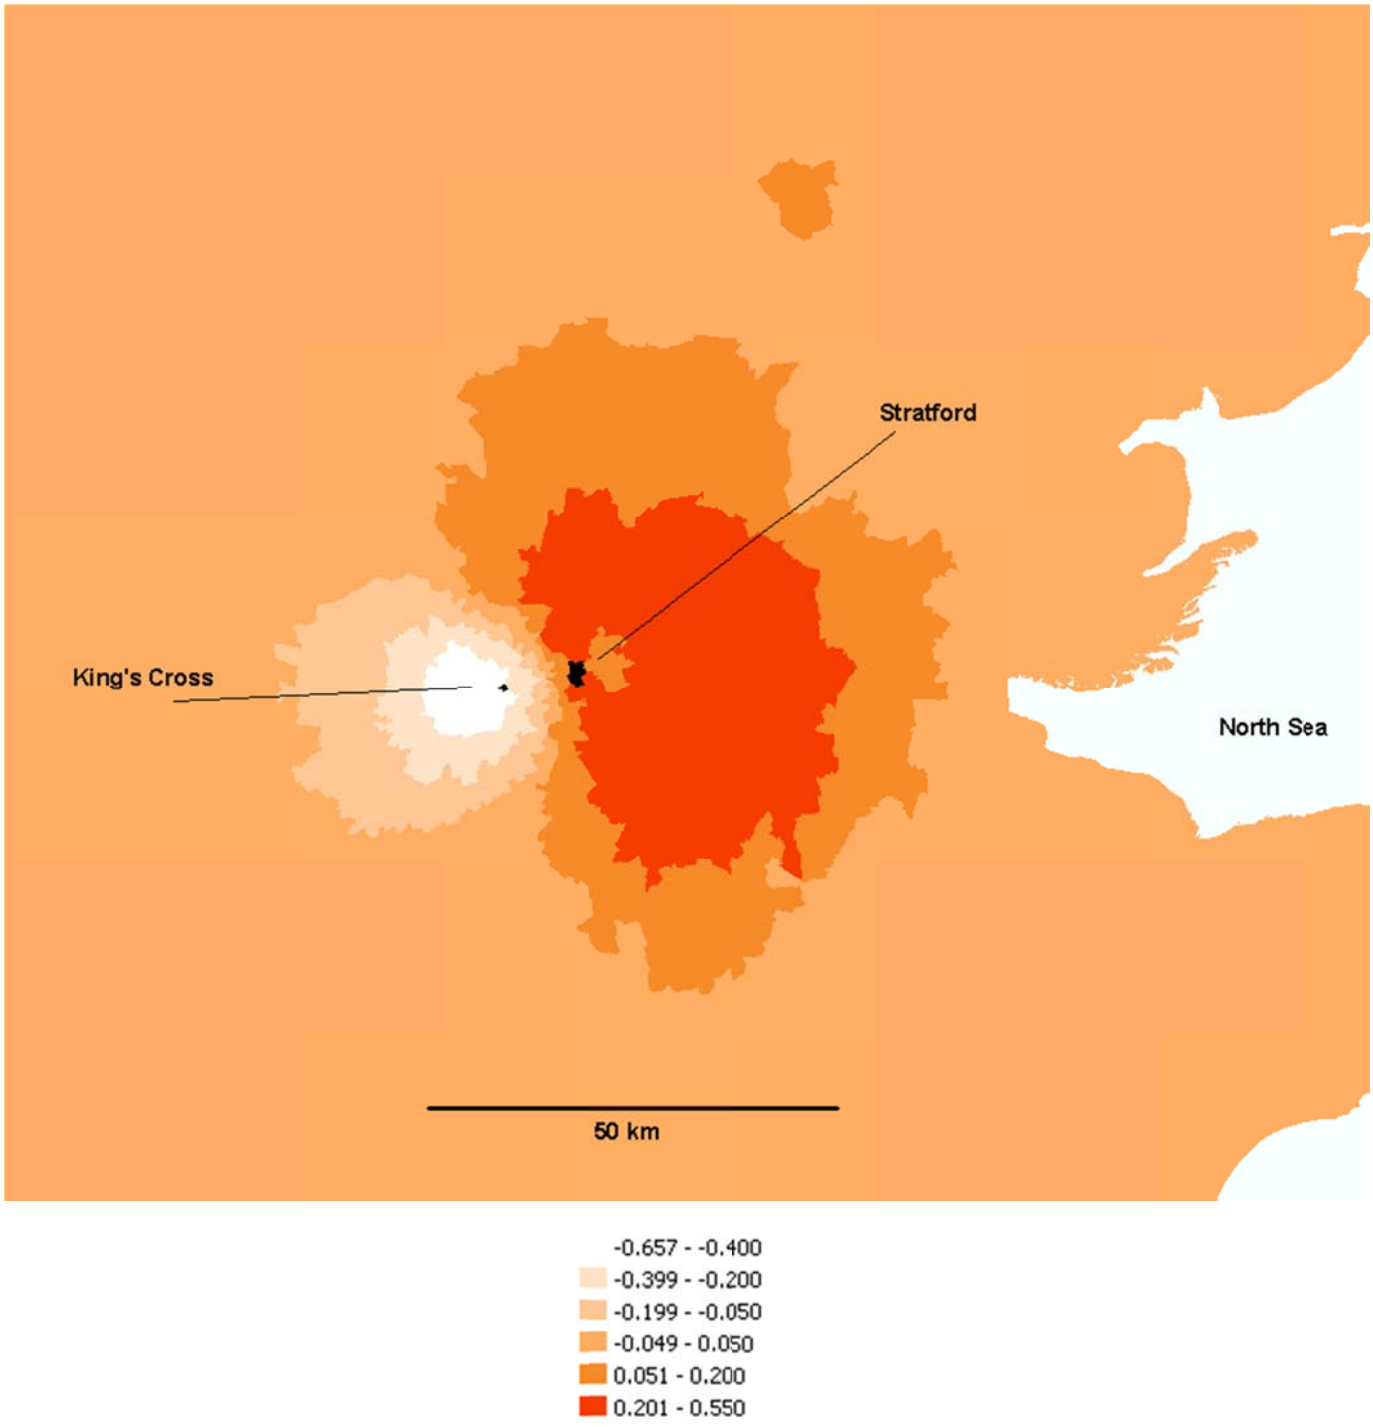

Table 1
Descriptive statistics on local labor markets

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)
$\log \widetilde{U}_{10b}$	0.201***	0.193***	0.200***	0.204***	0.0575***	0.0372
	(0.00441)	(0.00478)	(0.00272)	(0.00266)	(0.0190)	(0.0297)
$\log \tilde{V}_{10b}$	-0.224***	-0.214***	-0.227***	-0.230***	-0.194***	-0.187***
	(0.00528)	(0.00573)	(0.00336)	(0.00326)	(0.00924)	(0.0175)
$U_b / \widetilde{U}_{10b}$	0.713***	0.710***	0.623***	0.657***	-0.196**	-0.336**
	(0.0564)	(0.0560)	(0.0308)	(0.0326)	(0.0874)	(0.156)
$U_{5b} / \widetilde{U}_{10b}$	0.278***	0.273***	0.254***	0.258***	-0.0680	-0.0341
	(0.0268)	(0.0268)	(0.0146)	(0.0147)	(0.0909)	(0.143)
$U_{20b} / \widetilde{U}_{10b}$	-0.00135*	-0.00158**	0.000589	0.000747	-0.000775	-0.000131
	(0.000771)	(0.000756)	(0.000440)	(0.000460)	(0.000885)	(0.00275)
$U_{35b} / \widetilde{U}_{10b}$	0.000265**	0.000237**	0.000181**	0.000190**	0.000197	-8.53e-05
	(0.000108)	(0.000108)	(7.43e-05)	(8.05e-05)	(0.000136)	(0.000365)
$V_b / \tilde{V}_{10b}$	-0.856***	-0.846***	-0.762***	-0.771***	-1.351***	-0.896***
	(0.0444)	(0.0443)	(0.0222)	(0.0226)	(0.0438)	(0.0694)
$V_{5b} / \tilde{V}_{10b}$	-0.110***	-0.106***	-0.0975***	-0.101***	-0.0400	0.00774
	(0.0241)	(0.0241)	(0.0133)	(0.0133)	(0.0335)	(0.0529)
$V_{20b} / \tilde{V}_{10b}$	-0.000688	-0.000520	-0.00405***	-0.00421***	-0.000244	-0.00444***
	(0.00105)	(0.000999)	(0.000470)	(0.000481)	(0.000305)	(0.00146)
$V_{35b} / \tilde{V}_{10b}$	-2.42e-05	-4.40e-06	-6.00e-05	-8.20e-05	-9.47e-06	4.14e-05
	(0.000121)	(0.000116)	(0.000109)	(0.000112)	(7.71e-05)	(0.000363)
Observations	197579	197579	188648	188591	197579	188591
Time Effects	No	Yes	Yes	Yes	Yes	Yes
Ward Effects	No	No	No	No	Yes	Yes
Instruments	No	No	Vacancy variables	Vacancy and unemploym. variables	No	Vacancy and unemploym. variables

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

Table 3
Matching functions in log and level

	(1)	(2)	(3)	(4)	(5)
$\log \widetilde{U}_{10b}$	0.197***	0.202***	0.201***	0.196***	0.169***
	(-0.00484)	(0.0079)	(0.0080)	(-0.0114)	(0.0115)
$\log \tilde{V}_{10b}$	-0.206***	-0.213***	-0.202***	-0.199***	-0.155***
	(-0.00535)	(0.0076)	(0.0078)	(-0.0129)	(0.0132)
$U_b / \widetilde{U}_{10b}$	0.781***	0.307***	0.397***	0.642***	0.311***
	(-0.0575)	(0.0934)	(0.0943)	(-0.1450)	(0.1405)
$U_{5b} / \widetilde{U}_{10b}$	0.252***	0.092**	0.086*	0.192***	0.121***
	(-0.0269)	(0.0448)	(0.0456)	(-0.0648)	(0.0640)
$V_b / \tilde{V}_{10b}$	-0.962***	-0.420***	-0.349***	-0.853***	-0.521***
	(-0.0484)	(0.1043)	(0.1056)	(-0.1140)	(0.1155)
$V_{5b} / \tilde{V}_{10b}$	-0.0846***	0.018	0.0036	-0.0153	-0.062
	(-0.0238)	(0.0387)	(0.0394)	(-0.0575)	(0.0574)
Observations	197579	197579	208717	8282	8708
Funct. Form	Log	Level	Level	Log	Level
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). Standard errors are reported in brackets.

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
δ	0.300	0.064	0.209	0.470	24
β	0.745	0.035	0.652	0.793	24
γ	-0.068	0.075	-0.175	0.064	24
Constant (α)	-0.766	0.192	-1.055	-0.419	24

Notes. The Table reports mean estimates of the parameters δ, β, γ and α across the 24 months from May 2004-April 2006, together with standard deviations, minimum and maximum values. Monthly estimates are maximum likelihood estimates of equation (15), where the number of applications per job is given in equation (10).

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)	(8)
δ	0.259***	0.226***	0.241***	0.263***	0.204***	0.766***	0.260***	
	(0.058)	(0.070)	(0.069)	(0.059)	(0.029)	(0.117)	(0.038)	
β	0.747***	0.793***	0.770***	0.799***	0.758***	0.756***	0.745***	
	(0.064)	(0.076)	(0.087)	(0.050)	(0.039)	(0.042)	(0.044)	
γ	-0.043	-0.067	-0.009	-0.038	-0.060	-0.053	-0.048	
	(0.087)	(0.121)	(0.131)	(0.086)	(0.086)	(0.084)	(0.057)	
μ							0.783***	
							(0.182)	
Constant (α)	-0.771***	-0.770***	-0.777***	-1.004***	-0.754***	-0.762***	-0.758***	-0.899***
	(0.141)	(0.167)	(0.151)	(0.110)	(0.123)	(0.123)	(0.083)	(0.026)
U_b/V_b		0.044***	0.026					0.087***
		(0.008)	(0.025)					(0.013)
V_b			-0.020					
			(0.025)					
Observations	8709	8709	8709	8709	8709	8709	8709	8709
Adjusted R2	0.0395	0.0302	0.0305	0.0391	0.0377	0.0374	0.0410	0.0109
Distance concept	Geographic	Geographic	Geographic	Geographic	Time	Cost	Geographic	-
MF specification	Urn-ball	Urn-ball	Urn-ball	Log-linear	Urn-ball	Urn-ball	Urn-ball	Urn-ball

Notes. The table reports maximum likelihood estimates of the matching function (see equation (15) for the main specification), where the number of applications per job is given in equation (10). Standard error corrected for spatial correlation are reported in brackets.

Table 6
Estimates of a job application model: Controlling for commuting flows

Dependent variable: vacancy outflow rate			
	(1)	(2)	(3)
δ	0.350***	0.229***	0.877***
	(0.070)	(0.034)	(0.133)
β	0.752***	0.763***	0.762***
	(0.045)	(0.033)	(0.035)
γ	-0.036	-0.069	-0.063
	(0.048)	(0.068)	(0.063)
Commuters	-7.702***	-3.664**	-4.404**
	(2.536)	(1.648)	(1.738)
Constant (α)	-0.792***	-0.747***	-0.754***
	(0.061)	(0.086)	(0.080)
Observations	8709	8709	8709
Adjusted R2	0.0413	0.0383	0.0386
Distance concept	Geographic	Time	Cost
MF specification	Urn-ball	Urn-ball	Urn-ball

Notes. The table reports maximum likelihood estimates of the matching function (see equation (15)), where the number of applications per job is given in equation (10). The extra regressor included in the cost of distance measures the number of individuals resident in the origin ward a, who commute to the destination ward b. Standard error corrected for spatial correlation are reported in brackets.

Table 7
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 8
The distribution of actual and predicted commuting flows

	Percentage of commuters	
Distance	Actual	Predicted
0 km	23.7	10.1
(0,5] km	29.5	44.4
(5,10] km	18.4	25.5
(10,20] km	15.5	13.3
20+ km	12.9	6.7
Total	100	100

Notes. Actual commuting flows are obtained from the Census 2001, and predicted commuting flows are obtained from equation (21), evaluated at parameter values reported in column (1) of Table (5).

Table 9
The propagation of local shocks

	Percentage change		
Distance	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 $\gamma = 0$).

Table 10
The effect of reducing the cost of distance

	Percentage change		
Distance	Applications per job	Vacancy outflow	Unemployment outflow
King's Cross	4.38	1.09	-0.43
(0,3] km from King's Cross	1.39	0.36	-0.43
Stratford	1.73	0.43	0.41
(0,3] km from Stratford	-0.59	-0.15	0.20
(3,10] from both	0.13	0.03	-0.09
10+ km from both	-0.02	-0.00	0.00

Notes: The Table shows the simulated effect of halving the cost of distance between King's Cross Ward and Stratford and New Town Ward. The simulation uses estimates from column 1 of Table 5 (having set $\gamma = 0$).

Appendix Tables

Table A1
Estimates of a job application model: Average values for May 2004-April 2006

	Dependent variable: Vacancy outflow rate							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
δ	0.300	0.259	0.269	0.314	0.246	0.917	0.316	
	(0.064)	(0.036)	(0.044)	(0.091)	(0.097)	(0.323)	(0.074)	
β	0.749	0.793	0.774	0.798	0.765	0.762	0.766	
	(0.035)	(0.033)	(0.036)	(0.031)	(0.029)	(0.030)	(0.015)	
γ	-0.068	-0.101	-0.063	-0.060	-0.103	-0.097	-0.039	
	(0.075)	(0.083)	(0.110)	(0.081)	(0.075)	(0.074)	(0.074)	
μ							0.200	
							(0.396)	
Constant (α)	-0.766	-0.760	-0.756	-1.003	-0.737	-0.740	-0.813	-0.962
	(0.192)	(0.189)	(0.191)	(0.154)	(0.174)	(0.177)	(0.188)	(0.122)
U_b/V_b		0.046	0.032					0.091
		(0.016)	(0.018)					(0.017)
V_b			-0.015					
			(0.020)					
Observations	24	24	24	24	24	24	16	24
Distance concept	Geographic	Geographic	Geographic	Geographic	Time	Cost	Geographic	-
MF specification	Urn-ball	Urn-ball	Urn-ball	Log-linear	Urn-ball	Urn-ball	Urn-ball	Urn-ball

Notes. Model specifications are the same as in Table 5. Coefficients reported are averages across monthly estimates, with standard deviations reported in brackets. Specification (7) is only estimated on the months February 2005-May 2006, because unemployment data by occupation become available in January 2005.

Table A2
Estimates of a job application model, controlling for commuting flows: Average values for May 2004-April 2006

Dependent variable: vacancy outflow rate			
	(1)	(2)	(3)
δ	0.379	0.252	0.951
	(0.082)	(0.069)	(0.262)
β	0.748	0.767	0.764
	(0.036)	(0.030)	(0.031)
γ	-0.072	-0.106	-0.099
	(0.056)	(0.076)	(0.074)
Commuters	-10.786	-1.724	-2.458
	(7.251)	(6.114)	(6.755)
Constant (α)	-0.779	-0.739	-0.744
	(0.154)	(0.170)	(0.170)
Observations	24	24	24
Distance concept	Geographic	Time	Cost
MF specification	Urn-ball	Urn-ball	Urn-ball

Notes. Model specifications are the same as in Table 6. Coefficients reported are averages across monthly estimates, with standard deviations reported in brackets.