Search Across Local Labour Markets*

Panagiotis Nanos †
University of Sheffield

Christian Schluter‡
Aix-Marseille Université
(Aix Marseille School of Economics), CNRS & EHESS
and University of Southampton

April 2019

Please do not cite or circulate without the authors' permission

Keywords:
(spatial) job search, search frictions, local labour markets, geographic mobility, internal migration,

JEL Codes: J31, J61, J63, J64, R23

*Funding from ANR-DFG (contract ANR-15-FRAL-0007-01) is gratefully acknowledged. Thanks also to G. van den Berg, C. Dustmann, and L. Visschers for comments, as well as our seminar and conference participants.

†P.Nanos@sheffield.ac.uk, Department of Economics, The University of Sheffield, Sheffield, S1 4DT, UK.

‡christian.schluter@univ-amu.fr, C.Schluter@soton.ac.uk, Aix-Marseille School of Economics, 5 Boulevard Maurice Bourdet CS 50498, 13205 Marseille Cedex 01, France, and Economics Division, University of Southampton, Highfield, Southampton, SO17 1BJ, UK. https://sites.google.com/site/christianschluter1/

Abstract

Local labour markets exhibit substantial and persistent differences in terms of unemployment rates, nominal and net-of-housing-cost wages, as well as firm productivities. Yet the observed spatial mobility of workers searching for jobs, unemployed and on-the-job, is limited and the population response to localised labour demand shocks is very slow. In order to address this empirical puzzle, we propose a new dynamic structural and empirical model of workers' job search within and across many local labour markets that extends the frictional search paradigm by interpreting job search literally as spatial. The tractability of the model, which follows from the directness of search, allows us to accommodate many locations, worker heterogeneity and firm behaviour. The model enables us to quantify the underlying drivers of and barriers (such as relocation costs, classic search frictions, spatial differential flow of information about job opportunities, and their amplifying interaction) to the spatial mobility of workers; search frictions now also have a spatial dimension. We estimate this model structurally using individual transition data obtained from an administrative employer-employee panel from Germany (LIAB).

to do list

- check home productivity b_k
- check that $\lambda_u = 1$
- update identification discussion, specifically amenities vs moving costs
- empirical strategy focuses on specification being parsimonious, aiming to avoid over-fitting
- check/update productivity Appendix (discuss the strategy)
- new Appendix including proof of existence/uniqueness (see Menzio & Shi, JPE 2011)
- rewrite section Model fit using new results
- convincing/interesting/substantive counterfactual experiments
- provision of standard errors
- present fit for sth not targeted

1 Introduction

The literature has documented two principal empirical features characterising the spatial mobility of workers searching for jobs across local labour markets. One the one hand, this spatial mobility is limited,[1] while one the other hand, local labour markets exhibit substantial and persistent differences in terms of e.g. unemployment rates, wages, and firm productivities.[2] Why are the relocation rates of job seekers so low in the presence of such spatial disparities? In order answer this question and then to quantify the underlying drivers of and barriers to the spatial mobility of workers, we propose and then structurally estimate a new dynamic empirical model of spatial search, which essentially reinterprets and generalises the frictional job search paradigm as spatial search. In our spatial model, job searchers (both the unemployed and on-the-job)[3] search within and across local labour markets, trading off local differences in net salaries, unemployment probabilities and non-market aspects such as amenities in the face of moving costs. Such costs and job search frictions interact and lead to spatial search friction. This model is then structurally estimated using an administrative employee-employer panel (LIAB) for Germany.

This approach enables us to make several theoretical and empirical contributions to the search and matching literature. First, while much has been learned about the importance of search frictions, the literature tends to treat the process of job search rather abstractly. We innovate by interpreting job search literally as spatial, so workers search for jobs within and across local labour markets. Search frictions now also exhibit a spatial dimension, as barriers to mobility, spatially learning frictions about job opportunities, and the usual search frictions interact. Seen from an alternative vantage point, our dynamic model builds on the dynamic migration model of Kennan and Walker (2011), and adds a frictional job search perspective. In particular, the observed low mobility of workers is empirically rationalised by Kennan and Walker by very high moving costs. We add to this mechanism by considering search frictions; as in Schmutz and Sidibé (2018), we allow for informational frictions, captured by potentially higher job offer rates in the home location compared to alternative loca-

[1]For instance, Caliendo et al. (2017) report that yearly mobility rates in the US amount to about 3% and in European about 1%. For the sample of young workers (aged 18-27) in the US examined in Kennan and Walker (2011), the average ten-year interstate migration rate is 32%. Molloy et al. (2014) report for the period 2002-2012 an interstate migration rate of 3.3% for workers aged 20-24, 1.5% for workers aged 35-44, and rates of no more than 0.9% for older workers. For shorter movement, the within-county migration rate is 6.6% for workers aged 35-44. For France, Schmutz and Sidibé (2018) report yearly mobility rates across large regions (NUTS2) for the employed of no more than 1.5% and no more than 2.5% for smaller (NUTS3, departments) regions. The overall transition rates for all are only marginally larger. Amior and Manning (2018) also demonstrate convincingly that the population response to localised labour demand shocks is very slow.

[2] Such spatial variation is extensively documented in e.g. Moretti (2011), OECD (2005), or Overman and Puga (2002), Schmutz and Sidibé (2018) for France, and Section 2 below for evidence for Germany.

[3]Both unemployed and employed can search for new jobs in our model. This is in line with the data, as, for instance, Rupert and Wasmer (2012) have demonstrated that both groups are spatially mobile: using data from the 2000 US Census, they report that 17% of employed and 25% of the unemployed have changed residence, and that 42% of relocations are across counties. In our German administrative data, job-to-job transitions are of the same order of magnitude as out-of-job transitions. By contrast, some leading models in the literature restrict job search either to the employed or the unemployed (e.g. Beaudry et al., 2012).

tions. These additional frictions permit low mobility rates to be consistent with much lower moving costs. While Schmutz and Sidibé (2018) propose a dynamic random search model, our approach is complementary: We generalise and extend the directed search paradigm of Shi (2009) and Menzio and Shi (2011), which gives rise to several key differences that are outlined next.

As our second contribution, our dynamic directed spatial search model is very tractable, which allows us to overcome several key challenges that have affected the estimation of such dynamic location models. This tractability follows from the directedness of search as workers self-select into local labour markets: workers only apply to jobs they intend to accept and firms only meet workers willing to fill their vacancies. Contact probabilities (and hence value functions) are therefore independent of the distributions of workers across states/locations. Random search, by contrast, are less tractable and computationally more challenging because firms can meet workers unwilling to accept their job offers so contact probabilities depend on the distributions of workers across states (an exception is Lise and Robin (2017), who consider, as we do, the joint surplus of a match). This tractability of our model enables us to accommodates a large number of locations in our empirical analysis. Specifically, we consider travel-to-work areas (TTWAs) in West Germany, which, unlike administrative spatial units (such as municipalities or cities) reflect the spatial organisation of economic activity and the idea of a local labour market. Leading dynamic models in the literature often restrict attention to a very small number of locations.[4] Furthermore, the tractability of our model allows us to lift some restrictions Schmutz and Sidibé (2018) had to impose, such as considering workers as homogeneous, and abstracts from firm behaviour. By contrast, we model the firm's decision problem, estimate firm productivity, and in the empirical analysis allow for different worker types. Finally, the model permits the study of adjustment paths in the wake of a localised shock, and we conduct several such counterfactual experiments below.

Third, we estimate our model is structurally using individual-level transition data from an administrative employee-employer panel (LIAB) for Germany. @@expand@@

- need to discuss briefly / reference Ludo and Carlos.

The paper is organised as follows. Section 2 establishes stylised facts about the extent and persistence of spatial heterogeneity, as well as the limited mobility of workers based on our German data. These empirical features inform our model. The directed search model is described next. We start in Section 3 with a description of the environment within which agents interact. The workers' transitions within and across local labour markets by employment states are presented in Section 3.2, where we take as given the decisions of workers and firms. These decisions are presented for the decentralised economy in Section 4, while Appendix A presents in detail the Social Planner's Problem, characterises its solution, and demonstrates that the competitive

[4] For instance, Gould (2007) has two locations corresponding to a rural and an urban area, whereas Baum-Snow and Pavan (2012) consider three locations corresponding to small/medium/large cities). By contrast, Kennan and Walker (2011) consider inter-state moves, but have to restrict the information available to each individual. They observe that "(i)deally, locations would be defined as local labour markets; (...) even if J is the number of States, the model is computationally infeasible" (p.216).

equilibrium is also socially efficient. In Section 5, we describe our estimation strategy, discuss identification of our model's parameters, outline the parameterisation employed in our structural estimation, and present preliminary estimation results. The appendices provide further details and supplementary analyses.

2 Persistent Spatial Heterogeneity and Limited Mobility Across Local Labour Markets: Stylised Facts for Germany

We proceed to describe and quantify the persistent spatial heterogeneity across local labour markets and the associated limited spatial mobility of workers in the setting of our empirical application. Specifically, we have at our disposal individual-level administrative employee-employer data for Germany. Throughout, we will interpret the notion of a local labour market as a travel-to-work area (TTWA). These empirical observations set the scene for and inform our theoretical model presented in Section 3.

2.1 Data

Our individual-level transition data is drawn from a rich administrative employer-employee panel from the German Social Security system that has been assembled by the Research Data Centre (FDZ) of the German Federal Employment Agency into the LIAB data. The LIAB LM 9310 covers the years 1993 to 2010 (see Klosterhuber et al. (2013) for a recent description of the data, and Dustmann et al. (2009) and Card et al. (2013) for a recent use in the context of the German wage structure). We focus on workers in West Germany, given the persistent peculiarites of the labour market in East Germany.[5]

This dataset samples private sector workers and includes daily earnings and total days worked at each job in a year, the total length of unemployment spells, as well as information on occupation, industry and education. Since the data are based on administrative social security registers, individual information about labour market states and wage is of exceptional quality, and accurate to the day. While civil servants and the self-employed are not sampled, dependent private sector employees constitute about 80% of the workforce. The establishment identifiers link the workers employed in firms as of the 30th June to the annual waves of the IAB Establishment Panel. In each year, the data cover on average about 1.4 million individuals and 300,000 establishments.

The employer-employee panel enables the estimation of worker fixed effects in comprehensive wage regressions. Specifically, we use the worker fixed effects as obtained by Card et al. (2013) who apply the methodology of Abowd, Kramarz and Margolis (1999) using the universe of German private sector workers (assembled in the Integrated Employment Biographies), from which our data is drawn. These worker

[5]For instance, Heise and Porzio (2018) report a persistent unconditional wage gap of 27% between workers in the West and the East, persistent productivity differences, and attribute part of the wage gap to a preference of East Germans to live in the East.

Figure 1: Travel-to-work areas (TTWAs) and the spatial distribution of unemployment.

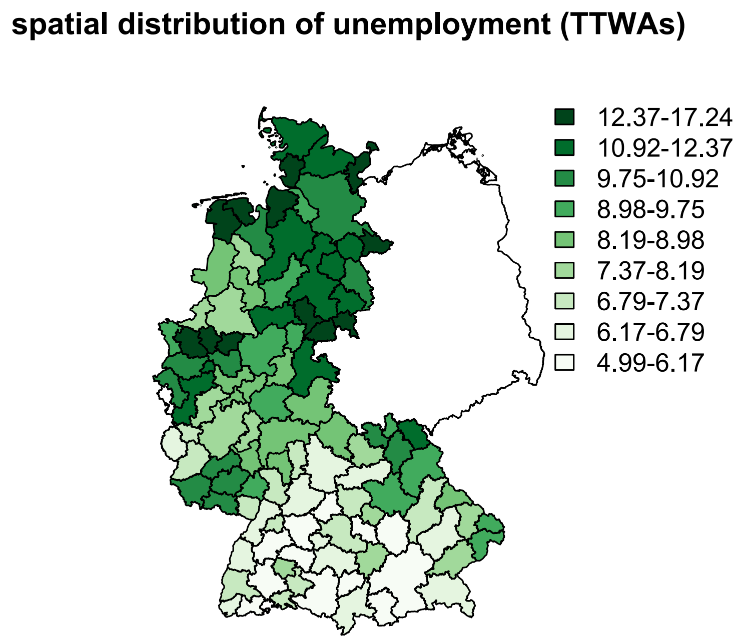

Notes: Depicted are 108 TTWAs in West Germany (we exclude Berlin), and mean local unemployment rates over time period 2002-2008, arranged into 9 quantile groups. The unemployment rate is obtained from www.regionalstatistik.de at the level of the district, and aggregated for TTWAs using weights given by district-level relative population size.

fixed effects will be used in the empirical analysis in order to allow for worker heterogeneity. Specifically, we consider 3 groups using the shorthand "low, midlle, and high" ability.

The employer side of the data also enables us to measure firm-level productivity.We therefore do not use the firm fixed effects that are typically included in the AKM approach, and whose interpretation has been challenged, see e.g. Eeckhout and Kircher (2011). We follow the established literature in order to estimate firm-level total factor productivity (tfp) using Cobb-Douglas production functions and recent empirical strategies that have originated in Olley and Pakes (1996). We also follow established LIAB-based strategies in the empirical implementation, which are detailed in Data Appendix D. Since it is customary to estimate production functions only for firms in manufacturing, our subsequent analysis is restricted to this sector.

In summary, we consider the time period 2002-2008, focussing our analysis on prime-aged males (20-60) who work in manufacturing and reside in West Germany.

2.2 Travel-to-work areas

Our analysis of local labour markets is enabled by the spatial information contained in the LIAB. The data at our disposal provides information about the place of residence and the place of work. The spatial unit is the district (consistently coded with respect to its status on 31.12.2010), West-Germany being partitioned into about 326 districts. As these spatial units are defined administratively, they do not necessarily reflect the spatial organisation of economic activity and the idea of a local labour market. We therefore aggregate these administrative spatial units into travel-to-work areas using the classification of Eckey et al. (2006), which is based on a detailed factor analysis of actual commuting flows within radii of up to 60 minutes travel time. Henceforth, we use the labels of travel-to-work areas (TTWA) and local labour markets interchangeably. This spatial aggregation of West German district results in 108 TTWAs, none of which is smaller than 50,000 inhabitants. We have excluded Berlin, given its special status as capital city and it being located in East Germany. The map of Figure 1 depicts these TTWAs. The federal constitutional and political structure of Germany also manifests itself in its urban structures, since Germany lacks a predominant center of gravity (such as Paris or London). Finally, in order to control for the spatial differences in the cost of housing and living, we also use a district-level house price index based on actual transactions recorded on the largest German online portal rendered comparable by hedonic price regressions.

2.3 Data Descriptives: Persistent Spatial Heterogeneity, and Transitions

Table 1: Heterogeneity across all local labour markets

	Percentiles		
	10	50	90
mean unemployment rate	6.04	8.53	12.55
relative house price index	0.59	0.74	0.87
mean daily log-wages	4.27	4.42	4.56
mean worker fixed effects	3.74	3.83	3.93
mean firm productivity (low manuf)	4.42	4.59	4.76
mean firm productivity (low services)	2.84	3.12	3.30
mean firm productivity (high manuf)	4.48	4.64	4.92

Notes: Period 2002-2008, the spatial units are 108 TTWAs. TTWA means computed using weights given by district-level relative population size. District-level data (population size) and mean unemployment rate obtained from www.regionalstatistik.de (Table 173-01-4 for year 2002, and Table 659-71-4 averaged over 2002-2008). The house price index is obtained from www.immobilienscout24.de, for year 2007, expressed relative to TTWA München. Worker fixed effect (FEs) obtained from log wage regression described in Card et al. (2013), and averaged across the districts of each TTWA using establishment/district employment levels as weights. Firm productivity (TFP) estimated using the methodology set out in Appendix C using firm-level LIAB data, based on Olley and Pakes (1996), and spatially aggregated.

We document the principal features of our data, first examining the evidence for persistent spatial heterogeneity across the all local labour markets, and then summarising the transition data. In order to provide greater detail and spatial resolution, in Data Appendix D, we consider explicitly 8 selected local labour markets.

Figure 1 depicts the spatial distribution of unemployment for 9 quantile groups. Southern Germany tends to have lower rates, while local unemployment in Lower Saxony and North Rhine-Westphalia is particulary elevated. It is clear that spatial heterogeneity is substantial. Table 1 reports the 10%/50%/90% deciles of the marginal distributions of unemployment, wages, relative house prices, as well as location specific mean worker fixed effects and firm TFP in manufacturing (low- and high-tech) and in low services. It is evident that spatial variations are pervasive. For instance, the 90/10 ratio of local unemployment rates is 2.1, for relative house prices 1.5, and log wages 1.07.

This heterogeneity across local labour markets is not only pervasive but also highly persistent. For instance, consider the year-to-year Spearman rank correlation of local unemployment rates. For all TTWAs, the smallest rank correlation is .968, for all districts it is .975. Even for a ten year lag, the rank correlation for all TTWAs is still .86.

Table 2: Spatial mobility and job transitions.

total $e \rightarrow u$ transitions	558,056	7.92%
total $u \rightarrow e$ transitions	547,823	7.77%
total $e \rightarrow e$ transitions	693,956	9.85%
total spells	7,046,710	
total relocations		
within TTWAs [%]	40.63	
across TTWAs [%]	59.37	
total relocations given transitions into employment $(u, e \rightarrow e)$		
within TTWAs [%]	72.70	
across TTWAs [%]	27.30	
total relocations given transitions into unemployment		
within TTWAs [%]	76.54	
across TTWAs [%]	23.46	

Notes: Based on LIAB. We report the share of spells by type for the window 2002-2008.

Turning to the transition data, Table 2 reports measures of worker transitions on the labour market and across locations. The incidence of job-to-job transitions is of the same magnitude as out-of-job transitions, which we interpret as strong indirect evidence of the importance of on-the-job search. Our model accommodates both types of transitions, and both type of searchers are permitted to change location. Empirically, most job-related mobility is short range, as, given a labour status and location change, only about 25% change TTWA while three-quarters change location within their TTWA. All these data features inform our model, which is presented next.

3 A Tractable Equilibrium Model of Directed Search across Local Labour Markets

We proceed to present our dynamic model of spatial search. Building on the directed search models of of Shi (2009) and Menzio and Shi (2011), we add a spatial perspective: Now workers can search within and across local labour markets, and make optimal location decisions in the face of moving costs. Job searchers might also be better informed about local job opportunities relative to more distant ones. Our objective is to build a very parsimonious model. To this end, we (could allow for but) refrain from allowing some frictional parameters to vary across locations. In what follows we present the generic model for one type of worker (in the empirical analysis we further introduce worker heterogeneity by stratifying the analysis by worker types, allowing some parameters to be segment specific whilst others are segment-invariant).

3.1 The Environment

Time is discrete and continues for ever. The economy is populated by a continuum of workers with measure 1. Each worker is endowed with an indivisible unit of labour and maximises the expected sum of periodical consumption discounted by the factor $\beta \in (0, 1)$.

Economic activity occurs within geographically defined markets or locations, indexed by $k \in K = \{1, ..., N(k)\}$ with $N(k) \geq 2$. Workers can move across locations, whereas firms cannot. Moving from source location l to a new destination location k is costly, and measured by a cost function $c_i(l, k) \geq 0$ that depends on the employment state of the individual i ($i \in \{u, e\}$). Hence we have ex ante heterogeneity of workers in terms of relocation costs; in our empirical application, we will introduce further sources of heterogeneity by segmenting the labour market by industry.

In addition to the endogenous mobility of workers, there are also exogenous relocations: every period a random sample of workers (employed or unemployed) leave the economy; for simplicity, these transitions are labelled as deaths. The "mortality rate" $\tau \in [0, 1)$ is exogenous. Deceased individuals are replaced by an equal measure of new-born workers. These new entrants are randomly allocated across locations, join the unemployment pool, and cannot search during their first period.

There is a continuum of firms with positive measure in every location k. Each firm uses a technology that turns one unit of labour into $\pi(y, \mu)+z$ units of output, where π is a constant returns to scale, increasing, and concave function. The first component of productivity, y, is common to all firms and its value lies in $Y = \{y_1, ..., y_{N(y)}\}$ with $N(y) \geq 2$. The second component of productivity, μ, is specific to the location of the firm $\mu \in \{\mu_1, ..., \mu_{N(k)}\}$. The third component of productivity, z, is specific to a firm-worker pair, and its value lies in $Z = \{z_1, ..., z_{N(z)}\}$ with $N(z) \geq 2$. The aggregate component of productivity y captures aggregate business cycle conditions, whereas μ captures local differences in productivity driven by e.g. agglomeration economies (as emphasised by e.g. Combes et al. (2012)). Hence our model exhibits ex ante heterogeneity on the side of the firm, since firms' productivities differ spatially (see also Kaas and Kircher (2015) on the importance of firm heterogeneity). Firms can enter freely a location, and inherit the common location-specific productivity component. In equlibrium, with free entry, firms will be indifferent in which location to produce since all firms will earn the same profits.

3.1.1 Timing

At the beginning of every period, the state of the economy can be summarised by $\psi = (y, u, g)$, where $y \in Y$ is the aggregate component of productivity, u denotes the measure of workers who are unemployed in the economy and is given by the sum of the measures of unemployed individuals in every location, $u = \sum_{k \in K}\{u_k\} \leq 1$ with $u_k \in [0, 1]$, and g is a function $g : Z \times K \to [0, 1]$ with $g(z, k)$ denoting the measure of workers who are employed in matches with idiosyncratic productivity z in location k. Every location k therefore consists of a collection of submarkets indexed by x.

Each period is divided into five stages: births and deaths, separation, search, matching, and production.

At the separation stage, within each submarket x, and each location k matches between firms and workers are destroyed with probability $d_e(z, k) \in [\delta, 1]$, where $\delta \in (0, 1)$ denotes the probability that a match is destroyed for exogenous reasons. Separated workers must spend one period in unemployment before searching. The unemployment pool consists of individuals searching for a match (labelled unemployed job-searchers) and individuals who cannot search for one period (labelled unemployed non-searchers) either because they are new entrants or recently separated. At the separation stage, within each location k unemployed job-searchers become unemployed non-searchers with probability $d_u(k) \in [0, 1]$. The measure of unemployed non-searchers in location k is denoted $ns_k \in [0, 1]$, and the corresponding measure of unemployed job-searchers is given by $u_k - ns_k$.

At the search stage, individuals can move and/or search for a job within or across locations. Unemployed non-searchers move from their current location, l, to a different location, k with probability $\eta_m(l, k) \in [0, 1]$. Unemployed job-searchers look for a job with probability $\lambda_u \in [0, 1]$, while employed workers search with probability $\lambda_e \in [0, 1]$. The probability that an unemployed individual in location l looks for a job in a different location, k, is $\eta_u(l, k) \in [0, 1]$. Similarly, the probability that a worker employed in a match of productivity z in location l looks for a match in a different location, k, is $\eta_e(z, l, k) \in [0, 1]$. At the search stage, firms decide how many vacancies to post. The cost of maintaining an open vacancy is $\xi > 0$ per period.

At the matching stage, individuals and vacancies searching in the same location and submarket meet. The meeting technology is constant returns to scale and can be expressed as a function of the submarket and location specific vacancy-to-searcher ratio, θ, i.e. the local labour market tightness. The probability that a job-seeker meets a vacancy in this submarket is $p(\theta)$ and the probability that a vacancy meets a worker is $q(\theta) = p(\theta)/\theta$. When a firm meets a job-seeker, nature draws z from the probability distribution $f(z)$.

As in Menzio and Shi (2011), we allow for learning frictions. The firm-worker pair do not directly observe their match-specific productivity: they observe s, which is a signal of z. With probability $\alpha \in [0, 1]$, the signal s is equal to z and with probability $(1 - \alpha)$ the signal s is drawn from f independently of z. At opposite ends of the spectrum stand the cases of experience goods ($\alpha = 0$) and inspection goods ($\alpha = 1$).

In the latter case, the quality of the match is known before forming it, in the former case no information is available. The informativeness of signals may differ across locations: $\alpha_{ll} \geq \alpha_{lk} \forall l, k \in K$.[6] Conditional on the signal, s, firms decide to hire the worker using a selection criterion r. A firm hires a worker if and only if the signal s about the quality of their match is greater than or equal to r. The probability that the signal about the quality of the match is above the selection cutoff r is given by $m(r) = \sum_{s \geq r} f(s)$.

At the production stage, an unemployed individual in location k, produces b_k units of output, where $b_k \in B = \{b_1, ..., b_{N(k)}\}$ and also enjoys flow utility A^u_k, where $A^u_k \in A^u = \{A^u_1, ..., A^u_{N(k)}\}$, from local amenities. A worker employed in a match with idiosyncratic productivity z in location k produces $\pi(y, \mu_k) + z$ units of output and also enjoys flow utility A^e_k, where $A^e_k \in A^e = \{A^e_1, ..., A^e_{N(k)}\}$, from local amenities. After production, the firm-worker pair observe z. At the end of this stage, nature draws next period's aggregate component of productivity, $\widehat{y}$, from the probability distribution $\phi(\widehat{y}|y)$, where $\phi : Y \times Y \to [0, 1]$.

3.1.2 The Labour Market

The labour market is organised in a continuum of submarkets indexed by (x, r, k), where x is the lifetime utility offered by a firm to a worker, r is the selection criterion, and k is the location of the submarket. As is usual in this type of directed search model, employment contracts are assumed to be complete in the sense that a contract can specify the wage, w, the separation probability, d_e, the probability of search in a different location, η_e, and the submarket where the worker searches while on the job, (x, r, k), as functions of the history of the aggregate state of the economy and the quality of the match, z. The firm maximizes its profits by choosing the contingencies for d_e, η_e, x, and r so as to maximize the joint value of the match, and by choosing the contingencies for w so as to deliver the promised value x. The assumption of complete contracts captures the view that firms and workers have an incentive to find ways in practice to maximise the joint gains from trade.

3.2 Transitions within and across Local Labour Markets

We consider in detail the possible transitions that could be experienced by a worker in a particular employment state, location, and submarket since these are the principal objects of the empirical investigation. We then establish for each location l next period's measures of non-searchers $\widehat{ns}_l$, of unemployed job-seekers $\widehat{u}_l$, and of employed workers of productivity z, $\widehat{g}(z, l)$. These transitions are, of course, based on the optimal choices of workers and firms. In this section, we take these as given, and defer their derivations to Section 4. For notational simplicity, the optimal policy functions are indicated by the max superscript.

[6] Kennan and Walker (2011) assume that workers only know the wage in their home location, and need to move to other locations to determine the local wage. Schmutz and Sidibé (2018) also assume the existence of informational frictions across locations.

3.2.1 The Unemployed: Non-searchers and Job-seekers

At the beginning of the period, a measure τ of individuals (employed and unemployed) leave the economy while an equal measure of new entrants are equally distributed across locations (τ/K per location) and join the unemployment pool. New entrants are not allowed to search for a match (hence the label non-searchers ns), but are allowed to move to a different location.

Consider such a non-searchers in location l. These consist of non-searchers who have originated in location k' and have optimally moved to location l, and of non-searchers who have decided to stay in l. We denote the 0-1 probability of the former event by $\eta_m^{max}(k', l)$, which equals 1 if the worker originating in location k' optimally chooses location l, and 0 otherwise. The staying probability is therefore $1 - \sum_{k' \neq l} \eta_m^{max}(l, k') \equiv 1 - \eta_m^{max}(l)$. Apart form these new entrants in l or k' who stayed in or moved to l, $\widehat{ns}_l$ includes: (a) individuals who entered the period as unemployed job-seekers in l or k', decided to stop searching with source location-specific probability $d_u(l)$ or $d_u(k')$, and stayed in or moved to l; and (b) individuals who entered the period as employed job seekers in l or k', were separated with source location-specific probability $d_e(z, l)$ or $d_e(z, k')$, and stayed in or moved to l. The measure of unemployed non-searchers in location l at the end of the search stage therefore equals

$$\widehat{ns}_l = \frac{\tau}{N(k)} \left\{ [1 - \eta_m^{max}(l)] + \sum_{k' \in K} \eta_m^{max}(k', l) \right\} \tag{1}$$

$$+ (1 - \tau) \left\{ [1 - \eta_m^{max}(l)] \tilde{d}(l) + \sum_{k' \in K} \eta_m^{max}(k', l) \tilde{d}(k') \right\}$$

where $\tilde{d}(k') = d_u(k') u_{k'} + \sum_{z \in Z} [d_e(z, k') g(z, k')]$. Consider next an individual who enters the period unemployed in location l. This individual does not leave the economy with probability $1 - \tau$, and enters the search stage as an unemployed job-searcher with probability $1 - d_u(l)$. At the beginning of the search stage, the individual looks for potential matches in location l with probability $1 - \eta_u^{max}(l)$. With complementary probability $\eta_u^{max}(l)$, she searches in a different location. At the matching stage, the unemployed job-seeker meets a firm with probability $\lambda_u \, p(\theta_u^{max}(l))$, and a match with idiosyncratic productivity $z = s$ is created with probability $h_u^{max}(s)[a + (1 - a) f(s)]$, while a match with $z' \neq s$ is created with probability $h_u^{max}(s)(1 - a) f(z')$. Therefore, at the production stage, the individual job-seeker is still unemployed with probability $1 - \lambda_u \, p(\theta_u^{max}(l)) \, m_u^{max}(l)$, where $m_u^{max}(l) = \sum_s [h_u^{max}(s, l) f(s)]$, while she is employed in a match of productivity z' with probability $\lambda_u \, p(\theta_u^{max}(l)) \, [a h_u^{max}(z', l) + (1 - a) m_u^{max}(l)] f(z')$.

Given these possible transitions, the measure of unemployed individuals in location l at the production stage $\widehat{u}_l$ includes the measure of unemployed non-searchers $\widehat{ns}_l$, and the measure of individuals who entered the period as unemployed job-seekers in location l, remained in the economy, searched for a job, but failed to find a match in l or in any other location k':

$$\widehat{u}_l = u_l \times (1 - \tau)(1 - d_u(l)) \times [1 - \lambda_u \, p(\theta_u^{max}(l)) \, m_u^{max}(l)] + \widehat{ns}_l \tag{2}$$

3.2.2 Employed Workers

An individual who enters the period employed in a match of productivity z in location l reaches the search stage with probability $(1-\tau)(1-d_e(z,l))$. On-the-job searchers decide optimally where to locate at the beginning of the search stage. The worker searches in a different location with 0-1 probability $\eta_e^{max}(z,l)$, and stays with probability $1-\eta_e^{max}(z,l)$. At the same time workers in matches of productivity $z' \neq z$ in other locations k decide to optimally relocate to l with probability $\eta_e^{max}(z',k',l)$, which equals 1 if this worker originating in location k' optimally chooses location l, and 0 otherwise.

At the matching stage, the worker meets a new firm with probability $\lambda_e\, p(\theta_e^{max}(z,l))$. A match of idiosyncratic productivity $z=s$ is created with probability $h_e^{max}(s,z,l)[a+(1-a)f(s)]$, while a match with productivity $z' \neq s$ is created with probability $h_e^{max}(s,z,l)(1-a)f(z')$. The worker stays in her original job with probability $1-\lambda_e\, p(\theta_e^{max}(z,l))\, m_e^{max}(z,l)$ where $m_e^{max}(z,l)=\sum_s[h_e^{max}(s,z,l)f(s)]$.

The production stage measure of individuals employed in matches of productivity z in location l, $\widehat{g}(z,l)$, therefore consists of workers who (a) entered the period employed in a match of productivity z in l and did not change employment status; or (b) entered the period employed in a match of productivity $z' \neq z$ in l or k' and found a match of productivity z in l; or (c) entered the period as unemployed job-seekers in l or k' and found a match of productivity z in l. In summary, we have

$$
\begin{aligned}
\widehat{g}(z,l) \quad &= \quad (1-\tau)\times\{g(z,l)[1-d_e(z,l)][1-\lambda_e\, p(\theta_e^{max}(z,l))\, m_e^{max}(z,l)] \\
&+ \quad \sum_{z'\in Z} g(z',l)[1-d_e(z',l)][1-\eta_e^{max}(z',l)]\tilde{\gamma}_e(l,l,z',z) \qquad (3)\\
&\qquad + \sum_{k'\in K}\sum_{z'\in Z} g(z',k')[1-d_e(z',k')]\eta_e^{max}(z',k',l)\tilde{\gamma}_e(k',l,z',z) \\
&+ \quad u_l[1-d_u(l)][1-\eta_u^{max}(l)]\tilde{\gamma}_u(l,z)+\sum_{k'\in K} u_{k'}[1-d_u(k')]\eta_m^{max}(k')\tilde{\gamma}_u(k',z)\ \}
\end{aligned}
$$

with

$$
\begin{aligned}
\tilde{\gamma}_u(k',z) \quad &\equiv \quad \lambda_u p(\theta_u^{max}(k'))[\alpha h_u^{max}(z,k')+(1-\alpha)m_u^{max}(k')]f(z) \\
\tilde{\gamma}_e(k',l,z',z) \quad &\equiv \quad \lambda_e p(\theta_e^{max}(z',k',l))[\alpha h_e^{max}(z,z',k')+(1-\alpha)m_e^{max}(z',k')]f(z)
\end{aligned}
$$

4 The Decentralised Economy

We proceed to discuss the optimal search behaviour of workers in the decentralised economy, the optimal behaviour of firms, and the resulting equilibrium. Appendix A presents in detail the associated problem of the central social planner. In line with the earlier literature on competitve and directed search, it will turn out that the decentralised equilibrium is socially efficient.

4.1 Value Functions

We proceed to consider in detail the value function of each worker by labour market status. The decision problem of where and in which submarket to search is broken down into two stages. In the first stage, a worker makes pairwise comparisons between the current location l and any possible destination k. In the second stage, the worker then picks the best alternative.

Consider an unemployed individual in location l at the beginning of the production stage. Her lifetime utility is denoted $U(l, y)$. In the current period, she produces b_l units of output and enjoys flow utility A_l^u from local amenities. With probability $(1-\tau)$, she survives until the next period. At the separation stage, with probability $d_u(l, \widehat{y})$ she decides to quit to a state of non-searching, which gives her lifetime utility $J_u^{max}(l, \widehat{y})$, or to enter the search stage looking for a match. During the search stage, she decides where to search for a match by comparing the potential net gains from searching in each submarket within every location. The unemployed individual's choice of destination submarket and location is made optimally in two stages: first, she chooses the submarket within each location that maximises her value, and then selects the destination location that maximises the net gain from search. Suppose the optimally chosen destination submarket is x in location k. At the matching stage, she meets a vacancy that leads to an acceptable match giving her expected lifetime utility x with probability $p(\theta(x, r, l, k, \widehat{y}))m(l, k, r)$. If source and destination locations differ ($k \neq l$), the individual has to incur the moving cost $c_u(l, k)$. With probability $1 - p(\theta(x, r, l, k, \widehat{y}))m(l, k, r)$, she fails to meet an acceptable vacancy, so her employment status does not change and her expected lifetime utility is $U(l, \widehat{y})$. The unemployed job searcher's value is given by:

$$\begin{aligned} U(l, y) = b_l + A_l^u + \beta(1-\tau)\mathbb{E}\max_{d_u}\{d_u J_u^{max}(l, \widehat{y}) \\ + (1-d_u)[U(l, \widehat{y}) + \lambda_u D_u^{max}(U(l, \widehat{y}), l, \widehat{y})]\}, \end{aligned} \tag{4}$$

where D_u^{max} is the total return to search function for unemployed job-searchers, giving the highest net gain from searching across all possible destination locations

$$D_u^{max}(U, l, y) = \max\{0, D_u(U, l, 1, y), D_u(U, l, 2, y), \ldots, D_u(U, l, K, y)\} \tag{5}$$

and D_u is the return to search function for unemployed job-searchers, giving the highest net return from searching across all possible submarkets within a destination location

$$\begin{aligned} D_u(U, l, k, y) = \max_{x, r, \eta_u}\{(1-\eta_u)\left[p(\theta(x, r, l, l, y))m(l, l, r)(x - U)\right] \\ + \eta_u\left[p(\theta(x, r, l, k, y))m(l, k, r)(x - U - c_u(l, k))\right]\} \end{aligned} \tag{6}$$

Consider now an individual who became an unemployed non-searcher during the separation stage. At the beginning of the search stage, the unemployed non-searcher decides optimally whether to relocate and where. This choice is made optimally in two stages: first the unemployed non-searcher makes pairwise comparisons between her value in the current location l and her value in any possible destination location k net of moving costs, $J(l, k, y)$. In the second stage, she chooses the location k^* that gives her the highest net value, $J^{max}(l,) = J(l, k^*,)$. For the remainder of the period, the unemployed non-searcher is inactive in the labour market, engaging solely in home production, b_{k^*}. Hence, the problem solved by an unemployed non-searcher in location l is given by

$$J_u^{max}(l, y) = \max\{J_u(l, 1, y), J_u(l, 2, y), \ldots, J_u(l, K, y)\} \tag{7}$$

where

$$J_u(l,k,y) = \max_{\eta_m}\{(1-\eta_m)U(l,y) + \eta_m[U(k,y) - c_u(l,k)]\}. \tag{8}$$

Finally, consider a firm-worker pair in location l with match-specific productivity z. At the beginning of the production stage, the joint value of this match, denoted $J_e(z,l,y)$, is given by the sum of the present discounted value of the worker's utility and the firm's profits. In the current period, the firm-worker pair produce $\pi(y,\mu_l)+z$ units of output, and the worker enjoys flow utility A_l^e from local amenities. With probability $(1-\tau)$, the worker survives until the next period. At the separation stage, with probability $d_e(z,l,\widehat{y})$ the match is destroyed: the worker joins the unemployment pool as a non-searcher with lifetime utility $J_u^{max}(l,\widehat{y})$, and the firm becomes idle with zero profit. With probability $(1-d_e(z,l,\widehat{y}))$ the match survives and the firm-worker pair enter the search stage. During the search stage, with probability λ_e, the worker searches for new matches: she decides where to search by comparing the potential net gains from searching in each submarket within every location. The worker's choice of destination submarket and location is made optimally in two stages: first, she chooses the submarket within each location that maximises the gain from search, and then selects the destination location that maximises the net gain from search.[7] Suppose the optimally chosen destination submarket is x in location k. At the matching stage, the worker meets a vacancy that leads to an acceptable match giving her expected lifetime utility x with probability $p(\theta(x,r,l,k,\widehat{y}))m(l,k,r)$. If source and destination locations differ ($k \neq l$), the worker has to incur the moving cost $c_e(l,k)$. With probability $1-p(\theta(x,r,l,k,\widehat{y}))m(l,k,r)$, the worker does not meet an acceptable vacancy, so she remains employed in a match of productivity z in location l with joint value $J_e(z,l,\widehat{y})$. The joint value of this match is given by:

$$\begin{aligned} J_e(z,l,y) = \pi(y,\mu_l) + z + A_l^e + \beta(1-\tau)\mathbb{E}\max_{d_e}\{&d_e J_u^{max}(l,\widehat{y}) \\ &+ (1-d_e)\left[J_e(z,l,\widehat{y}) + \lambda_e D_e^{max}\left(J_e,z,l,\widehat{y}\right)\right]\} \end{aligned} \tag{9}$$

where D_e^{max} is the total return to search function for employed job-searchers, giving the highest net gain from searching across all possible destination locations

$$D_e^{max}(J_e,z,l,y) = \max\{0, D_e(J_e,z,l,1,y), D_e(J_e,z,l,2,y), .., D_e(J_e,z,l,K,y)\} \tag{10}$$

and D_e is the return to search function for employed job-searchers, giving the highest net return from searching across all possible submarkets within a destination location

$$\begin{aligned} D_e(J_e,z,l,k,y) = &\max_{x,r,\eta_e}\{(1-\eta_e)\left[p(\theta(x,r,l,l,y))m(l,l,r)(x - J_e(z,l,y))\right] \\ &+\eta_e\left[p(\theta(x,r,l,k,y))m(l,k,r)(x - J_e(z,l,y) - c_e(l,k))\right]\} \end{aligned} \tag{11}$$

[7] Complete contracts imply that workers internalise the effect of their search decisions on the profits of the firm. Therefore, the solution to the search problem of employed workers should lead to a match that yields the maximum joint value net of search costs.

4.2 Vacancy Creation

The decision of firms to post vacancies in a submarket (x, k, r) is made optimally by weighing up the costs and benefits of vacancy creation at the margin. The cost of posting a vacancy is ξ. The expected benefit of posting a vacancy in submarket (x, k, r) is given by the product between the probability that the firm fills the vacancy, $q(\theta(x, r, l, y))$, and the value to the firm from filling the vacancy, $\sum_{s \geq r} \{(\alpha J_e(s, k, \widehat{y})+$ $(1-\alpha)\mathbb{E}_z J_e(z, k, y) - x) f(s)\}$. A firm never creates a vacancy in any submarket where the cost of posting the vacancy ξ exceeds the expected benefit of matching with a worker. By contrast, if the expected benefit exceeds the cost of posting a vacancy, then the firm would seek to open as many vacancies as possible in that market. The free entry of firms guarantees that any profits are competed away, so submarket tightness, θ, is such that

$$\xi \geq q(\theta(x, r, k, y)) \sum_{s \geq r} \{(\alpha J_e(s, k, y) + (1-\alpha)\mathbb{E}_z J_e(z, k, y) - x) f(s)\} \tag{12}$$

and $\theta(x, r, k, y) \geq 0$ with complementary slackness. Condition (12) ensures that the market tightness function, θ is consistent with the incentives of firms to create vacancies.

4.3 Policy Functions

All policy functions reflect the two-stage optimisation strategy of workers. In particular, for unemployed job-searchers, the policy functions pertaining to the pairwise location comparison of stage one, given by equation (6), are $x_u(l, k, \widehat{y})$, $r_u(l, k, \widehat{y})$, $\eta_u(l, k, \widehat{y})$, and for the optimal location choice in stage two, given by equation (5), $x_u^{max}(l, \widehat{y})$, $r_u^{max}(l, \widehat{y})$, $\eta_u^{max}(l, \widehat{y})$. $d_u(l, \widehat{y})$ is the decision of unemployed job-searchers to quit to a state of non-searching and solves equation (4). Analogously, for unemployed non-searchers, we have the first stage policy function $\eta_m(l, k, \widehat{y})$, and the second stage policy function $\eta_m^{max}(l, \widehat{y})$. The policy functions for employed job-searchers follow the same logic: for the first stage we have $x_e(z, l, k, \widehat{y})$, $r_e(z, l, k, \widehat{y})$, $\eta_e(z, l, k, \widehat{y})$, for the second stage these are $x_e^{max}(z, l, \widehat{y})$, $r_e^{max}(z, l, \widehat{y})$, $\eta_e^{max}(z, l, \widehat{y})$. $d_e(z, l, \widehat{y})$ is the decision of the firm-worker pair to destroy the match. Finally, we observe that the policy functions can be expressed in terms of θ by solving (12): for the first stage problem we thus obtain $\theta_e(z, l, k, y)$ and $\theta_u(l, k, y)$, and for the second stage problem $\theta_e^{max}(z, l, y)$ and $\theta_u^{max}(l, y)$.

4.4 Equilibrium: Definition, Tractability and Efficiency

An equilibrium consists of a set of market tightness functions for unemployed job-searchers, $(\theta_u, \theta_u^{max})$, a set of market tightness functions for employed workers $(\theta_e, \theta_e^{max})$, a value function for unemployed non-searchers, J_u^{max}, a set of policy functions $(\eta_m,$ $\eta_m^{max})$ for unemployed non-searchers, a value function for unemployed job-searchers, U, a set of policy functions for unemployed job-searchers, $(x_u(l, k, y)$, $r_u(l, k, y)$, $\eta_u(l, k, y)$, $x_u^{max}(l, y)$, $r_u^{max}(l, y)$, $\eta_u^{max}(l, y))$, a joint value function for firm-worker matches, J_e, and a set of policy functions for firm-worker matches, $(x_e(z, l, k, y)$, $r_e(z, l, k, y)$, $\eta_e(z, l, k, y)$, $x_e^{max}(z, l, y)$, $r_e^{max}(z, l, y)$, $\eta_e^{max}(z, l, y)$, $d_e(z, l, y))$. These functions satisfy the following conditions:

i. $(\theta_u(l,k,y),\ \theta_u^{max}(l,y))$ satisfy (12), (5), and (6);

ii. $(\theta_e(z,l,k,y),\ \theta_e^{max}(z,l,y))$ satisfy (12), (10), and (11);

iii. J_u^{max} satisfies (7) and (8), and $(\eta_m(l,k,y),\ \eta_m^{max}(l,y))$ are the associated policy functions;

iv. U satisfies (4), (5), and (6), and $(x_u(l,k,\widehat{y}),\ r_u(l,k,\widehat{y}),\ \eta_u(l,k,\widehat{y}),\ x_u^{max}(l,\widehat{y}),\ r_u^{max}(l,\widehat{y}),\ \eta_u^{max}(l,\widehat{y}),\ d_u(l,y))$ are the associated policy functions;

v. J_e satisfies (12), (10), and (11), and $(x_e(z,l,k,y), r_e(z,l,k,y), \eta_e(z,l,k,y), x_e^{max}(z,l,y), r_e^{max}(z,l,y),\ \eta_e^{max}(z,l,y),\ d_e(z,l,y))$ are the associated policy functions.

These conditions ensure that the strategies of each agent are optimal given the strategies of other agents.

In line with the earlier literature on competitive and directed (non-spatial) search we observe that the equilibrium agents' value functions and policy functions depend on the aggregate state of the economy only through aggregate productivity and not through the distribution of workers across employment states or locations. The equilibrium is therefore block recursive, and computations tractable. Furthermore, the decentralised equilibrium coincides with the solution of the social planner's problem, and is therefore socially efficient. Intuitively, these properties follow from the directedness of job search: a worker self-selects into the submarket that maximises her expected gains from search, by trading off employment probabilities and the value of moving from their current position to a new job/location. In particular, by equation (12), unemployed/low value workers search in submarkets where the probability of entering is high and the gain is low, while high value workers search in submarkets where the probability of entering is low and the gain is high. Firms in a submarket therefore know who they will meet and that their job offer will not be rejected. Therefore, a firm's value from meeting a worker in a particular submarket is independent of the distribution of workers and so is the tightness in this submarket.

We proceed to estimate structurally this search model using individual level data on transitions across labour market states and geographical transitions within and across local labour markets.

5 Empirical Implementation

The model is estimated using our individual level transition data from the LIAB employer-employee panel. In order to maintain the spatial representativeness of this administrative dataset, we consider only the largest 30 TTWAs in West Germany (excluding Berlin). In our baseline model, we abstract from spatial learning frictions, $\alpha_{l,l} = \alpha_{l,k} = 1$, so close and distant jobs are inspection goods. In order to accommodate worker heterogeneity, we stratify workers into groups by skill level, age, and industry. We call a particular skill/age/industry cell a segment, and we have 12 such segments. Our objective is to estimate a parsimonious model. To this end, we do not allow the frictional parameters $(\delta, \lambda_e, \lambda_u)$ and the vacancy posting cost ξ to vary across locations,[8] but they are allowed to vary by segment. The parameters in our amenity and moving costs function are segment invariant. Overall, we will estimate 20 parameters.[9]

- Give details / definitions about stratification. It might also be better to replace Appendix D with a new one for the data at hand: Some summary stats for the 30 TTWAs, perhaps by segment.

5.1 Specifications

As regards the production function, we combine an aggregate component (y), a location-specific component (μ_l for location l), and a match-specific component (z). Location and match specific components are allowed to differ by subindustry. y follows a two-state Markov process with unconditional mean 1. z is given, as in Menzio and Shi (2011), by a discrete approximation of a Weibull distribution with mean μ_z, scale σ_z, and shape ν_z. Its dispersion is set as to reflect the global dispersion of the firm fixed effects. The location-specific component μ_l has been set to the rescaled deviation of the firm TFP in that location for the specific subindustry from the overall mean, i.e. by subindustry $\mu_l = \beta_1 + \beta_2(FFE_l - F\bar{F}E)$ where $F\bar{F}E = K^{-1}\sum_l FFE_l$. We combine these 3 components using a production function given by $\pi(y, \mu_l) + z$, where $\pi = 2(0.5y^{1/2} + 0.5\mu_l^{1/2})^2$. Our production technology combines a CES production function and an additive match specific component, implying substitutability between the match specific component and the aggregate/location specific components, and, at the same time, some complementarity between the aggregate and the location specific components.

Amenities are assumed to be a summary measure of all pull factors that explain population distributions over and above the distributions implied by productivity differences. Specifically, we use the following parametric specification for amenities:

$$A_k^e = \alpha_1 + \alpha_2 \times Dec(pop, k) + \alpha_3 \times \sum_{i \in \{sector\}} [Dec(pop_i, k) - Dec(\mu_i, k)] \tag{13}$$

where $Dec(pop, k)$ is the position of region k (i.e. the decile) in the distribution of employed workers across regions; similarly, $Dec(pop_i, k)$ is is the position of region k (i.e. the decile) in the distribution of sector i workers across regions.[10]

We assume that the flow utility from amenities in location k enjoyed by the unemployed, A_k^u, is a constant share of the corresponding flow utility enjoyed by employed workers, A_k^e, and that home productivity in this location, b_k, is also a constant share of average productivity across all sectors in this location. We set these constant shares

[8] Observe, though, that the job finding probability $\lambda_e p(\theta(.))$ is location dependent because $\theta(.)$ is.

[9] Schmutz and Sidibé (2018) estimate 913 parameters, having reduced the dimensionality of their problem by also parametrising the functions determining moving costs and informational frictions.

[10] Diamond (2015) takes the complementary approach of enumerating explicitly specific dimensions of amenities. By contrast, Kennan and Walker (2011) capture amenities by estimating their model including fixed effects for different locations/regions.

at 50%, a proportion approximately equal to the average unemployment replacement rate in Germany between 2002 and 2008 (DICE Database 2013).

Turning to the moving cost function, we take into account several factors. Moving costs (monetary and psychological) might be a function of the physical distance between origin (l) and potential destination (k), which might also reflect the distinct regional identities and the federal structure of Germany. Also taken into account are difference between housing costs. In order to capture the idea that adjusting to life in a big city is more costly than settling in a smaller place, we also include an indicator for whether two TTWAs l and k include both one of the five largest German cities. We also allow moving costs to differ by segment since some groups may be more mobile than others (e.g. the young age group). Since moving costs (psychological and direct) might differ between employed and unemployed, we also include an indicator for labour market status (a dummy equal to one if the individual is unemployed). In summary, the moving cost function is for current location l and potential destination k:

$$c(l,k) = \left(\alpha^u \times \mathbb{1}_{\{u\}} + 1\right) \times \left\{ \sum_{i \in \{segment\}} \alpha_i^s \times \mathbb{1}_{\{i=segment\}} + \alpha^c + \alpha^{bc} \times \mathbb{1}_{\{\text{big city}\}} + \alpha^{hp} \times \Delta hp(l,k) + \alpha^d \times distance(l,k) \right\} \tag{14}$$

We fix some global parameters at plausible values. The discount factor β is set to 0.984, which corresponds to an annual discount rate of 5%. The elasticity of the matching function γ is set to 0.35, a value similar to the elasticity of the matching function estimated by Kohlbrecher et al. (2016) using the LIAB dataset.

5.2 Estimation Strategy

The model is evaluated, for a given set of parameters, by value function iterations. The parameters themselves are estimated by minimising the distance between a set of empirical moments of our transition data, and the corresponding simulated model-based moments. Specifically, collect all parameters to be estimated in a vector Θ, denote the i$^{\text{th}}$ moment calculated from the data by $\widehat{m}_i$ and the corresponding model-simulated moment by $m_i(\Theta)$. The GMM criterion to be minimised is

$$GMM(\Theta) = \sum_i w_i \left(\frac{\widehat{m}_i - m_i(\Theta)}{\widehat{m}_i} \right)^2 \tag{15}$$

where w_i is a weight (set to unity currently, $w_i \equiv 1$); we standardize by the moments themselves instead of their standard errors (which are less precisely estimated). We take into account the local level of unemployment rates, the transition between labour markets states in each location (i.e. $e \to u$, $e \to e$, $u \to e$) as well as moving rates between any two locations. These moments are stated explicitly in Table B1 of Appendix B. In particular, we use 330 moments: local unemployment (30 moments, since $K = 30$), relocations into each travel to work area (30 moments), local sector-specific job-to-job transitions (3 segments $\times$ 30 moments), local sector-specific job-to-unemployment transitions (3 segments $\times$ 30 moments), and local sector-specific employment shares (3 segments $\times$ 30 moments). These moments are used to estimate the sector-specific frictional parameters (δ, λ_e), the vacancy posting cost ξ, and the parameters of the moving cost and amenities function, a total of $8 + 4 \times S = 20$ parameters where $S = 3$ is the number of sectors. The objective is to estimate a parsimonious empirical model that avoids the danger of over-fitting that arises were the frictional parameters and cells of a moving cost matrix allowed to vary across locations without restrictions.

For a large number of locations the computations are time-consuming. The minimum distance criterion is therefore minimised using an evolutionary / genetic algorithm which is designed to locate the global minimum of the objective function relatively rapidly because computations are parallelised across processors. Computational details as well as validation experiments, demonstrating the good performance of our estimation strategy, are collected in Appendix B.

Statistical inference for Θ in this GMM framework is standard, see e.g. Wooldridge (2002, chapter 14). In practice, in order to estimate the derivative of each moment with respect to each parameter, $\partial m(\Theta)/\partial\Theta$, we follow Lise and Robin (2017), since the moments are not necessarily smooth functions of the parameters: In particular, we compute a partial derivative by simulating the model, taking the parameter in question from a grid centered about the estimated value while keeping all other parameters at their estimated values, computing the new moment function, fitting a high order polynomial for each moment, and estimating the partial derivatives by the derivatives of such fitted polynomials evaluated at $\hat{\Theta}$.

5.3 Identification

We discuss how transition data within and across labour market states and locations and their spatial variation, as well as the observed spatial variation in productivities and populations, identify the parameters of our model. Since the policy functions of the model are not available in closed form, it is impossible to obtain formal classic identification demonstrations. Instead, we follow the literature and consider the behaviour of the estimation objective function in the context of a simulation in order to provide a numerical illustration of identification. Appendix B.4 presents the results of two experiments, which demonstrate that our model is identified and that our estimation algorithm works very well.

In what follows, we consider informally the sources of identification that reflect both incentives to relocate (productivity and amenity gains) and barriers (moving costs). Observed high productivity in location k, a high population share, and a high inflow rate are jointly indicative of sufficiently low moving costs. By contrast, high productivity and low a population share is consistent with both high moving costs and low amenity values, but the two scenarios could be distinguished by sufficiently high population outflow rates. Similarly, observed low local productivity and a high population share is consistent with either high mobility costs or a high local amentiy value, but sufficiently high population inflow rates would point to the latter. In order to discipline moving costs further, we use the convergence of the model to the steady state starting from arbitrary initial population distributions: in the presence of excessively high moving costs that suffocates relocations, the model's steady state distribution would then look very different from the observed population distribution, as would the relocation rates. Job separation and job search probabilities as well as vacancy posting costs are principally identified by transitions within local labour markets. Finally, we observe that our rich transition data contribute complementary sources of identifying varation; in particular, for workers employed in two consecutive periods, we have three subgroups, namely location stayers and firm non-switchers, (location) stayers and switchers, and movers (who therefore are also switchers). As regards the unemployed, we have both stayers and movers.

5.4 Results (Preliminary)

Tables 4 and 3 report the results of our GMM estimation. Table 4 presents the sector-specific estimated model parameters for manufacturing and services. Table 3 reports global (sector-invariant) and sector-specific moments aggregated across travel-to-work-areas. The model fits all sector specific data moments very well, indicating that our strategy of introducing heterogeneity by partitioning the labour market into segments works. Our model-simulated global (non-sector specific) moments are also close to the corresponding data moments indicating that our model captures the main features of the aggregate labour market, and, more importantly, the relocation patterns observed in the data.

The sector-specific data moments suggest that the incidence of labour market transitions in the service sector is higher than in the manufacturing sector: both job-to-job transition and job separation rates in the service sector are approximately two times higher than the corresponding rates observed in the manufacturing sector. This implies larger transition parameter estimates (λ_e, δ) in the service sector: our estimated sector-specific parameters, reported in Table 4, are indeed larger in services than in manufacturing. Vacancy costs are estimated to be lower in the service sector than in manufacturing (approximately 1/3 of manufacturing costs), as one would expect. Finally, the estimated sector specific parameters of the cost function suggest that moving costs are highest for workers in the service sector and lowest in the manufacturing sector.

Table 3 reports that our model fits the data well by comparing global and sector-specific moments aggregated across travel-to-work areas. To illustrate how well our model performs in matching spatial heterogeneity in the data, we present Figure (2), which demonstrates that our model describes the spatial variation of employment in the aggregate. Our model's performance in matching spatial heterogeneity in employment in the different sectors is further evidenced by the high Spearman rank correlations between data- and model-computed local employment shares: 0.61 in manufacturing, 0.58 in services. Figure (3) demonstrates that our model does a very good job in matching aggregate (across sectors) relocation rates into travel-to-work areas: the Spearman rank correlation between data- and model-computed relocation rates is 0.726.

5.5 Counterfactual Experiments (Preliminary)

In Section 5.5.1, we show how a reduction in moving costs can affect worker mobility. Further experiments are in progress.

Table 3: Model Fit

	data	model		
unemployment	2.760%	1.850%		
relocations	0.336%	0.329%		
	Manufacturing		Services	
	data	model	data	model
e $\rightarrow$ e transitions	2.14%	1.74%	4.83%	3.20%
e $\rightarrow$ u transitions	1.06%	1.10%	2.22%	1.91%

Notes: Time unit is a quarter. Based on LIAB for years 2002-08. Estimation by GMM, using $8 \times K = 240$ moments (where $K = 30$ denotes locations), to estimate $8 + 4 \times S = 16$ parameters (where $S = 2$ denotes sectors). Specifically, the global (non-sector specific) estimated parameters in the amenities and moving cost functions are: $\alpha_1 - \alpha_3$ from (13) and α^c, α^{bc}, α^{hp}, α^d, α^u from (14); the remaining 8 sector-specific parameters are presented in Table 4. The GMM criterion (15) is minimised by our evolutionary algorithm (see Appendix B.3).

Table 4: Sector Specific Parameter Estimates

	Manufacturing	Services
ξ	4.7263	1.6324
δ	0.0088	0.0176
λ_e	0.5251	0.7919
α^s	4.0515	6.3535

Notes: Time unit is a quarter. Based on LIAB for years 2002-08. Estimation by GMM, using $8 \times K = 330$ moments (where $K = 30$ denotes locations). The GMM criterion (15) is minimised by our evolutionary algorithm (see Appendix B.3).

5.5.1 Reducing Moving Costs

To illustrate how our model can provide useful insights into the motivations of workers to move across local labour markets, we conduct a counterfactual experiment: using our estimated model parameters, see Table 4, and the corresponding moments, see Table 3, as a benchmark, we gradually decrease moving costs and recalculate the global (sector-invariant) and sector specific moments. The results are reported in Table 5.

A decrease in moving costs by 25% increases relocations by a factor of 1.6. Decreasing moving costs by 75%, further increases relocations. In both cases, unemployment falls in response to these changes, but the decrease is small. Setting moving costs to

Figure 2: Employment across 30 locations

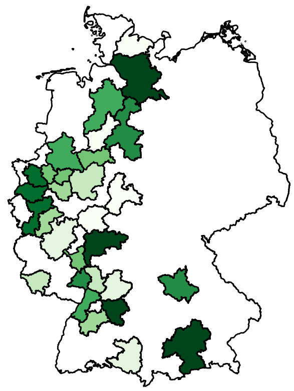

Notes: Local employment in manufacturing, services, and financial services in the LIAB data (left) and in the model (right), arranged into 9 quantile groups. Using LIAB 2005 data and model generated data, we express the population of employed workers in these three sectors in every TTWA as a share of the population of all workers employed in these sectors across the 30 TTWAs.

Figure 3: Relocation rates

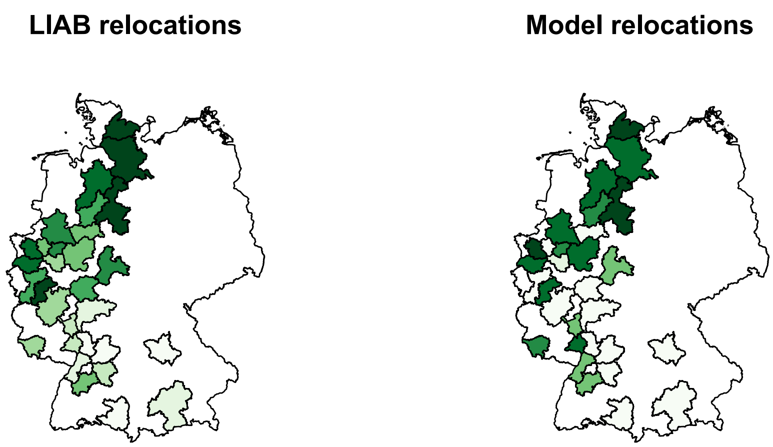

Notes: Relocation rates into TTWAs in the LIAB data (left) and in the model (right).

zero makes relocations shoot up to 3.7% of total spells, suggesting that even moderate moving costs play a significant role in the allocation of workers across local labour markets. The effect of zero moving costs on unemployment is not significant, suggesting that lower moving costs lead to a re-allocation of workers across TTWAs: workers are more likely to move to high-productivity TTWAs even if job-queue lengths are longer.

An important observation is the sectoral heterogeneity in the responses to lower moving costs, captured entirely by varying job-to-job transition rates; job separations are almost unaffected by moving costs. In the service sector, job-to-job transitions respond only to sizeable decreases in moving costs (75%) and more than triple relative to the baseline if moving costs are eliminated. By contrast, the manufacturing sector exhibits the lowest sensitivity to moving costs: job-to-job transition rates remain almost unaltered until moving costs are set to zero.

Table 5: Counterfactual Experiments: Reducing moving costs

	Baseline	75% cost	25% cost	no cost
relocations	0.329%	0.510%	0.674%	3.659%
unemployment	1.850%	1.826%	1.769%	1.651%
Manufacturing				
e $\rightarrow$ e transitions	1.74%	1.75%	1.86%	5.65%
e $\rightarrow$ u transitions	1.10%	1.10%	1.07%	1.12%
Services				
e $\rightarrow$ e transitions	3.20%	3.21%	5.06%	10.15%
e $\rightarrow$ u transitions	1.91%	1.90%	1.93%	1.96%

6 Conclusion

We have built a general equilibrium model of directed search on-the-job, where workers are allowed to search within and across regional labour markets that differ in terms of firms' productivities. Workers' job search yields an equilibrium characterised by a spatial distribution of wages and unemployment, and rich dynamics as workers experience transitions between different labour market states and between regional labour markets.

References

Abowd, J. M., F. Kramarz, and D. N. Margolis. (1999) "High Wage Workers and High Wage Firms," *Econometrica*, 67, 2, 251-334.

Amior, Michael, and Alan Manning. (2016) "The Persistence of Local Joblessness", *CEP working paper*.

Bagger, Jesper, Francois Fontaine, Fabien Postel-Vinay, and Jean-Marc Robin. (2014) "Tenure, experience, human capital, and wages: A tractable equilibrium search model of wage dynamics", *The American Economic Review*, 104, 6, 1551-1596.

Beaudry, Paul, David A. Green, and Benjamin M. Sand.(2012) "Does Industrial Composition Matter for Wages? A Test of Search and Bargaining Theory." *Econometrica*, 80(3): 1063-1104.

Beaudry, Paul, David A. Green, and Benjamin M. Sand. (2014) "Spatial Equilibrium with Unemployment and Wage Bargaining: Theory and Estimation." *Journal of Urban Economics*, 79: 2-19.

Baum-Snow, Nathaniel and Ronni Pavan. (2012) "Understanding the city size wage gap." *Review of Economic Studies*, 79, 1, 88-127.

Card, D., J. Heining, and P. Kline. (2013) "Workplace Heterogeneity and the Rise of West German Wage Inequality," *The Quarterly Journal of Economics*, 128, 3, 967-1015.

Combes, Pierre-Philippe, Gilles Duranton, Laurent Gobillon, Diego Puga, and Sebastien Roux. (2012), "The productivity advantages of large cities: Distinguishing agglomeration from firm selection." *Econometrica*, Vol. 80, No. 6, 2543-2594.

Cornelissen, Thomas, and Olaf Huebler. (2011) "Unobserved Individual and Firm Heterogeneity in Wage and JobDuration Functions: Evidence from German Linked Employer-Employee Data", *German Economic Review*, 12, 4, 469-489.

Diamond, Rebecca. (2016) "The determinants and welfare implications of US workers' diverging location choices by skill: 1980-2000", *American Economic Review*, 106(3), 479-524.

Dustmann, Christian, Johannes Ludsteck, and Uta Schonberg. (2009), "Revisiting the German Wage Structure," *Quarterly Journal of Economics*, 124, 2, 843-881.

Eckey, Hans-Friedrich, Reinhold Kosfeld und Matthias Tuerck. (2006) "Delineation of German labour market regions", *Raumforschung und Raumordnung*, 4, 299-309.

Ehrl, Philipp. (2013) “Agglomeration economies with consistent productivity estimates”, *Regional Science and Urban Economics*, 43, 5, 751-763.

Gould, Eric. (2007) “Cities, Workers and Wages: A Structural Analysis of the Urban Wage Premium.” *Review of Economic Studies*, 74,2, 477-506.

Hall, R.E. and A.B. Krueger. (2010) “Evidence on the determinants of the choice between wage posting and wage bargaining”, NBER Working Paper 16033.

Kennan, John and John R. Walker. (2011) “The effect of expected income on individual migration decisions.” *Econometrica*, 79(1), 211-251.

Klosterhuber, Wolfram, Jorg Heining and Stefan Seth. (2013) “Linked-Employer-Employee-Data from the IAB: LIAB Longitudinal Model 1993-2010 (LIAB LM 9310)”, FDZ-Datenreport 08/2013.

Kohlbrecher, Britta, Christian Merkl, and Daniela Nordmeier. (2016) “Revisiting the matching function”, *Journal of Economic Dynamics and Control*, 69, 350-374.

Lehmer, Florian, and Joachim Moeller. (2010) “Interrelations between the urban wage premium and firm-size wage differentials: a microdata cohort analysis for Germany”, *The Annals of Regional Science*, 45, 1, 31-53.

Lise, Jeremy, and Jean-Marc Robin. (2017). “The Macrodynamics of Sorting between Workers and Firms”, *American Economic Review*, 107(4): 1104-35.

Menzio, Guido, and Shouyong Shi. (2011) “Efficient Search on the Job and the Business Cycle.” *Journal of Political Economy*, 119, 3, 468-510.

Mion, Giordano, and Paolo Naticchioni. (2009) “The spatial sorting and matching of skills and firms”, *Canadian Journal of Economics*, 42, 1, 28-55.

Molloy, R. , C. L. Smith, and A. Wozniak. (2014) “Declining Migration within the US: The Role of the Labor Market.” Finance and Economics Discussion Series, Federal Reserve Board, Washington, D.C.

Moretti, Enrico. (2011) Local Labor Markets. In *Handbook of Labor Economics.* Vol. 4B, ed. David Card and Orley Ashenfelter, 12371313. New York: Elsevier.

OECD, (2005) “How Persistent Are Regional Disparities in Employment? The Role of Geographic Mobility”, *OECD Employment Outlook.*

Olley, G. Steven, and Ariel Pakes. (1996) “The dynamics of productivity in the telecommunications equipment industry”, *Econometrica*, 64, 6, 1263.

Overman, Henry G., and Diego Puga. (2002) Unemployment Clusters across Europes Regions and Countries. *Economic Policy*, 17, 34, 115-148.

Rupert, Peter and Etienne Wasmer. (2012) “Housing and the labor market: Time to move and aggregate unemployment.” *Journal of Monetary Economics*, 59, 24-36.

Schmutz, Benoit, and Modibo Sidibe. (2018) “Frictional Labor Mobility.” *The Review of Economic Studies*, .

Shi, Shouyong. (2009) “Directed Search for Equilibrium Wage-Tenure Contracts.” *Econometrica*, 77, 561-84.

Appendices

The structure of this Appendix is as follows:

- Appendix A presents in detail the Social Planner's Problem. We state the Planner's value function, and demonstrate that it can be decomposed into a set of smaller problems, labelled the component value functions. We then characterise the solution to the Planner's problem that enables us to establish who moves, when and where.
- In Appendix B we provide details of our estimation algorithm and report the results of a validation exercise.
- In Appendix C we provide details of our estimation of firm productivities.
- In Data Appendix D, we illustrate in some greater detail the heterogeneity among local labour markets and the geographic mobility among them by focussing on 8 selected travel-to-work areas.

A The Social Planner's Problem

At the beginning of every period, the social planner observes the aggregate state of the economy $\psi = (y, u, g)$. Births and deaths occur exogenously. At the separation stage, the planner chooses the probability $d_e(z, l)$ of destroying a match with productivity z in location l, $Z \times K \to [\delta, 1]$, and the probability $d_u(l)$ with which an unemployed job-searcher in location l stops searching for a job, $K \to [0, 1]$.

The social planner takes decisions in two steps. In the first step, the planner makes pairwise comparisons between an individual's current value in her location/employment state/submarket and her value in all possible destination locations/employment states/submarkets. In the second step, the planner chooses the destination location/employment state/submarket that maximises the individual's value. To maintain notational transparency, we denote the policy function chosen by the planner in the second step using a *max* superscript. Therefore, at the search stage, the social planner makes the following choices in two steps:

- the planner chooses the probability with which an unemployed non-searcher in l would relocate to any possible destination location $k \in K$, $\eta_m(l, k)$: $K \times K \to [0, 1]$; given this set of choices, the planner chooses the probability that an unemployed non-searcher in location l moves to a different location, $\eta_m^{max}(l)$: $K \to [0, 1]$;
- the planner chooses the probability with which an unemployed job-seeker in location l would search for a job in any possible destination location $k \in K$, $\eta_u(l, k)$: $K \times K \to [0, 1]$; given this set of choices, the planner chooses the probability that an unemployed job-seeker in l searches in a different location, $\eta_u^{max}(l)$: $K \to [0, 1]$;
- the planner chooses the probability with which a job-seeker currently employed in a match of productivity z in location l, would search for a job in any possible destination location $k \in K$, $\eta_e(z, l, k)$: $Z \times K \times K \to [0, 1]$; given this set of choices, the planner chooses the probability that a worker in a match with productivity z in location l searches for a job in a different location, $\eta_e^{max}(z, l)$: $Z \times K \to [0, 1]$;
- for unemployed job-seekers in location l, the planner chooses the tightness at any possible destination submarket $k \in K$, $\theta_u(l, k)$: $K \times K \to \mathbb{R}_+$; given this set of choices, the planner chooses the tightness at the submarket where unemployed job-seekers in l look for a match, $\theta_u^{max}(l)$: $K \to \mathbb{R}_+$;
- for job-seekers employed in matches of productivity z in location l, the planner chooses the tightness at any possible destination submarket $k \in K$, $\theta_e(z, l, k)$: $Z \times K \times K \to \mathbb{R}_+$; given this set of choices, the planner chooses the tightness at the submarket where workers employed in matches of productivity z in location l search for a job, $\theta_e^{max}(z, l)$: $Z \times K \to \mathbb{R}_+$.

At the matching stage, the social planner makes the following choices in two steps:

- the planner chooses the probability with which a meeting between an unemployed job-seeker in location l and a firm in any possible destination location $k \in K$ is turned into a match, given the signal s, $h_u(s, l, k)$: $Z \times K \times K \to [0, 1]$; given this set of choices, the planner chooses the probability that an unemployed job-seeker in l will match with a firm conditional on s, $h_u^{max}(s, l)$: $Z \times K \to [0, 1]$;

- the planner chooses the probability with which a meeting between a worker employed in a match with productivity z in location l and a firm in any possible destination location $k \in K$ is turned into a match given the signal s, $h_e(s, z, l, k)$: $Z \times Z \times K \times K \to [0, 1]$; given this set of choices, the planner chooses the probability that a worker employed in a match of productivity z in location l will match with a firm conditional on s: $h_e^{max}(s, z, l)$: $Z \times K \times K \to [0, 1]$.

Given the choices of the social planner, $\Omega = \{d_e,\, d_u,\, \eta_m^{max},\, \eta_u^{max},\, \eta_e^{max},\, \theta_u^{max},\, \theta_e^{max},\, h_u^{max},\, h_e^{max}\}$, aggregate consumption is given by total production minus relocation costs and search costs:

$$
\begin{aligned}
F(\Omega|\psi) = & \sum_{k \in K} \{b_k \, \widehat{u}_k\} + \sum_{k \in K} \sum_{z \in Z} \{(\pi(y, \mu_k) + z) \, \widehat{g}(z, k)\} \\
& + \sum_{k \in K} \left\{ A_k^u \times \widehat{u}_k + A_k^e \times \sum_{z \in Z} \widehat{g}(z, k) \right\} \\
& - \sum_{k \in K} \left\{ \eta_m^{max}(k) \left(\frac{\tau}{N(k)} + d_u(k) u_k + \sum_{z \in Z} [d_e(z, k) g(z, k)] \right) \times c_u(k, k^*) \right\} \\
& - (1 - \tau) \lambda_u \sum_{k \in K} \Bigg\{ \eta_u^{max}(k)(1 - d_u(k)) p(\theta_u^{max}(k)) \times \\
& \qquad \times \mathbb{E}_s[\alpha_{kk^*} h_u^{max}(s, k) + (1 - \alpha_{kk^*}) m_u^{max}(k)](u_k - \widehat{ns}_k) c_u(k, k^*) \Bigg\} \\
& - (1 - \tau) \lambda_e \sum_{k \in K} \sum_{z \in Z} \Bigg\{ \eta_e^{max}(z, k) \, [1 - d_e(z, k)] \, p(\theta_e^{max}(z, k)) \times \\
& \qquad \times \mathbb{E}_s[\alpha_{kk^*} h_e^{max}(s, z, k) + (1 - \alpha_{kk^*}) m_e^{max}(z, k)] g(z, k) c_e(k, k^*) \Bigg\} \\
& - (1 - \tau) \xi \lambda_u \sum_{k \in K} \{[(1 - d_u(k)) \theta_u^{max}(k)] \, (u_k - \widehat{ns}_k)\} \\
& - (1 - \tau) \xi \lambda_e \sum_{k \in K} \sum_{z \in Z} \{[1 - d_e(z, k)] \; \theta_e^{max}(z, k) g(z, k)\},
\end{aligned}
\tag{A1}
$$

where k^* denotes the destination location for any source location k, and $\widehat{u}_k$, $\widehat{g}(z, k)$ denote the distribution of individuals across employment states, locations, and submarkets at the production stage and at the beginning of the next period.

A.1 The Social Planner's Value Function

The social planner maximises the sum of current and future aggregate consumption discounted at the factor β. Hence, the planner's value function, $W(\psi)$, solves the following Bellman equation:

$$W(\psi) = \max_{\Omega} \{F(\Omega|\psi) + \beta \mathbb{E} W(\widehat{\psi})\} \tag{A2}$$

subject to (1), (2), (3), and

$$
\begin{array}{lll}
d_e : Z \times K \to [\delta, 1], & d_u(l) : K \to [0,1], & \eta_m^{max}(l) : K \to [0,1], \\
\eta_u^{max}(l) : K \to [0,1], & \eta_e^{max}(z,l) : Z \times K \to [0,1], & \theta_u^{max}(l) : K \to \mathbb{R}_+, \\
\theta_e^{max}(z,l) : Z \times K \to \mathbb{R}_+, & h_u^{max}(s,l) : Z \times K \to [0,1], & h_e^{max}(s,z,l) : Z \times K \times K \to [0,1]
\end{array}
$$

A.1.1 Separability of the Social Planner's Problem

The social planner's value function, $W(\psi)$, depends on the aggregate productivity, y, the measure of workers who are unemployed across $N(k)$ locations, u, and the measure of workers who are employed in $N(z)$ submarkets across $N(k)$ locations, g. Directed search (and the self-selection it implies) enables this decomposition of the Planner's problem into worker-specific problems.

Consider the planner's value function $W(\psi)$, which solves (A2); it is possible to express $W(\psi)$ as follows:

$$
\begin{aligned}
W(\psi) = \sum_k \{Q_u^{max}(k,y) \times ns_k\} + \sum_k \{W_u(k,y) \times (u_k - ns_k)\} \\
+ \sum_k \sum_z \{W_e(z,k,y) g(z,k)\}, \qquad \text{(A3)}
\end{aligned}
$$

where $Q_u^{max}(k,y)$, $W_u(k,y)$, $W_e(z,k,y)$ are the component value functions for the unemployed non-searchers in location k, the unemployed job-seekers in location k, and the workers employed in matches of productivity z in location k, respectively.

A.1.2 Component Value Functions

Inspection of (A3) suggests that the social planner's value function, $W(\psi)$, is linear in u and g. This implies that the social planner's problem is equivalent to solving $(N(z)+2) \times N(k)$ smaller problems, each one of which is associated with workers in a particular submarket, and/or employment state, and/or location. The planner's problem is equivalent to the optimisation of the following component value functions subject to the constraints given in equation (A2).

The component value function for the unemployed is:[11]

$$
W_u(l,y) = \max_{d_u}\{d_u \times Q_u^{max}(l,y) + (1-d_u) \times S_u^{max}(l,y)\} \qquad \text{(A4)}
$$

where $S_u^{max}(l,y)$ is the component value function for the unemployed job-seekers in location l and $Q_u^{max}(l,y)$ is the component value function for the unemployed non-searchers. $S_u^{max}(l,y)$ solves

$$
S_u^{max}(l,y) = \max\left\{S_u(l,1,y), S_u(l,2,y), \ldots, S_u(l,K),y)\right\}, \qquad \text{(A5)}
$$

[11] To keep the notation manageable, $\{\theta_u, h_u, m_u\}$ is used to denote both $\{\theta_u(l,k), h_u(l,k), m_u(l,k)\}$ and $\{\theta_u(l,l), h_u(l,l), m_u(l,l)\}$.

where $S_u(l,k,y)$ is the value of an unemployed job-seeker in l searching for a match in $k \in K$

$$\begin{aligned} S_u(l,k,y) = \max_{\eta_u,\theta_u,h_u} \{ & -(1-\eta_u)\,\xi\lambda_u\theta_u - \eta_u\,\xi\lambda_u\theta_u \\ & + [1-\eta_u]([1-\lambda_u p(\theta_u) m_u][b_l + A_l^u + \beta \mathbb{E} W_u(l,\widehat{y})] \\ & + \lambda_u p(\theta_u)\mathbb{E}_s[\alpha_{ll} h_u(s) + (1-\alpha_{ll}) m_u][\pi(y,\mu_l) + s + A_l^e + \beta \mathbb{E} W_e(s,l,\widehat{y})]) \\ & + \eta_u([1-\lambda_u p(\theta_u) m_u][b_l + A_l^u + \beta \mathbb{E} W_u(l,\widehat{y})] \\ & + \lambda_u p(\theta_u)\mathbb{E}_s[\alpha_{lk} h_u(s) + (1-\alpha_{lk}) m_u][\pi(y,\mu_k) + s + A_k^e - c_u(l,k) + \beta \mathbb{E} W_e(s,k,\widehat{y})]) \\ & \} \end{aligned} \tag{A6}$$

Similarly, $Q_u^{max}(l,y)$ solves

$$Q_u^{max}(l,y) = \max\{Q_u(l,1,y), Q_u(l,2,y), \ldots, Q_u(l,K,y)\}, \tag{A7}$$

where $Q_u(l,k,y)$ is the value of an unemployed non-searcher in l who examines the possibility of relocating to $k \in K$

$$\begin{aligned} Q_u(l,k,y) = \max_{\eta_m}\{ & (1-\eta_m)(b_l + A_l^u + \beta \mathbb{E} W_u(l,\widehat{y})) \\ & + \eta_m(b_k + A_k^u - c_u(l,k) + \beta \mathbb{E} W_u(k,\widehat{y}))\} \end{aligned} \tag{A8}$$

The component value function for the employed is:[12]

$$W_e(z,l,y) = \max_{d_e}\{d_e \times Q_u^{max}(l,y) + (1-d_e) \times S_e^{max}(z,l,y)\}, \tag{A9}$$

where $S_e^{max}(z,l,y)$ is the component value function for workers employed in matches of productivity z in location l and $Q_u^{max}(l,y)$ is the component value function for the unemployed non-searchers, given by (A7). $S_e^{max}(z,l,y)$ solves

$$S_e^{max}(z,l,y) = \max\{S_e(z,l,1,y), S_e(z,l,2,y), \ldots, S_e(z,l,K,y)\}, \tag{A10}$$

where $S_e(z,l,k,y)$ is the value of a worker employed in a match of productivity z in l searching for a match in $k \in K$

$$\begin{aligned} S_e(z,l,k,y) = \max_{\eta_e,\theta_e,h_e} \{ & \\ & -[1-\eta_e(z)][1-d_e]\xi\lambda_e\theta_e - \eta_e(z)[1-d_e]\xi\lambda_e\theta_e \\ & + [1-\eta_e(z)][1-d_e][1-\lambda_e p(\theta_e) m_e][\pi(y,\mu_l) + z + A_l^e + \beta \mathbb{E} W_e(z,l,\widehat{y})] \\ & + [1-\eta_e(z)][1-d_e]\lambda_e p(\theta_e)\mathbb{E}_s[\alpha_{ll} h_e(s) + (1-\alpha_{ll}) m_e][\pi(y,\mu_l) + s + A_l^e + \beta \mathbb{E} W_e(s,l,\widehat{y})] \\ & + \eta_e(z)[1-d_e][1-\lambda_e p(\theta_e) m_e][\pi(y,\mu_l) + z + A_l^e + \beta \mathbb{E} W_e(z,l,\widehat{y})] \\ & + \eta_e(z)[1-d_e]\lambda_e p(\theta_e)\mathbb{E}_s[\alpha_{lk} h_e(s) + (1-\alpha_{lk}) m_e][\pi(y,\mu_k) + s + A_k^e - c_e(l,k) + \beta \mathbb{E} W_e(s,k,\widehat{y})] \\ & \} \end{aligned} \tag{A11}$$

[12] As in the previous footnote, $\{\theta_e, h_e, m_e\}$ is used to denote both $\{\theta_e(z,l,k,y), h_e(z,l,k,y), m_e(z,l,k,y)\}$ and $\{\theta_e(z,l,l,y), h_e(z,l,l,y), m_e(z,l,l,y)\}$.

A.2 Who Moves, When and Where? The Solution to the Social Planner's Problem

The planner's problem can be decomposed into worker-specific problems that depend only on the aggregate productivity because the search process is directed. In this section, we provide a description of the solution. The planner solves the following $(N(z)+2) \times N(k)$ problems, each one of which corresponds to individuals in a particular location, employment state, and submarket.

A.2.1 Unemployed Non-Searchers

There are $N(k)$ problems for unemployed non-searchers. The planner solves each one of these problems in two steps. First, conditional on the non-searchers' current location, l, the planner makes pairwise comparisons of non-searchers' lifetime utility in l and their corresponding utility in every possible destination location k (accounting for moving costs), and chooses $\eta_m^*(l,k,y)$, which determines whether an unemployed non-searcher is better-off staying in her current location or moving to location k. In particular, the efficient choice of $\eta_m(l,k,y)$ is $\eta_m^*(l,k,y)=1$ if

$$b_l + A_l^u + \beta \mathbb{E} W_u(l,\widehat{y}) \leq b_k + A_k^u + \beta \mathbb{E} W_u(k,\widehat{y}) - c_u(l,k) \tag{A12}$$

and $\eta_m^*(l,k,y)=0$ otherwise.

Having made all possible pairwise comparisons between the current location and destination locations, the planner then chooses $\eta_m^{max}(l,y)=\eta_m^*(l,k^*,y)$ where $k^* \in K$ is the destination location that maximises the present value of the lifetime utility of unemployed non-searchers, as given by equation (A7).

A.2.2 Unemployed Job-Seekers

There are $N(k)$ problems for unemployed job-seekers. The planner solves each one of these problems in three steps. First, conditional on the unemployed job-seekers' current location, l, the planner chooses $d_u^*(l,y)$, which determines whether job-seekers are better-off stopping their search and becoming non-searchers or continuing their search for matches. Specifically, the efficient choice of $d_u(l,y)$ is $d_u^*(l,y)=1$ if

$$Q_u^{max}(l,y) \geq S_u^{max}(l,y) \tag{A13}$$

and $d_u^*(z,l,y)=\delta$ otherwise (see (A4)).

In the second step the planner makes pairwise comparisons of job seekers' lifetime utility in l and their corresponding utility in every possible destination location k (accounting for search and moving costs), and chooses $\eta_u^*(l,k,y)$, $\theta_u^*(l,k,y)$, and $h_u^*(s,l,k,y)$, which determine whether an unemployed job-seeker is better-off searching for a match in her current location or in location k. In particular, the efficient choice of $\eta_u(l,k,y)$ is $\eta_u^*(l,k,y) = 1$ if

$$
\begin{aligned}
&-\xi\lambda_u\theta_u^*(l,l,y) + [1-\lambda_u p(\theta_u^*(l,l,y))m_u^*][b_l + A_l^u + \beta\mathbb{E}W_u(l,\widehat{y})] \\
&+ \lambda_u p(\theta_u^*(l,l,y))\mathbb{E}_s[\alpha_{ll}h_u^*(s) + (1-\alpha_{ll})m_u^*][\pi(y,\mu_l) + s + A_l^e \\
&+ \beta\mathbb{E}W_e(s,l,\widehat{y})] \leq \\
&-\xi\lambda_u\theta_u^*(l,k,y) + [1-\lambda_u p(\theta_u^*(l,k,y))m_u^*][b_l + A_l^u + \beta\mathbb{E}W_u(l,\widehat{y})] \\
&+ \lambda_u p(\theta_u^*(l,k,y))\mathbb{E}_s[\alpha_{lk}h_u^*(s) + (1-\alpha_{lk})m_u^*][\pi(y,\mu_k) + s + A_k^e - c_u(l,k) \\
&+ \beta\mathbb{E}W_e(s,k,\widehat{y})]
\end{aligned}
\tag{A14}
$$

and $\eta_u^*(l,k,y) = 0$ otherwise (see (A6)).
The efficient choice of $\theta_u(l,k,y)$ solves:

$$
\begin{aligned}
\xi \geq p'(\theta_u^*(l,k,y)) \sum_{s\geq r_u^*(l,k)} (&\alpha_{lk}\{\pi(y,\mu_k) + s + A_k^e - b_l - A_l^u - c_u(l,k) \\
&\quad + \beta\mathbb{E}[W_e(s,k,\widehat{y}) - W_u(l,\widehat{y})]\} + \\
&+ (1-\alpha_{lk})\mathbb{E}_{z'}\{\pi(y,\mu_k) + z' + A_k^e - b_l - A_l^u - c_u(l,k) \\
&\quad + \beta\mathbb{E}[W_e(z',k,\widehat{y}) - W_u(l,\widehat{y})]\})f(s)
\end{aligned}
\tag{A15}
$$

Finally, the efficient choice of $h_u(s,l,k,y)$ is $h_u^*(s,l,k,y) = 1$ if

$$
\begin{aligned}
b_l + A_l^u + \beta\mathbb{E}W_u(l,\widehat{y})] + c_u(l,k) \leq \alpha_{lk}[\pi(y,\mu_k) + s + A_k^e + \beta\mathbb{E}W_e(s,k,\widehat{y})] + \\
+ (1-\alpha_{lk})\mathbb{E}_{z'}[\pi(y,\mu_k) + z' + A_k^e + \beta\mathbb{E}W_e(z',k,\widehat{y})]
\end{aligned}
\tag{A16}
$$

and $h_u^*(s,l,k,y) = 0$ otherwise

Having made all possible pairwise comparisons between the current location and destination locations, the planner then chooses in step 3 $\eta_u^{max}(l,y) = \eta_u^*(l,k^*,y)$, $\theta_u^{max}(l,y) = \theta_u^*(l,k^*,y)$, and $h_u^{max}(s,l,y) = h_u^*(s,l,k^*,y)$ where $k^* \in K$ is the destination location that maximises the present value of the lifetime utility of unemployed job-seekers (as given by equation (A5)).

A.2.3 Employed Workers

There are $N(z) \times N(k)$ problems for employed workers. The planner solves each one of these problems in three steps. First, conditional on the idiosyncratic productivity of the workers' current match, z, and conditional on the workers' current location, l, the planner chooses $d_e^*(z,l,y)$, which determines whether workers are better-off separating or remaining employed in this type of match. Specifically, the efficient choice of $d_e(z,l,y)$ is $d_e^*(z,l,y) = 1$ if

$$
Q_u^{max}(l,y) \geq S_e^{max}(z,l,y) \tag{A17}
$$

and $d_e^*(z,l,y) = \delta$ otherwise. (See (A9))

In the second step the planner makes pairwise comparisons of workers' lifetime utility in l and their corresponding utility in every possible destination location k (accounting for search and moving costs), and chooses $\eta_e^*(z,l,k,y)$, $\theta_e^*(z,l,k,y)$, and $h_e^*(s,z,l,k,y)$, which determine whether a worker employed in a match of productivity z is better-off searching for a match in her current location, l, or in location k.

In particular, the efficient choice of $\eta_e(z,l,k,y)$ is $\eta_e^*(z,l,k,y) = 1$ if

$$
\begin{aligned}
& -\xi\lambda_e\theta_e^*(z,l,l,y) + [1-\lambda_e p(\theta_e^*(z,l,l,y))m_e^*][\pi(y,\mu_l) + z + A_l^e + \beta\mathbb{E}W_e(z,l,\widehat{y})] \\
& + \lambda_e p(\theta_e^*(z,l,l,y))\mathbb{E}_s[\alpha_{ll}h_e^*(s) + (1-\alpha_{ll})m_e^*][\pi(y,\mu_l) + s + A_l^e \\
& + \beta\mathbb{E}W_e(s,l,\widehat{y})] \leq \\
& -\xi\lambda_e\theta_e^*(z,l,k,y) + [1-\lambda_e p(\theta_e^*(z,l,k,y))m_e^*][\pi(y,\mu_l) + z + A_l^e + \beta\mathbb{E}W_e(z,l,\widehat{y})] \\
& + \lambda_e p(\theta_e^*(z,l,k,y))\mathbb{E}_s[\alpha_{lk}h_e^*(s) + (1-\alpha_{lk})m_e^*][\pi(y,\mu_k) + s + A_k^e - c_e(l,k) \\
& + \beta\mathbb{E}W_e(s,k,\widehat{y})]
\end{aligned}
\tag{A18}
$$

and $\eta_e^*(z,l,k,y) = 0$ otherwise (see (A11)).

The efficient choice of $\theta_e(z,l,k,y)$ solves:

$$
\begin{aligned}
\xi \geq p'(\theta_e^*(z,l,k,y)) \times \sum_{s \geq r_e^*(z,l,k,y)} & (\alpha_{lk}\{\pi(y,\mu_k) - \pi(y,\mu_l) + s - z + A_k^e - A_l^e - c_e(l,k) \\
& + \beta\mathbb{E}[W_e(s,k,\widehat{y}) - W_e(z,l,\widehat{y})]\} \\
& + (1-\alpha_{lk})\mathbb{E}_{z'}\{\pi(y,\mu_k) - \pi(y,\mu_l) + z' - z + A_k^e - A_l^e - c_e(l,k) \\
& + \beta\mathbb{E}[W_e(z',k,\widehat{y}) - W_e(z,l,\widehat{y})]\}\})f(s)
\end{aligned}
\tag{A19}
$$

Finally, the efficient choice of $h_e(s,z,l,k,y)$ is $h_e^*(s,z,l,k,y) = 1$ if

$$
\begin{aligned}
\pi(y,\mu_l) + z + A_l^e + \beta\mathbb{E}W_e(z,l,\widehat{y}) + c_e(l,k) \leq\ & \alpha_{lk}[\pi(y,\mu_k) + s + A_k^e + \beta\mathbb{E}W_e(s,k,\widehat{y})] \\
& + (1-\alpha_{lk})\mathbb{E}_{z'}[\pi(y,\mu_k) + z' + A_k + \beta\mathbb{E}W_e(z',k,\widehat{y})]
\end{aligned}
\tag{A20}
$$

and $h_e^*(s,z,l,k,y) = 0$ otherwise.

Having made all possible pairwise comparisons between the current location and destination locations, the planner then chooses in step 3 $\eta_e^{max}(z,l,y) = \eta_e^*(z,l,k^*,y)$, $\theta_e^{max}(z,l,y) = \theta_e^*(z,l,k^*,y)$, and $h_e^{max}(s,z,l,y) = h_e^*(s,z,l,k^*,y)$ where $k^* \in K$ is the destination location that maximises the present value of the lifetime utility of employed job-seekers (as given by equation (A10)).

B Computational Details and Validation Studies

B.1 Moments used in the GMM criterion

The GMM objective function, see (13), calculates the distance between data moments and model generated moments. Table B1 presents the data and model simulated moments considered in our estimation. Labour market transitions are considered by segment and local labour market, while regional transitions and unemployment are aggregated across segments at the local labour market level.[13]

Table B1: GMM moments

data moments	model-based moments	
local $e^{sec} \rightarrow e^{sec}$ transition rate	ee_l^{sec}	see eq. (B1)
local $e^{sec} \rightarrow u$ transition rate	eu_l^{sec}	see eq. (B2)
local sectoral employment share	g_l^{sec}	see eq. (B3)
local unemployment	$unemp_l$	see eq. (B4)
relocations into l	$reloc_l$	see eq. (B5)

Notes. Data moments based on LIAB 2002-2008, time-averaged. Subscript l denotes (destination) travel to work area, superscript sec denotes sector.

To calculate the job-to-job transition rate, we examine flows into matches of productivity z in segment sec at TTWA l between the beginning of the period and the production stage. The job-to-job transition rate is then determined by dividing the total number of flows into $z \forall z \in Z$ by total employment in segment sec at TTWA l at the beginning of the period, see (B1): for each segment, we first compute the number of workers who, at the beginning of the period, were employed in a match of productivity z' in location l and moved into a match of productivity z in TTWA l, and the number of workers who, at the beginning of the period, were employed in a match of productivity z' in location k' and moved into a match of productivity z in TTWA l; we then aggregate across productivity levels ($\forall z \in Z$) and divide by the local sectoral employment level at the beginning of the period.

$$\begin{aligned} ee_l^{sec} = & \frac{1}{\sum\limits_{z \in Z} g(z,l)} \times (1-\tau) \times \sum_{z \in Z} \Bigg\{ \sum_{z' \in Z} \Bigg\{ g(z',l)[1-d_e^*(z',l)] \times \\ & \qquad \times [1-\eta_e^{max}(z',l)] \lambda_e p(\theta_e^{max}(z',l)) h_e^{max}(z,z',l) f(z) \Bigg\} + \\ & \quad + \sum_{k' \in K} \Bigg\{ \sum_{z' \in Z} \Bigg\{ g(z',k')[1-d_e^*(z',k')] \eta_e^{max}(z',k',l) \times \\ & \qquad \times \lambda_e p(\theta_e^{max}(z',k',l)) h_e^{max}(z,z',k') f(z) \Bigg\} \Bigg\} \quad \Bigg\} \end{aligned} \tag{B1}$$

[13] To keep the notation manageable, we suppress the sectoral superscript from the right-hand-sides of equations (B1)-(B6).

Similarly, to calculate the job-to-unemployment transition rate, we compute all the job separations in segment sec at TTWA l that occurred between the beginning of the period and the production stage, and divide by the local sectoral employment level at the beginning of the period.

$$eu_l^{sec} = \frac{1}{\sum_{z \in Z} g(z,l)} \times \left\{ \tau \times \sum_{z \in Z} g(z',l) + (1-\tau) \times \sum_{z \in Z} \left\{ d_e^*(z,l) g(z,l) \right\} \right\} \tag{B2}$$

Local employment in every sector is expressed as a share of the total employment in this sector across all local labour markets

$$g_l^{sec} = \frac{\sum_{z \in Z} g(z,l)}{\sum_{k \in K} \sum_{z \in Z} g(z,k)} \tag{B3}$$

The remaining moments used in our GMM objective function are aggregated across segments at the TTWA level. We consider local unemployment levels at TTWA l and relocation flows into l at the production stage of the period. Local unemployment is given by

$$unemp_l = \sum_{sec} u_l^{sec} \tag{B4}$$

where sec denotes segment, and u_l^{sec} is given by (2).[14]

Similarly, relocation flows into l are given by

$$reloc_l = \sum_{sec} reloc_l^{sec} \tag{B5}$$

where $reloc_l^{sec}$ denotes relocation flows into l for segment sec. To calculate segment specific relocation flows into l, we first compute the number of regional flows into matches of productivity z in TTWA l aggregated across origin labour market state (unemployment or employment in a match of productivity z'), and also aggregated across origin TTWA ($\forall k' \in K_{-l}$); we then aggregate across productivity levels ($\forall z \in Z$), and add regional flows into the unemployment pool of TTWA l, which include

[14] Note that in the presentation of our model, we only consider one segment, so (2) presents u_l as total unemployment in travel to work area l.

unemployed non-searchers or new entrants.

$$
\begin{aligned}
reloc_l^{sec} =& (1-\tau) \times \sum_{z\in Z} \Bigg\{ \sum_{k'\in K} \Bigg\{ u_{k'}[1-d_u(k')]\eta_m^{max}(k')\lambda_u p(\theta_u^{max}(k')) \times h_u^{max}(z,k')f(z) + \\
& \qquad + \sum_{z'\in Z} \Bigg\{ g(z',k')[1-d_e^*(z',k')]\eta_e^{max}(z',k',l) \times \\
& \qquad \lambda_e p(\theta_e^{max}(z',k',l))h_e^{max}(z,z',k')f(z) \Bigg\} \quad \Bigg\}\Bigg\} + \\
& + (1-\tau) \times \Bigg\{ \sum_{k'\in K} \eta_m^{max}(k',l) \times \Bigg(d_u(k') \times u_{k'} + \sum_{z\in Z}[d_e(z,k')g(z,k')] \Bigg) \Bigg\} + \\
& + \frac{\tau}{N(k)} \times \Bigg\{ [1-\eta_m^{max}(l)] + \sum_{k'\in K} \eta_m^{max}(k',l) \Bigg\} \qquad \text{(B6)}
\end{aligned}
$$

B.2 Statistical inference

Statistical inference for the vector of parameters in this GMM framework is standard, see e.g. Wooldridge (2002, chapter 14). To be precise, let $\hat{m}$ denote the $N \times 1$ vector of data moments, and $m(\Theta)$ the corresponding vector of theoretical moments, where all parameters are collected in the $P \times 1$ vector Θ. Under standard regularity conditions, assuming that the moment vector follows asymptotically a normal law, $\hat{m} \sim N(m(\Theta_0), \Sigma)$ where Θ_0 denotes the population value, then

$$\hat{\Theta} \sim N(\Theta_0, \hat{J}^{-1}\hat{I}\hat{J}^{-1}),$$

with $\hat{J} = \hat{M}^t\hat{\Omega}\hat{M}$, $\hat{\Omega} = \text{diag}(w_1/\hat{m}_1^2, \ldots, w_N/\hat{m}_N^2)$, $\hat{M} = \partial m(\hat{\Theta})/\partial\Theta^t$, and $\hat{I} = \hat{M}^t\hat{\Omega}\hat{\Sigma}\hat{\Omega}\hat{M}$. The variance estimator of the moment function is $\hat{\Sigma} = (\hat{m} - m(\hat{\Theta}))(\hat{m} - m(\hat{\Theta}))^t$.

In practice, in order to estimate the derivative of each moment with respect to each parameter, $\partial m(\Theta)/\partial\Theta$, we follow Lise and Robin (2017), since the moments are not necessarily smooth functions of the parameters: In particular, we compute a partial derivative by simulating the model, taking the parameter in question from a grid centered about the estimated value while keeping all other parameters at their estimated values, computing the new moment function, fitting a high order polynomial for each moment, and estimating the partial derivatives by the derivatives of such fitted polynomials evaluated at $\hat{\Theta}$.

B.3 Computational Details: An Evolutionary Algorithm

In the empirical implementation of the model, we segment the economy by worker characteristics (i.e. age and skill level) and by industry in order to accommodate better worker and firm heterogeneity. Our estimation algorithm reflects that our objective function depends on a set of global and a set of segment-specific parameters. This is achieved by switching successively between estimating segment-specific parameters given global parameters, and then estimating global parameters given the estimates of the segment-specific parameters. Each such inner and outer loop features a minimisation of the GMM criterion, see (13), which is a function of the distance between the empirical and simulated model-based moments. We use the moments presented in Table (B1).

In principle, the minimisation problem could be solved by an application of the Nelder Mead algorithm. In practice, this is very slow given the high dimensionality of our set-up, so we develop an evolutionary optimisation algorithm to estimate our model. The comparative strength of our algorithm is that it is paralleliseable.[15] Our evolutionary optimisation algorithm is as follows:

- step (i) generate a "population", equal to N, of randomly drawn parameter vectors within the lower and upper parameter bounds;
- step (ii) evaluate the objective function N times and sort parameter vectors by the corresponding "fitness" (i.e. objective function) value;
- step (iii) pick the S (where $S < N$) best fitness parameter vectors and store them, discard the remaining parameter vectors;
- step (iv) generate R (where $R < N$) randomly drawn parameter vectors as in step (i); generate M (where $M+R \leq N$) new parameter vectors that correspond to linear combinations of the stored S parameter vectors; generate B (where $B+M+R \leq N$) new parameter vectors, such that $B = (1+p) \times S$, where $p \in [-0.1, 0.1]$; use the minimum and maximum values of each parameter in the stored S parameter vectors as the new lower and upper bounds and generate $N - R - B - M$ randomly drawn parameter vectors within these new (contracted) lower and upper bounds;
- step (v) using the new population of $R+M+B+S = N$ parameter vectors repeat steps (ii)-(iv)
- step (vi) repeat until I different populations of N parameter vectors have been generated and evaluated.

B.4 Validation Experiments

We present several validation experiments that, in different settings, enable us to illustrate numerically identification, and the performance of the estimation algorithms. Throughout, we consider models with 30 locations, which correspond to the 30 largest TTWAs in our German data. The time period is a quarter.

B.4.1 Experiment I: A Simple Cost Function

In this experiment, we focus on the frictional parameters and the moving cost function by considering only one sector and segment. Amenities are assumed to play no role.

[15] In Appendix B.4, we present the results of a simple experiment in the context of which our algorithm yields results that are equivalent to Nelder Mead.

The moving cost function is parsimonious, and only physical distance and differences in house prices play a role. In particular, the moving cost function is given by

$$c(l,k) = \alpha_0 + \alpha_1 \times \Delta hp(l,k) + \alpha_2 \times distance(l,k) \tag{B7}$$

where $distance(l,k)$ measures the geographic distance between two locations, and $\Delta hp(l,k)$ is the difference between the relative house price indices. The objective is to estimate λ_e, δ, ξ as well as $\alpha = (\alpha_0, \alpha_1, \alpha_2)$. The population parameters are set at $\lambda_e = 0.85$, $\delta = 0.026$, $\xi = 3.65$, $\alpha_1 = 5$, and $\alpha_2 = 0.004$.

Identification

We begin by studying numerically identification by evaluating the contours of the estimation objective function given by the GMM criterion, given by equation (15) in the main text. We do so by varying pairs of parameters, holding the remainder at their population values. The criterion function evaluated at the population values equals, of course, zero.

Figure B1 reports the results. It is evident that the contours are well behaved. The population values give the global minimum (depicted by the red circle in the centre of the plot), and in view of the observed gradients, estimators should converge to the population values. We conclude that the examined model is indeed identified.

Estimation

The shape of the contours depicted in Figure B1 suggest that our estimation, which involves minimising the GMM criterion function (13) with respect to the parameters considered, should converge eventually. The specification of the simple model permits estimation by the "slow" Nelder Mead algorithm, which we use here as a benchmark. We also examine the performance of our genetic algorithm. The left panel of Table B2 reports the estimates obtained using Nelder Mead (column N-M) and using our genetic algorithm (column GA). In both cases, the estimates are very close to the population values. The right panel of Table B2 reports the fit of the model evaluated in terms of spatially aggregated (transition) measures; given the closeness of the estimates to the population values, the model fit is perfect. We conclude that our "fast" genetic algorithm works very well.

B.4.2 Experiment II: Multiple industries, age and skill groups

In this experiment, we consider the setting of the empirical application. We now introduce amenities, as given by equation (14) in the main text, and consider three sectors (manufacturing, services, financial services) and four segments within each sector (two age- and two skill-groups). Recall that the frictional parameters and one parameter of the cost function, a_i^s in (15), are segment-specific, while the remaining parameters of the cost and the amenity function are invariant across segments.

Identification

As before, we illustrate numerically identification by considering the contours of the GMM objective function, varying pairs of parameters holding the remainder at the

Figure B1: Contours: Experiment 1

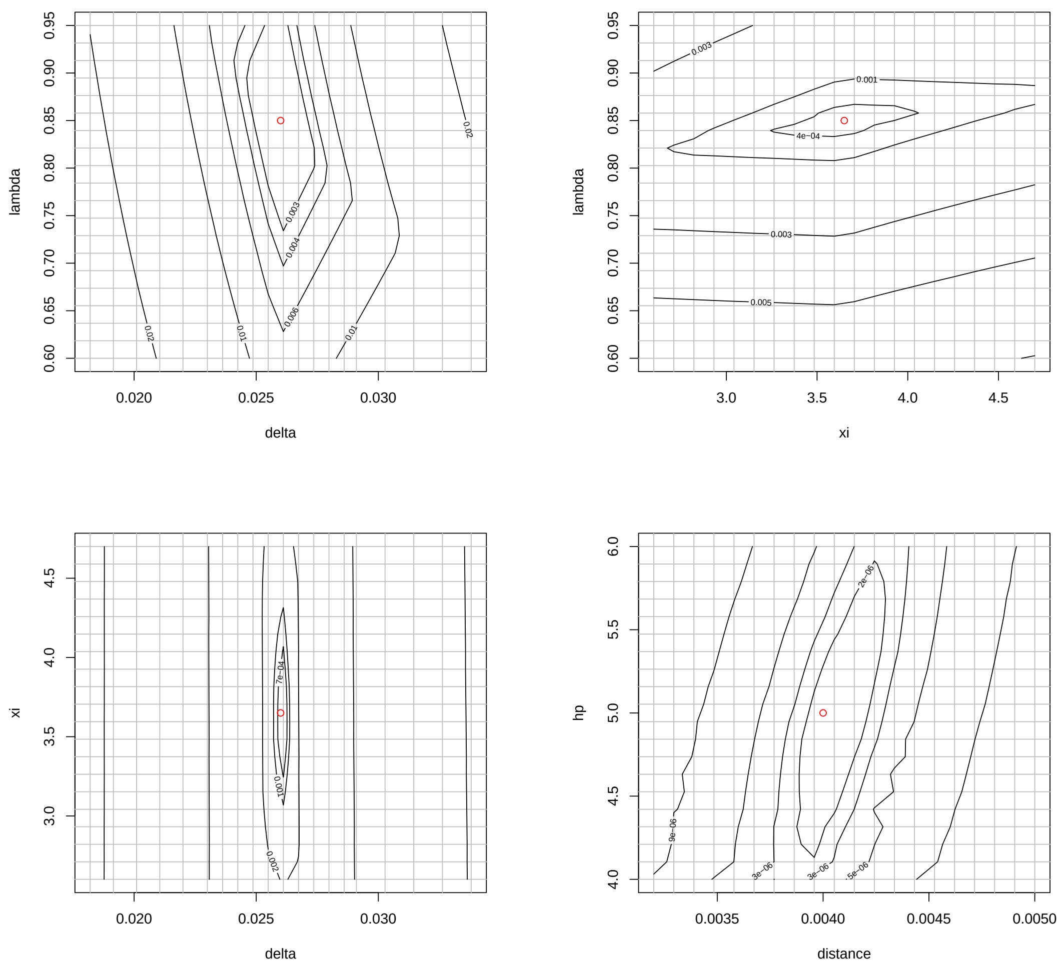

Table B2: Validation experiment I (simple cost function)

	Estimates			Model Fit			
	pop. value	N-M	GA		true	N-M	GA
λ_e	0.85	0.8500	0.8499	unemp [%]	.04798	.04798	.04798
δ	0.026	0.0260	0.0260	reloc [%]	.01447	.01447	.01447
ξ	3.65	3.6507	3.6479	u $\rightarrow$ e [%]	.03899	.03899	.03899
α_0	1.0	0.9857	0.9916	e $\rightarrow$ e [%]	.05426	.05426	.05426
α_1	0.004	0.0041	0.0041				
α_2	5.0	5.2995	4.9662				
obj. fct		5.376e-07	7.827e-07				

Notes: Optimisation using the Nelder-Mead algorithm (N-M) and the genetic algorithm (GA) described in Appendix B.3.

population values. In the first set of experiments, we focus on one segment and sector (young skilled workers in manufacturing), and vary sector/segment specific parameters. This generalises Experiment I by embedding it in the more complicated multi sector/segment model.

Figure B2 reports the results. As in the preceding Figure B1, it is evident that the contours are well behaved. The population values give the global minimum, and in view of the observed gradients, estimators should converge to the population values. We conclude that the examined model is indeed identified.

In Figure B3, we consider another experiment in which we vary at least one segment invariant parameter, such as a parameter in the amenities or the moving cost function. As previously, we conclude that the examined model is identified.

At the same time, this experiment illustrates one particular feature of the model, which further supports our approach of estimating the model using an inner and an outer loop, as described in Appendix B.3: our GMM objective function is more sensitive to variation in segment invariant parameters than in segment specific parameters. This is evident in Figure B3, where we present how variation in a segment-specific and in a segment invariant parameter affects the GMM objective function: if we vary sufficiently one sector invariant parameter, e.g. one of the parameters of the amenities function as in the left panel of Figure B3, then this can induce big population reallocations across local labour markets, and the resulting change in the model-implied moments can lead to jumps in the GMM objective function. Figure B3 illustrates this clearly: in both panels, we observe a critical ridge at which the GMM criterion increases by a jump. By contrast, the contours in the neighbourhood of the population values are well behaved. This behaviour of the objective function implies that any attempt to estimate all the parameters (segment-specific and segment invariant) of the model simultaneously would be very time-consuming, and justifies our approach of switching successively between estimating segment-specific parameters given global parameters, and then estimating global parameters given the estimates of the segment-specific parameters.

Figure B2: Contours for Experiment 2, variation of within-segment parameters

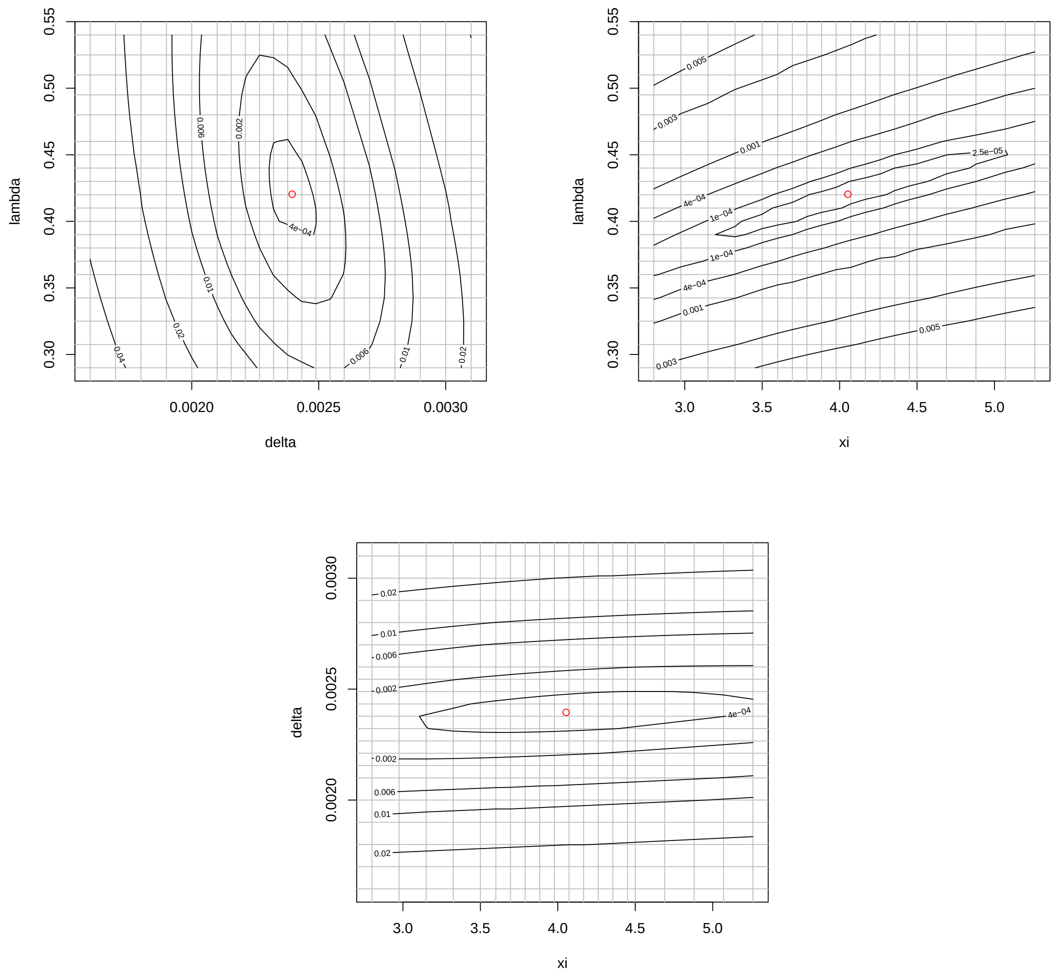

Figure B3: Contours for Experiment 2, variation of across-segment parameters

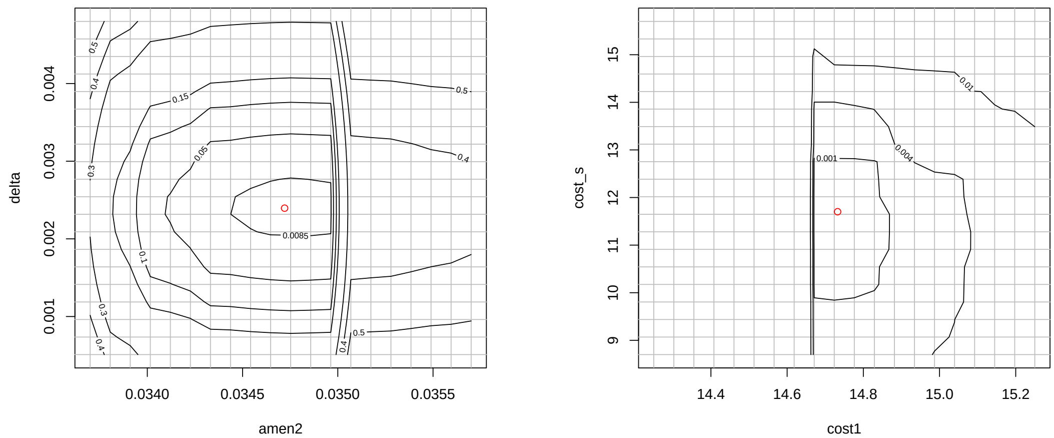

Estimation

The model with multiple sectors, age and skill groups is used to examine the performance of our estimation algorithm: we attempt to recover the population values of all global and segment-specific parameters of the model by minimising the GMM criterion function (13). In this experiment, we estimate eight global (segment-invariant) parameters (five parameters of the cost function and three parameters of the amenities function), as well as fourty-eight segment-specific parameters: we consider twelve segments (three sectors and four segments per sector) each one of which has four segment-specific parameters. These parameters are estimated using the genetic algorithm described in Appendix B.3. Figure B4 reports the performance of our algorithm: the top panel presents the average value of the GMM criterion function (13) in each iteration, while the bottom panel presents the best/minimum value of the GMM criterion per iteration. It is evident that our algorithm achieves convergence fairly fast, despite the large number of parameters estimated. More importantly, the behaviour of the average and the "best" fitness value in iterations 1-6 suggests that our algorithm can flexibly explore the parameter space and focus on the most promising regions without completely eliminating seemingly irrelevant regions.

Figure B4: Convergence plot

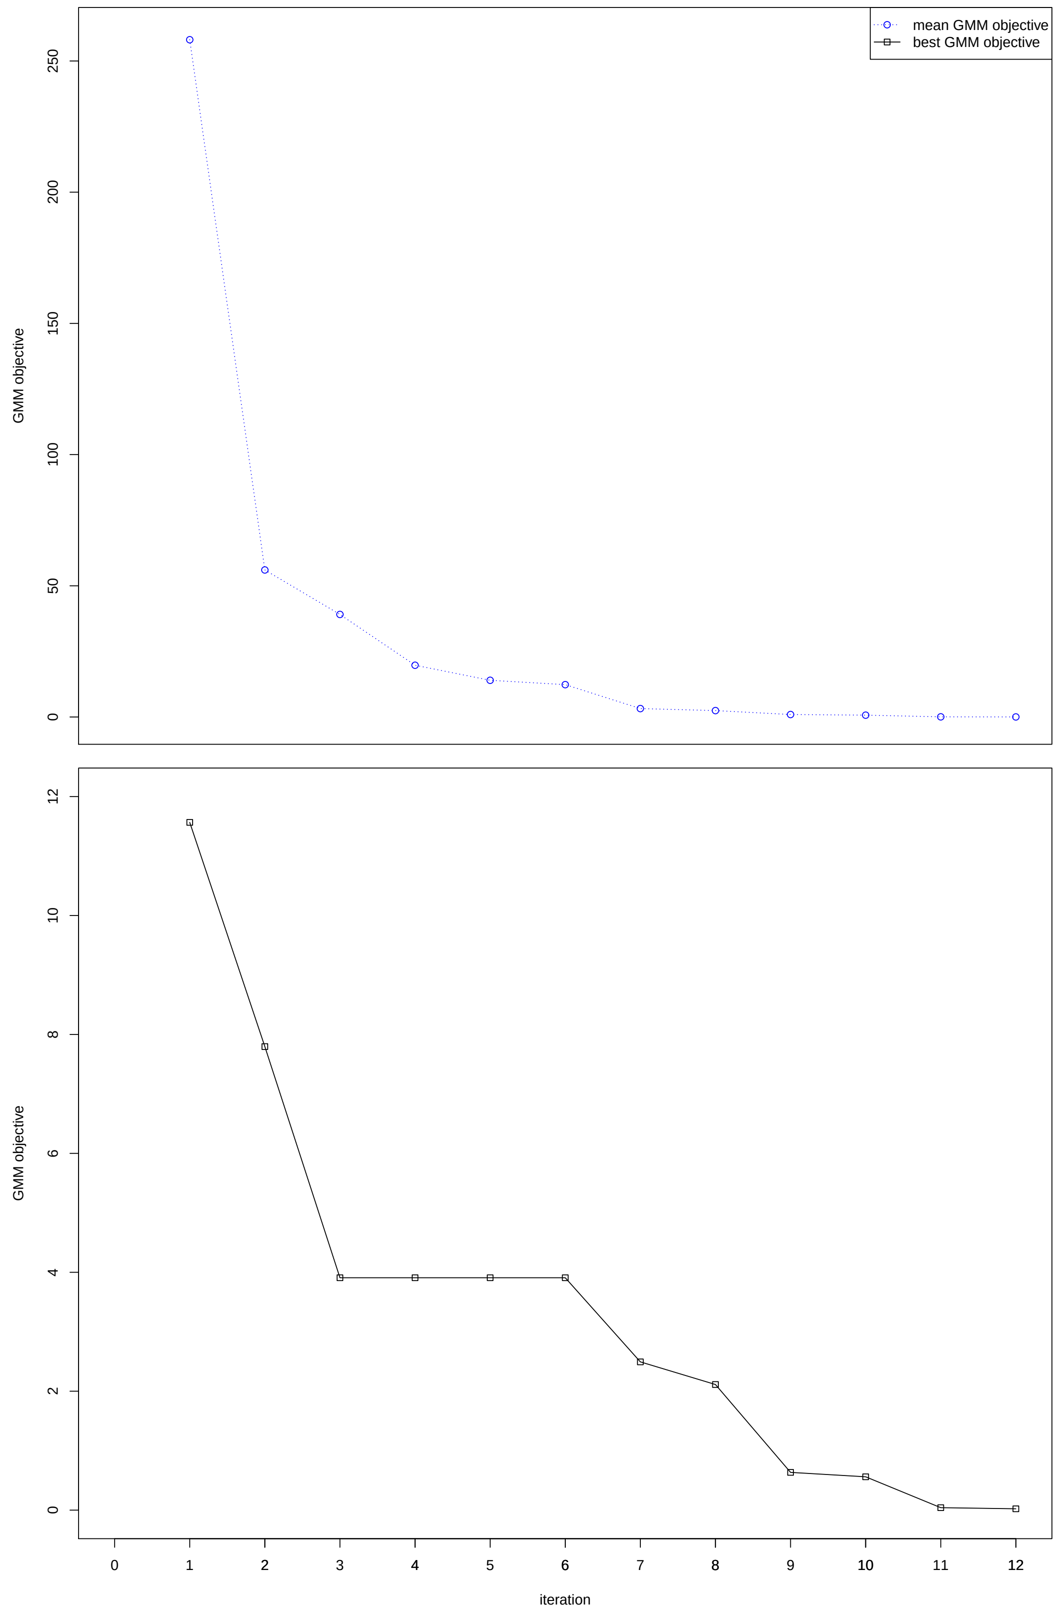

C Productivity Estimation

We follow the Olley and Pakes (1996) procedure to obtain estimates of total factor productivity (TFP) using sector-specific plant-level production functions. We consider three sectors: Low-Tech Manufacturing, Low Services, and High-Tech Manufacturing. As is standard in the literature, we assume a Cobb-Douglas production function with labour, capital, and materials as inputs. A control for the share of high skilled workers is also included in the production function. Table C1 reports the coefficients.

Table C1: Production function coefficients from OP estimation

	Low Manufacturing	Low Services	High Manufacturing
labour	0.3312	0.2793	0.3774
capital	0.0692	0.0441	0.0471
high skill share	0.0137	0.0358	0.0785
materials	0.5903	0.7151	0.5924

Notes: Based on LIAB for years 2002-08.

The estimates reported in Table C1 are similar to the coefficients estimated by Ehrl (2016) using LIAB data for a period similar to the period we consider (2000-2007).

D Data Appendix

D.1 8 selected local labour markets

Figure D1: Eight selected travel-to-work areas.

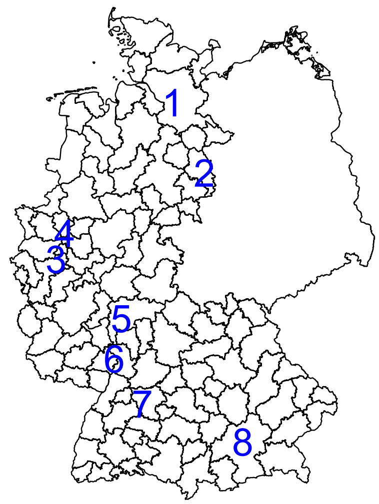

Notes: 8 Local labour markets (TTWAs): 1: Hamburg, 2: Wolfsburg, 3: Leverkusen, 4: Bochum, 5: Frankfurt, 6: Mannheim, 7: Stuttgart, 8: München.

In order to illustrate further and in more specific detail the spatial heterogeneity of local labour markets, and the observed patterns of geographic mobility, we report summary statistics for 8 selected local labour markets. These include the largest and economically most important cities (Hamburg, Frankfurt, Munich) and a selection of smaller ones (such as Bochum or Wolfsburg). These 8 local labour markets are depicted on the map of Figure D1.

Table D1 reports some summary statistics. It is evident that these local labour markets differ markedly in terms of unemployment rates, wages, living costs, as well as productivities. For instance, the mean unemployment rate in the TTWA of Bochum is 2.4 times higher than that of Munich, while the mean daily wage in the former is 91% of the latter. The largest mean daily wage is paid in the TTWA of Frankfurt, which also exhibits the largest mean worker fixed effect and the largest firm fixed effect in services, the latter being expected given the area's central role in banking and finance. Wolfsburg is a comparatively small TTWA and, as the seat of car manufacturer Volkswagen, being the principal employer, exhibits the largest firm fixed effect in manufacturing; since its workforce is mainly employed in blue collar occupations, it is no contradiction that the mean daily wage and the mean worker fixed effect are fairly low, as are housing costs. Stuttgart, by comparison, is much larger and much more diversified industrially, which is manifested in a smaller mean firm fixed effect in manufacturing but a larger worker fixed effect. Overall, the TTWA of Munich is the most expensive to inhabit.

Table D2 describes the geographic mobility among these 8 selected local labour markets (and all others aggregated in the cells labelled 'Rest'). The table reveals that, conditional on being located in a big urban zone (such as Hamburg, Frankfurt, or Munich), the worker is more likely to relocate to another such urban zone than to a smaller urban zone. For instance, originating in the TTWA of Munich, moves to Wolfsburg, Bochum or Mannheim are extremely unlikely. However, spatial moves are not exhausted by the 8 selected TTWA, as the residual category exceeds by at least one order of magnitude all other conditional probabilities.

Table D1: Eight local labour markets

no.	TTWA name	# of districts	inhabitants [$\times 10^6$]	mean u rate	rel. house price index	rel. mean daily wage	rel. mean worker FE
1	Hamburg	8	3.19	10.02	0.75	0.94	0.987
2	Wolfsburg	3	0.39	10.37	0.58	0.85	0.979
3	Leverkusen	5	2.04	11.49	0.74	1.04	0.983
4	Bochum	4	1.05	14.90	0.60	0.91	0.975
5	Frankfurt	8	2.49	8.55	0.80	1.05	1.022
6	Mannheim	8	0.98	9.53	0.67	0.82	0.963
7	Stuttgart	4	1.99	6.12	0.82	0.94	0.983
8	Munchen	11	2.65	6.20	1	1	1

Notes: See Table 1. Columns 4 and 5: District-level data obtained from www.regionalstatistik.de: Col. 4 (inhabitants) Table 173-01-4 for year 2002, Col. 5 (mean unemployment rate) Table 659-71-4 averaged over 2002-2008.

Table D2: Spatial mobility across TTWAs

	1 [H]	2 [Wo]	3 [Lev]	4 [Bo]	5 [Fr]	6 [Man]	7 [Stut]	8 [Mun]	Rest
1 [H]		0.037	0.286	0.106	0.550	0.057	0.203	0.475	10.363
2 [Wo]	0.052		0.018	0.016	0.024	0.001	0.011	0.011	1.194
3 [Lev]	0.243	0.019		0.143	0.240	0.050	0.104	0.174	5.603
4 [Bo]	0.098	0.004	0.132		0.082	0.014	0.048	0.073	3.780
5 [Fr]	0.415	0.018	0.250	0.114		0.169	0.224	0.342	6.401
6 [Man]	0.048	0.006	0.042	0.016	0.200		0.097	0.057	2.603
7 [St]	0.220	0.010	0.101	0.050	0.190	0.092		0.196	5.715
8 [Mun]	0.466	0.013	0.229	0.090	0.437	0.070	0.267		6.455
Rest	12.252	1.070	7.481	3.714	8.419	2.918	6.422	7.911	

Notes: Bi-stochastic transition matrix for moves between selected TTWAs. The category labelled "Rest" aggregates all other TTWA. Based on LIAB, time period 2002-2008. Reported is $s_{l,k}$ = number of relocations from l to k divided by the total number of relocations in Germany. By definition $\sum_l \sum_k s_{l,k} = 100$.