Job Search and Migration



Maria Balgova

Corpus Christi College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity 2019

Acknowledgements

Thank you to my supervisors, Margaret Stevens and Abi Adams, for their endless patience, wisdom, and support. It has been a great privilege to learn from you. You have made it possible for me to see myself as an economist and an academic, and have given me the best possible example.

I am also grateful to all the inspiring and caring professors, supervisors, and teachers I have met in Oxford, Cambridge, and Slovakia, for all your help and encouragement on this long journey.

To my almost-husband Jonathan, whose kindness, patience, and generosity deserve a degree on its own.

To my family and friends, in Slovakia and the UK, for standing by me and giving my life and this thesis the right perspective. Ďakujem vám, Ivan, Jarka, Ivanka, Viki a Zorko, za to, že ste pri mne vždy stáli. To my DPhil friends, Eileen, Andrew, Alessio, Jonas and Greg, without whom these past six years would not be nearly as much fun. Thank you for putting up with me when times where tough, and for inspiring me with your amazing work.

Last but not least, I would like to thank the ESRC, Corpus Christi College, and the Dulverton and Wolf scholarships for their incredibly generous financial support during my time at Oxford.

Abstract

Standard models of within-country mobility assume that all migration is speculative: workers move to search for jobs in other labour markets. I establish a new stylised fact showing that the majority of cross-regional migration in the US is in fact with a job in hand, as a result of a job offer from another region. The fact that a significant fraction of migration is a result of cross-regional hiring, and that this hiring is subject to search and matching frictions, has important implications for the level of regional mobility and hence regional differences within a country.

In the first two chapters, I show empirically that the availability of cross-regional employment opportunities does matter for workers' migration behaviour. Using US panel data, I demonstrate that the observed gap in migration propensity between the more and less educated workers can be partly explained by differences in labour market frictions. The less educated workers find job search in more distant regions much more difficult, which limits their options to move for a specific job and reduces their overall mobility. This result opens a new policy channel in addressing regional differences and those left behind: the importance of the ability to find a job before moving suggests a large social return to improving regional search and matching for less educated groups.

In the third chapter, I explore the theoretical implications of an imperfect search and matching process on the existence and persistence of regional disparities. First, I show that if the efficiency of the search and matching process increases in the size of the labour market, migration will serve to exacerbate, not equilibrate, the differences in regional economies. A bigger labour market will experience higher wages and lower unemployment rates, attracting migrants from other regions and thus accelerating the process of regional divergence. At the same time, larger labour market frictions in the smaller regions reduce firms' incentives to hire there, curtailing cross-regional hiring and hence the opportunity to move for a specific job. The interaction of these two forces give rise to multiple equilibria and explain how, in the presence of labour market frictions, regional disparities can be large and persistent at the same time.

Contents

List of Figures **vii**

List of Tables **ix**

1 Introduction **1**

2 Why don't less educated workers move? A discrete choice model of cross-regional job search and migration **9**

2.1 Introduction 9
2.2 Data 14
2.2.1 The SIPP 14
2.2.2 Data on recruitment strategies 17
2.3 Stylised facts about regional migration and education 18
2.3.1 The propensity to migrate 18
2.3.2 Speculative migration vs. migration for a specific job 19
2.4 A discrete choice model of employment and location 22
2.4.1 General determinants of migration 22
2.4.2 Discrete choice model of location and employment 26
2.4.3 Estimation 33
2.4.4 Results 35
2.5 Conclusion 44

3 The role of job search in migration decisions: a structural approach **46**

3.1 Introduction 46
3.2 A theoretical model of cross-regional job search and migration . . . 51
3.2.1 The setup 51
3.2.2 Decisions and choice sets 55
3.2.3 Values of employment and unemployment in different regions 57
3.2.4 Equilibrium 60
3.3 Structural estimation 65
3.3.1 Data 65
3.3.2 Identification of parameters 65

3.3.3 Method: estimation by simulated moments 69
3.3.4 Results 71
3.4 Explaining the patterns of internal migration 73
3.4.1 Decomposing the migration propensity gap between the more and the less educated 73
3.4.2 Policy implications: the role of job search frictions 75
3.5 Conclusion 79

4 Labour market frictions and regional disparities 81
4.1 Introduction 81
4.2 Motivation 86
4.3 Local labour markets under autarky 90
4.3.1 The model 90
4.3.2 The autarkic equilibrium 93
4.3.3 Comparative statics: the role of regional population size . . 96
4.4 Local labour markets with migration 99
4.4.1 The model 100
4.4.2 Speculative migration 105
4.4.3 Migration for a specific job 106
4.4.4 Steady-state equilibria with endogenous population 111
4.4.5 The adjustment process 113
4.5 Comparative statics: labour market frictions and regional disparities 116
4.5.1 Comparative statics: the role of search frictions and migration costs in regional divergence 116
4.5.2 Welfare comparisons 124
4.6 Conclusion 126

5 Conclusion 129

Appendices

A Appendix to Chapter 2 133
A.1 Data 133
A.2 Stylised facts about regional migration and education 138
A.3 Post-migration outcomes 141
A.4 Imputing wage offers 141

B Appendix to Chapter 3 145
B.1 The role of initial population distribution 145
B.2 Closed-form solutions for model moments 148

C Appendix to Chapter 4 **151**
C.1 Motivation 151
C.2 Local labour markets under autarky 159
C.2.1 The model 159
C.3 Local labour markets with migration 160
C.3.1 The model 160
C.3.2 Equilibria with endogenous population 163
C.3.3 Derivatives of the migration conditions with respect to regional population levels 164
C.4 Modelling regional matching spillovers 165

Bibliography **167**

List of Figures

2.1 The 4 large census regions of the US 16
2.2 Propensity to migrate into another state, US, 1996-1999 19
2.3 Probability of moving into a specific job (conditional on moving), US, 1996-1999 21
2.4 Recruitment proxies and migration 31
2.5 Estimated distribution of the probability of choosing from a full choice set (the probability of receiving a cross-regional job offer), by education 38
2.6 Marginal effects of income and migration costs on the probability that an average high school graduate in the Northeast chooses employment in the South (option 5) 39
2.7 Impact of choice set variation: counterfactual migration behaviour if all workers received cross-regional job offers vs. if none of them did. 40
2.8 Marginal effects of income on the probability that an average worker chooses employment in the South (option 5), by education. 42

3.1 Choices and possible outcomes for an unemployed worker 55
3.2 Choices and possible outcomes for an employed worker 56
3.3 Equilibrium in a 2-region model 61
3.4 Equilibrium in a 2-region model: constant stock of employed workers in region A 62
3.5 Decomposition of the education differences in migration propensity 74
3.6 Decomposition of the education differences in migration for a specific job 75
3.7 Counterfactuals comparing the role of migration costs and cross-regional search frictions in migration behaviour 76
3.8 Counterfactuals evaluating the role of cross-regional search frictions on migration behaviour 77

4.1 Distribution of regional income and unemployment rates 87
4.2 Correlations between regional population size, income, and employment rate 88
4.3 Persistence of regional income and employment rate over time . . . 88

4.4 Correlation between regional population and the size of net migration flows . . . 89
4.5 Existence of equilibria under different matching functions . . . 95
4.6 Local labour market equilibria as regional population increases . . . 98
4.7 Migration conditions and the location of steady-state equilibria . . . 116
4.8 Migration conditions and the location of equilibrium types . . . 118
4.9 Comparative statistics for the migration cost . . . 120
4.10 Comparative statistics for the extra cost of search for migrants . . . 121
4.11 The effect of matching spillovers on regional disparities . . . 123
4.12 Comparative statistics for matching function spillovers . . . 124
4.13 The effect of matching spillovers on regional disparities when allowing for agglomeration in production . . . 125
4.14 Comparative statistics for the matching function spillovers when allowing for agglomeration in production . . . 126

A.1 Propensity to migrate, US, 1996-1999 . . . 138
A.2 Probability of moving into a specific job (conditional on moving), US, 1996-1999 . . . 138
A.3 Probability of moving into a specific job, conditional on moving and employment state before migration, US, 1996-1999 . . . 139
A.4 Unemployment rates for stayers and migrants, US, 1996-1999 . . . 139
A.5 Migration for a specific job by education as a function of months after moving, US, 1996-1999 . . . 140
A.6 Speculative mirgation vs. unemployment duration, by education and sector, US, 1996-1999 . . . 140

C.1 Distribution of regional income . . . 152
C.2 Distribution of regional unemployment rates relative to aggregate unemployment rate over time . . . 153
C.3 Correlation between regional population size and employment rate . 154
C.4 Correlation between regional population size and income . . . 155
C.5 Correlation between regional income and employment (wage curve) 156
C.6 Correlation between regional income at the start and the end of the period . . . 157
C.7 Correlation between regional employment rate at the start and the end of the period . . . 158

List of Tables

2.1 Binary logit models of propensity to migrate and the type of migration, between regions . . . 20
2.2 Regional variation in monthly wages and non-employment rates, by education . . . 23
2.3 Across-state variation in monthly wages and non-employment rates, by education . . . 24
2.4 The list of options in the logit model of the employment-location choice 26
2.5 Full and restricted choice sets; example for a worker resident in the West . . . 29
2.6 Multinomial logit model of the choice of employment or unemployment in one of the 4 large US regions, allowing for choice set heterogeneity. . . . 37
2.7 Multinomial logit model of the choice of employment or unemployment in one of the 4 large US regions, allowing for choice set heterogeneity, by education . . . 43

3.1 Data moments: matrix of transition probabilities for the more educated (in %) . . . 67
3.2 Data moments: matrix of transition probabilities for the less educated (in %) . . . 67
3.3 The SMM estimates of the structural model: goodness of fit . . . 71
3.4 The SMM estimates of the model parameters, by education . . . 72

4.1 Migration conditions and types of equilibria . . . 114
4.2 Parameter values in simulations . . . 117

A.1 Descriptive statistics for the sample of working men between the age of 25 and 50 . . . 135
A.2 Observed choices . . . 136
A.3 Definitions and sources of variables. . . . 137
A.4 Impact of migration on wages, within-person comparison . . . 141
A.5 Impact of migration on unemployment . . . 141
A.6 Heckman selection model for monthly regional wages. . . . 143
A.7 Dahl selection model for monthly regional wages . . . 144

1 Introduction

Location matters. In the US, an average worker moving from a city in the bottom decile of the wage distribution to a city in the top decile will see her wage increase by a third (Moretti (2011)). In the UK, the average full time earnings in London were a quarter higher than in the Northeast of the country[1]. Regional differences are similarly large in Europe: the unemployment rate in the poorest region of Italy in 1996 (Campania) was 4.4 times higher than in the most well-off region (Overman and Puga (2002)). Moreover, these differences are not just a consequence of temporary local shocks. Their persistence is strong: Campania was the Italian region with the highest unemployment in the late 1990s as well as in 2015, while the gaps in regional per capita output in the UK have persisted for more than 30 years (Rice and Venables, 2003). And these differences are not limited to economic indicators: there is also considerable inequality in mortality, education, and intergenerational mobility (Chetty and Hendren (2018)). In the UK, for example, the probability of achieving at least AAB in A-levels is 4 percentage points higher in the Southeast compared to the Northwest, with up to three years difference in life expectancy[2].

The question that underpins this thesis is simple: if one's region of residence determines income and employment, the life chances of their children, and even

[1] ONS data for 2018.

[2] Based on Department of Education data for 2014, and ONS data for 2012-2014.

one's life expectancy, why don't people move? The most recent annual migration rate between US states is 2.3% [3], and it has been in decline since the 1980s. Mobility in Europe is even lower, as are the migration rates for specific groups of the population, such as those without a college degree, for whom the benefits of migration may potentially be the largest.

The existing literature explains the decision to migrate as a careful trade-off between the costs and the benefits. Moving into another location is fundamentally equivalent to the decision to attend college, or to make a financial investment. They all require an immediate cost in exchange for an uncertain gain, and individuals move only if the expected return of moving exceeds its costs. While earnings and the quality of local schools and hospitals matter, there are other factors that may tip the balance: the value of local social networks, preferences for certain amenities, house prices and congestion, the cost of moving itself. For most individuals the combination of these factors outweighs the benefits of moving, which is why large regional disparities don't lead to large-scale within-country migration.

In this thesis, I argue that there is another crucial determinant of the migration decision: the labour market. Earnings and the return to skills have long been considered important factors in the decision to move, but I demonstrate that the *process* of job search matters too. The standard model of migration assumes that employers and workers can only match when they reside in the same region. All migration is speculative: workers move first and search only later. This implies that migration is a decision made under (income) uncertainty, where the costs have to be paid upfront but gains are risky. However, if hiring can happen across regions, remotely, a worker can search for jobs in other regions *before* she has to decide whether to move, and her income in another region will be known at the time of making the migration decision. Unlike speculative migration, she has the option to move only if the net benefit is positive. As a result, if the process of job search allows it, migration becomes more like a purchase of a risk-free asset rather than a risky investment.

[3] US Census Bureau data for 2017.

To make this point further, consider the following example of a two-region economy. Let's say that the expected income of staying in one's home region is \$600. If the worker moves to the other region, she may either be unemployed and earn \$0, or get a job that pays \$1,000. With a 50% probability of finding a job, the expected income of moving into the other region is \$500. If only speculative migration is possible, the worker will never move (\$600 > \$500). If, on the other hand, she is able to search for jobs in the other region before she moves, there is a 50% chance of migration: she would move if she was offered a job in the other region (\$600 < \$1,000) but not if she didn't (\$600 > \$0). This is why job search frictions matter for migration.

More broadly, via its impact on migration, cross-regional job search affects the existence and persistence of regional disparities. Because migration for a specific job allows for cost-sharing of the migration expense between the worker and her employer, and brings a higher net gain on average, it will happen even if the size of regional differences does not justify speculative migration. Depending on whether moves to the more prosperous regions increase or decrease the level of regional disparities, this will fuel or alleviate regional inequality. It also opens up a new set of policies to address regional divergence.

The link between migration and job search is supported by data. First, I present a new stylised fact that the majority of inter-state moves in the US are for a specific job – the result of remote job search. This is in contrast with the standard assumption in the literature that job search follows *after* moving. Second, drawing on empirical IO models, I show that migration increases when there are more cross-regional employment opportunities. Finally, I discuss how the lack of incentives for cross-regional hiring could be in part responsible for the observed persistence of large regional differences.

The empirical evidence on the role of job search in migration is established in the first chapter of this thesis. I use SIPP, a panel dataset of a representative sample of US workers which captures their monthly employment status and location. The high-frequency nature of the observations, and the fact that the survey follows

respondents when they move, allows me to distinguish between migration into employment and unemployment. The former implies that the worker moved for a specific job, while the latter means she moved speculatively. The data reveals that 65% of migration is with a job lined up. I provide empirical evidence that this is not a coincidence – that job search has a significant impact on the migration decision. I build a discrete choice model of migration and employment, and augment it to include the recruitment practices of workers' employers. The fact that these have a statistically significant impact on workers' outcomes, even after controlling for standard migration variables, suggests that the availability of employment opportunities elsewhere influences the migration decision.

I apply this idea to explain the patterns of migration behaviour of workers of different education in the US. It has been long known in the literature that the less educated individuals move less (Greenwood (1997)). My dataset allows me add a second stylised fact: not only do the less educated move less, they are less likely to move for a specific job. I show that these two observations are related. Those without a college degree are more likely to be employed in sectors and occupations where cross-regional hiring is less frequent, which forces them to move speculatively, reducing their overall mobility.

In Chapter 3, I model the interaction between job search and migration formally. I extend a model of job search (McCall (1970)) so that workers now may receive job offers from other regions, and decide to move speculatively to search elsewhere. As a result, the decision about accepting a job depends on local labour market conditions across regions, and at the same time, the availability of such cross-regional job offers shapes the location choice.

I then estimate the structural model to revisit the question of why the less educated don't move much. My model allows me to explicitly control for differences in wage distributions, migration costs, and location preferences. It also allows me to understand the search frictions faced by the less and the more educated. The estimates of the model parameters show that it is harder to find a job across regions than locally, and that these cross-regional search frictions are greater for

the less educated. I demonstrate that, while the standard variables such as wages and migration costs play a significant role in explaining the observed migration patters, the availability of cross-regional employment opportunities matters too. Up to 50% of the gap in migration propensity and 75% of the difference in the type of migration can be attributed to the relative difficulty the less educated workers face when looking for jobs remotely. This lends further empirical support to the suggestion that job search is an important driver of migration decisions.

In the third chapter of this thesis, I analyse the consequences of this new model of migration on regional outcomes, assuming that regional differences are driven by increasing returns to matching in local labour markets. I model both the migration decisions of workers and optimal vacancy creation by firms in order to analyse what makes speculative migration, and migration for a specific job, optimal. Both types of migration depend on the size of regional differences, and in turn these differences are shaped by the extent and direction of migration flows. I describe when the regional economies are in a steady state, and what the process of divergence looks like. Finally, linking the theory back to the stylised facts that motivate this thesis, I show how different types of cross-regional search frictions increase or decrease the probability that an economy might get stuck in an equilibrium with large and persistent regional disparities.

In terms of the subject topic, I draw on two fields of literature: new economic geography and local labour markets. Together, this work has shaped our understanding of the spatial distribution of economic activity and internal migration. These papers build a comprehensive model of spatial equilibrium within a country, encompassing production, regional trade links, knowledge spillovers, amenities, congestion, and the housing market. They describe the determinants of the distribution of economic activity over time and across space, and provide a framework to evaluate place-based policy. Migration serves as the equilibrating force. Population flows facilitate adjustment in output, agglomeration and congestion so that in equilibrium, the utility of the marginal worker across space is equalised. The determinants of

migration used in this thesis, and the concept of spatial steady state equilibrium, are based on this literature.

The theoretical framework underpinning my models comes from the literature on search and matching in the labour markets. These models analyse individual worker and firm behaviour to understand how output, wages, and employment depend on the imperfections in the labour market. Adding a spatial dimension then allows me to study how search frictions within and across regions affect the decision to hire remotely (by firms), and to accept these offers (by workers), leading to migration for a specific job. By also allowing workers to move speculatively, I nest my model of search and matching across regions in the wider class of models of migration behaviour and regional economics.

The interactions between employment and location choice are modelled using two different frameworks. In Chapter 3, I present a partial equilibrium model of worker's search in an economy with multiple regions and a range of exogenous, fixed job-finding probabilities. Workers differ in their education and skills, and their idiosyncratic preferences over locations. In this world, firms' behaviour is taken as given, and their actions are summarised by the observed wage distribution and the estimated job offer arrival rates. On the other hand, in Chapter 4, I model both workers and firms as they decide how long to stay unemployed, and whether to post any vacancies. Unlike in Chapter 3, the setup assumes away any heterogeneity; instead, their decisions are shaped by the changing size of their local labour market as workers move across the regions. The partial-equilibrium approach in Chapter 3 allows me to capture the rich set of job offers and speculative moves a worker is choosing from every period, and link it to her employment-location outcomes. This is not possible in the model in Chapter 4, where only one type of migration and hiring (local vs. cross-regional) is optimal at any given point in time. On the other hand, the advantage of this model lies in its ability to capture the feedback loop from migration to vacancy creation and local labour market conditions. Together, the two models capture the different implications of the role of job search in migration.

Finally, the main empirical part of this thesis contributes to the growing literature on discrete choice models with heterogeneous choice sets. I draw on the concept of unobserved variation in choice sets, which was developed in the empirical IO literature (Abaluck and Adams (2017); Goeree (2008)). I adapt the method to labour market behaviour: instead of capturing the set of products a consumer is aware of, I use it to model the sequencing of migration and job search decisions. Applying this method to the labour setting generates several specific issues, such as the treatment of unobserved option attributes, that are discussed in Chapter 2.

The core result of this thesis – that job search matters for migration – has the potential to change the narrative about what drives individuals' responsiveness to economic incentives. This is particularly pertinent to the analysis of lower mobility of the less educated in the US. The results suggest that the observed mobility differences are not only because the return to migration is higher for college graduates, but also due to the inequality in employment opportunities between the two groups. In other words, the less educated move less in part because they lack the option to move for a specific job that are ample in the labour market for college graduates. While there may be economic reasons why employers do not hire less educated workers cross-regionally, this does not diminish the fact that moving for opportunity is harder as a result.

The results of this thesis have several implications for policy. Understanding why and how the less educated move is vital for addressing regional inequalities and worker mobility. The augmented model of migration in Chapter 2 demonstrates that failing to control for differences in employment opportunities leads to overestimating migration costs, and underestimating individuals' responsiveness to economic incentives. This helps to reconcile the low observed migration rate with the work on local labour market shocks which shows that the less educated, low-income households are also most exposed to regional shocks. It opens a new policy channel in addressing the left behind communities: the importance of the ability to find a job before moving suggests a large social return to improving regional search and matching for less educated groups. Finally, the theoretical model of labour

markets and regional disparities provides a framework to think about how the type and magnitude of search frictions influences the process of regional differences and divergence in aggregate.

2

Why don't less educated workers move? A discrete choice model of cross-regional job search and migration

2.1 Introduction

Less educated individuals move less: compared to their college-educated counterparts, they are up to three times less likely to move into another region of the same country[1] (Greenwood (1997)). This is despite the evidence that benefits to migration can be large. In the short run, moving to another region makes it possible for an individual to weather negative local shocks, while in the long run, migration can significantly improve the lifetime earnings of the movers and their offspring alike (Bound and Holzer (1996), Kennan and Walker (2011), Chetty and Hendren (2018)). Both of these effects have been found to be larger for the less educated: they are more exposed to regional economic fluctuations, and the impact of where a child grows up on his future earnings is particularly strong for low-income households (Gregg et al. (2004), Hoynes (1999), Bartik (2018)). This makes the low migration propensity of the less educated a puzzle.

The existing literature explains this low mobility as a reflection of small net

[1] Based on inter-state migration of a sample of US adults, 1996-1999.

expected gains to moving. When deciding whether to move, the worker weighs up expected utility in different regions, which depends on factors such as expected income, location preferences, and migration costs. The existing empirical evidence shows that these are less favourable for the less educated: they tend to have lower and less dispersed wages, different amenity preferences, and larger migration costs. As a result, their expected returns to moving are relatively small, and they move relatively infrequently.

In this chapter, I argue that the literature has overlooked a crucial mechanism: cross-regional job search frictions. They not only change the expected returns to moving, but also the way migration decisions are made. I show that college graduates usually move for a specific job, while the less educated tend to move speculatively, without a job offer in hand. I estimate a discrete choice model of the joint employment and location decision to argue that the less educated workers find job search in more distant regions much more difficult. This diminishes their options to move for a specific job, forcing them to move under greater uncertainty, and thereby reducing their overall mobility.

The analysis is based on a monthly panel of a sample of US inhabitants from 1996 to 1999 (SIPP). The monthly data on employment and location, combined with the fact that the survey follows respondents when they move, allows me to infer whether a particular move was for a specific job or speculative. I analyse these outcomes by drawing on other labour market information about the workers and matching it with proxy data on recruitment strategies of firms.

To test whether the less educated have fewer opportunities to migrate with a job in hand, I draw on the growing literature on discrete choice models with heterogeneous choice sets. The key feature of my empirical model is that unlike speculative migration, which is independent of employers' decisions, migration for a specific job is impossible without first receiving a job offer. This means that an individual's choice set – and his outcomes – is determined by the recruitment behaviour of his employers. Because these choice sets are unobserved, I estimate their impact by embedding them in a model of joint location and employment choice.

The estimates of the importance of employer behaviour on worker outcomes are then informative about the role of cross-regional job search in migration decisions.

I start by estimating a conditional logit in which the worker weighs up employment and unemployment options in different regions of the same country. The utility of an option depends on its attributes, such as income, local rent, amenities, and migration costs; these may vary depending on whether the option is for employment or unemployment. The worker chooses the option that gives him the highest utility.

Following the consumer choice literature (Goeree (2008), Abaluck and Adams (2017)), I augment this model by allowing for choice set heterogeneity. Namely, not all options are equally likely to be available for the worker to choose. I assume that a worker can be unemployed anywhere, but employment in "away" regions is only possible if potential employers advertise their vacancies outside of their local labour market. In the absence of a large-scale dataset on hiring strategies, I draw on the County Business Patterns data to construct measures of spatial concentration and firm size (by sector). I also use data on the representativeness of online vacancies by occupation. All of these proxy for the likelihood that a firm advertises its vacancies more widely, making it easier for a worker in a distant region to apply. The model is estimated as a conditional logit with unobserved heterogeneous choice sets using maximum likelihood.

I find that the less and the more educated face significantly different opportunities to move for a specific job. The more educated are more likely to work for large and more spatially concentrated companies, which tend to advertise their vacancies more broadly; as a result, the more educated are more likely to be offered a job in an away region. Overall, less than a half of all high school dropouts can choose to migrate for a specific job, but this increases to 95% for college graduates. This translates into large differences in migration propensity: high school dropouts would be about 40% more likely to move if they were given the same job search opportunities as the more educated.

Moreover, controlling for differences in opportunity is crucial for correct interpretation of the estimates of sensitivity to income variation and migration costs. A

standard logit implicitly assumes that all workers have the same choice set, which does not allow us to distinguish between workers' preferences and their options. As a result, its estimates suggest that the less educated are not as responsive to income as college graduates; letting choice sets vary virtually eliminates this gap. Intuitively, a model with heterogeneous choice sets allows for the possibility that the choice probability of employment away is low not because the less educated don't care for the better paid employment in other regions, but because that option is not available.

Related literature This chapter makes several contributions.

First, I introduce a new stylised fact about the differences in migration behaviour between the less and the more educated. The two ways to move – speculatively or for a specific job – have been discussed before. Silvers (1977) and Molho (2001) have included both types of migration in their theoretical models. In terms of empirical work, a similar stylised fact was first reported by Basker (2018), who uses CPS data in which respondents are asked for their main motivation to move. He finds that those without a college degree are more like to report to be moving to search for work, as opposed to for a new job. Using the same dataset, Amior (2015) notes that only a small fraction of the job-motivated movers that are the subject of his paper report doing so without a job lined up. However, to my best knowledge, this is the first study to construct a statistic on the type of migration using actual observed moves rather than self-reported motivation. Using high-frequency observational data, as opposed to annual self-reported data, sidesteps possibly unreliable narrative and make a more accurate inference from individuals' behaviour. This adds to the work on education differences in internal migration and general trends in within-country mobility (Malamud and Wozniak (2012)).

Second, I estimate an empirical model to investigate the links between the two stylised facts – the size and the type of migration flows. There is a long history of descriptive discrete choice models of migration, but none of them distinguishes between migration into unemployment and employment[2]. They either explain the

[2]The exception to this is Basker (2018), who uses discrete choice models to understand correlations with the stated reasons to move.

binary migration decision as a function of the worker's demographic characteristics, or study the multinomial location choice in relation to the economic features of the regions of destination and origin (Davies et al. (2001)). The logit model I outline in this chapter nests these two approaches and adds the employment-unemployment margin to present a more complete model of job search and migration.

More broadly, I contribute to the literature on dynamic discrete choice models. Starting with the seminal work of Kennan and Walker (2011), this literature explains an individual's migration choice in the context of a rich set of his past choices, including past moves, education and employment history (Bayer and Juessen (2012), Gemici (2011)). This chapter borrows the core of this approach – combining dynamic considerations with a discrete choice framework – but is considerably simpler. I assume the world is stationary and that the individual's choice, while forward-looking, does not depend on his past decisions. This allows me to extend the model to choice set heterogeneity and a rich set of option attributes that may vary across individuals.

The treatment of choice set variation is built on insights from consumer choice literature. Similarly to this study, this literature asks whether people choose as they do because they like the option, or simply because they don't pay attention to the alternatives. There are several competing approaches to identification of these two parts, drawing on specific assumptions about what goods the consumers pay attention to (Conlon and Mortimer (2008)), observed asymmetries in their responsiveness to changes in option attributes (Abaluck and Adams (2017)), and external data on workers' choice sets (Goeree (2008), Ho et al. (2017)). Applying choice set variation to a labour market model is simpler because, in contrast to the universe of health insurance plans or personal computers, the set of possible options is small. On the other hand, the fact that migration is a relatively rare event, and the lack of any direct data on firms' hiring strategies, makes estimating my model more difficult, because some of the identification strategies based on external information and data on many repeated purchases are not applicable.

2.2 Data

One of the main points of this chapter is to distinguish between two types of migration: before and after job search. This is possible by using monthly panel data that follows the individuals when they move. This section describes this dataset, and defines the education categories and regions used in the analysis. I also describe the data on choice set variation.

My main dataset comes from the 1996 wave of the Survey of Income and Program Participation (for a discussion of why I use this particular wave, see Appendix A.1). It tracks a nationally-representative sample of US inhabitants over the period of 4 years, providing monthly data on their income, employment, and residence, among others. I also use four supplementary datasets (County Business Patterns, IPUMS, the Bureau of Labor Statistics, FRED) to capture firms' recruitment behaviour and regional variables such as housing costs and unemployment rates. The full list of variables and their sources is summarised in Table A.3 in Appendix A.1.

2.2.1 The SIPP

I use the Survey of Income and Program Participation (SIPP) because of its three key features: it is (i) a panel with (ii) monthly data that (iii) follows its respondents when they move. These three features are necessary to analyse the relationship between cross-regional job search and migration. I need panel data in order to observe the outcomes of the worker's job search, and I need information on his geographic movement to be able to study migration. Finally, the data needs to be of sufficiently high frequency to allow me to link the two, i.e. to be able to see the worker's labour market outcomes just after the move. SIPP is, to my best knowledge, the only major dataset that combines these three features (for a detailed comparison of the different potential data sources, see Hernández-Murillo et al. (2011)).

I restrict my attention to a subsample of adult working men. The aim of this paper is to understand the role of job search in migration decisions, which is why I exclude individuals below the age of 24 (and college-related migration) and above 50 (to avoid migration motivated by retirement), as well as those adults with no

employment spell over the duration of the panel. I focus on men to reflect the fact that 75% of respondents in the data live in households where the primary earner is male and is thus likely to have a disproportionate influence on the migration decision.

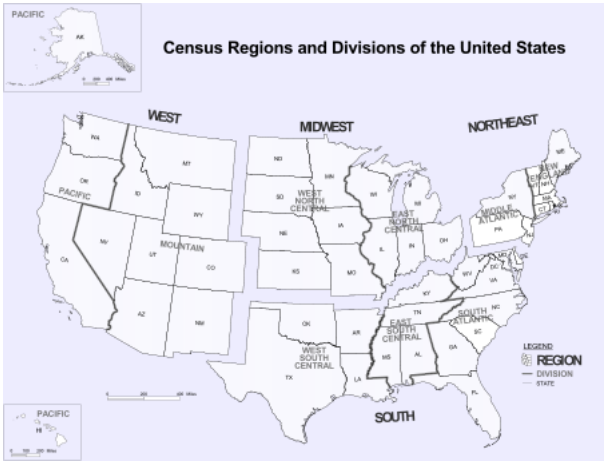
The descriptive statistics on the subsample used in my analysis are summarised in Table A.1 in Appendix A.1. There are 16,720 working-age males in my sample, about 13% of which did not finish high school, while a quarter of the sample has completed a four-year college degree or more. The table shows that there is considerable variation between the education groups in terms of individual and household income, employment, and migration propensity. At the same time, however, the three groups are similar in terms of their age, racial and urban profile, making it easier to abstract from these variables.

Throughout the paper, the outcome of primary interest is the location-employment combination (based on the four large US regions) chosen by the workers each month. Unsurprisingly, employment is more frequently observed than unemployment, and the most frequent location choice corresponds to the most populous region of the US (the South). Cross-regional migration is a relatively rare event, with 2.5% of the workers choosing to become employed in an away region and 0.5% of opting for unemployment away, during the 4 years of the sample. A full breakdown of the distribution of individuals' choices can be seen in Table A.2 in Appendix A.1.

Defining education categories This paper works with three education categories: high school dropouts, high school graduates, and college graduates. They are based on SIPP's *eeducate* variable which captures the highest degree or grade completed. The mapping is the following:

- high school dropouts: individuals that have not received a high school diploma or equivalent
- high school graduates: individuals who have graduated from high school and may have vocational degrees or some college

Figure 2.1: The 4 large census regions of the US



Source: US Census Bureau, Geography Division.

- college graduates: individuals that hold a bachelor's degree (graduated from a four-year college) or more

In some parts of the paper, I merge the first two categories to only distinguish between the less educated (i.e. less than a four-year college degree) and the more educated (i.e. college graduates).

Defining migration The SIPP contains several levels of geographic information. Each household is located in a state and metropolitan area (if applicable), alongside a within-household address index which changes every time a household moves, even if it is within the same neighborhood. The majority of moves fall into the latter category - 63 % of moves are within the same urban or rural area[3].

For the purposes of my analysis, I work with two definitions of migration: inter-state and cross-regional. Inter-state migration is widely used in literature when more detailed commuting zones are not available. In my dataset, about 13% of all moves are between states. Most of the analysis is conducted at the level of regions (Figure 2.1). Because of the larger geographic areas, this type of migration is somewhat rarer: just below 50% of all the inter-state moves are moves between regions.

[3]The dataset distinguishes between urban and rural households, and identifies the metropolitan area of the urban households. A "rural area" is thus defined as the area of each state that is not urban. The 63% hence refers to moves within the same metropolitan area, plus moves between different parts of the same state which are rural.

2.2.2 Data on recruitment strategies

In the absence of detailed information on the geographic radius of recruitment, I draw on the findings of the human resources literature to construct proxy variables.

The literature suggests that firms that hire across a wide geographic radius are likely to fall into one of two categories. They may be looking for workers of specific skill that may be hard to find in their local labour market. The studies by Barron et al. (1985) and Russo et al. (1996) have demonstrated that firms that are looking for more educated workers search across greater distance, using more resources. Second, the search radius is also going to increase with the size of the firm itself. Tardos and Pedersen (2011) and Russo et al. (1996) show that larger firms spend more money on recruitment, interview more applicants, and reach out into more distant regions. Moreover, these studies point out that larger companies are more likely to use formal recruitment channels, such as job board or newspaper postings, and professional recruitment agencies, that are more likely to reach outside of the local labour market. Small companies, on the other hand, are much more likely to rely on informal channels, such as referrals and word of mouth, as well as using only one recruitment channel (Barber et al. (1999)).

Based on these findings, I construct three proxies for the geographic radius of recruitment. Using firm data from the County Business Patterns dataset for the years 1996 – 1999, I calculate the average company size (as measured by the number of employees) by sector. I use the firm location information in this dataset to calculate an index of spatial concentration of each sector. The index (based on Ellison and Glaeser (1997)) compares the total employment in a sector in a county against the total population in each county of the US. The higher the number, the more clustered the given sector is. Finally, I use more recent data on online vacancy posting that looks at what proportion of all vacancies, within each occupation, are advertised on the Internet (Hershbein and Kahn (2018)). Even though the data in question is from a more recent period (2007, 2010-2014), this variable provides a more direct evidence on the willingness of a firm to disseminate their vacancy information as widely as possible.

I match these proxies on firm recruitment with the worker dataset. The SIPP provides detailed data on sector and occupation for each employed worker, distinguishing between 409 occupation categories and 233 different sectors. I aggregate these into 21 occupation and 13 sector groups, using the internal grouping in the census codes for the occupation and sector variables. These are then matched with the occupation and sector data used in the CBP and online vacancy data. The result is a merged dataset which records recruitment proxies for each worker, based on his sector and occupation group.

2.3 Stylised facts about regional migration and education

This paper is motivated by two stylised facts: the less educated move less, and they are less likely to move for a specific job. The education differences in migration propensity are well-known in the literature (Greenwood (1997)), but the pattern in the type of migration has not been established before[4]. In this section, I document both stylised facts in my dataset and explain the categorisation of migration into speculative and job-specific. I also discuss the wider implications of the different migration behaviour.

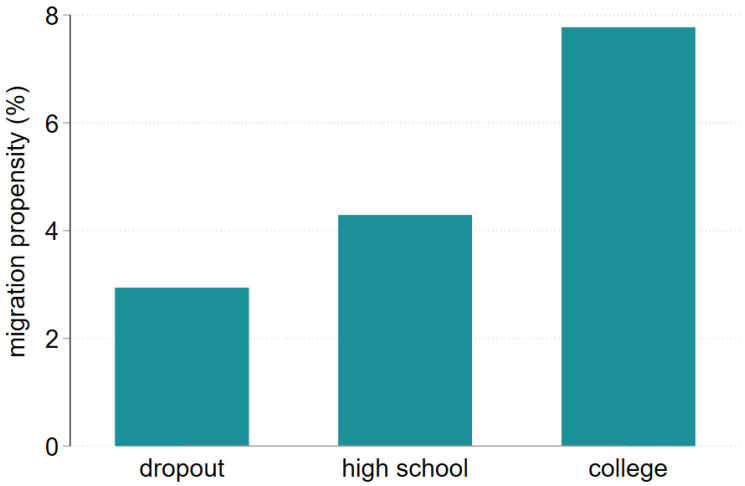
2.3.1 The propensity to migrate

Migration is a rare event. The probability that an adult in the SIPP sample of US inhabitants moves into a different state in the time period 1996-1999 is about 4.5%. In general, the inter-state migration propensity reported in the literature ranges between 3 and 7%, and has been in decline since the 1980s (Molloy et al. (2011)).

Despite that, a permanent feature of the migration statistics is the large gap in migration propensity between the more and the less educated (Greenwood (1997)). In my sample, college graduates are up to 3 times more likely to move into another states than high-school dropouts (Figure 2.2). These numbers are

[4]Using revealed preference data. Basker (2018) looks at workers' stated motivation to move, and his findings corroborate the patterns presented in this chapter.

Figure 2.2: Propensity to migrate into another state, US, 1996-1999



Calculated from the Survey of Income and Program Participation, 1996-1999 panel. Sample: adults (anyone over the age of 18). States refer to the 50 US states.

similar to the estimates in other studies: for example, Hernández-Murillo et al. (2011) report the migration propensity to be 2.6% for high school dropouts and 5.7% for college graduates.

This effect is robust to demographic characteristics, household composition, employment status, and distance migrated. The first two columns in Table 2.1 present the results of a binary logit model of the migration decision, in which the coefficient on education remains positive and statistically significant even after including a set of controls; further robustness checks can be found in Figure A.1 in Appendix A.1.

2.3.2 Speculative migration vs. migration for a specific job

The second stylised fact motivating this paper is about *how* individuals move. My panel data allows me to distinguish between migration followed by employment (*migration for a specific job*) and migration followed by unemployment (*speculative migration*) by looking at the timing of their migration and employment. I categorise all individuals that became employed within one month of their move as having migrated with a job in hand, while all migrants that were unemployed for at least a month after moving are assumed to have moved speculatively and had to search

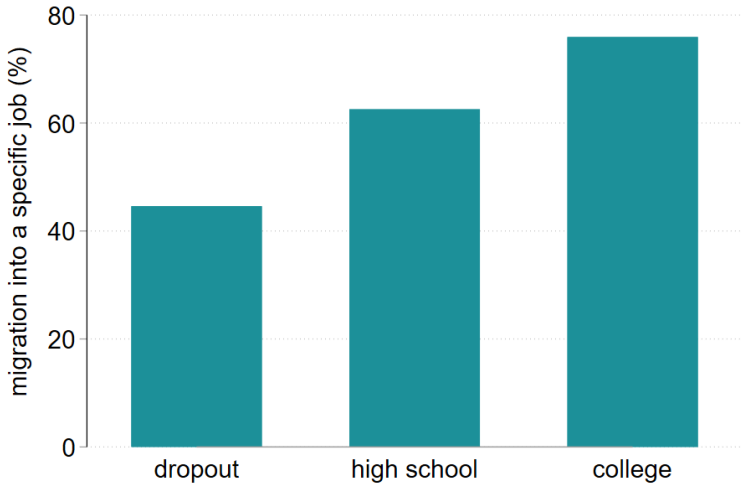
Table 2.1: Binary logit models of propensity to migrate and the type of migration, between regions

	(1) migration propensity	(2) migration propensity	(3) speculative migration	(4) speculative migration
college	1.1003***	1.1191***	-0.7856**	-0.7035**
	(0.0969)	(0.0974)	(0.3222)	(0.3325)
employed in previous month	-1.0738***	-0.9871***	-1.7744***	-1.8898***
	(0.1368)	(0.1382)	(0.3894)	(0.4073)
married		-0.1835		-0.5011
		(0.1216)		(0.3824)
kids		-0.1064		-0.5243
		(0.1177)		(0.4043)
young		-0.5989***		0.4678
		(0.0980)		(0.3247)
unemployment rate			1.6547***	1.6317**
			(0.6373)	(0.6549)
constant	-6.4955***	-6.0552***	-0.4140	-0.1427
	(0.1297)	(0.1417)	(0.4101)	(0.5125)
N	16720	16720	428	428

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$
The sample is restricted to working men between the age of 25 and 50. Columns (1) and (2) model the probability of moving into a different census region. Columns (4) and (5) explain the likelihood that, conditional on moving, the migration is speculative (as opposed to for a specific job). "Employed in previous month" is a dummy variable equal to 1 if the worker was in employment the month before migration. "Unemployment rate" refers to individual-specific probability of being unemployed in any given month.

upon arrival. Inevitably, this categorisation does not capture the two types of migration perfectly. Some workers that move speculatively and find a job within the first four weeks of their arrival will be mis-classified as moving with a job in hand, while those that do move with a job in hand but take several months to commence employment will be mis-classified as speculative migrants. Figures A.5 and A.6 in Appendix A.2 present robustness checks in the data, showing that these misclassifications are unlikely to be systematic.

Figure 2.3: Probability of moving into a specific job (conditional on moving), US, 1996-1999



Calculated from the Survey of Income and Program Participation, 1996-1999 panel. Sample: adults (anyone over the age of 18). Migration for a specific job refers to migration followed by employment.

The data shows that conditional on migration, college graduates are up to 2 times more likely to move into a specific job than high school dropouts (Figure 2.3). While more than half of high school dropouts that move across states have no job lined up, the number is only about a fifth for college graduates.

Table 2.1 establishes this stylised fact formally. I run a binary logit model of whether a migrant moves speculatively or not. The estimates show that the negative relationship between education and speculative migration is a robust one (see also Figures A.2, A.3 and A.4 in Appendix A.2). In particular, the less educated are less likely to move into a specific job (i.e. be employed after migration) even after controlling for their greater unemployment rate overall[5].

[5]In fact, there are two different effects that I have to control for: employment before moving, and overall unemployment rate. The unemployed move more (Herzog et al. (1993)), which is why the binary regressions in Table 2.1 include employment status in previous month. Moreover, since the unemployment rate of the less educated tends to be higher in general, I also have to control for this when estimating their propensity to migrate speculatively. This is captured by the individual-specific unemployment rate before moving ("unemployment rate").

2.4 A discrete choice model of employment and location

Are the two stylised facts – the relatively low migration propensity and the low probability of migrating into a specific job – describing the migration of the less educated related? To answer this question, I estimate a multinomial model of migration and employment choice. Each worker makes optimal decisions about where and whether to work, but his options are limited by whether the firms in his occupation and industry hire cross-regionally. At the same time, employment-location options may differ in the offered income, living costs, and migration costs. This setup allows me to decompose the observed behaviour into variation in the attributes of the different employment-location options, and variation in the availability of these options. The results show that, among other reasons, the less educated move less because they receive relatively fewer opportunities to do so. Moreover, I demonstrate that controlling for choice set heterogeneity is important for producing unbiased estimates of workers' preferences.

2.4.1 General determinants of migration

The first step in explaining the education differences in migration propensity is to understand why people migrate in general.

Income, and especially earnings differences across regions, is one of the most important drivers of migration both in theoretical models and in terms of the focus of the existing literature. Most the existing models of migration follow in the vein of Sjaastad (1962) and Harris and Todaro (1970), in which mobility is the response to income and job differences between regions. This has been confirmed in the empirical literature: a seminal paper by Kennan and Walker (2011) demonstrates the importance of expected income for migration decisions of young men, while Borjas et al. (1992) highlights the role of state-level differences in the return to skill.

Given than earnings vary significantly with education, income is one of the leading explanations for the education migration gap. In general, average earnings rise with education, making it more likely that the more educated can afford to pay

Table 2.2: Regional variation in monthly wages and non-employment rates, by education

	region			
education	Northeast	Midwest	South	West
	nominal wages			
dropouts	1,704	1,637	1,504	1,549
high school graduates	2,351	2,247	2,079	2,390
college graduates	4,145	3,691	3,629	4,066
	non-employment rate (%)			
dropouts	19.86	20.53	18.05	23.34
high school graduates	10.71	9.41	10.83	12.27
college graduates	6.27	5.97	5.72	6.95

Calculated from the Survey of Income and Program Participation, 1996-1999 panel. The sample are working males between the age of 25 and 50.

migration costs. Higher returns to skill (Borjas et al. (1992)) and larger regional dispersion of wages of the more educated (Amior (2015)) have also been suggested as explanations. However, several studies have pointed out that regional differences in employment rates in particular are in fact larger for the less educated (Hoynes (1999), Gregg et al. (2004)), making their small migration propensity something of a puzzle.

The differences in wages and unemployment rates across US regions are considerable (Moretti (2011)). Table 2.2 documents these differences for the individuals in my data. A high school graduate may gain on average $300 in his monthly nominal wage by moving from the lowest-wage to the highest-wage region (the South and the West, respectively). The average gain for a college graduate is even higher at $500 a month (for a move from the South to the Northeast). Differences in non-employment rates[6] are also non-negligible. The difference between the best and worst performing region is more than 5 percentage points for a high school dropout and almost 3 percentage points for a high school graduate.

At the same time, the results of direct comparisons of income variation by education are mixed. While the absolute wage differences between regions are larger for the more educated, Table 2.3 shows that this difference relative to nominal wage

[6] I define non-employment rate as 100 − employment rate.

Table 2.3: Across-state variation in monthly wages and non-employment rates, by education

	average wage		maximum wage		non-employ. rate (%)	
	dropouts	college	dropouts	college	dropouts	college
mean	1,913	4,361	6,307	22,469	19.3%	5.8%
standard deviation	338	824	4,164	10,097	0.07	0.03
coeff. of variation	0.1769	0.1891	0.6603	0.4494	0.37	0.51

Calculated from the Survey of Income and Program Participation, 1996-1999 panel. The sample are working males between the age of 25 and 50.

(as measured by the coefficient of variation) is comparable, if not greater for the less educated. The opposite is true for non-employment rates: the regional variation (as measured by standard deviation) is larger for the less educated, although the coefficient of variation is actually larger for college graduates.

However, determinants of migration go beyond expected income: amenities (climate, culture, quality of schools, etc.), living and housing costs, retirement and other lifecycle considerations, and proximity to friends and family. Kennan and Walker (2011), about half of the moves in their sample are subsequently reversed by return migration, highlighting the importance of residence history and social networks. Moving for college is another widely-studied motive for migration: Gregg et al. (2004) hypothesise that this experience of moving makes college graduates more mobile for the rest of their lives. Clark and Hunter (1992) and Whisler et al. (2008) show how individuals' changing preference for amenities over their lifetime shapes migration behaviour. Rappaport (2007) estimates that 44% of the variation in population growth in metropolitan areas over the past 50 years is due to winter weather alone.

Similarly, the education gap in migration may exist because individuals with less education face higher migration costs (Greenwood (1997), Moretti (2011)), or because they have different preferences over locations (Diamond (2016)). Researchers hypothesised that the stylised fact is simply a spurious correlation arising from those that self-select into college education also self-selecting into migration. Malamud

and Wozniak (2012) bring evidence that the effect from education to migration propensity is causal, but the exact mechanism remains unresolved.

Since the seminal work of Roback (1982), models of local labour markets assumed that workers' utility depends on a combination of all of these factors. In this framework, the worker weighs up expected income, local amenities, migration costs, social networks, etc. to maximise his utility. Relevant for this paper is the fact that, under these assumptions, the migration decision depends on labour markets even if employment is not the primary reason for moving.

Nevertheless, survey data suggests that job-related reasons are one of the main factors behind migration decisions. Using the Current Population Survey, which asks respondents about their reason to migrate, Amior (2015) documents that employment is a significant reason for moving across education and age categories. Almost a third of all cross-county moves, and just over a half of all cross-state moves were primarily for job-related reasons. This number increases in the distance migrated, which is encouraging given that this chapter uses data on migration across the four census regions of the US.

Does migration pay off? Data on post-migration outcomes confirms that a large majority of migrants benefit from their move to a new labour market. I calculate the impact of migration on wages by comparing an individual's wage before and after migration. The average increase in nominal monthly wages varies between 19% to 46%, with only about 30% of workers taking a pay cut. In this calculation, I do not control for self-selection into migration, but the existing literature has in general documented a positive effect of migration on worker outcomes. Rodgers and Rodgers (2000) find positive returns to moving for every postmove year, estimating the wage premium at 20% on average. The data on unemployment incidence and duration paints a similar picture: migrants are more likely to be unemployed following migration, but they stay unemployed for shorter, compared to stayers. The full results of the ex-post benefits to migration are summarised in Tables A.4 and A.5 in Appendix A.3.

Table 2.4: The list of options in the logit model of the employment-location choice

option	employment status	region
1	employed	Northeast
2	unemployed	Northeast
3	employed	Midwest
4	unemployed	Midwest
5	employed	South
6	unemployed	South
7	employed	West
8	unemployed	West

An individual can decide to be employed or unemployed in each of the census regions of the US (conditional on that option being in his choice set). The dependent variable in the model is a dummy vector, the elements of which are equal to 1 if a particular option was selected, and 0 otherwise.

2.4.2 Discrete choice model of location and employment

Typically, empirical literature on migration focuses on two decisions: whether to migrate, and where to migrate to. Conditional on location decision, models of the labour market then estimate the worker's decision whether to accept a job or remain unemployed.

In this paper, I nest these three choices within a single model of employment-location choice. Every period, the worker decides where to live and whether to work, which determines his employment status and may result in migration. Using the four census regions of the US as locations, the model works with eight outcome variables, listed in Table 2.4. The dependent variable is a 8×1 vector, the elements of which are equal to 1 if the given option was selected, and 0 otherwise.

This discrete choice model seeks to decompose the observed choice probabilities of these eight options into variation in the attributes of these options, and differences in the choice sets available to individuals. The interpretation of the model is based on the random utility model (RUM) framework. RUM assumes that the decision-maker is a utility-maximiser, so that the observed choice probabilities are proportional to the utility differences between options, once choice set heterogeneity is controlled for.

As a result, the estimated coefficients are informative both about workers preferences over income, migration costs, etc., and how likely a given option is available to them.

The theoretical interpretation of this model is that the worker is continuously searching (including on the job) and every period re-optimises his decision about whether and where to be employed. Every period, he receives at least one job offer from his home region (if he is employed, he is offered to stay in his current job, but he also may be offered a new job in his home region). Depending on the recruitment behaviour in his occupation and sector, he may or may not also receive job offers from away regions. The worker can also always choose to become unemployed at home or in any of the other three regions (this would constitute speculative migration). Every period, he weighs up the utilities of all available options, and chooses the optimal location-employment combination. This then results in the employment-location outcomes analysed in this model.

2.4.2.1 Modelling utility: conditional choice probability

First, I define the choice probability conditional on the option being available to the individual.

The starting point is to decompose the utility of each option into an observed and unobserved part:

$$T_{it} = O_{it} + \epsilon_{it} \tag{2.1}$$

where T_{it} denotes total utility of option t for worker i; O_{it} denotes the observable part of utility, and ϵ_{it} captures the unobservable component, which is treated as random by the econometrician.

Our specification of the observed part of utility, O_{it}, is based on the seminal model of migration by Sjaastad (1962). In this standard migration model, migration decision is a form of human capital investment: the individual chooses to move if the benefits (increased income, better amenities) outweigh the costs (living costs in the new location and migration costs, which can include both direct financial losses and the social and psychological costs). I define X_{it}, the components of O_{it},

as the monthly income from option t, local rent (to capture living costs), and a dummy for migration, which will be my catch-all variable for all types of costs associated with moving. I also include a dummy for unemployment, to reflect the fact that just like migration, unemployment also often costs the worker more than just the pure financial loss of forgone income.

X_{it} varies both across the eight options and across individuals. For example, if $t = \{\text{employment, South}\}$, the income will capture the current wage offer for a job in the South, rent will correspond to the local rent, the migration dummy will equal one if the worker currently does not reside in the South, and the unemployment dummy will be equal to 0 since this option involves employment. On the other hand, if $t = \{\text{unemployment, West}\}$, O_{it} will consist of the worker's reservation wage in the West, the local rent, unemployment dummy equal to 1, and the migration dummy either 0 or 1 depending on whether worker i currently resides in the West or not.

The part of the utility unobserved by the econometrician, ϵ_{it}, corresponds to the worker's idiosyncratic preferences over regions and employment options: individual-specific family ties, location history, particular amenity preferences (love of skiing vs. love of sandy beaches), the worker's disutility of work.

Worker chooses the option that gives him the highest utility. As a result, the probability of choosing a particular option s, P_{is}, is equal to the probability that its utility, T_{is} is greater than that of the other options:

$$\begin{aligned} P_{is} &= \text{Prob}(T_{is} > T_{it'} \ \forall \ t' \neq s) \\ P_{is} &= \text{Prob}(O_{is} + \epsilon_{is} > O_{it'} + \epsilon_{it'} \ \forall \ t' \neq s) \\ P_{is} &= \text{Prob}(\epsilon_{it'} - \epsilon_{is} < O_{is} - O_{it'} \ \forall \ t' \neq s) \\ P_{is} &= \text{Prob}(\epsilon_{it'} - \epsilon_{is} < \beta X_{is} - \beta X_{it'} \ \forall \ t' \neq s) \end{aligned} \tag{2.2}$$

where β denotes the worker's preferences over the components of X_{it}.

As equation (2.2) shows, P_{is} is a function of the probability distribution of the unobserved utility component ϵ_{is}. Assuming that ϵ_{is} is type I extreme value

Table 2.5: Full and restricted choice sets; example for a worker resident in the West

options in...	
full choice set	restricted choice set
1 [employed, Northeast]	-
2 [unemployed, Northeast]	2 [unemployed, Northeast]
3 [employed, Midwest]	-
4 [unemployed, Midwest]	4 [unemployed, Midwest]
5 [employed, South]	-
6 [unemployed, South]	6 [unemployed, South]
7 [employed, West]	7 [employed, West]
8 [unemployed, West]	8 [unemployed, West]

The left hand column lists all the employment-location options available to a worker who does receive job offers from other regions. The right hand side column lists the options available to a worker who only receives job offers from his home region (is not capable of cross-regional job search).

distributed, the discrete choice model is a logit and the choice probability P_{is} has a closed-form solution (McFadden et al. (1973)):

$$P_{is} = \frac{exp(\beta X_{is})}{\sum_{t=1}^{8} exp(\beta X_{it})} \tag{2.3}$$

2.4.2.2 Modelling choice sets: recruitment behaviour

The second part of the total (unconditional) choice probability is the likelihood that a given option lies in the individual's choice set. This addition to the standard multinomial logit draws on the empirical IO literature (Abaluck and Adams (2017), Goeree (2008)), which explicitly recognises that the consumer's optimal choice may be skewed by the set of goods he pays attention to. Allowing for choice set heterogeneity makes sense in my model because a worker can only choose to be employed if he has received a job offer. If some workers are more likely to receive job offers than others, they will be able to choose employment options more often. And if workers differ in their probability of receiving a job offer from an away region, this will impact their ability to migrate for a specific job.

I impose a specific structure on the choice set heterogeneity in this model. A worker may face one of two choice sets: a full choice set with all eight options, or a restricted one which does not include employment in regions other than his home one. This means that a worker can always choose to be unemployed in any region, because that decision does not depend on the behaviour of employers. He can also always choose to be employed in his home region, because I hypothesise that it is easier to find a job in a familiar labour market compared to searching in a distant region. Table 2.5 lists an example of the options in these two choice sets for a worker who resides in the West.

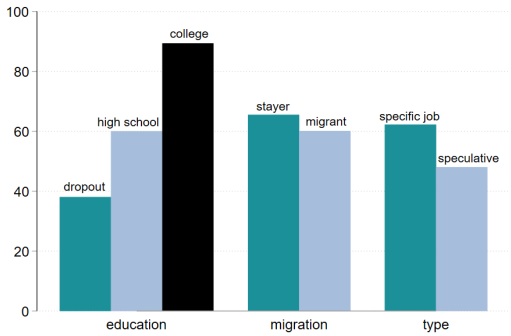
A worker's choice set depends on the recruitment practices in his sector and occupation, Z_i. I use three variables to proxy for these: average firm size and spatial concentration (by sector), and future online vacancy posting (by occupation). The first two variables[7] are selected based on findings in the human resources literature, which show that firms hire within a broader geographic radius if they are looking for skill or quantity of labour that is relatively harder to find in their local labour market. This holds true for firms that are large or in spatially concentrated sectors. The data on online vacancy posting (Hershbein and Kahn (2018)) from the late 2000s is used to pin down occupations for which the employers are willing to search more broadly. As not all vacancies are posted online, and there is significant variation across education groups and occupations, this variable is likely to capture some of the relevant variation in recruitment strategies.

The empirical link between migration and recruitment proxies are documented in Figure 2.4. Its three panels plot the average values of firm size, spatial concentration, and online vacancy posting, separately for migrants and stayers, and for the two different types of migration. It shows that migrants on average work in more spatially concentrated sectors, and in occupations that are more likely to be advertised online. However, this relationship is mostly driven by those that move into a specific job: moving with a job in hand is more likely for workers working for larger and more spatially concentrated companies that tend to post their vacancies online.

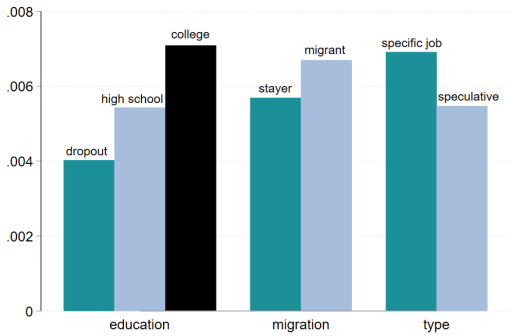
[7]They are constructed from the Country Business Patterns dataset; see section 2.2 for more detail.

Figure 2.4: Recruitment proxies and migration

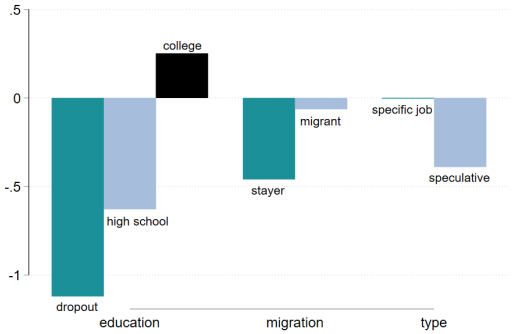
(a) Average firm size



(b) Spatial concentration, by industry



(c) Online vacancy posting



Recruitment behaviour and worker characteristics. Panel (a) summarises average firm size by education, migration propensity, and type of migration (conditional on migration). Panel (b) plots the average index of spatial concentration by education, migration propensity, and type of migration (conditional on migration). The higher the value of the index, the more spatially concentrated the sector. Panel (c) summarises the gap between all vacancies and those posted online. The variable is calculated as log odds ratio: it is the log of odds ratio of the probability density mass of each occupation category in the two datasets. Negative number means that the vacancies are under-represented in the online data, and vice versa. Data sources: SIPP (1996-2000), County Business Patterns (1996-1999), JOLTS (2007-2014), Burning Glass (2007-2014).

I model choice set heterogeneity as a binary variable equal to 1 if a worker can choose from all options, and 0 if his choice set is restricted. Assuming that the unobserved part of the choice set variation is also type I extreme value distributed, the probability that worker i has a full choice set, μ_i, is a logit function[8] of Z_i:

$$\text{Prob (full choice set} = 1) = \mu_i = \frac{exp(\gamma Z_i)}{1 + exp(\gamma Z_i)} \tag{2.4}$$

where γ stands for the coefficients of Z_i.

[8]The RUM interpretation does not apply in this case. Choice set is not something the worker can choose, it is given by factors exogenous to him.

2.4.2.3 Unconditional choice probability

The overall probability that the worker chooses option s is a combination of the conditional choice probability P_{is} and the probability that it is in his choice set, which is a function of μ_i. The choice set itself depends on the worker's home region, and so the exact unconditional probabilities will be individual-specific. Let's return to the example of a worker residing in the West of the US, whose two possible choice sets were described in Table 2.5. He can only choose to be employed in the Northeast (option 1) if he receives cross-regional job offers, that is, if he has a full choice set. This will happen with probability μ_i. Conditional on option 1 being in his choice set, he will choose it with the probability P_{i1}. The unconditional probability of choosing employment in the Northeast for this worker, R_{i1}, is thus:

$$R_{i1} = \mu_i P_{i1} = \frac{exp(\gamma Z_i)}{1 + exp(\gamma Z_i)} \frac{exp(\beta X_{i1})}{\sum_{t=1}^{8} exp(\beta X_{it})} \tag{2.5}$$

The unconditional choice probability for the other employment-away choices (in this case, options 3 and 5) can be derived analogously.

In contrast, the option to be unemployed anywhere, or employed locally, is available under both choice sets. The unconditional choice probability for options $s = \{2, 4, 6, 7, 8\}$ will be thus a weighted average of the conditional choice probabilities under those two scenarios. If the worker faces a full choice set (with probability μ_i), he will choose to stay home as unemployed with the probability P_{i8}, where P_{i8} is the conditional probability of choosing $s = 8$ from among all the eight options. With probability $1-\mu_i$, the worker receives no away employment offers, and the conditional probability of choosing unemployment at home is S_{i8}. S_{i8} is different from P_{i8}, because it calculates how optimal option 8 is among five, not eight, options. For the worker in our example, the unconditional choice probability for option 8, R_{i8}, is:

$$R_{i8} = \mu_i P_{i8} + (1 - \mu_i) S_{i8} \tag{2.6}$$

where

$$S_{i8} = \frac{exp(\beta X_{i8})}{\sum_{t=\{2,4,6,7,8\}} exp(\beta X_{it})} \tag{2.7}$$

The unconditional choice probabilities for the other options in the restricted set, $s = \{2, 4, 6, 7\}$ can be derived similarly.

2.4.3 Estimation

2.4.3.1 Identification

The goal of the this model is to decompose the variation in workers' choices into variation in the utility of the options, and variation in their availability.

The workers' preferences are identified from the variation in attributes across options. The coefficients on the various attributes reveal how much does the probability of choosing an option change when the value of an attribute changes. I use seven option attributes: income, regional unemployment rate, unemployment dummy, migration dummy, switching cost dummy, regional rent, and regional house price. The data sources on these variables are described in section 2.2 and Table A.3 in Appendix A.1.

Because the estimated model is a conditional logit, normalisation of coefficients with respect to some arbitrary option is not necessary. Normalisation of the scale of utility is done implicitly by assuming that the unobserved part of utility, ϵ_{it}, is type I extreme value distributed with variance equal to $\pi^2/6$. The coefficients in the utility part of the model should be interpreted as indicating preferences about an attribute relative to the unobserved part of utility.

The challenge in estimating the variation of the choice sets lies in the fact that they are unobservable to the researcher. In the previous section, I built a model where the probability of full choice set was denoted as the latent variable μ_i. We learn about it indirectly, by embedding it in the unconditional probabilities (2.5) and (2.6).

Identification of μ_i rests on two components. It is identified from the variation in cross-regional recruitment practices, as proxied by variables Z_i. It also depends on the specific choice set structure I imposed in the model, which manifests itself in the functional forms of the unconditional probabilities. In particular, I assume that unemployment anywhere and employment at home are always in the worker's choice set, but employment away may not be.

An important assumption for the identification of both parts of the model is that the errors are drawn from a type I extreme value distribution. On a practical level, logit errors determine the specific, closed-form expression for the conditional

probabilities P_{it}, S_{it}, and the choice set probability μ_i. In terms of identification, assuming logit errors means assuming that the unobserved part of utility, ϵ_{it}, is independent across alternatives and over time, which in my case means that the worker's idiosyncratic preferences about working in the Midwest are independent from his idiosyncratic preferences about being unemployed in the Midwest. It also means that liking employment in the Midwest in the previous period has no impact on his preferences about this option the next period. I relax the implications of the logit error assumption by allowing for option-specific constants[9], so that certain regions and employment states can be more attractive for everybody in all time periods.

There are two exclusion restrictions for the identification of choice sets. First, that firms recruit all the workers in the same sector and occupation the same way. Second, that workers do not choose their sector and occupation because of the migration opportunities it offers – worker's choice of industry and occupation is taken as given at the point of his employment-location choice.

2.4.3.2 Imputing income

The model includes monthly income because it is an important part of the worker's migration and employment decision. However, some parts of the income variable have to be imputed before the full model can be estimated.

For options that include employment, the income variable corresponds to the worker's monthly wage offer. This will be unobserved unless it is accepted. For unemployment options, income equals the worker's reservation wage. It reflects non-labour income (social security, savings) and discounted expected future earnings, and it is also unobserved.

Reservation wages are calculated as the smallest observed wage in each worker-region bin. Workers are sorted into bins by occupation category, experience (captured as 5-year age intervals), education, year, and region of residence. The smallest observed wage in each bin is a super-consistent estimate of the reservation

[9]In discrete choice models, only differences in utility matter, which is why I cannot add a constant to the utility of every option. Identification of option-specific constants requires normalisation. I set the constant of option 8 (unemployment in the West) equal to 0.

wage (Keane et al. (2011)), assuming that the reservation wages are uncorrelated with any relevant unobservables, such as differences in propensity to migrate or location preferences.

Monthly wage offers are extrapolated from the observed monthly earnings. I need to control for two sources of selection bias: the probability that a worker would choose to be employed, and that he would choose the given region. To do so, I follow the method proposed by Dahl (2002), which expands the seminal work by Heckman (1979) to multi-regional labour markets (see Appendix A.4 for more detail). I estimate a wage offer equation separately for each of the four regions, which is then used to impute the individual-specific average monthly wage offer.

2.4.3.3 The likelihood function

The likelihood function that rationalises the data follows the standard form for a discrete choice model:

$$L(\beta, \gamma) = \prod_i \prod_t (R_{it})^{d_{it}} \tag{2.8}$$

the log-likelihood of which is

$$logL(\beta, \gamma) = \sum_i \sum_t d_{it} log(R_{it}) \tag{2.9}$$

R_{it} denotes the unconditional probability of worker i choosing option t, and d_{it} is a dummy variable which equals 1 for the selected t and 0 otherwise. This expression assumes that the unobserved components of utility are independent across options, and that the workers' decisions are independent of each other.

The assumption that both parts of R_{it}, conditional choice probability and the choice set probability, are logit, means that the log-likelihood has a closed-form solution. The model is estimated by maximum likelihood.

2.4.4 Results

This section presents estimates of the discrete choice model. It explains the worker's choice over employment and location as a function of the attributes of these options and the worker's choice set, proxied by data on firm recruitment. The

dependent variable is a dummy signalling the workers' choice over eight options, employment or unemployment in each of the four large US Census regions. There are two main results. First, the estimates show that more educated workers have more migration opportunities; second, controlling for these differences is vital for obtaining unbiased preference estimates.

2.4.4.1 The role of job search in migration decisions: analysing the impact of choice set heterogeneity

The results of the model are summarised in Table 2.6. The first column is a standard conditional logit. It serves as a benchmark for the augmented model estimated in the second column. It explains the observed choices as a function of the option attributes, implicitly assuming that all workers always have full choice sets (so there is no choice set heterogeneity). It shows that workers prefer options that give them higher income (although the income coefficient is not statistically significant), better amenities (as captured by house prices), and with lower living costs (rents). Employment options are preferred to unemployment. The coefficients on the education-specific migration dummies show that migration costs are greatest for the least educated. The model also estimates a relatively large and significant switching cost. Such strong persistence is expected when looking at month-to-month decisions.

The second column in Table 2.6 allows choice sets to vary. These results are presented in two panels: the upper summarises workers' preferences, and the lower the variables that determine their choice set.

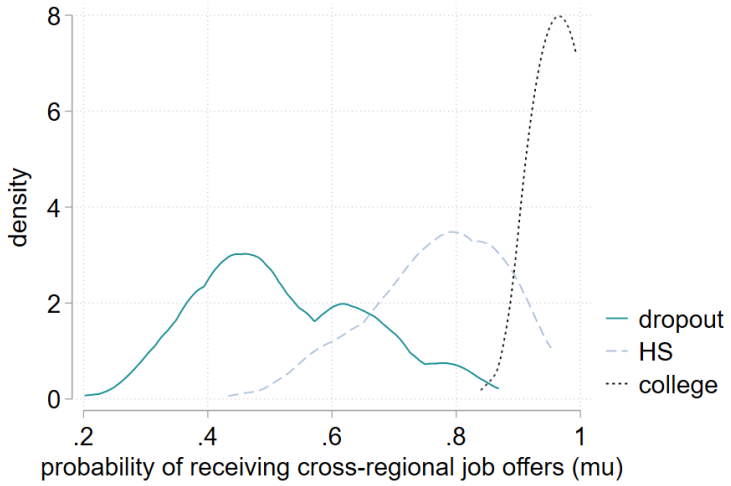
Starting from the bottom, the statistical significance of the recruitment proxies tells us that choice set variation matters. Workers coming from industries and occupations that are more spatially concentrated, with larger firms, and where vacancies are more likely to appear online, are more likely to receive a cross-regional job offer, and choose to be employed away. At the same time, the education-specific constants in the second column show that the average probability of having access to the full choice set increases in education, suggesting there are other factors behind choice set variation not covered by my model.

Table 2.6: Multinomial logit model of the choice of employment or unemployment in one of the 4 large US regions, allowing for choice set heterogeneity.

	full choice set	heterogenous choice set
utility		
Dependent variable: probability of choosing one of the 8 options		
income	0.032	0.144***
	(0.054)	(0.029)
regional unemp. rate	0.941***	-0.611***
	(0.336)	(0.061)
unemployment dummy	-0.821***	-1.253***
	(0.192)	(0.076)
regional house price	8.836	9.251***
	(8.098)	(1.445)
regional rent	-0.008**	-0.011***
	(0.004)	(0.001)
mig. cost (dropouts)	-2.333***	-2.058***
	(0.300)	(0.258)
mig. cost (high school)	-1.390***	-1.313***
	(0.129)	(0.083)
mig. cost (college)	-0.369***	-0.351***
	(0.130)	(0.083)
switching cost	-2.742***	-2.786***
	(0.103)	(0.052)
choice set		
Dependent variable (latent): probability of receiving cross-regional job offer		
employees per firm		0.801**
		(0.4226)
spatial index (industry)		1.002
		(0.6746)
online vacancies (occupation)		0.692**
		(0.2770)
dropout		-2.536
		(0.727)
HS		-1.040
		(0.895)
college		0.274
		(2.146)
N	3997	3997

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ *Both models also contain option-specific constants. The first column assumes all workers face the same choice set.*

Figure 2.5: Estimated distribution of the probability of choosing from a full choice set (the probability of receiving a cross-regional job offer), by education

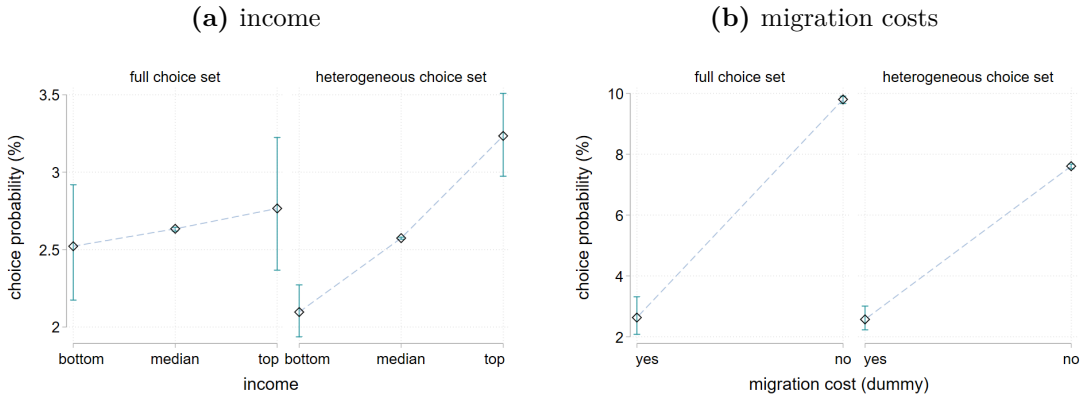


Kernel density plot of estimated μ_i, by education. Solid line: dropouts. Dashed line: high school graduates. Dotted line: college graduates. Based on the second column in Table 2.6.

Based on these estimates, Figure 2.5 plots the probability that a worker may receive cross-regional job offers, μ_i. As education increases, the average probability of receiving cross-regional job offers increases, and the variance of this probability falls. The median dropout worker has a 48% chance of being able to select employment away. The median μ_i rises to 78% for high school graduates, and to almost 97% for college graduates. Fewer than 1% of college graduates face a less than 90% chance of being able to choose from the full choice set.

The upper panel of Table 2.6 shows that allowing for choice set heterogeneity changes the estimates of workers' preferences. This acts as further evidence that choice set heterogeneity matters. Comparing coefficients in the two columns, we can see that income and unemployment dummy increase in magnitude, while migration costs become smaller, and the sign on regional unemployment rate changes. The coefficients also become more statistically significant, because varying choice sets explain some of the behaviour a standard conditional logit could not explain. The magnitude and sign of the variables changes because allowing for choice set heterogeneity is similar to including an omitted variable, which reduces the omitted

Figure 2.6: Marginal effects of income and migration costs on the probability that an average high school graduate in the Northeast chooses employment in the South (option 5)



The figures plot the estimated probability that an average high school graduate in the Northeast chooses employment in the South. Full choice set plots the probabilities for a standard logit model, which assumes that all workers always choose from all options. Heterogeneous choice set plots the probabilities for a conditional logit in which the probability that a worker receives a job in another region varies with his education, sector, and occupation. Left hand panel shows how this choice probability changes as the income offered in this option increases. Bottom = bottom 1% of the wage distribution in the South. Median = median wage. Top = top 1% of the wage distribution. Right hand side panel shows how this choice probability changes between a world with 0 migration costs, and a world where migration costs exist (the average migration costs estimated in Table 2.6). The diamond represent the point estimates; the solid lines depict 95% confidence intervals.

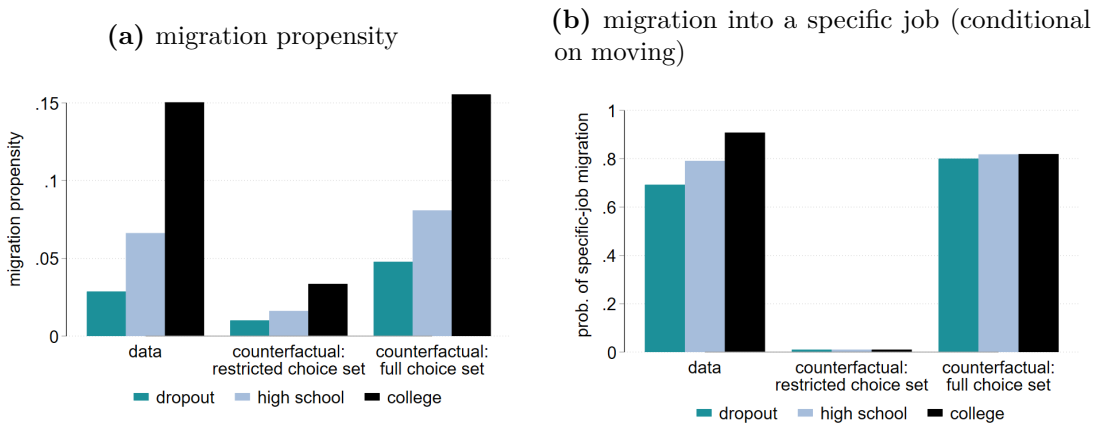
variable bias in the estimated coefficients[10].

The coefficient on regional unemployment changes from positive to negative, indicating that workers indeed do prefer locations where the labour market is tighter. It suggests that workers stay in regions with relatively higher unemployment rate because speculative migration is not optimal, and they lack the option to move for a specific job.

Take the example of income and migration costs, the estimates of which figure prominently in the migration literature. Figure 2.6 plots how the predicted choice probability of a particular option changes when income and migration costs change. Each panel contains two lines: one representing the standard conditional logit (where all workers have a full choice set), the other based on the model with heterogeneous choice sets. Looking at income, we can see that workers are estimated to be much

[10]The bias discussed here pertains to my interpretation of the logit coefficients as estimating workers' *preferences*; the coefficients are not biased as such. If we ignore the RUM framework and instead interpret the logit coefficients as the workers' *responsiveness* to option attributes, then controlling for choice set heterogeneity provides more detail on this responsiveness.

Figure 2.7: Impact of choice set variation: counterfactual migration behaviour if all workers received cross-regional job offers vs. if none of them did.



Left hand side: probability of regional migration. Right hand side: probability of moving for a specific job, conditional on migration. The first cluster of bars plots the observed migration behaviour. The other clusters capture counterfactuals calculated from the estimated reduced-form model. The middle cluster plots migration behaviour when all workers must choose from the restricted choice set, i.e. never receive cross-regional job offers. The values of job-specific migration are 0 for all education groups. The last cluster plots migration behaviour when all workers can always choose from the full choice set, i.e. employment in other regions is always available. All predicted probabilities are based on the coefficients from my preferred specification, column (3) in Table 2.6.

more sensitive to income variation when I allow their choice sets to vary. Intuitively, in order for the standard logit to explain the low probability of selecting employment in distant regions, workers cannot care much for the higher income these options offer. By letting choice sets to vary, I bring in another explanation: workers do not choose employment away because they may not have received a job offer from another region. For migration costs, the effect is the opposite: allowing for choice set heterogeneity actually decreases the estimated migration costs, resulting in a flatter line on the right hand side panel of Figure 2.6. However, the intuition is the same: some of the low choice probability of employment away is now explained by the lack of cross-regional job offers rather than high migration costs. Looking at the second column in Table 2.6, we can see that this effect is greatest for the high school dropouts, who also face the most severe choice set restrictions.

How much does choice set heterogeneity matter? To answer this question, I use the model estimates from Table 2.6 to construct two counterfactuals: a world where all workers have a restricted choice set, and a world where all workers have a

full choice set. The corresponding predicted migration propensities and predicted probabilities of migration into a specific job are plotted in Figure 2.7. The differences in migration behaviour are stark. Effectively ending any cross-regional hiring, as I do in the middle set of bars, shows that removing the possibility to move into a specific job would have a large and negative impact on migration propensity of the workers. This effect is the strongest for the most educated group, whose migration propensity would decline by more than 10 percentage points, compared to about 2 percentage points decline for high school dropouts. On the other hand, allowing all workers to always choose from among all options doubles the migration propensity of the less educated. It would also almost completely eliminate any education-based differences in the type of migration, suggesting that the relatively low share of migration for a specific job among the less educated is more due to their restricted choice sets rather than their preferences or the attributes of their options.

To summarize, the estimates of the empirical model complement the existing knowledge about internal migration in several ways. They re-affirm previous findings from the literature about the importance of income, amenities, and living costs in the migration decision. They also confirm the existence of considerable migration costs which fall with education. Adding choice set heterogeneity refines these estimates: income becomes more important and migration costs less important for the utility of the worker. My main finding is that job search matters for migration, as evidenced by the coefficients on recruitment proxies. This is further highlighted by the counterfactual exercise that shows that a large proportion of the college graduates' migration propensity, and almost all of the education gap in migration type, can be explained by the relative lack of migration opportunities for the less educated.

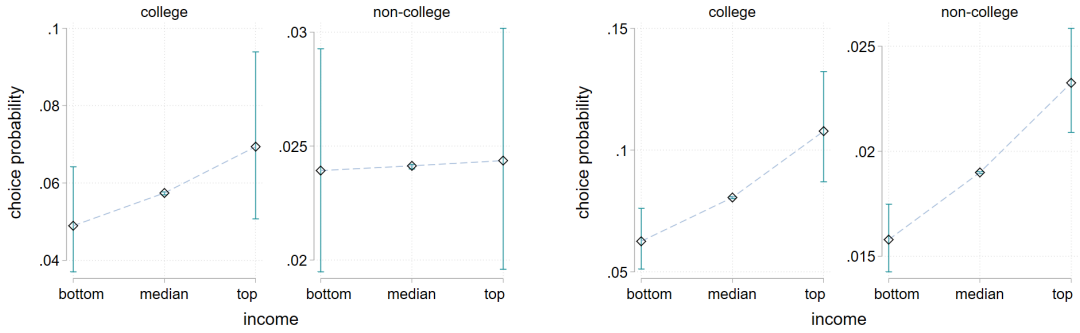
2.4.4.2 Explaining education differences in migration behaviour

To investigate the differences between education groups further, I estimate the discrete choice model of migration separately for those with and without a college degree[11]. In the model estimated in Table 2.6, I allow migration costs to vary

[11] I pool the high school graduates and dropouts in this way because the dropout subsample is not large enough to identify the model.

Figure 2.8: Marginal effects of income on the probability that an average worker chooses employment in the South (option 5), by education.

(a) full choice set (standard logit) **(b)** heterogeneous choice set



The figures plot the estimated probability that average worker chooses employment in the South, if he currently resides in the Northeast, as the income offered in this option increases. The probabilities are calculated separately for workers with and without college education, using estimates from Table 2.7. Left hand side panel: full choice set plots the probabilities for a standard logit model, which assumes that all workers always choose from all options. Right hand side panel: heterogeneous choice set plots the probabilities for a conditional logit in which the probability that a worker receives a job in another region varies with his education, sector, and occupation. Bottom = bottom 1% of the wage distribution in the South. Median = median wage. Top = top 1% of the wage distribution.The diamond represent the point estimates; the solid lines depict 95% confidence intervals.

by education; here I allow other preferences to vary, too. If the differences in education are primarily driven by variation in workers' tastes, rather than different job search opportunities and option attributes, we should see different preference coefficients for the less and the more educated, while the choice set variables should lose their significance.

The results in Table 2.7 show that this is not the case. If we compare the coefficients in the two rightmost columns, we can see that some differences in preferences exist – the more educated seem to be more responsive in general to the different option attributes – but they are small. Furthermore, the variables capturing choice set variation remain statistically significant. They are different from the coefficients estimated in Table 2.6, but that is because here they relate to within-group variation, whereas the general model captures between-group variation.

Moreover, controlling for choice set heterogeneity reduces the bias in the education-specific preference estimates. Similarly to the general results, I estimated the standard conditional logit alongside the logit with heterogeneous choice sets.

Table 2.7: Multinomial logit model of the choice of employment or unemployment in one of the 4 large US regions, allowing for choice set heterogeneity, by education

	full choice set		heterogenous choice set	
	less than college	college	less than college	college
preferences				
Dependent variable: probability of choosing one of the 8 options				
income	0.006	0.095	0.127***	0.159***
	(0.072)	(0.081)	(0.036)	(0.061)
regional unemp. rate	0.669	0.767	-0.320***	-0.475***
	(0.442)	(0.488)	(0.085)	(0.095)
unemployment dummy	-1.313***	-1.542***	-1.093***	-2.130***
	(0.230)	(0.312)	(0.087)	(0.191)
regional house price	4.815	12.600	8.902***	9.218***
	(12.025)	(12.500)	(2.071)	(2.270)
regional rent	-0.006	-0.009	-0.003***	-0.003***
	(0.005)	(0.006)	(0.001)	(0.001)
migration cost	-1.575***	0.057	-1.426***	1.399***
	(0.132)	(0.285)	(0.0803)	(0.097)
switching cost	-2.668***	-3.185***	-2.636***	-3.671***
	(0.110)	(0.281)	(0.067)	(0.094)
choice set				
Dependent variable (implied): probability of receiving cross-regional job offer				
employees per firm			5.262***	0.785**
			(1.777)	(0.436)
spatial index (industry)			0.164	0.747
			(15.033)	(4.789)
online vacancies (occupation)			0.744***	0.506***
			(0.271)	(0.202)
constant			0.546	-0.342**
			(0.410)	(0.147)
N	2840	1157	2840	1157

Standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ *All models contain option-specific constants.*

These results are given in the first two columns of Table 2.7. I use these estimates to predict the choice probability of a particular option as the income offered in that option increases. As Figure 2.8 shows, the standard logit estimates the less educated to be much less responsive to income changes than the more educated. This difference virtually disappears when I control for differences in choice sets faced by these two groups.

Overall, education-specific estimates of the model confirm that heterogeneity in cross-regional employment opportunities is a significant driver of the observed migration patterns. Furthermore, the comparison of standard and augmented logits highlights the need to distinguish between preferences and opportunities to arrive at a more nuanced explanation of the lower migration propensity of those without college education.

2.5 Conclusion

The main message of this chapter is that job opportunities matter for understanding migration behaviour. I build an empirical model of job search and migration, and apply it to the question of differing migration behaviour between the less and the more educated in the US. I show that the less educated are more likely to move speculatively, and move less overall, because they face significantly fewer employment opportunities in other regions before they move.

I motivate the paper by establishing a new stylised fact about the relative timing of migration and job search. The less educated are not only less likely to move, they are significantly less likely to move into a specific job. I then analyse what drives this behaviour, first by estimating a conditional logit of employment-location decision with unobserved choice set heterogeneity. By drawing on the empirical IO literature, I am able to decompose the observed migration patterns into the differences in the *attributes* of the different options, and the differences in their *availability*.

The estimates of the model show that job search matters for migration behaviour. Moreover, controlling for the difference in opportunities is necessary for a correct interpretation of our estimates of worker preferences and migration costs. Not taking

them into account leads to an overestimation of migration costs and underestimation of the sensitivity to income and amenities, especially among the less educated.

These findings have important implications for the way we model migration. The new stylised fact shows that 65% of migration is into a specific job, which is at odds with the standard assumption that all migration is speculative, under uncertain income. In the model presented here, I propose one possible way to incorporate the different types of migration into the existing framework. The estimates of this model, especially the significant link between recruitment behaviour and migration outcomes, suggests that this topic is a potentially fruitful area for future research.

3

The role of job search in migration decisions: a structural approach

3.1 Introduction

The decisions about employment and location are interlinked. The wage distribution and labour demand conditions in a worker's region drive her decision to search, and what jobs to accept. At the same time, labour market conditions are an important determinant of the worker's location. There is evidence that factors such as expected income and unemployment rate significantly influence migration (Kennan and Walker (2011)). Despite this, the employment and location choices are usually modelled sequentially: the worker first chooses where to live, and only then thinks about the process of job search.

In the first chapter of this thesis, I present empirical evidence that these two decisions should in fact be modelled jointly. Using data on inter-regional mobility in the US, I demonstrate that an employment opportunity in another region significantly influences location decision. In other words, worker's job search does not only determine her employment status, but also where she decides to settle. In the case of internal migration patterns in the US, this means that the less educated workers are less likely to move because they find job search in more distant regions much more difficult.

In this chapter, I build a theoretical model of the joint location-employment decision. I describe how the worker's decision about where to live and whether to work interact with each other and with the search frictions in the labour market. I then estimate the structural parameters of the model, which allows me to quantify the importance of job search in the migration decision. Finally, I return to the question of differential migration rates between the less and the more educated workers in the US. The structural estimates pin down the cross-regional search frictions that drive the choice set heterogeneity documented in Chapter 2. As a result, these estimates provide another test of the hypothesis that the lower migration propensity of the less educated in the US is partly due to their difficulty in moving for a specific job.

The point of departure for the theoretical model is the seminal partial equilibrium model of job search (McCall (1970)), in which workers wait to receive random wage offers, and follow an optimal stopping rule in deciding whether to accept them or not. I add a location dimension, so that workers may receive job offers from other regions, as well as migrate speculatively. Job-finding probabilities may vary with the worker's region of origin and the location of the vacancy, as well as with the worker's employment status and education. The fact that the worker considers her employment options in a spatial context generates a rich set of possible choice sets and outcomes, and allows us to think about the various interactions between labour markets and geography.

I estimate the model using the method of simulated moments. The parameters are identified from the transition rates in and out of unemployment, both within and between regional labour markets. I derive the theoretical expressions for these transition rates; the parameter estimates come from minimising the squared distance between the empirical and model-generated moments. Because the data allows me to distinguish between migration into employment vs. unemployment (migration for a specific job vs. speculative migration), as well as capturing the overall migration flows, I can separately estimate migration costs and cross-regional search frictions, alongside local labour market parameters.

I use the same dataset as in Chapter 2. It comes from the SIPP, a monthly panel of a sample of US inhabitants from 1996 to 1999. It captures, among other characteristics, workers' employment, wages and location. This is used to construct the data moments for structural estimation.

The results of the structural estimation show that employment opportunities have a significant impact on an individual's migration/location choice. A series of counterfactuals constructed from the structural model suggest that almost 50% of the gap in migration propensity, and three quarters of the difference in the type of migration observed between the less and the more educated workers in the US can be attributed to the relative difficulty of searching across regions.

Second, these cross-regional search frictions are considerable. When searching on the job, the probability that a college graduate would receive a job offer from another region is about 30% smaller than the probability of receiving a local job offer. This gap is more than 60% for those without a college degree. Cross-regional job search when unemployed is even harder, with individuals of any education seeing less than 1% of the vacancies in distant regions.

These numbers also highlight the extent of the disparities in cross-regional employment opportunities between the less and the more educated workers. The counterfactuals show that, while the lower migration propensity of those without a college degree partly reflects their lower wages, higher migration costs, and different location preferences, it is also a consequence of a labour market that makes it much more difficult for them to find a job in another region without having to move there first.

The structural estimates hence provide additional evidence on the link between job search and the migration patterns of the less and the more educated presented in Chapter 2. In the previous chapter, the estimates of choice set heterogeneity were based on the variation in recruitment behaviour between the less and the more educated. The results presented in this chapter, on the other hand, do not make use of this supplementary data and are instead derived from the structure of the model. The fact that the qualitative findings of both chapters are similar

lends support to the idea that the employment and location choices and deeply inter-linked, and that the way local labour markets interact has an important impact on mobility patterns within a country.

Related literature The chapter's main contribution is in modifying the way we model within-country migration. The standard assumption in the literature is that the question of employment is only resolved after the worker has decided on her location: she moves speculatively, for an uncertain income. In Chapter 2, I present evidence that speculative migration coexists with migration for a specific job. In fact, the majority of long-distance moves in the US are the result of job search. The theoretical model presented in this chapter is the first to combine both search-move and move-search strategies in one framework[1]. The existing literature assumes that individuals either only search after they migrate (Kennan and Walker (2011), Epifani and Gancia (2005), Kline and Moretti (2013a)), or that migration can only happen as a result of a specific employer-employee match (Beaudry et al. (2014), Lutgen and der Linden (2015), Amior (2015)). Incorporating both enables a better understanding of the the interaction between job search and migration decisions.

The findings of this paper add to the literature on local labour markets (Moretti (2011)) and within-country mobility (Malamud and Wozniak (2012); Ransom (2019)) by highlighting the role of job search in migration decisions, and the impact of cross-regional search frictions on the overall migration propensity. At the same time, the migration focus of the theoretical model complements the local labour markets literature. Work by Roback (1982); Moretti (2011) and Kline and Moretti (2013a) shows how regional equilibria – output, employment, wages, house prices, congestion – are determined, using migration as a black box equilibrating mechanism. Here, I take the regional equilibrium as given and explicitly model migration flows.

[1]Molho (2001) does include both, but his work focuses on analysing spatial equilibrium instead of migration behaviour, and he assumes the migration strategy is a choice made by the individual as opposed to something partially determined by the labour market. Schmutz and Sidibé (2018) also present a spatial extension to McCall (1970), and allow for both type of migration. However, in their setup speculative migration is not an equilibrium outcome.

The results of the chapter help to bridge the literatures on local labour demand shocks and regional adjustment. There is substantial evidence on the role of mobility as a buffer to local demand shocks (Bartik (2018), Blanchard and Katz (1992)), but research also shows that the less educated, low-income households who are more exposed to these shocks are also least likely to move, exacerbating regional disparities (Bound and Holzer (1996), Wozniak (2010), Hoynes (1999)). This paper explains that this is partly because the frictions between regional labour markets make it difficult for the less educated to migrate.

Understanding why and how the less educated move is essential for policy. Governments intervene in regions that have experienced prolonged periods of high unemployment and low growth, often by implementing job-creation policies such as hiring subsidies and special tax arrangements for new businesses. As the literature shows (Kline and Moretti (2013a), Glaeser and Gottlieb (2008)), the welfare impact of such an intervention largely depends on workers' migration response. This chapter deepens our understanding of that response, and the estimated structural model allows us to analyse the migration impact of various policy measures. My results also suggest that, if the aim of place-based policies is in fact to improve welfare of the most vulnerable groups, improving the cross-regional labour market is a viable alternative to subsidies and taxation (Schmutz and Sidibé (2018)).

The rest of this chapter is organised as follows. In section 3.2, I outline the theoretical model of the joint location-employment decision. In section 3.3 I describe how this model is estimated, and present the main results. In section 3.4, these structural estimates are then used to analyse the role of job search in migration decisions, and why the less educated US workers don't move much. I also discuss the policy implications of the model. Section 3.5 concludes.

3.2 A theoretical model of cross-regional job search and migration

The empirical results in Chapter 2 demonstrated that different workers have different opportunities to migrate into a specific job, and that these differences matter for their overall migration behaviour. To analyse the role of cross-regional job search formally, I build a model of frictional regional labour markets. Every period, a worker decides where to live and whether to work there, responding to job offers from various regions and to the utility of the different options. The probability of receiving an offer depends both on the location of the vacancy and of the worker, which generates a rich structure of choice sets and outcomes. The model pins down equilibrium transition rates in and out of unemployment and within and across regions. In section 3.3, I use these to estimate the model parameters.

3.2.1 The setup

The point of departure is a standard partial equilibrium model of job search (McCall (1970)), where workers wait to receive random wage offers and follow an optimal stopping rule in deciding whether to accept them or not. The same mechanism lies at the heart of this model, but I add a location dimension, so that workers may receive wage offers from different regions, and may choose to move to another region even in the absence of a specific job offer. This means there are 3 endogenous variables in this model: a worker's employment status, her wage, and her location.

The focus of the model, as well as the whole chapter, is equilibrium analysis. In this model, I assume a steady-state equilibrium exists, and then describe and determine its properties. Although the model analyses dynamic and forward-looking behavior, the environment is stationary. All the model parameters and distributions (of location preferences and wage offers) are constant over time, and so are the values of employment and unemployment in different regions.

I define region to correspond to the local labour market, i.e. a commuting zone. This means that a move to a different region necessarily implies a move to a new labour market. Furthermore, all the workers within the same region

are assumed to be subject to the same labour market conditions. There are J regions in the model, denoted $j = 1, 2, ..., J$.

Regions can differ in their labour market characteristics, as well as in features unrelated to the labour market. Local (regional) labour markets may differ in wage offer distributions ($F_j(z)$), job offer arrival rates (θ_j, λ_j), and the probability of exogenous job destruction (δ_j). The attractiveness of a region also depends on weather, landscape, local amenities, geographic links to other regions, etc. This is captured by a vector of individual location preference parameters $\boldsymbol{\gamma} = \begin{bmatrix}\gamma_1 & \gamma_2 & ... & \gamma_J\end{bmatrix}$. Each γ_j is a random variable drawn from a fixed multivariate distribution G with variance g^2 and region-specific means $\bar{\gamma}_j$.

Migration is costly. Every time a worker moves, she pays a fixed cost K. It captures the financial, social and psychological burden of moving.

Job search happens both when unemployed and on the job. All workers draw from the same wage distributions, so the employed are not more or less likely to receive high wage offers than the unemployed. The job offer arrival rate follows a Poisson process, in which the probability of receiving more than one offer is equal to 0. I denote θ as the probability of receiving a job offer when unemployed, and λ as the job offer rate for the employed. There is no explicit job search decision: all workers search.

The probability of receiving a job offer may not only vary by employment status but also by geography. First, the job offer rates may be higher in one region than another. Second, it may be easier to receive job offers from one's current region of residence than from other regions. To capture this variation, I differentiate between θ's and λ's by adding superscripts to denote what regional labour market the worker currently resides in, and subscripts to denote the region of the job offer. θ_j^k thus stands for the probability of receiving an offer to work in region j when residing (unemployed) in region k, and λ_k^k denotes the probability of receiving a home job offer in k (when employed). The empirical results from Chapter 2 suggest that the probability of home job offer is higher than the probability of job offer from another region ($\theta_k^k > \theta_j^k$ and $\lambda_k^k > \lambda_j^k$), but that this difference

falls with education. The structural estimates presented in section 3.3 will provide another test of this hypothesis.

The full set of job-finding probabilities for the unemployed and employed in a four-region economy ($J = 4$) are summarised in matrices Θ and Λ, respectively.

$$\Theta = \begin{bmatrix} \theta_1^1 & \theta_2^1 & \theta_3^1 & \theta_4^1 \\ \theta_1^2 & \theta_2^2 & \theta_3^2 & \theta_4^2 \\ \theta_1^3 & \theta_2^3 & \theta_3^3 & \theta_4^3 \\ \theta_1^4 & \theta_2^4 & \theta_3^4 & \theta_4^4 \end{bmatrix} \qquad \Lambda = \begin{bmatrix} \lambda_1^1 & \lambda_2^1 & \lambda_3^1 & \lambda_4^1 \\ \lambda_1^2 & \lambda_2^2 & \lambda_3^2 & \lambda_4^2 \\ \lambda_1^3 & \lambda_2^3 & \lambda_3^3 & \lambda_4^3 \\ \lambda_1^4 & \lambda_2^4 & \lambda_3^4 & \lambda_4^4 \end{bmatrix} \tag{3.1}$$

where the diagonal elements are probabilities of receiving a domestic job offer, and the off-diagonals are cross-regional probabilities.

For the purposes of estimation in section 3.3, I impose a particular structure on this general model. In particular, I assume that the cross-regional job-finding probability is a fixed proportion of the domestic offer probability in a given region. I define two job search wedges, ζ_1 and ζ_2, for search on the job and when unemployed, respectively. They denote how much more difficult it is to find a vacancy in another region compared to searching locally. The probability of receiving a job offer from region k when the worker resides in region m can then be defined as

$$\theta_k^m = \zeta_2 * \theta_m \tag{3.2}$$

$$\lambda_k^m = \zeta_1 * \lambda_m \tag{3.3}$$

I assume that the wedge between local and cross-regional job search is the same in all regions. Under this simplifying assumption, the job-finding probabilities Θ and Λ are:

$$\Theta = \begin{bmatrix} \theta_1^1 & \zeta_2\theta_1^1 & \zeta_2\theta_1^1 & \zeta_2\theta_1^1 \\ \zeta_2\theta_2^2 & \theta_2^2 & \zeta_2\theta_2^2 & \zeta_2\theta_2^2 \\ \zeta_2\theta_3^3 & \zeta_2\theta_3^3 & \theta_3^3 & \zeta_2\theta_3^3 \\ \zeta_2\theta_4^4 & \zeta_2\theta_4^4 & \zeta_2\theta_4^4 & \theta_4^4 \end{bmatrix} \qquad \Lambda = \begin{bmatrix} \lambda_1^1 & \zeta_1\lambda_1^1 & \zeta_1\lambda_1^1 & \zeta_1\lambda_1^1 \\ \zeta_1\lambda_2^2 & \lambda_2^2 & \zeta_1\lambda_2^2 & \zeta_1\lambda_2^2 \\ \zeta_1\lambda_3^3 & \zeta_1\lambda_3^3 & \lambda_3^3 & \zeta_1\lambda_3^3 \\ \zeta_1\lambda_4^4 & \zeta_1\lambda_4^4 & \zeta_1\lambda_4^4 & \lambda_4^4 \end{bmatrix} \tag{3.4}$$

The worker is a rational utility maximiser. She decides where to live and which job offers to accept based on all the information available to her so that her expected lifetime utility is maximised. All workers are infinitely-lived, risk-neutral, and homogenous up to their idiosyncratic location preferences. The worker has perfect information about the aggregate characteristics of the local labour markets and of

her own location preferences, although she does not know what wage offers she will receive in a given period (if any), and what her draw of location preferences will be.

There are three reasons why a worker would change her location or employment. Each period, an employed worker may be fired; this happens with region-specific probability δ_j. Each period, a worker (unless she has just been fired) may receive a job offer. The probability of such an event depends on θ_j and λ_j, and the wage will be drawn from the relevant region-specific wage offer distribution $F_j(z)$. Third, the worker makes a draw of her location preferences $\boldsymbol{\gamma}$ from distribution G. As a result of these stochastic events, the worker re-optimizes her location and employment choice and may choose to migrate.

3.2.1.1 Timing of events and decisions

Upon learning about her job options and location preferences, the worker makes the joint decision about where to live and whether to work (if available). Her decisions then come into force, and she makes a new draw of the random variables.

The timing of the model is as follows:

1. The worker lives in region j. She may be employed or unemployed.

2. The worker draws her location preferences.

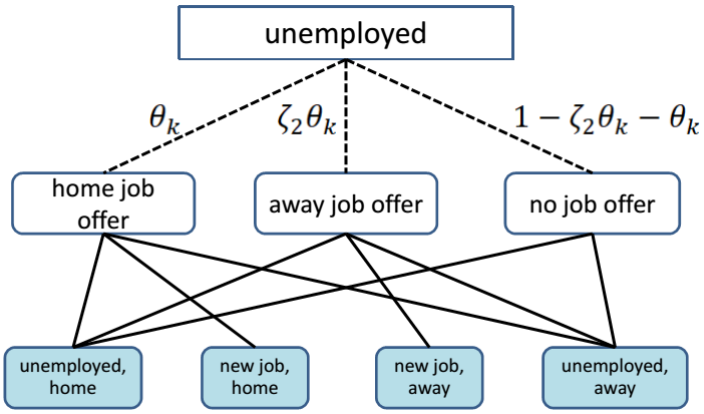
3. If the worker is employed, her job may be exogenously destroyed, or she may receive at most one job offer from one of the J regions of the economy.

4. If the worker is unemployed, she may receive at most one job offer from one of the J regions of the economy.

5. The worker makes her choice. She compares the value of her current state with all her other options, taking into account her location preferences, and choose that location-employment combination that maximises her expected lifetime utility.

6. The worker's decisions come into force.

Figure 3.1: Choices and possible outcomes for an unemployed worker



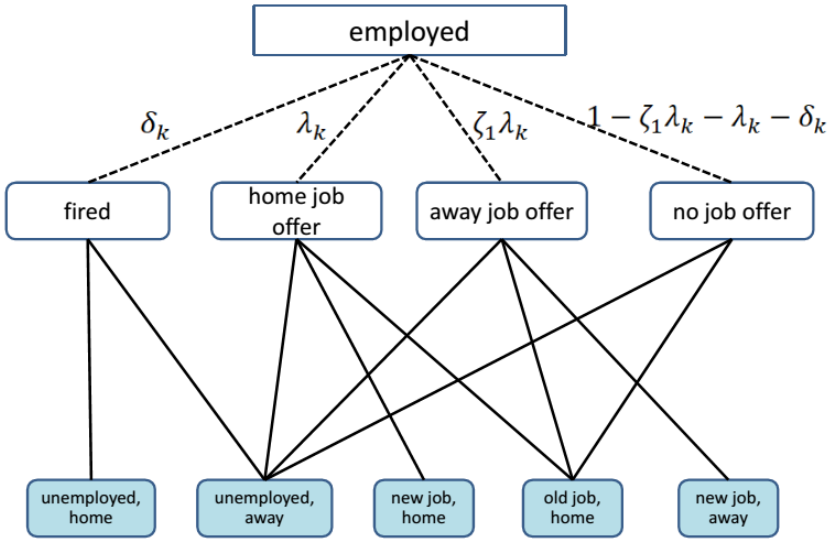
The scheme links the functioning of the labour market with workers' options and their potential outcomes. For instance, with the probability θ_k, the worker receives a job offer in her home region. In that case, she can decide between three outcomes: accept the offer, reject it and stay unemployed at home, or reject it and move into unemployment in another region. On the other hand, migration for a specific job is only possible if she receives a relevant job offer first.

3.2.2 Decisions and choice sets

As shown in Chapter 2, workers' outcomes depend on their choice sets. The options available to a worker vary by her employment status and whether she has received any job offers. Here I look closely at the choice set patterns and the outcomes they may lead to.

Figure 3.1 summarizes the links between an *unemployed* worker's job offers and her potential choices, for a stylized example of a 2-region economy. One of three things can happen to an unemployed worker: she may receive a home job offer (with the probability θ_k); a job offer in an away region (with the probability $\zeta_2\theta_k$); or receive no offers at all (with the probability $1-\zeta_2\theta_k-\theta_k$). As a result, she has four potential outcomes: to accept a new job in her home region, to accept a new job in another region (which results in migration for a specific job), to stay unemployment in her home region, or to move to search as unemployed in another region (which implies speculative migration). The diagram shows that the optimisation looks differently depending on which of the three situations the worker ends up in. If she receives no offers at all, her optimisation choice involves choosing the best region to be unemployed in. If, on the other hand, she receives a home job offer, she

Figure 3.2: Choices and possible outcomes for an employed worker



The scheme links the functioning of the labour market with workers' options and their potential outcomes. For instance, with the probability λ_k, the worker receives a job offer in her home region. In that case, she can decide between three outcomes: accept the offer, reject it and stay employed at home, or reject it and move into unemployment in another region.

compares its value against that of being unemployed in either region. She has to make a similar choice when she receives an offer of a job in another region.

The choice set structure of an employed worker is more complicated, because she also has the option to remain in her existing job. As Figure 3.2 shows, an employed worker has four possible routes and five different potential outcomes. If the worker is fired (with probability δ_k), she weighs up the values of home and away unemployment (speculative migration). With probability λ_k, she receives a new job offer in her home region, which she then compares against the value of her current job and the values of unemployment away (i.e. speculative migration) in different locations. If she receives an away job offer instead (with the probability $\zeta_1\lambda_k$), her decision process is similar, only now its outcome can also include specific-job migration along with staying employed at home or moving speculatively. If no new job offers arrive (which happens $1-\zeta_1\lambda_k-\lambda_k-\delta_k$ of the time), she reconsiders whether her current job is still better than unemployment away, given her new location preferences.

In Figure 3.2 I claim that an employed worker will never voluntarily quit to stay unemployed in her home region. This is clearly at odds with reality, but is implied

by the assumptions of the model. By accepting a job offer in region k, the worker reveals that the value of such employment exceeds the value of unemployment in the same location. Because I assume the environment of this model is stationary, this will be true in all periods. The worker will draw different location preferences, which may change her opinion about how employment in k compares with unemployment (and employment) in some other region j, but location preferences do not enter into the comparison of employment and unemployment within the same location. As a result, the transition from employment to unemployment within the same region is only possible as the result of exogenous job destruction.

Figures 3.1 and 3.2 make it clear that regardless of her employment status, a worker can always choose to be unemployed in any of the J regions, which also means that she can always choose to migrate speculatively. Migration for a specific job, on the other hand, is linked to receiving distant region job offers, and is hence less likely to be a part of a worker's choice set. That does not necessarily imply that speculative migration *will* be chosen more often than specific-job migration. The flows of utility of an option matters too, an employment options are likely to be of higher value than unemployment. They are derived in the next section.

3.2.3 Values of employment and unemployment in different regions

A worker's optimal decision is based on the values of employment and unemployment in different regions. These depend on the parameters of the model (job-finding and job-destruction probabilities, migration costs, wage offer distributions) and reflect the discounted expected utility of making a particular choice. There are J Bellman equations that describe the discounted expected value of employment, $\mathbf{V_j}$, and J Bellman equations for unemployment, U_j.

The values of employment and unemployment are dynamic, because the current choice has an impact on what options the worker will have in the future. The values of these options form a part of the values of employment and unemployment, creating a recursive structure.

In the previous section, I argued that the inclusion of particular options in the worker's choice set depends on her employment status and is partly stochastic. However, all choice sets have a common part: unemployment in any of the J regions. To capture this part of the choice set, I introduce the concept of *best unemployment option*. U_m^* is the highest net utility a worker can receive as unemployed, given her home region. U_m^* is equal to the highest value of unemployment less migration costs, given the vector of idiosyncratic preference draws $\boldsymbol{\gamma}$:

$$U_m^*(\boldsymbol{\gamma}) = \max[U_j + \gamma_j - K] \tag{3.5}$$

U_m^* may change every period because it is a function of $\boldsymbol{\gamma}$.

The value of unemployment U_m (equation (3.6)) has two parts: the current period payoff, and the expected utility of future optimisation. For every period the worker spends as unemployed in region m, she receives non-labour income (net of living costs) denoted b_m. The second part of (3.6) describes the future payoffs. With probability $\zeta_2\theta_m$, the worker receives a wage offer z in region j, which she then compares against her best unemployment option U_m^*, taking into account her location preferences and migration costs. Alternatively, with probability $1-(J-1)\zeta_2\theta_m-\theta_m$, she receives no job offers and picks her best unemployment option. The expectations of the future payoffs are taken over both the values of wages z and idiosyncratic location preferences $\boldsymbol{\gamma}$, since these are the elements unknown to the worker when making a choice today. The utility flow is discounted at the rate i.

$$\begin{aligned} U_m = \frac{1}{1+i}\left[b_m + \theta_m E_{z,\gamma}\max[V_m(z)+\gamma_m, U_m^*] + \sum_{j\neq m}^{J}\zeta_2\theta_m E_{z,\gamma}\max[V_j(z)+\gamma_j-K, U_m^*]\right] \\ + \frac{1}{1+i}\left[(1-(J-1)\zeta_2\theta_m-\theta_m)\,E_\gamma U_m^*\right] \end{aligned} \tag{3.6}$$

The value of employment in region m (equation (3.7)) also has two parts, although the future payoff is more complex. As described in Figure 3.2, there are 4 different events that can happen in the future. With probability δ_m, she gets fired and takes her best unemployment option. With probability $\zeta_1\lambda_m$, she receives a wage offer z in region j, in which case she selects the highest payoff from

between the new offer, the existing employment $V_m(w)$, and her best unemployment option. If she receives no new job offers, she can still choose to either keep her existing job or quit and take her best unemployment offer. The per-period payoff is a particular wage (net of living costs) w.

$$\begin{aligned} V_m(w) = & \frac{1}{1+i}\left[w + \lambda_m E_{z,\gamma}\max[V_m(w)+\gamma_m, V_m(z)+\gamma_m, U_m^*]\right] \\ & + \frac{1}{1+i}\left[\sum_{j\neq m}^{J} \zeta_1 \lambda_m E_{z,\gamma}\max[V_m(w)+\gamma_m, V_j(z)+\gamma_j - K, U_m^*]\right] \\ & + \frac{1}{1+i}\left[\delta_m E_\gamma U_m^* + (1-(J-1)\zeta_1\lambda_m - \lambda_m - \delta_m)\, E_{z,\gamma}\max[V_m(w)+\gamma_m, U_m^*]\right] \end{aligned} \tag{3.7}$$

Vector $\mathbf{V_j}$ stacks all the possible values of employment, as a function of wage, for a given region j. Assuming that the regional wage offer distribution $F_j(z)$ is discrete, the length of $\mathbf{V_j}$ corresponds to the length of the support of $F_j(z)$.

When making her optimal choice, the worker compares the values of employment and unemployment available to her plus her current draw of idiosyncratic location preferences $\boldsymbol{\gamma}$, and chooses the one that gives her the highest utility. This choice pins down her employment status, wage, and location. For example, if the worker is unemployed and has not received any job offers, her optimal choice will be unemployment in her best region:

$$U_m^*(\boldsymbol{\gamma}) = \max_j[U_j + \gamma_j - K]$$

On the other hand, if she is employed and receives a job offer from another region, her best choice is defined as:

$$\max[V_m(w)+\gamma_m, V_k(z)+\gamma_k - K, U_m^*(\boldsymbol{\gamma})]$$

where m stands for home region and z is the new job offer from region k.

3.2.4 Equilibrium

3.2.4.1 Defining the equilibrium

To close the model, I define equilibrium as one where the endogenous variables of the model, regional unemployment rates ($\{\mu_j\}_1^J$), population shares ($\{\alpha_j\}_1^J$), and accepted wage distributions $\{H_j(z)\}_1^J$, are constant. Throughout the paper, the total population is normalised to 1, so per-person and aggregate outcomes are the same[2].

A general J-region model requires $J(2+\omega)$ equilibrium conditions to pin down the J local unemployment rates μ_j, the J regional population shares α_j, and the $J\omega$ equilibrium wage distributions $H_j(z)$, where ω denotes the size of the discrete wage support. They are all based on imposing an equilibrium of flows between different parts of the labour market or different regions.

Equilibrium in local labour markets Equilibrium in the local labour market requires a steady unemployment rate, which means that the flows in and out of the stock of unemployed must be balanced.

$$\textit{flows into local unemployment in } k = \textit{flows out of local unemployment in } k \tag{3.8}$$

Equilibrium in regional population shares To find equilibrium regional population shares α_j, the net migration into every region has to be zero. I can set up a similar balance-of-flows expression as for the local unemployment rate that pins down a particular regional population α_k.

$$\textit{migration into region } k = \textit{migration out of region } k \tag{3.9}$$

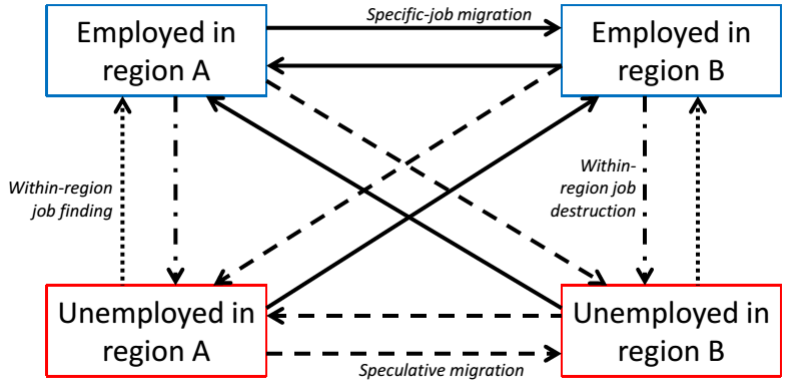
Given that the population shares α_j must sum up to one, the Jth condition is

$$\sum_1^J \alpha_j = 1 \tag{3.10}$$

which implicitly pins down α_J.

[2]This holds true as long as population stays constant.

Figure 3.3: Equilibrium in a 2-region model



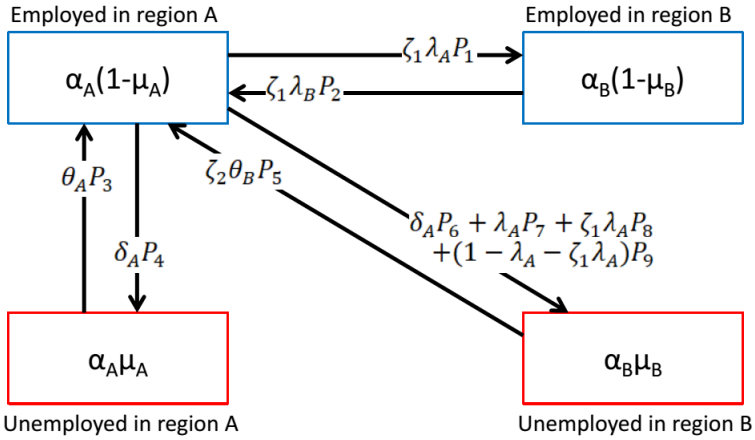
The arrows represent flows between the different employment-region stocks. Solid line: migration for a specific job. Dashed line: speculative migration. Dotted line: flows from unemployment to employment withing the same region. Dot-dashed line: flows from employment to unemployment within the same region.

Wage distribution $H_j(z)$ is the equilibrium distribution of accepted wages. It has the same support as the distribution of wage offers, $F_j(z)$. It is derived by setting equal the flows of workers paid a particular wage w, which is why the equilibrium condition needs to be specified for each point of the wage support, giving us $J\omega$ equations in total.

$$\textit{new workers that are paid } w \textit{ in region } k = \textit{workers that used to be paid } w \textit{ in } k \tag{3.11}$$

To understand how these flows equations link to the model developed in the previous section, consider a two-region example with a single wage. Figure 3.3 illustrates all the possible stocks and flows in a $J = 2$ economy. The boxes represent stocks of workers, employed and unemployed, in each region, while arrows depict the flows between these stocks, in and out of employment and between and within regions. The economy is in equilibrium when the size of the flows is such that the size of the stocks stays the same. For example, the equilibrium condition for regional population shares would specify that the magnitude of the flows in and out of each region must perfectly offset each other. Figure 3.3 makes it clear that this implies equality between the flows from employment to employment, employment to

Figure 3.4: Equilibrium in a 2-region model: constant stock of employed workers in region A



The four boxes represent the 4 different stocks of workers. The arrows demonstrate the possible flows in and out of the stock of employed in A. In equilibrium, these flows must perfectly offset each other. There are 3 inflows (from unemployment in A, unemployment in B, and employment in B) and 3 outflows (into unemployment in A, unemployment in B, and employment in B); the expressions on each arrow denote the share of the corresponding stock that flows in the given direction.

unemployment, unemployment to employment, and unemployment to unemployment from B to A, and the same flows in the opposite direction, from A to B.

Figure 3.4 shows how the model pins down the size of these flows. For example, the flow of employed in A to employment in B is $\alpha_A(1-\mu_A)\zeta_1\lambda_A P_1$. There are $\alpha_A(1-\mu_A)$ employed workers in region A. Fraction $\zeta_1\lambda_A$ of them receive a job offer in the other region, which is necessary in order to move for a specific job. Conditional on that, P_1 denotes the probability that the worker accepts the offer:

$$\begin{aligned} P_1 &= \Pr(V_B \text{ is the optimal choice out of } \{V_A, U_A, V_B, U_B\}, \text{ given preferences } \boldsymbol{\gamma}) \\ &= f_B(z)\left[\Pr(V_B(z)+\gamma_B-K > \{V_A(w)+\gamma_A, U_A+\gamma_A, U_B+\gamma_B-K\})\right] h_A(w) \end{aligned} \tag{3.12}$$

where $f_j(z)$ and $h_j(z)$ are probability density functions of $F_j(z)$ and $H_j(z)$, respectively. The expression of the overall flows thus depends on worker's optimisation, which is in turn the function of employment and unemployment values in different locations, and her choice set.

3.2.4.2 The role of initial population distribution

The equilibrium in this model does not depend on the initial population distribution across regions. This is a useful feature because it means that two economies with the same parameters will converge to the same equilibrium, regardless of their initial conditions. This result is summarised in the proposition below.

Proposition 1. *If some migration is optimal, the equilibrium distribution of regional population* $\{\alpha_j\}_1^J$ *will be independent of the initial regional population distribution* $\{\alpha_{0j}\}_1^J$.

Intuitively, the initial population distribution does not matter for equilibrium because the "attractiveness" of a region does not depend on its size[3]. As I demonstrated in the earlier part of this section, the size of the flows in and out of regions and employment depends on three factors: the size of the given stock (e.g. how many employed workers reside in region A), the job-finding and job-destruction probabilities, and workers' choice probabilities. The latter two of these are fixed and exogenous, which means that the fraction of workers leaving a particular stock for another destination will always be the same, regardless of the size of that stock. In other words, there is no feedback loop from a region's size to the flows in and out, which is why its initial size does not impact on its population in equilibrium.

A direct consequence of this is that an unattractive region that starts out large will diminish over time, and vice versa. For example, a region that is populous initially but has low wages will experience a large emigration rate but a small immigration rate, and these rates are fixed (in this toy example, they would depend on relative wage distributions). As a result, the size of the region will fall until the two flows just offset each other (so that net migration is zero). Regions do not stay large simply because they were populous in the past.

One exception to the above is when migration cost K is so large as to make any migration suboptimal. In the absence of any population flows, there would

[3]I relax this assumption in Chapter 4.

be no equilibrium adjustment, and the initial population distribution would be also the equilibrium population distribution.

Saying that equilibrium is independent of initial population distribution does not mean that migration costs do not matter. K, as one of the parameters of the model, directly determines the equilibrium by making migration more or less attractive. A world with high migration costs will converge to a different equilibrium than a world where mobility is cheap. However, holding K fixed, the point to which different initial population distributions will converge is the same. A formal proof of Proposition 1 can be found in Appendix B.1.

3.2.4.3 Solving for the equilibrium

The model is solved numerically.

I use value function iteration to calculate the values of employment and unemployment across regions, $\mathbf{V_j}$ and U_j. The challenging part of this procedure is to find the expected maximum value of future decisions, such as $E_{z,\gamma}\max[V_j(z)+\gamma_j-K, U_m^*]$. This is a core part of the Bellman equations, in which expectations need to be taken over both wage offer distributions and location preferences. The assumption that γ_j are type I extreme value distributed simplifies this step considerably. Rust (1987) demonstrates that under this assumption the expected utility has a closed-form solution, which I can plug into the Bellman equation:

$$E_{z,\gamma}\max[V_j(z)+\gamma_j-K, U_m^*] = f_j(z)^T[ln(exp(\mathbf{V_j}(z)-K)+exp(U_m^*))] \quad (3.13)$$

Having solved for values of employment and unemployment as a function of the model parameters, I can solve for conditional choice probabilities such as P_1 in equation (3.12). These probabilities are necessary in order to construct the equlibrium conditions of the model.

The final step is solving the system of flow equations for the equilibrium endogenous variables $\{\mu_j\}_1^J$, $\{\alpha_j\}_1^J$, and $\{H_j(z)\}_1^J$. I do this using value function iterations in two stages. First, I use the wage offer distribution $F_j(z)$ instead of the endogenous distribution of wages $H_j(z)$ and solve for equilibrium population shares

and unemployment rates. I then plug in these first-stage solutions to calculate the equilibrium wage distribution. I repeat the iteration until the values from the two stages converge.

3.3 Structural estimation

Having specified the interdependence between location and employment choice, I estimate the model parameters by matching the migration and employment flows generated by the theoretical model with those observed in the US between 1996 and 1999. The results show that employment considerations play an important role in the location decision, and that cross-regional search frictions are substantial. Receiving a job offer from another region is especially difficult for the less educated, and for those searching as unemployed.

3.3.1 Data

The data used to estimate the structural model is SIPP, a US panel dataset providing monthly information on individuals' location, employment status, wages, and a host of demographic variables. I use the 1996-1999 wave. Working with a sample of men between the age of 25 and 50 that have been employed for at least one month during the 4-year period covered by the panel, I calculate average monthly transition rates between employment and unemployment, within and across the four large census regions of the US[4]. As I further explain below, these data moments are used to identify the model parameters. For more details on the properties of the dataset, see section 2.2.

3.3.2 Identification of parameters

Using the four large Census regions of the US as my unit of analysis, the structural model contains 20 parameters to be estimated:

- 4 domestic job offer probabilities for search when unemployed, θ_j

[4]There are 16,720 men in the dataset who satisfy this criteria. See Table A.1 in the Appendix for the descriptive statistics.

- 4 domestic job offer probabilities for search on the job, λ_j
- 2 job offer wedges, ζ_1 and ζ_2, capturing the cross-regional job-finding frictions relative to home job-finding frictions
- 4 job-destruction probabilities δ_j
- 1 migration cost parameter K
- 4 mean location preference $\bar{\gamma_j}$
- 1 variance of the location preference distribution g^2

I set non-labour earnings b_j to 0 for all regions. The variation in non-labour income and savings will be instead estimated as a part of the region-specific location preferences, $\bar{\gamma_j}$. I set the distribution of these preferences, γ_j, to be type I extreme value distributed. As in other parts of this chapter, it makes it possible to find closed-form solutions for workers' choice probabilities, as well as the moments I base my estimation on. The parameters of the regional wage distributions are taken from the data, given that the partial equilibrium model presented here does not explain how wages are set.

The model is identified from data moments on the worker movement between regions and in and out of employment. In this section, I show how the theoretical expressions for these flows link to data to help uncover the model parameters.

A worker can be in one of eight states (employment or unemployment in one of the 4 census regions of the US), which gives me $8 \times 8 = 64$ transition probabilities. However, the probabilities in a transition matrix must sum up to 1, giving us 63 free moments. These moments are calculated from the 1996 SIPP panel. They are presented in Tables 3.1 and 3.2, for those with and without college education, respectively.

The model is estimated separately for the less and the more educated. The "more educated" category includes workers with a college degree or more; the "less educated" are high school graduates, high school dropouts, and workers with less than four years of college education (section 2.2 for more detail on this categorisation).

Table 3.1: Data moments: matrix of transition probabilities for the more educated (in %)

	e,1	u,1	e,2	u,2	e,3	u,3	e,4	u,4
e,1	19.24984	0.16173	0.00682	0.00038	0.00909	0.00076	0.00492	0.00038
u,1	0.13370	1.05521	0.00000	0.00152	0.00076	0.00189	0.00000	0.00076
e,2	0.00682	0.00000	24.32591	0.20150	0.01477	0.00038	0.00720	0.00076
u,2	0.00000	0.00114	0.15112	1.13361	0.00038	0.00152	0.00076	0.00114
e,3	0.00682	0.00000	0.00985	0.00076	28.57478	0.26361	0.01212	0.00038
u,3	0.00000	0.00038	0.00114	0.00076	0.20188	1.71690	0.00038	0.00152
e,4	0.00682	0.00000	0.00909	0.00076	0.00909	0.00000	20.78645	0.24695
u,4	0.00038	0.00038	0.00000	0.00038	0.00038	0.00227	0.19127	1.47980

The transition probabilities are calculated from the 1996-1999 SIPP data. The sample is working adult males between the age 25 and 50. e = employment. u = unemployment. The numbers denote one of the four census regions of the US (see section 2.2).

Table 3.2: Data moments: matrix of transition probabilities for the less educated (in %)

	e,1	u,1	e,2	u,2	e,3	u,3	e,4	u,4
e,1	14.90083	0.23907	0.00048	0.00000	0.00556	0.00097	0.00145	0.00000
u,1	0.20692	1.91083	0.00000	0.00024	0.00000	0.00073	0.00000	0.00000
e,2	0.00073	0.00024	21.64856	0.33382	0.00580	0.00193	0.00338	0.00048
u,2	0.00000	0.00000	0.27508	2.34255	0.00073	0.00169	0.00000	0.00097
e,3	0.00193	0.00024	0.00411	0.00048	31.29649	0.53929	0.00604	0.00024
u,3	0.00000	0.00024	0.00024	0.00097	0.42640	3.87870	0.00097	0.00024
e,4	0.00193	0.00000	0.00218	0.00024	0.00508	0.00121	18.86438	0.36549
u,4	0.00024	0.00000	0.00024	0.00193	0.00000	0.00193	0.32391	2.39162

The transition probabilities are calculated from the 1996-1999 SIPP data. The sample is working adult males between the age 25 and 50. e = employment. u = unemployment. The numbers denote one of the four census regions of the US (see section 2 in Chapter 1).

The first step in the identification is to normalise the model. Because in the worker's optimisation only differences in utility matter, I set the average location preference for the West to 0, which reduces the number of parameters to be estimated to 19.

The theory-consistent expressions for the moments are backed out by considering the circumstances under which a worker would find it optimal to make a particular transition. Take the transition rate from unemployment to employment from region m to region k, $T_{ue,mk}$. Using the notation from section 3.2, there are $\alpha_m \mu_m$

unemployed residing in region m. In order for some of these unemployed to move into employment in region k, they need to receive a relevant offer. This happens with probability $\zeta_2\theta_m$. However, the move is only going to happen if it is better than all other options available to the worker, in this case unemployment in one of the 4 regions of the US. This conditional probability depends on the utility of the employment and unemployment options, the worker's idiosyncratic location preferences, and the relevant migration costs. The probability has to be calculated for all possible wage offers, and the overall expression needs to be weighted by the probability of receiving the particular wage offer, $f_k(z)$. This is expressed in the second line of the transition rate, equation (3.14). Assuming that the idiosyncratic location preferences are type I extreme value distributed, this probability has a closed-form, logit solution. The resulting transition rate is given in expression (3.14); note that τ denotes the transpose of a matrix.

$$
\begin{aligned}
T_{ue,mk} &= \text{job-specific migration from unemployment in } m \text{ to } k \\
&= \alpha_m \mu_m \zeta_2 \theta_m f_k(z)^\tau \begin{bmatrix} \Pr(V_k(z_1) + \gamma_k - K > \{U_m + \gamma_m, ..., U_J + \gamma_J - K\}) \\ \Pr(V_k(z_2) + \gamma_k - K > \{U_m + \gamma_m, ..., U_J + \gamma_J - K\}) \\ ... \\ \Pr(V_k(z') + \gamma_k - K > \{U_m + \gamma_m, ..., U_J + \gamma_J - K\}) \end{bmatrix} \\
&= \alpha_m \mu_m \zeta_2 \theta_m \int \frac{exp(V_k(z) - K)}{exp(V_k(z) - K) + exp(U_m) + \sum_{j \neq m}^{J} exp(U_j - K)} \mathrm{d}F_k(z)
\end{aligned}
\tag{3.14}
$$

The other moments are derived analogously (see Appendix B.2).

The 19 model parameters are determined jointly by all the data moments, but moment expressions such as (3.14) enable us to see the links between the model and the data explicitly. In this particular case, knowing the equilibrium population distribution and unemployment rates, α_m and μ_m, along with the rest of the model parameters, would allow me to uniquely identify the unemployed job search wedge ζ_2.

Identification of other parameters can be argued similarly. ζ_1 is identified from the between-region employment-to-employment transition rates. The region-specific domestic job-finding probabilities for the employed and the unemployed are

backed out from the within-region employment-to-employment and unemployment-to-employment transitions, respectively. δ_j, the regional job-destruction probability, is determined by the local employment-to-unemployment moves. Overall migration rates pin down migration costs, while the region-specific populations identify average location preferences.

With 19 unknowns and 63 equations, the order condition is satisfied. This is particularly helpful in my case because migration is a relatively rare event. Some off-diagonal elements of the moment matrices in Tables 3.1 and 3.2 are small and close to 0, so having multiple observations on between-region unemployment-to-employment flows allows me to estimate parameters such as ζ_2 more accurately.

3.3.3 Method: estimation by simulated moments

The model parameters are recovered using the method of simulated moments (SMM) (McFadden (1989)). An algorithm finds the set of parameter values such that the moments generated by the data match those generated by the model.

In practice, SMM estimation involves minimising the squared distance between the data moments (as given in Tables 3.1 and 3.2) and the corresponding moment expressions as derived from the structural model (an example of such expression is given in equation (3.14)). Define $\mathbf{D}$ to be a 63×1 vector of data moments and $\mathbf{M}(\rho)$ the 63×1 vector of moment expressions, where ρ is the set of the 19 unknown parameters:

$$\rho = \{\delta_1, \delta_2, ..., \delta_4, \theta_1, ..., \theta_4, \lambda_1, ..., \lambda_4, \zeta_1, \zeta_2, K, \bar{\gamma_1}, ..., \bar{\gamma_3}, g\} \tag{3.15}$$

The SMM estimate of ρ, $\hat{\rho}$, is then a 19×1 vector of variables that can be defined as:

$$\hat{\rho} = \text{argmin}(\mathbf{M}(\rho) - \mathbf{D})^{\tau}\mathbf{W}(\mathbf{M}(\rho) - \mathbf{D}) \tag{3.16}$$

$\mathbf{W}$ is the weighting matrix. It determines how much "attention" to pay to individual moments in the minimisation process. $\mathbf{W}$ is often the inverse of the variance-covariance matrix of the data moments. I use an identity matrix, giving an equal weight to all moments. Altonji and Segal (1996) argue that variance-covariance

matrices suffer from a serious downward bias both in small samples and when the distribution is poorly behaved. In this case, even though each sample contains at least 4,000 individuals, the moments that describe migration are often based on a much smaller number of observations. As a result, a variance-covariance matrix would offer relatively little benefit over a simple identity matrix.

The first step in the estimation procedure is to randomly set the parameters and use them to evaluate employment and unemployment Bellman equations and workers' conditional choice probabilities.

Apart from conditional choice probabilities, the model moment expressions also depend on the equilibrium population distribution and local unemployment rate, α_j and μ_j. I take these directly from the data. Ideally, I would solve the model for every set of parameter values and use the corresponding α_j and μ_j; this would considerably increase the computational burden of estimation. Given that the analysis in this model is based on the assumption that the economy is in equilibrium, the observed population shares and unemployment rates should be equal to these values.

I make a similar argument about the equilibrium wage distribution, which is also taken from the data. I approximate it using a 2-point support separately for each region: log wages at 25th and 75th percentile, attributing a 50% probability density to both. To recover the distribution of wage offers, I plug the initial model parameters, along with the endogenous variables calculated above, into the equilibrium condition for the accepted wage distribution. Inverting the set of equations that define accepted wages in equilibrium allows me to solve for the distribution of wage offers, which can then be plugged into the moments. In the baseline specification, results of which are presented in the next section, I skip this step and use the approximation for accepted wage distribution as approximation for wage offers.

Once the model moments for a particular guess of parameter values are calculated (vector $\mathbf{M}(\rho)$), I compare this against data moments (vector $\mathbf{D}$). I repeat these steps until the squared difference of the model and data moments is sufficiently close to 0. The corresponding parameter values are the SMM estimate $\hat{\rho}$.

Table 3.3: The SMM estimates of the structural model: goodness of fit

	less educated		more educated	
	data	model	data	model
employed to				
employed home	98.27%	98.31%	98.95%	98.98%
employed away	0.04%	0.04%	0.11%	0.11%
unemployed home	1.67%	1.65%	0.93%	0.91%
unemployed away	0.01%	0.00%	0.00%	0.00%
unemployed to				
employed home	10.47%	10.22%	11.15%	10.81%
employed away	0.02%	0.00%	0.07%	0.03%
unemployed home	89.43%	89.68%	88.56%	88.93%
unemployed away	0.08%	0.10%	0.22%	0.23%

3.3.4 Results

The goodness of the model's fit is summarised in Table 3.3. Instead of comparing the model predictions for all 63 moments, they are aggregated into 8 transition rates, between employment and unemployment and home and away regions. The fit is good: for most moments, the difference between data and the model prediction is less than two decimal places.

The results of the estimation are presented in Table 3.4, separately for the less and the more educated. The main parameters of interest are the cross-regional job search wedges, ζ_1 and ζ_2. These variables capture the proportion between away and home job offers. A number close to 1 means that receiving a cross-regional job offer is as easy as finding a job locally, while a number close to 0 means that cross-regional job offers are virtually impossible compared to local ones.

I find that the wedges are sizable and vary by education. An employed college graduate is 35% less likely to receive a job offer from another region, compared to local offers. The wedge is even greater for those without a college degree, who receive on average four cross-regional offers for each ten local ones. There are also differences between search on-the-job and when unemployed. The latter is virtually infeasible: workers of either education receive less than one cross-regional offer for a hundred of local offers.

Table 3.4: The SMM estimates of the model parameters, by education

description	parameter	less educated	more educated
job offer arrival rate, on-the-job, Northeast	λ_1	0.857	0.660
job offer arrival rate, on-the-job,Midwest	λ_2	0.907	0.514
job offer arrival rate, on-the-job, South	λ_3	0.869	0.644
job offer arrival rate, on-the-job, West	λ_4	0.998	0.619
job offer arrival rate, unemployed, Northeast	θ_1	0.096	0.110
job offer arrival rate, unemployed, Midwest	θ_2	0.102	0.113
job offer arrival rate, unemployed, South	θ_3	0.096	0.102
job offer arrival rate, unemployed, West	θ_4	0.116	0.110
job destruction probability, Northeast	δ_1	0.016	0.008
job destruction probability, Midwest	δ_2	0.015	0.008
job destruction probability, South	δ_3	0.017	0.009
job destruction probability, West	δ_4	0.019	0.012
job search wedge, on-the-job	ζ_1	0.399	0.674
job search wedge, unemployed	ζ_2	0.0001	0.001
migration cost (log)	K	12.626	11.485
std. dev. of location preferences	g	1.609	1.627
mean location preference, Northeast	$\bar{\gamma_1}$	-0.162	-0.037
mean location preference, Midwest	$\bar{\gamma_2}$	0.004	0.108
mean location preference, South	$\bar{\gamma_3}$	0.289	0.240

The differences in the estimated migration costs are similarly large. Converted to US dollars, the total moving costs for a college graduate are almost $100,000; they rise to $300,000 for those without a college degree. These estimates are in line with the large range found in the literature: Kennan and Walker (2011) estimate the migration cost of high-school graduates at more than $300,000, while Amior (2015) arrives at a much more modest $13,000.

The education variation in the rest of the model parameters are much smaller. In general, in line with the greater turnover in the labour markets for less skilled labour, the job-finding and job-destruction variables are somewhat larger for those without a college degree. On the other hand, the distributions of their location preferences are very similar, suggesting that variation in location preferences is

unlikely to account for much of the migration differences. The next section explores this question in greater depth.

3.4 Explaining the patterns of internal migration

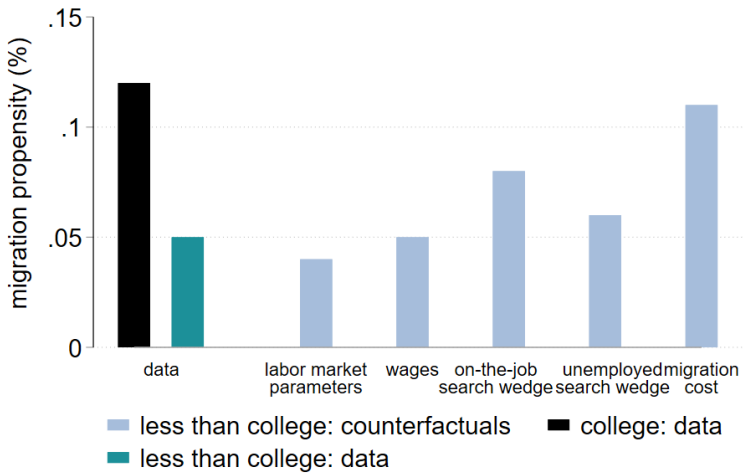
Estimates of the structural model in section 3.3 suggest that cross-regional job search is more frictional than local search, and that these frictions are larger for the less educated. In this section, I revisit the question of why the less educated workers in the US move less than those with a college degree. I use the structural estimates to compare the role of search frictions against other factors that shape migration behaviour, such as wages, migration costs, and preferences. I show that job search frictions can account for almost half of the observed gap in migration propensity between the less and the more educated, and 75% of the differences in the type of migration. This provides additional support for the findings from Chapter 2, and allows me to draw more general conclusions about the importance of the interactions between the employment and location choices of a worker, in particular how much job search matters in the migration decision.

3.4.1 Decomposing the migration propensity gap between the more and the less educated

The less and the more educated vary along many dimensions: labour market variables, wage distribution, location preferences, migration costs, and cross-regional job search frictions. These differences combine to give rise to the migration propensity and migration type gaps described in Chapter 2: first, that the less educated move less than those with a college degree; and that conditional on moving, the less educated are less likely to move for a specific job. I use the structural estimates (Table 3.4) to decompose the observed behaviour into these individual components.

The decomposition of the relative difference in migration propensity by education is presented in Figure 3.5. It plots the changes to the migration propensity of the less educated when I change specific parameters so that they equal those of the

Figure 3.5: Decomposition of the education differences in migration propensity

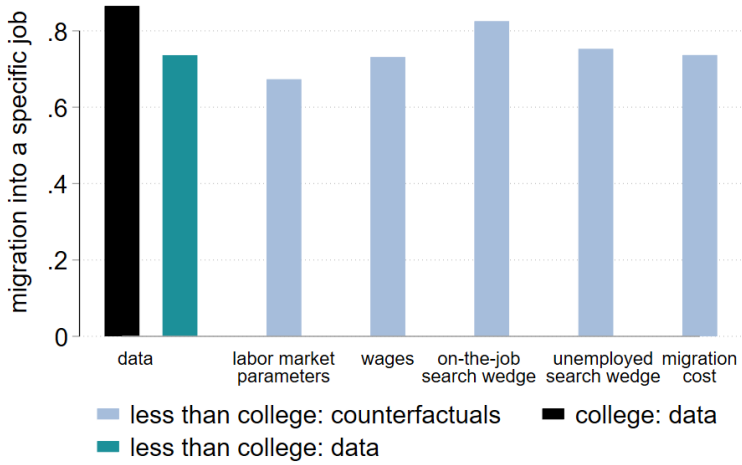


Calculated using the structural estimates in Table 3.4. The two leftmost bars represent observed behaviour. The other five bars plot the change in the migration behaviour of the less educated if the given parameters were equal to that of the more educated. For example, the rightmost bar in panel (a) shows that if the less educated faced the same migration costs as the more educated, their propensity to migrate would be almost as high as that of the more educated.

more educated. For example, the "labour market parameters" bar shows that the less educated would move less if their job-finding and job-destruction probabilities were the same as that of the more educated, keeping all the other parameters fixed. In general, the decomposition shows that the combined job search frictions are as important as migration costs in explaining the migration propensity differences. On the other hand, wage differences play an insignificant role. This is partly because higher mean and higher variance have opposing effects on migration propensity, effectively cancelling each other out. However, it also reflects that the observed regional wage differences are not sufficiently large to motivate migration flows.

Figure 3.6 repeats the exercise for the propensity to migrate for a specific job, conditional on moving. Here, cross-regional job search frictions are the primary explanation, combining to wipe out most of the higher share of job-specific migration among the more educated. Migration costs, a prominent factor behind the lower migration propensity, play a minor role, because they have to be paid regardless of which way the worker migrates. The education differences in wages are virtually insignificant, for the same reason as above.

Figure 3.6: Decomposition of the education differences in migration for a specific job



Calculated using the structural estimates in Table 3.4. The two leftmost bars represent observed behaviour. The other five bars plot the change in the migration behaviour of the less educated if the given parameters were equal to that of the more educated. For example, the rightmost bar in panel (a) shows that if the less educated faced the same migration costs as the more educated, their propensity to migrate would be almost as high as that of the more educated.

The impact of search frictions is large because they change the sequencing of workers' migration and employment decisions. In the absence of cross-regional job offers, all migration is by definition speculative – a decision in which income is uncertain. Receiving job offers in other regions makes it possible for the worker to decide about migration *after* the income uncertainty has been resolved[5]. These two strategies may deliver very different utility, resulting in different migration patterns. This is why the estimated cross-regional job search frictions result in stark differences in the migration propensity of college graduates versus those without a degree. It also demonstrates the importance of the interactions between the employment and the location choice of an individual[6].

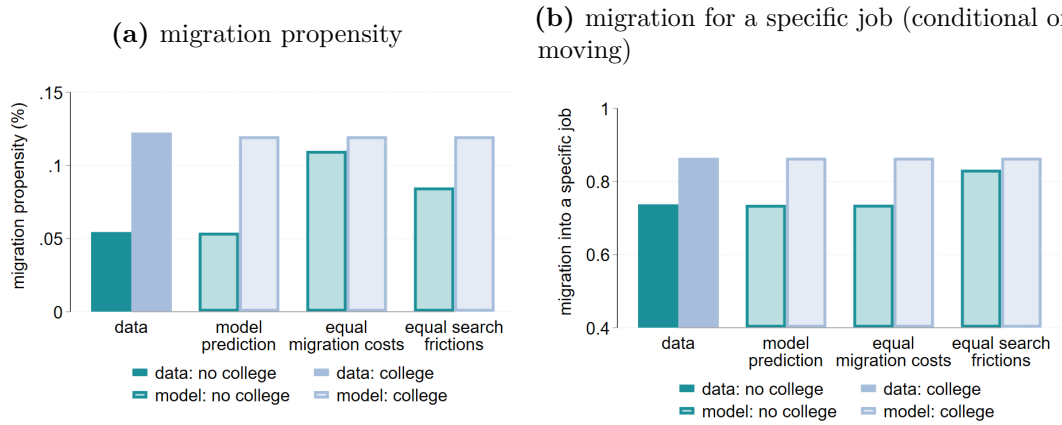
3.4.2 Policy implications: the role of job search frictions

Understanding why and how the less educated move is vital for addressing both regional inequalities and worker mobility. There is evidence that migration is

[5]However, this does not mean that the decision is made under certainty. Factors such as location preferences and future employment opportunities are still unknown to the migrant.

[6]This mechanism does not rely on risk aversion. The theoretical model presented in this chapter assumes all workers are risk-neutral.

Figure 3.7: Counterfactuals comparing the role of migration costs and cross-regional search frictions in migration behaviour



Calculated using the structural estimates in Table 3.4. The two leftmost bars represent observed behaviour. The next set from the left plot the values predicted by the model. The other bars plot the counterfactual migration behaviour of the less educated if the given parameters were equal to that of the more educated. For example, the rigthmost bars show that if the less educated faced the same search frictions in their cross-regional job search as the more educated, about half of the migration propensity gap, and almost all of the migration type gap, would disappear. The search frictions scenario changes the search wedge for both on-the-job and unemployed search.

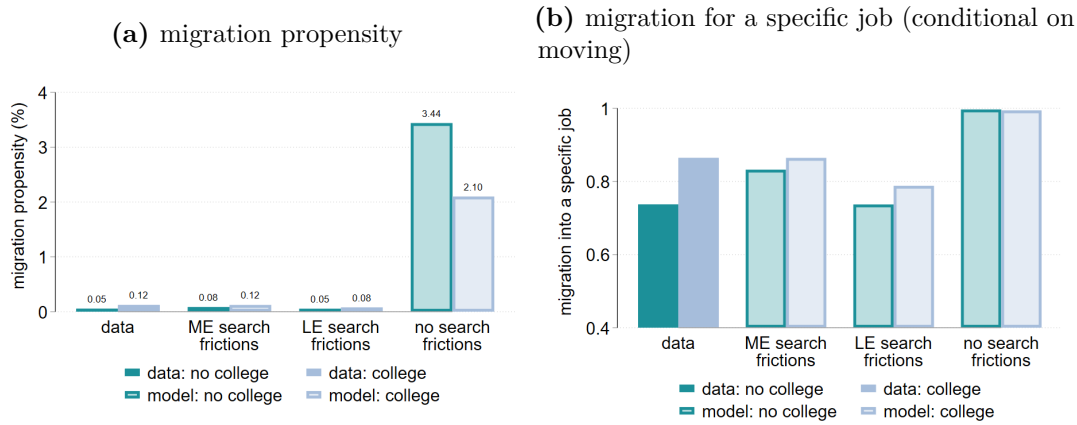
an important buffer to local demand shocks (Bartik (2018), Blanchard and Katz (1992)), but research also shows that the less educated, low-income households that are more exposed to these shocks are also least likely to move (Bound and Holzer (1996), Wozniak (2010), Hoynes (1999)). The implications are potentially large: Amior and Manning (2018) argue that this lack of mobility leads to the persistent and substantial regional disparities in unemployment and welfare. Governments do intervene in these regions, but the impact of their policies is often ambiguous and largely depends on workers' migration response[7] (Kline and Moretti (2013a), Glaeser and Gottlieb (2008)).

In this section, I construct a series of counterfactuals evaluating the possible impact of reducing job search frictions on migration behaviour. I compare it to the effect of reduction in migration costs, since subsidies to moving are often suggested as a solution to low mobility. I then discuss the potential policy lessons of this paper.

First, I look at what would happen if cross-regional search frictions were equalised across education, and compare it to an alternative policy in which I eliminate the

[7]For example, high worker mobility in the surrounding regions reduces the local impact of job-creation subsidies.

Figure 3.8: Counterfactuals evaluating the role of cross-regional search frictions on migration behaviour



Calculated using the structural estimates in Table 3.4. The two leftmost bars represent observed behaviour. The other bars plot migration behaviour when cross-regional search frictions change. The "ME search frictions" cluster represents the migration behaviour if both education groups faced the cross-regional search frictions of the more educated group. The "LE search frictions" do the same but with search frictions of the less educated. The rightmost bars consider what would happen if there were no cross-regional frictions, i.e. if it was as easy to find a job elsewhere as it is to find a job locally. The job-finding probabilities (and other labour market variables) remain education-specific.

differences in migration costs. The first set of bars in both panels in Figure 3.7 plots the observed migration behaviour (migration propensity and the share of migration that is into a specific job). The last set depicts a scenario in which I make cross-regional job search as easy for the less educated as it is for the college graduates (changing the job search wedges for both the employed and the unemployed). The central set does the same, but with migration costs, lowering them to the level of the college graduates. Lowering migration costs has a bigger impact on closing the education gap in migration propensity, but a large portion of the less educated are still forced to move speculatively. Improving cross-regional job search, on the other, closes about 75% of the gap in the type of migration, and almost a half of the difference in migration propensity.

What would happen if we could remove cross-regional search frictions entirely? I answer this question with a second set of counterfactuals, depicted in Figure 3.8. It allows me to see how would migration behaviour of the less and the more educated change under search frictions of the other group, and if cross-regional job search became as easy as the local one. The main takeaway, in line with

the decomposition in the previous section, is that migration behaviour would be radically different if the cross-regional search wedge was eliminated. Not only would virtually no workers move speculatively, but the less educated would in fact be more mobile than the more educated. This has a deep implication for the role of labour markets on mobility: the less educated don't move because they get few opportunities to move for a specific job.

Figure 3.8 also shows how much the migration behaviour of the college graduates depends on their relatively frictionless cross-regional job search. Their migration propensity would fall by a third, and the rate of their speculative migration would be almost in line with the less educated, if they faced the same cross-regional search frictions as the less educated. The difference in the overall migration propensity between the two education groups would fall to 3 percentage points.

Of course, these counterfactuals estimate the upper bound of the search friction effect. The analysis ignores any negative spillover and general equilibrium effects, such as equilibration of rents and wages, changes in firms' vacancy posting strategies, and negative search externalities among the workers. All of these would likely reduce the overall effect estimated here.

There are several reasons why a policymaker wishing to improve mobility may want to consider addressing the job search frictions rather than simply subsidising migration costs. First, even though levelling migration costs between the education groups has a bigger impact on the propensity to migrate than equalising cross-regional search frictions, the latter has a bigger impact on the differences in the type of migration (Figure 3.7). This matters because the discrete choice model in Chapter 2 shows that migration for a specific job is preferred to speculative migration, so allowing workers to move more *and* move more into a specific job would result in greater welfare gains. By changing migration costs, the policymaker is changing the cost-benefit balance of migration, but by addressing search frictions the policymaker improves the functioning and the efficiency of the market, increasing

rather merely re-distributing welfare[8]. Second, the difference in the migration costs between the more and the less educated estimated in the structural model is large, in the ballpark of $200,000, making investments into frictionless cross-regional job search a potentially cheaper alternative.

Improving cross-regional job search for those without a college degree would allow economic growth to spread more evenly. A policymaker may want to focus on sustaining strong growth in the better-off regions, rather than trying to relocate business into the depressed regions. In that case, a well-working inter-state job market would be a key part of the equilibrating channel.

An active labour market policy aimed at improving posting vacancies across regions would probably focus on two mechanisms. First, there is the cost of vacancy posting across large geographic areas. It may seem that this would be lowered sufficiently with the advancement of the internet, but the actual impact on the labour market is ambiguous. Data on online vacancy posting (Hershbein and Kahn (2018)) reveals there still are significant differences in the likelihood of different jobs being posted. However, advertisement of vacancies is only a small portion of the "black box" frictions discussed in this paper. The human resources literature points out that larger, more specialised firms not only advertise more widely, but they also interview more candidates and use different selection procedures. The second part of labour market intervention may thus look at the factors that determine the firm's recruitment strategy. For example, interview transportation subsidies for the unemployed could make it feasible for workers to not only apply, but actively pursue job search in distant regions (for empirical evidence on the effectiveness of this channel, see Franklin (2018)).

3.5 Conclusion

The decisions about employment and location are not made in isolation. In this chapter, I propose an extension to the standard model of job search to capture

[8]Of course, the matter is not as simple. Kline and Moretti (2013a) demonstrate that in a second-best world and/or in the presence of agglomeration forces, place-based policies may in fact improve overall welfare.

their interaction. The model summarises the links between job search across and within region and a worker's location choice. It allows for a rich set of parameters describing the local labour markets, worker's location preferences, and migration costs. This then allows me to map these parameters onto the observed transition rates in and out of employment and between and within local labour markets.

The estimates of the structural model build on the results in Chapter 2. I show that cross-regional search suffers from greater frictions than searching at home, and that these frictions decrease in education. The subsequent impact on migration behaviour is significant and highlights the need to take into account the workers' job search patterns in order to understand their mobility. Because the results from this chapter rely on the structure of the model rather than on proxies of recruitment behaviour, I provide complementary evidence on the drivers of the migration propensity between the less and the more educated workers in the US. The counterfactuals of the structural model suggest that almost a half of this gap in mobility can be attributed to the fact that those without a college degree find remote job search more difficult, both because they spend more time unemployed and because they face greater search frictions.

The policy implications of this analysis are two-fold. They help to explain why the less educated move less, even though they are more exposed to local labour market shocks. The counterfactuals suggest that improving the efficiency of cross-regional labour market could not only have a potentially large impact on overall mobility, but also improve individuals' welfare by allowing them to migrate after the job search uncertainty has been resolved.

In general, the relatively simple structure of the model opens the scope to study the impact of active labour market policies on mobility, and to deepen our understand of why people do, or do not, move for opportunity.

4

Labour market frictions and regional disparities

4.1 Introduction

Location matters. Even within the same country, where laws, language, and culture should present minimal obstacles, there are stark regional differences in determinants of welfare such as incomes, unemployment rates, education attainment and life expectancy (Moretti (2011); Kline and Moretti (2013b); Chetty et al. (2016)). Perhaps more importantly, these differences persist over time. The regions that suffer from the highest unemployment rates today are very likely to be the same regions that suffered from high unemployment decades ago (Overman and Puga (2002)).

Search frictions play a potentially large role in generating and perpetuating these regional disparities. First, local labour markets vary in how effectively they match workers and prospective employers (Gautier and Teulings (2009)). This translates into differences in the type of local jobs and their productivity. We would expect these regional differences to be arbitraged away by people moving from the low-welfare to the high-welfare regions, but search frictions *across* regions can significantly slow down this process. As I showed in Chapter 2, about two thirds of inter-state migration in the US is for a specific job. Search frictions make it

more difficult to find a job before moving, reducing overall migration flows and slowing down the equilibrating process.

In this chapter, I build a theoretical model of search and matching within and across regional labour markets that formalises this argument. Its key assumption is that the matching function exhibits increasing returns to scale. I show that if matching efficiency increases in the size of labour market, more populous regions will enjoy lower unemployment rates and higher wages than the smaller regions, even if their production functions are identical. I then introduce two types of cross-regional search frictions: cost of searching in a new labour market and the degree of regional matching spillovers. Equilibrium migration flows depend on the type of these frictions and their magnitude. The relatively easier it becomes for a migrant to search in her new region of residence, the less attractive migration for a specific job becomes for employers, and this reduces total mobility. On the other hand, as regional matching functions become more inter-linked, hiring in the smaller regions where wages are lower becomes relatively easier. This encourages migration for a specific job, and boosts mobility.

My point of departure is a standard search and matching model with homogenous workers and firms (Pissarides (2000)). In this model, workers decide how many vacancies to post, and unemployed workers search. When a potential employer and employee meet, they bargain over how to split the match output and form an employment contract. The meetings between workers and their potential employers take time and resources, and this friction translates into positive equilibrium unemployment. Following Ellison et al. (2014), I extend this model by assuming that the matching function exhibits increasing returns to scale. As a result, the size of the worker population becomes another model parameter. I prove that under certain conditions a unique labour market equilibrium exists, and there will be a one-to-one mapping between labour market conditions and regional population. The comparative statics then show how population differences lead to divergence in local labour market outcomes.

In the second part of the chapter, I make regional population endogenous by modelling migration. Workers can either move speculatively, arriving at their new destination unemployed and searching for jobs there, or move for a specific job they found while in their home region. Workers move speculatively whenever the utility gap between regions outweighs the migration cost. Migration for a specific job, on the other hand, also depends on whether any firms find it profitable to advertise their vacancies in another region. While the direct costs of vacancy creation are the same everywhere, the expected benefit is determined by the wages and the efficiency of matching in the local labour market. Cross-regional vacancies are profitable if lower wages outweigh the cost of hiring in smaller, less efficient labour markets. I show that both types of migration flows will be from the smaller region to the larger one, in line with the differences in the efficiency of the labour markets and hence unemployment rates and income. I also identify the intervals in which the economy is in a steady state, and describe certain aspects of the adjustment process between the steady-state equilbria.

Finally, I link the model to the stylised facts about regional disparities that motivate this chapter. The presence of increasing returns to scale in matching can replicate the pattern found in most developed countries of large regional disparities in income and the unemployment rate, and their positive correlation with each other and the regional population. Due to migration costs and search frictions, migrants move from smaller, less prosperous regions into the larger, more prosperous ones, which exacerbates regional differences. These differences persist over time because the regional economies are prone to fall into a high-divergence steady state: migration costs and cross-regional search frictions are too large to encourage speculative migration, but not so small as to prevent the migration for a specific job that drove the economy into a high-divergence steady state in the first place. As a result, labour market frictions can both generate and perpetuate regional differences.

Contribution to the literature This chapter builds on the existing literature on local labour markets and search and matching models. It contributes to this literature in three ways.

First, it fills a gap in the way we model regional variation in labour market outcomes. The key assumption lies in increasing returns to matching as the source of agglomeration forces: I show how differences in labour market frictions can theoretically cause the observed regional disparities. This is not a new idea: the possibility of thick labour markets or labour market pooling effects originate in Marshall (Krugman (1991)), and have been explored in other fields of labour economics[1]. However, the existing literature on regional disparities either assumes exogenous productivity differences across cities/regions (e.g. Moretti (2011)), models agglomeration as the result of transportation costs and forward and backward linkages (Epifani and Gancia (2005), Francis (2009), Lutgen and der Linden (2015)), or of knowledge spillovers. In these models, the observed differences in local labour markets are the consequence of the underlying agglomeration forces. In the model presented here, on the other hand, it is the differences in matching efficiency in the local labour markets that drive regional differences.

The assumption of increasing returns to matching is empirically plausible, but the existing research is so far inconclusive. Pissarides et al. (1986) and Blanchard and Diamond (1989) run regressions of number of matches on the number of vacancies and unemployed for the aggregate economy; both are unable to reject the null hypothesis of constant returns to scale, although Blanchard and Diamond (1989) find evidence of economies of scale in the industry-specific matching functions. On the other hand, similar research by Yashiv (2000) and Warren (1996) finds evidence of economies of scale. Petrongolo and Pissarides (2006) point out that the hiring rate depends not only on the rate of meetings between firms and workers, but also on wages. In particular, the upward wage adjustment may completely offset the impact of the higher meeting rate, so that increasing returns would not be detected from hiring data only. The authors thus separately estimate a wage

[1] Ellison et al. (2014) analyse the behaviour of a standard search and matching model as returns to scale change from decreasing to constant to increasing.

equation and a workers' hazard rate. They find an economies of scale effect in wages, but not in the matching function. Indirect evidence of localised returns to scale can also be found in Wheeler (2008) and Bleakley and Lin (2012). Perhaps the most supportive evidence comes from Baker et al. (1996). The authors focus on the potential mis-measurement of the number of searching workers which arises when studies don't include on-the-job search. Using a Canadian dataset, they demonstrate that correcting for this underestimate reveals significant economies of scale in the matching function. Their paper also reveals considerable variation of the efficiency of matching across Canadian provinces.

Second, this chapter extends how we think about the interactions between local labour markets, by modelling cross-regional hiring alongside speculative migration. Most of the existing literature assumes only the latter: in effect, regional labour markets are modelled the same way as countries, only with relaxed migration restrictions. Notable exceptions are Schmutz and Sidibé (2018), Lutgen and der Linden (2015) and Beaudry et al. (2014), who allow workers and firms from different regions to meet in other ways than migration. In Beaudry et al. (2014), workers' migration is conditional on first forming a match with an employer; and both Schmutz and Sidibé (2018) and Lutgen and der Linden (2015) allow for both move-search and search-move migration. However, in none of these papers does the distinction matter for the equilibrium result.

In the existing literature, these interactions between local labour markets do not affect regional outcomes. This is in sharp contrast with the model presented here, where workers and firms make optimal decisions about whether to migrate and in which way, and these decisions shape regional divergence. This is the third contribution of this chapter. I demonstrate how migration increases regional disparities in labour market outcomes and utility, and how different types of cross-regional labour market frictions make migration more or less attractive. Overall, the state of local labour markets not only drives regional differences at a given point in time, but also shapes the dynamic process of regional divergence.

This chapter is organised as follows. I start by presenting the five stylised facts that motivate this paper. In section 4.3, I outline the basic model of search and matching in local labour markets when regional population levels are fixed; in section 4.4, I endogenise this by allowing for migration. In section 4.5, I simulate the model to illustrate the process of regional divergence and how it links to migration costs and cross-regional search frictions. I also discuss the welfare implications of regional divergence. Section 4.6 concludes.

4.2 Motivation

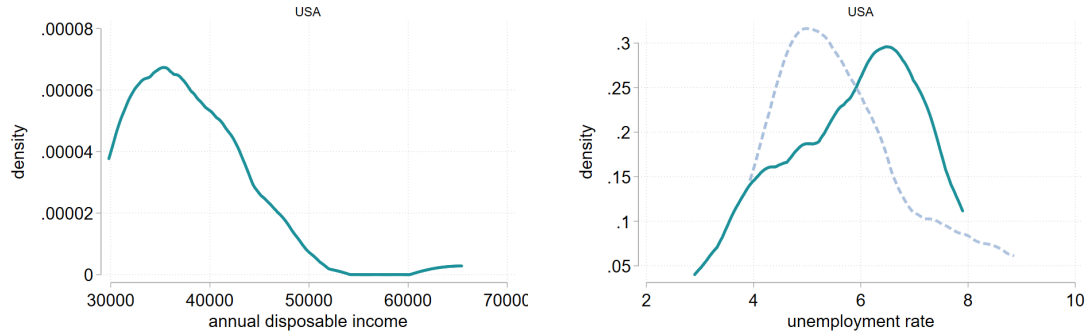
This chapter is motivated by five stylised facts about regional labour markets in developed countries. In the remainder of the chapter, I present a model of local and cross-regional labour market frictions that can generate these stylised facts.

First, regional differences are large. Second, the regions with high wages also tend to enjoy high employment rates. Third, both the regional employment rate and income are positively associated with the population size of the region. Four, these regional differences persist strongly over time. Finally, workers migrate from the smaller, poorer regions to the large ones, exacerbating the regional disparity.

The extent of regional differences in income and unemployment rates for the case of US states is plotted in Figure 4.1. The left panel shows that regional income in the richest US states is more than twice as high as that in their poorer neighbours. The pattern for regional unemployment rate is similar. To understand the magnitude of these regional differences, the right panel also plots the variation in aggregate unemployment over time. As Moretti (2011) points out, regional differences are often comparable, if not larger, than the business cycle variations that receive a lot of attention from economists and policymakers alike.

In Figure 4.2, I show that these regional differences in income and employment are associated with regional population size. US states with higher population density enjoy on average higher incomes and lower unemployment rates than the more sparsely populated regions. As a consequence, regional differences display

Figure 4.1: Distribution of regional income and unemployment rates



The left panel plots the kernel density of annual average regional disposable income in the year 2011. The right panel plots the kernel density of annual average regional unemployment rate in 2014 (solid line), and the kernel density of the average aggregate unemployment rate during the years 1990-2014 (dashed line). Data source: OECD regional database for TL2 – large regions.

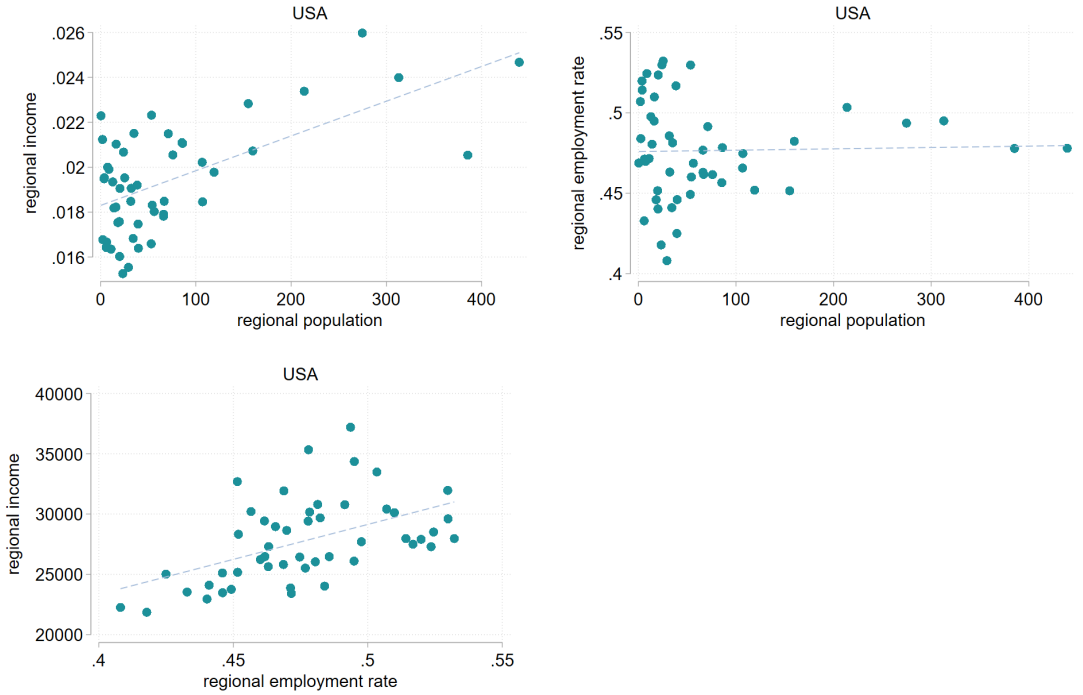
a "wage curve" relationship: a negative correlation between regional income and unemployment (Blanchflower and Oswald (1995)).

Regional differences are not only large, but often strongly persistent as well. As shown in Figure 4.3, the level of regional income and employment in the past has a significant impact on their values today: the correlation coefficient between 1990 and 2014 values is 0.8 and 0.7 for regional income and employment rate, respectively.

Finally, the observed migration flows seem to add to these differences. Workers tend to move away from the high-unemployment, low-wage, low-density regions to the more populated ones. This gives rise to a positive relationship between net emigration from a region, and its size (Figure 4.4).

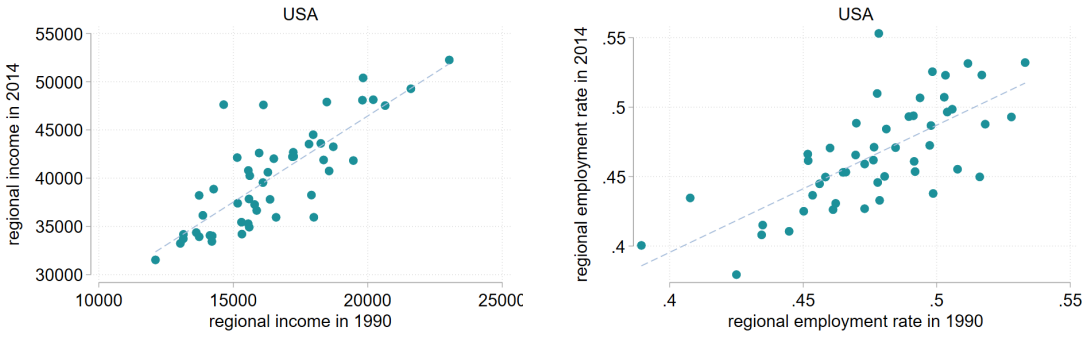
These stylised facts extend to other developed countries (see Appendix C.1), and have been widely documented in the literature. Moretti (2011) and Kline and Moretti (2013b), and Overman and Puga (2002) document the extent and persistence of disparities across US metropolitan areas and transnational EU regions, respectively. Moretti (2011) gives evidence on the sizeable differences across metropolitan wages, productivity and innovation, while Overman and Puga (2002) highlight the persistence of regional differences in wages and unemployment. Kline and Moretti (2013b) show that a substantial regional dispersion in wages and unemployment rates persists even after controlling for local human capital; and Moretti (2011) and Overman and Puga (2002) argue that the gap is in fact widening.

Figure 4.2: Correlations between regional population size, income, and employment rate



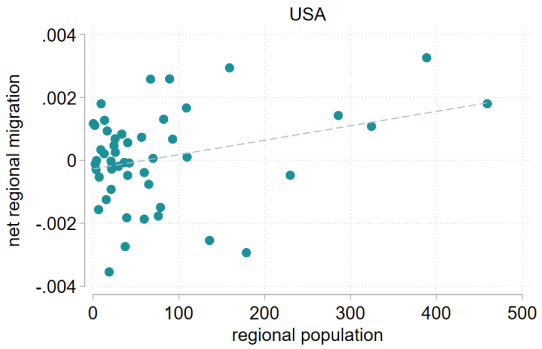
The top left panel: regional population density in the US against relative regional disposable income. Top right panel: regional population density in the US against employment relative to total population. Bottom left panel: employment rate (based on population) against regional disposable income in the US (the wage curve). Data corresponds to averages over the years 1990-2014, based on availability. Data source: OECD regional database for TL2 – large regions.

Figure 4.3: Persistence of regional income and employment rate over time



The graphs plot the correlation between regional income (left panel) and regional employment rate (right panel) in the first year (1990) and last year (2104) of data availability, in the US. Data source: OECD regional database for TL2 – large regions.

Figure 4.4: Correlation between regional population and the size of net migration flows



The graph plots the correlation between regional population density in the US and net regional immigration in 2011. Data source: OECD regional database for TL2 – large regions.

This is partly due to the above mentioned direction of migration into regions with higher wages and lower unemployment: regions with greater agglomeration experience positive net inflows of labour (Beaudry et al. (2014), Pissarides and Wadsworth (1989) for the UK).

4.3 Local labour markets under autarky

4.3.1 The model

4.3.1.1 The matching technology

The labour market is stochastic: workers and employers meet at random and decide whether to form an employment contract. The number of meetings per unit of time, M, depends on the stock of unemployed u and the stock of existing vacancies v. Following Ellison et al. (2014), I assume that $M(u, v)$ is quasi-concave and homothetic. This allows me to describe the matching on the labour market as a function of the labour market composition (vacancies relative to unemployed) and size.

First, the stock of unemployed and vacancies combine to determine the size of the labour market $m(u, v)$. This function is increasing in both arguments, concave and homogenous of degree 1. As a result, market size (or activity) per unemployed worker, μ, as well as the elasticities of m with respect to u and v, depend only on labour market tightness $\theta = v/u$. θ captures the composition of the labour market: the ratio of vacancies to unemployed.

$$\mu(\theta) \equiv \frac{m}{u} = m(1, \theta) \tag{4.1}$$

$$\frac{\mathrm{d}m}{\mathrm{d}v}\frac{v}{m} \equiv \alpha(\theta) \tag{4.2}$$

$$\frac{\mathrm{d}m}{\mathrm{d}u}\frac{u}{m} \equiv 1 - \alpha(\theta) \tag{4.3}$$

The matching function Φ then converts the market size m into matches: $M = \Phi(m)$, where $\Phi(m)$ is a strictly increasing function of m. The key assumption in this paper is that this process exhibits increasing returns to scale. In particular, I assume that the elasticity of the matching function Φ is globally increasing in market size m.

$$\frac{\mathrm{d}\Phi}{\mathrm{d}m}\frac{m}{\Phi} \equiv \eta(m) > 1 \ \forall \ m \tag{4.4}$$

I also define the average matching rate $\phi(m)$, which will be increasing when Φ exhibits increasing returns to scale.

$$\phi(m) \equiv \frac{\Phi}{m} \tag{4.5}$$

As Ellison et al. (2014) point out, the homotheticity assumption allows me to express the standard matching function (with constant returns to scale) as a nested version of the more general M. In the general model, the probability that a worker finds a job is product of labour market tightness and size:

$$\frac{M}{u} = \frac{M}{m}\frac{m}{u} = \phi(m)\mu(\theta) \tag{4.6}$$

When $\eta(m) = 1$, this probability collapses to a function of labour market tightness only, $m/u = \mu(\theta)$, as is the case in standard models of constant returns to matching (e.g Pissarides (2000)).

4.3.1.2 Employment and unemployment

Workers live infinitely, discounting the future at the rate r. They can be either employed, or unemployed and searching for a job. There is no labour market participation decision, and no on-the-job search. The income of an unemployed worker is normalised to 0.

Each firm has at most one potential job. It will create a vacancy if the expected profit of doing so is weakly greater than the cost of advertising, c. Under the standard assumption of no barrier to entry and constant returns to labour in the production function, the expected profit of a vacancy will always be pushed down to 0. If the vacancy is filled, the worker and the firm jointly generate a per-period output y until the match is dissolved.

Given the size and tightness of the labour market, the firms decide whether to create a vacancy, and the workers decide whether to accept the job offer. At the same time, δ of the existing matches are exogenously destroyed. These actions determine the labour market size and tightness in the next period.

4.3.1.3 The match surplus and surplus sharing

The matched firms and workers bargain over how to split their joint output y. Since all firms, workers, and matches are homogenous, all matches have the same surplus, and all pay the same wages and profits. In a steady state environment, the match surplus and the way it is shared is constant over time. Because of these

assumptions, all meetings will be converted into employment[2], and a match once created will not be endogenously terminated.

The surplus of a match is its product less the outside option of the firm and the worker. It will be shared according to their bargaining powers: the standard assumption is that this is a fixed proportion, with workers receiving a share β of the surplus and firms taking the remaining $1 - \beta$.

The outside option of a worker is the per-period return to unemployment, rU. It equals her share of the expected present discounted value of match surplus:

$$rU = \phi(m)\mu(\theta)\beta\frac{y - rU}{r + \delta} \tag{4.7}$$

The free entry condition implies that the outside option of a firm is 0. The firm's share of the expected present discounted value of match surplus will be just equal to the per-period cost of advertising a vacancy.

$$c = \frac{\phi(m)\mu(\theta)}{\theta}(1 - \beta)\frac{y - rU}{r + \delta} \tag{4.8}$$

Equations (4.7) and (4.8) describe the behaviour of the workers and the firms on the labour market, respectively. For the full derivation, including the Bellman equations that underpin these two equations, the Nash bargaining and the resulting wages (which are not explicitly necessary for this analysis), see Appendix C.2.

4.3.1.4 Local labour market

The total worker population of a region is P. The regional (local) labour market consists of the vacancies posted in this region, and all the unemployed workers residing there. I assume initially that regional population is fixed (this assumption will be relaxed in section 4.4). Firms are free to enter and exit the local labour market, so that the expected profit of a vacancy never exceeds the cost of posting it.

In standard search and matching models with constant returns to matching (such as Pissarides (2000)), the mass of workers is normalised at 1. Setting the

[2]In models with heterogeneous workers and/or firms, the profitability of a match is only revealed when the two parties meet. This may mean that not all meetings will result in employment. In this model, all matches are the same, so the surplus of a match is a priori known to both sides, and they would not search if it was not sufficiently profitable.

mass at P instead allows me to analyse how the labour market size and tightness, vacancy creation, unemployment rates and wages change with the size of the regional population in a labour market with increasing returns to matching.

4.3.2 The autarkic equilibrium

I define an autarkic local labour market steady state equilibrium as one where, given some level of regional population P, labour market size and tightness (m, θ) are constant.

The equilibrium is described by the worker and firm equations (4.7) and (4.8). Combining these two equations gives us a complete description of job creation in the local labour market as a function of two variables only, labour market tightness θ and labour market size m.

$$c = \frac{\phi(m)\mu(\theta)}{\theta} \frac{(1-\beta)y}{r+\delta+\beta\phi(m)\mu(\theta)} \tag{4.9}$$

An equilibrium requires θ, m to be constant, which in turn implies a constant stock of unemployed workers in the regional economy. Equation (4.10) below specifies the set of m, θ at which the flows into unemployment (the fraction δ of matches $P-u$ which were exogenously terminated) are just balanced by the outflow from unemployment as new matches are created.

$$\delta\left(P - \frac{m}{\mu(\theta)}\right) = \phi(m)m \tag{4.10}$$

An equilibrium in the local labour market is given by the labur market size and tightness pair (m, θ) that satisfy equations (4.9) and (4.10). The existence and uniqueness of such an equilibrium depends on the properties of the matching function (for a discussion and an extended proof of existence of equilibria, see Ellison et al. (2014)). In particular, under the assumption of increasing returns to matching ($\eta > 1$), a unique equilibrium at non-zero labour market activity will exist if there are decreasing returns to vacancy creation ($\alpha\eta < 1$). Intuitively, if the returns to vacancy posting were increasing, any positive stock of vacancies would lead to an explosion in matches, making it impossible for the flows in and out of unemployment to balance out.

Proposition 1. *Assume globally increasing returns to scale in matching:* $\eta(m) > 1 \ \forall \ m$. *An equilibrium* (m^*, θ^*) *exists at zero activity,* $m = \theta = 0$. *If* $\alpha\eta(m) < 1 \ \forall \ m$, *there is also a second equilibrium where* $m > 0$. *This equilibrium will be unique for all* $m > 0$. *If* $\alpha\eta(m) \geq 1 \ \forall \ m$, *there is no equilibrium where* $m > 0$. *(The model has a unique equilibrium at* $m > 0$ *if the returns to matching are constant or decreasing.)*

Proof. Equations (4.9) and (4.10) give us two implicit definitions of θ as a function of m: $\theta = \gamma_1(m)$ and $\theta = \gamma_2(m)$. An equilibrium is such that $\gamma_1(m) = \gamma_2(m)$.

$\theta = \gamma_1(m)$ is defined by equation (4.9), re-arranged below.

$$\phi(m) = \frac{\theta c(r+\delta)}{\mu(\theta)[(1-\beta)y - \beta\theta c]} \tag{9A}$$

The left hand side of this equation is some function of the market activity m, bounded by 0 and ∞. The right-hand side is a continuous, differentiable and increasing function of θ. Because $\phi(m)$ cannot be negative, the denominator of the expression implicitly sets an upper limit on θ:

$$\bar{\theta} = \frac{y}{c}\frac{1-\beta}{\beta} \tag{9B}$$

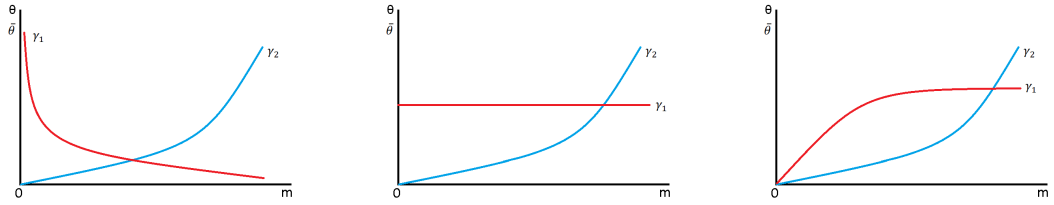
As $\theta \to \bar{\theta}$, the right hand side expression approaches ∞, while its lower limit, when $\theta \to 0$, is 0. The solution for $\gamma_1(m)$ thus lies on the $[0, \bar{\theta}]$ interval.

Differentiating (9A) as an implicit function of m allows us to derive the slope of γ_1:

$$\frac{\mathrm{d}\gamma_1(m)}{\mathrm{d}m} = [\eta(m) - 1]\frac{\theta^2}{m}\frac{c(r+\delta)}{\phi(m)\mu(\theta)(1-\beta)y - \alpha\theta c(r+\delta)} \tag{9C}$$

The sign of this derivative is given by $\eta(m) - 1$ and will be positive when the matching function exhibits globally increasing returns to scale. In that case, $\gamma_1(0) = 0$, and $\gamma_1(m) \to \bar{\theta}$ as $m \to \infty$. If the returns to scale of the matching function were decreasing, i.e $\eta(m) < 1$, $\gamma_1(m)$ would be a decreasing function of m. For the case of constant returns, $\eta(m) = 1$, $\gamma_1(m)$ would be a horizontal line, making labour market tightness independent of m.

Figure 4.5: Existence of equilibria under different matching functions



(a) decreasing returns to scale **(b)** constant returns to scale **(c)** increasing returns to scale with decreasing returns to vacancies

$\gamma_2(m) = \theta$ is defined by equation (4.10), re-arranged as:

$$\mu(\theta) = \frac{\delta m}{\delta P - \phi(m)m} \tag{10A}$$

The left hand side is a positive function of θ, equal to 0 when $\theta = 0$. The right hand side has a lower bound limit at 0 as $m \to 0$; as for its upper bound, there must be some positive $m = \bar{m}$ for which $\gamma_2 = \bar{\theta}$. The slope of $\gamma_2(m)$ is always positive:

$$\frac{\mathrm{d}\gamma_2(m)}{\mathrm{d}m} = \frac{\theta\delta + \mu(\theta)\eta(m)\phi(m)\theta}{\alpha m\delta} \tag{10B}$$

so $\gamma_2(m)$ is upward-sloping in the θ, m space, with $\gamma_2(0) = 0$ and $\gamma_2(\bar{m}) = \bar{\theta}$.

The relative position of γ_1, γ_2 in the (θ, m) space for different types of returns to scale in the matching function is plotted in Figure 4.5. When the matching function exhibits decreasing returns to scale, $\gamma_1(m)$ will be downward-sloping, and there will be a unique equilibrium at $m > 0$. A similar scenario occurs under constant returns to scale, where $\gamma_1(m)$ is constant as labour market tightness is independent of market size. When the returns to market size are increasing, as in the third panel of Figure 4.5, multiple equilibria can occur. There is an equilibrium at $(0, 0)$; and under certain conditions, there will be at least one equilibrium such that $m > 0$.

The existence condition for an equilibrium at a positive level of market activity is $\eta(0)\alpha < 1$. We know that, in general, $\gamma_1(\bar{m}) < \gamma_2(\bar{m}) = \bar{\theta}$; hence, for an intersection point to exist, $\frac{m}{\gamma_1}\frac{\mathrm{d}\gamma_1(m)}{\mathrm{d}m} \Big/ \frac{m}{\gamma_2}\frac{\mathrm{d}\gamma_1(m)}{\mathrm{d}m} < 1$ for some m close to 0. (The curves in panel (c) of Figure 4.5 give one such example.) The expression for the ratio of elasticities can be derived from the expressions for the slope of the two functions, γ_1' and γ_2'. The

limit of this ratio as $m \to 0$ is $\frac{\alpha(\eta(0)-1)}{1-\alpha}$. This will be smaller than 1 if $\alpha\eta(0) < 1$, i.e. there are decreasing returns to matching as vacancies increase. If, on the contrary, matching became more efficient with more vacancies, $\gamma_1(m)$ would be flatter than $\gamma_2(m)$, and the two curves would never cross. Multiple equilibria at $m > 0$ would arise if the matching function exhibited local, as opposed to global, returns to scale.

$$\left[\frac{m}{\gamma_1}\frac{\mathrm{d}\gamma_1(m)}{\mathrm{d}m}\Big/\frac{m}{\gamma_2}\frac{\mathrm{d}\gamma_1(m)}{\mathrm{d}m}\right] = [\eta(m)-1]\frac{\theta c(r+\delta)}{\phi(m)\mu(\theta)(1-\beta)y - \alpha\theta c(r+\delta)}\frac{\alpha\delta}{\delta+\mu(\theta)\eta(m)\phi(m)}$$

$$\lim_{m\to 0}\left[\frac{m}{\gamma_1}\frac{\mathrm{d}\gamma_1(m)}{\mathrm{d}m}\Big/\frac{m}{\gamma_2}\frac{\mathrm{d}\gamma_1(m)}{\mathrm{d}m}\right] = \frac{\eta(0)-1}{1-\alpha}\alpha = \frac{\alpha\eta(0)-\alpha}{1-\alpha}$$

$$\lim_{m\to 0}\left[\frac{m}{\gamma_1}\frac{\mathrm{d}\gamma_1(m)}{\mathrm{d}m}\Big/\frac{m}{\gamma_2}\frac{\mathrm{d}\gamma_1(m)}{\mathrm{d}m}\right] < 1 \;\; if \;\; \alpha\eta(0) < 1$$

$\alpha\eta(0) < 1$ is also a sufficient condition for the uniqueness of the equilibrium. Given the values of $\gamma_1(m)$ and $\gamma_2(m)$ at $\bar{m}$, the single crossing property will be satisfied if

$$\frac{\mathrm{d}(\gamma_1-\gamma_2)}{\mathrm{d}m} < 0$$

$$\begin{aligned}
\gamma_1' - \gamma_2' &\overset{\text{sgn}}{=} \frac{(\eta-1)(r+\delta)}{\phi(m)\mu(\theta)\beta+(1-\alpha)(r+\delta)} - \frac{\delta+\mu(\theta)\eta\phi(m)}{\alpha\delta} \\
&\overset{\text{sgn}}{=} \alpha\delta(\eta-1)(r+\delta) - [\phi(m)\mu(\theta)\beta+(1-\alpha)(r+\delta)][\delta+\mu(\theta)\eta\phi(m)] \\
&\overset{\text{sgn}}{=} (r+\delta)\delta[\alpha\eta(m)-1] - (1-\alpha)(r+\delta)\mu(\theta)\eta(m)\phi(m) - \phi(m)\mu(\theta)\beta[\delta+\mu(\theta)\eta(m)\phi(m)] \\
&< 0 \text{ if } \alpha\eta(m)-1 < 0 \;\forall m
\end{aligned}$$

□

4.3.3 Comparative statics: the role of regional population size

One of the implications of the existence of a unique equilibrium at $m > 0$ is that there is a one-to-one mapping between regional population level and local labour market conditions. As Proposition 2 states below, both labour market size and tightness are increasing in P: regions with a larger population will also have a larger labour market, which under increasing returns to matching makes the matching technology

more efficient. Smaller labour market frictions encourage vacancy creation, which in turn increases labour market tightness, with positive effects for the workers: a higher probability of finding a job (both thanks to a larger and tighter labour market) will translate into higher wages and a lower unemployment rate.

Proposition 2. *If $\eta(0) > 1$ and $\alpha\eta(m) < 1 \ \forall \ m$, labour market tightness θ and size m are increasing in regional population P.*

Proof. In the absence of a closed-form solution for equilibrium θ and m, I instead implicitly differentiate equilibrium equations (4.9) and (4.10).

The implicit derivative with respect to P of the worker-employer behaviour relationship, equation (4.9), is:

$$c(r+\delta)\left[\frac{\mathrm{d}\theta}{\mathrm{d}P}m + \frac{\mathrm{d}m}{\mathrm{d}P}\theta\right] = [\beta + (1-\beta)y]\left[\mu(\theta)\frac{\mathrm{d}\Phi}{\mathrm{d}m}\frac{\mathrm{d}m}{\mathrm{d}P} + \Phi(m)\frac{\mathrm{d}\theta}{\mathrm{d}\theta}\frac{\mathrm{d}\theta}{\mathrm{d}P}\right] \tag{9D}$$

$$\frac{\mathrm{d}m}{\mathrm{d}P}\Big/\frac{\mathrm{d}\theta}{\mathrm{d}P} = \frac{1-\alpha}{\eta-1}\frac{m}{\theta} \tag{9E}$$

$\frac{\mathrm{d}m}{\mathrm{d}P}/\frac{\mathrm{d}\theta}{\mathrm{d}P}$ will be positive, i.e the two derivatives will be of equal sign, if $\eta > 1$.

The implicit derivative with respect to P of the equilibrium employment-unemployment flows, equation (4.10), is:

$$\delta\mu(\theta) + \delta P\frac{\mathrm{d}\mu}{\mathrm{d}\theta}\frac{\mathrm{d}\theta}{\mathrm{d}P} - \delta\frac{\mathrm{d}m}{\mathrm{d}P} = \frac{\mathrm{d}\Phi}{\mathrm{d}m}\frac{\mathrm{d}m}{\mathrm{d}P}\mu(\theta) + \Phi(m)\frac{\mathrm{d}\mu}{\mathrm{d}\theta}\frac{\mathrm{d}\theta}{\mathrm{d}P} \tag{10C}$$

$$\delta\mu(\theta) = \frac{\mathrm{d}m}{\mathrm{d}P}\frac{\mu(\theta)}{m}[\delta P + (\eta-1)\Phi] - \frac{\mathrm{d}\theta}{\mathrm{d}P}\alpha\delta\frac{m}{\theta} \tag{10D}$$

Combining these two equations to solve for $\frac{\mathrm{d}m}{\mathrm{d}P}$:

$$\delta\mu(\theta) = \frac{\mathrm{d}m}{\mathrm{d}P}\left[\frac{\mu(\theta)}{m}(\delta P + (\eta-1)\Phi(m)) - \frac{\eta-1}{1-\alpha}\alpha\delta\right] \tag{4.11}$$

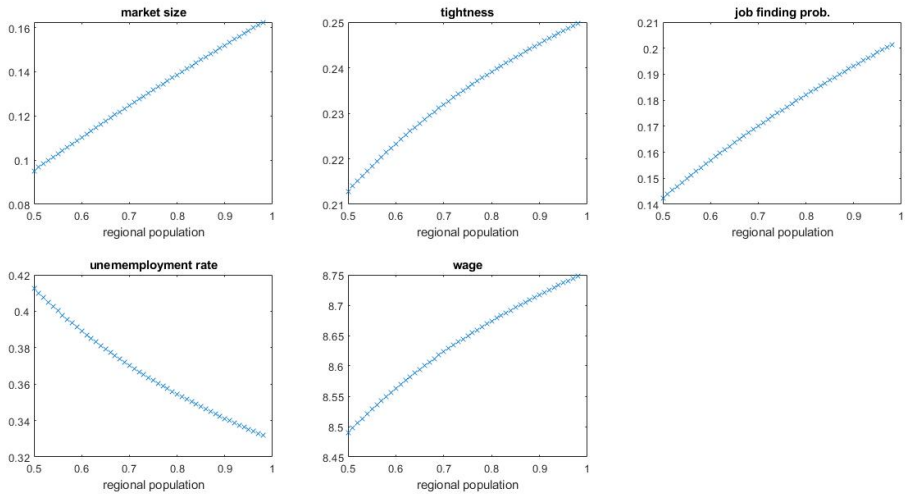
$\frac{\mathrm{d}m}{\mathrm{d}P}$ is positive if $1 > \alpha\eta$:

$$\frac{\mu(\theta)}{m}(\delta P + (\eta-1)\Phi(m)) - \frac{\eta-1}{1-\alpha}\alpha\delta > 0$$

$$\text{if } \ \delta m(1-\alpha\eta) > -\mu(\theta)\Phi(m)(1-\alpha)\eta \ \ \Rightarrow \ \ 1 > \alpha\eta$$

This result shows that under decreasing returns to vacancy creation, $\frac{\mathrm{d}m}{\mathrm{d}P}$ will be positive. The earlier result showed that under increasing returns to scale to the matching function, $\frac{\mathrm{d}m}{\mathrm{d}P} and \frac{\mathrm{d}\theta}{\mathrm{d}P}$ will be of equal signs. The necessary and sufficient conditions for a positive impact of population on regional labour market size and tightness are hence $\eta > 1$ and $1 > \alpha\eta$. □

Figure 4.6: Local labour market equilibria as regional population increases



A numerical simulation of a regional labour market as population exogenously increases. The labour market equilibrium is re-calculated for each value of regional population level and then plotted against its value. Top row, left to right: market size m, labour market tightness θ, job finding probability. Bottom row, left to right: unemployment rate, wage.

4.3.3.1 A numerical simulation

To illustrate how the labour market changes with the level of regional population, I run a numerical simulation of a local labour market. As its population increases, labour market size, tightness, and a host of other labour market variables change as a result. I pick a simple matching function $M = (\sqrt{uv})^{3/2}$. The parameter values are chosen following a calibration exercise in Ellison et al. (2014)(see Table 4.2 in the next section), but the example here is purely illustrative.

The result of this exercise is plotted in Figure 4.6. As population exogenously increases, the labour market expands in size, and the labour market tightness grows. The more efficient matching technology results in a lower unemployment rate, a higher job-finding rate, and higher wages. These patterns mimic the relationship between regional population and labour market conditions observed empirically: Figure 4.2 shows that regions with relatively larger population enjoy higher income and a lower unemployment rate than the relatively smaller regions. Moreover, the model also gives rise to a wage curve relationship: wages are higher in the region with the lower unemployment rate. In this case, this relationship is causal.

4.4 Local labour markets with migration

In section 4.3, I presented a model of local labour markets with increasing returns to scale in matching. In this section, I extend the basic model by allowing the regional population to adjust endogenously to the differences in utility across regions. There will be two different types of migration: workers can make an individual decision to move into another region to search for employment (*speculative migration*), or accept a job in another region that has been advertised in her local labour market (*migration for a specific job*). Unlike the former, migration for a specific job is linked to the behaviour of firms. Both types of migration will depend not only on migration costs, but also on the level of labour market frictions within and across local labour markets. The combination of these factors gives rise to multiple equilibria with regional differences in wages and unemployment rates, despite the fact that the regions are a priori identical.

I start with the model description, which explains how the framework introduced in section 4.3 changes when I allow for migration. In sections 4.4.2 and 4.4.3, I describe the conditions under which speculative migration and migration for a specific job are optimal. In section 4.4.4, I define the model's steady-state equilibria, and describe their location and properties. Finally, I sketch the nature of the regional adjustment between the steady states, although the full description of the dynamic process is beyond the scope of this chapter.

The analysis presented in this section focuses on steady-state equilibria. A steady-state equilibrium is one where regional populations are constant, i.e. there is no migration. Despite the fact that migration is an out-of-equilibrium mechanism, understanding the conditions under which it happens is essential for pinning down the location of steady states. It also allows me to qualitatively describe the adjustment mechanism between the multiple steady-state equilibria.

4.4.1 The model

4.4.1.1 The matching technology, employment and unemployment

As in the model in section 4.3, workers and firms meet, as governed by some matching function. They bargain over how to split the surplus of the job, and then remain in a match until it is exogenously destroyed. The matching technology exhibits increasing returns to scale, which is why the size of the labour market matters.

The environment of the labour market is static. All equilibria are steady states, so endogenous variables, such as the values of unemployment and vacancies, and regional population levels, will also be constant in equilibrium. Throughout the analysis, I assume the economy is in a steady state equilibrium. I explore the out-of-equilibrium states at the end of this section.

The economy has two regions, A and B. They have the same labour market parameters and matching technology. The only driver of labour market differences is their population levels. I denote A the weakly larger region, and B as the region with a weakly smaller population. I normalise the total population of the economy at 1.

Workers can freely move between the two regions. Because they are utility maximisers, they will move to the region that offers higher utility, and their utility function depends only on the labour market features of the given region. However, migration is costly. First, every time a worker relocates, she has to pay a one-off fixed cost K, which represents the financial costs of moving, as well as the utility cost of having to arrange new housing, losing one's social network, etc. I assume that the migration cost exceeds the difference between the value of employment and unemployment in any given region. This ensures that regions are sufficiently "separated" by geography. If the gain from employment was greater than the migration cost, workers may find it optimal to move into another region for employment even if the labour market conditions in the two regions were identical. The following condition, that the migration cost is greater than the worker's share (β) of the maximum match surplus, prevents such an outcome.

$$K > \frac{\beta y}{r + \delta} \tag{4.12}$$

Second, an unemployed migrant searching in the new local labour market has to pay a search cost b every period until she is matched with a firm. This reflects the fact that job search requires local labour market knowledge which takes time to acquire, and that migrants may be at a disadvantage in a new market compared to locals. The model is memory-less in that the worker's migration history does not matter for their current migration status. In general, a worker's home region is the one she currently resides in; an unemployed migrant becomes "local" when she meets her first potential employer.

Firms are free to enter either regional labour market without any barriers. Moreover, they can decide to settle in one region, but hire workers from another one (*cross-regional hiring*), taking advantage of the fact that the two regions form a single country. If they do decide to hire cross-regionally, the hired worker will move for the specific job into the region of the firm's residence. Once the contract is exogenously terminated, she will become a local in this region. The cost of vacancy creation is the same in both regions.

Each local labour market thus consists of all the unemployed workers currently residing there, u_A and u_B for regions A and B respectively, and all the vacancies posted there, v^A and v^B. The option of cross-regional hiring means that not all the vacancies posted in a given labour market will be for local jobs. To distinguish between local and cross-regional vacancies, I use the subscript to denote where the job is located, and superscript to record the region of posting. v_B^A hence denotes vacancies advertised in A for jobs in B, and v_A^A are local vacancies in A. By construction:

$$v^A = v_A^A + v_B^A \tag{4.13}$$

$$v^B = v_B^B + v_A^B \tag{4.14}$$

The matching technology does not distinguish between local and cross-regional vacancies; they enter the matching function additively:

$$M^A = \Phi(m(u_A, v_A^A + v_B^A)) \tag{4.15}$$

$$M^B = \Phi(m(u_B, v_B^B + v_A^B)) \tag{4.16}$$

The probability that a cross-regional vacancy is filled is equal to the probability that any vacancy is filled in the given labour market. Similarly, the probability that a worker is matched with a cross-regional vacancy depends on the ratio of matches to unemployed, and on the number of cross-regional vacancies advertised.

$$\text{prob. of filling a cross-regional vacancy} = \frac{M_B^A}{v_B^A} = \frac{M^A}{v_B^A}\frac{v_B^A}{v^A} = \frac{M^A}{v^A} \tag{4.17}$$

$$\text{prob. of receiving a cross-regional offer} = \frac{M_B^A}{u_A} = \frac{M^A}{u_A}\frac{v_B^A}{v^A} \equiv \lambda_B^A \tag{4.18}$$

Given the size and tightness of the regional labour markets, the firms decide in which region to create a vacancy, and where to advertise it. Workers are in their home market where they search if unemployed, and choose optimally between local and cross-regional job offers. If they receive no offers - or no offer is good enough - workers may decide to move speculatively into another region. δ of the existing matches are exogenously destroyed: these workers become unemployed in their current region of residence. These actions determine the regional population size, labour market size and tightness for both regions at every point in time.

4.4.1.2 The match surplus and surplus sharing

I assume that the economy is in an equilibrium, so the values of employment and unemployment, and of filled and unfilled vacancy, are constant over time.

The matched firms and workers bargain over how to split their joint output y. Workers and firms are homogenous, as are all the match products, but the surplus of a match will depend on where is the job located, and where was the vacancy advertised. In a two-region economy with cross-regional hiring, there will be four different types of match surplus: local match in regions A and B, and cross-regional match in either region. The surplus will be shared between the worker and the employer proportionally to their bargaining power β and $1-\beta$. Because the values of unemployment and vacancies (and of jobs and employment) are constant in an equilibrium, the surplus itself, as well as the way it is shared, will not change over the duration of the match.

The expression for the surplus of *local* matches S_A^L, S_B^L is equivalent to that of a match under autarky:

$$S_A^L = \frac{y - rU_A^L}{r + \delta} \tag{4.19}$$

$$S_B^L = \frac{y - rU_B^L}{r + \delta} \tag{4.20}$$

where U_A^L, U_B^L denote the value of unemployment for the local unemployed in region A and B, respectively.

The surplus of a *cross-regional* match has to reflect not only the forgone value of unemployment, but also the migration cost that a migrant worker has to pay in order to take up the job, and the change in utility that comes from living in a different region. The expression for the surplus of a cross-regional job located in region A and B respectively, S_A^M, S_B^M, is:

$$S_A^M = \frac{y - rU_A^M + \delta(U_A^L - U_A^M)}{r + \delta} - K \tag{4.21}$$

$$S_B^M = \frac{y - rU_B^M + \delta(U_B^L - U_B^M)}{r + \delta} - K \tag{4.22}$$

The relevant outside option (the forgone value of unemployment) is now the value of being unemployed as a migrant in the destination region, denoted as U_A^M, U_B^M for regions A and B, respectively. Even though the cross-regional hiring takes place in the worker's home region, once she makes the move, her outside option is that of a migrant worker[3], and that is what is used in the bargaining process. However, when fired, the worker becomes a local in the new labour market, and her outside option increases to U_A^L (U_B^L). This change in utility is captured by the term $\delta(U_A^L - U_A^M)$ in the surplus expression. The final term is the one-off fixed cost of moving K. It enters the surplus because the worker has to pay it in order to accept the cross-regional job offer.

[3]This is in line with the assumption that searching in a local labour market requires some (local) experience. Given that a migrant worker matched with her employer in her home market, she would have to pay the search cost b were she quit her cross-regional job voluntarily. After the match is dissolved exogenously, however, the worker becomes a local, perhaps because she has established local networks that make search easier. This assumption is necessary to prevent employers from discriminating between locals and migrants for a long time after the move.

The worker's outside options, the return to unemployment, depend on the worker's share of the expected discounted surplus for each of the four matches. They are given by the four equations below. For example, the value of being an unemployed migrant in region A (U_A^M) equals the probability of receiving a local job offer (λ_A^A) times the extra utility from such a job, plus the probability of receiving a cross-regional job offer (λ_B^A) multiplied by the surplus of that job. Employment of either type also includes a change in status into a local the next time the worker is unemployed, which is captured by the $U_i^L - U_i^M$ terms. The last term of the equation captures the fact that if the worker does not receive any job offers (or if none of them offer a positive surplus), she can always choose to relocate into region B and search for employment there. The equation also reflects the extra cost of search in a new region b.

$$\begin{aligned} rU_A^L &= \lambda_A^A \beta \frac{y - rU_A^L}{r+\delta} + \lambda_B^A \beta \frac{y - rU_B^M + \delta(U_B^L - U_B^M) - (r+\delta)K}{r+\delta} \\ &\quad + [1 - \lambda_A^A - \lambda_B^A]\max[0, U_B^M - K - U_A^L] \end{aligned} \tag{4.23}$$

$$\begin{aligned} rU_A^M &= \lambda_A^A \beta \frac{y - rU_A^L}{r+\delta} + \lambda_B^A \beta \frac{y - rU_B^M + \delta(U_B^L - U_B^M) - (r+\delta)K}{r+\delta} - b \\ &\quad + \lambda_A^A(U_A^L - U_A^M) + \lambda_B^A(U_B^L - U_B^M) + [1 - \lambda_A^A - \lambda_B^A]\max[0, U_B^M - K - U_A^M] \end{aligned} \tag{4.24}$$

$$\begin{aligned} rU_B^L &= \lambda_B^B \beta \frac{y - rU_B^L}{r+\delta} + \lambda_A^B \beta \frac{y - rU_A^M + \delta(U_A^L - U_A^M) - (r+\delta)K}{r+\delta} \\ &\quad + [1 - \lambda_B^B - \lambda_A^B]\max[0, U_A^M - K - U_B^L] \end{aligned} \tag{4.25}$$

$$\begin{aligned} rU_B^M &= \lambda_B^B \beta \frac{y - rU_B^L}{r+\delta} + \lambda_A^B \beta \frac{y - rU_A^M + \delta(U_A^L - U_A^M) - (r+\delta)K}{r+\delta} - b \\ &\quad + \lambda_B^B(U_B^L - U_B^M) + \lambda_A^B(U_A^L - U_A^M) + [1 - \lambda_B^B - \lambda_A^B]\max[0, U_A^M - K - U_B^M] \end{aligned} \tag{4.26}$$

The vacancy creation conditions for the four types of vacancies compare the cost of posting a vacancy against the firm's share of the expected discounted surplus. The four conditions below describe, respectively, local hiring in region A, hiring in

region B for a job in A, local hiring in B, and hiring in region A for a job in B.

$$c \geq \frac{M^A}{v^A}(1-\beta)\frac{y - rU_A^L}{r+\delta} \tag{4.27}$$

$$c \geq \frac{M^B}{v^B}(1-\beta)\frac{y - rU_A^M + \delta(U_A^L - U_A^M) - (r+\delta)K}{r+\delta} \tag{4.28}$$

$$c \geq \frac{M^B}{v^B}(1-\beta)\frac{y - rU_B^L}{r+\delta} \tag{4.29}$$

$$c \geq \frac{M^A}{v^A}(1-\beta)\frac{y - rU_B^M + \delta(U_B^L - U_B^M) - (r+\delta)K}{r+\delta} \tag{4.30}$$

The free entry condition implies that the cost will weakly exceed the benefit. The given vacancy is created if this condition holds with equality.

These two sets of four equations describe the behaviour of the workers and the firms on the labour market. For an alternative derivation using Bellman equations and Nash bargaining, see Appendix C.3.1.

In a steady-state equilibrium, not all the four types of matches may exist at the same time. In fact, my definition of steady-state equilibrium precludes any equilibrium migration. However, analysing migration behaviour is necessary to pin down the existence and location of steady-state equilibria. I do this in the next two sections.

4.4.2 Speculative migration

Speculative migration refers to migration of unemployed workers. It is optimal if the increase in utility associated with the relocation offsets the cost of moving, which consists of the fixed cost of moving and the extra cost of search in a new local labour market. In my setting, the region with the higher level of regional population has a more efficient labour market, which implies higher utility (Proposition 2). As a result, speculative migration is always only optimal from the smaller (lower utility) region into the larger region. The optimality conditions are summarised in the proposition below.

Proposition 3. *Speculative migration is optimal if the value of being an unemployed migrant, less the migration cost, exceeds the value of being an unemployed local in the worker's region of origin. Denote the weakly smaller region B, and the weakly*

larger region A. It is then never optimal for unemployed workers to move from region A to region B. Speculative migration from region B to A will be optimal if

$$U_A^M - K > U_B^L \tag{M1}$$

Proof. First, if A is the region with weakly larger population, Proposition 2 implies that labour market size and tightness, and hence the value of local unemployment, will be larger in region A:

$$U_A^L \geq U_B^L \tag{4.31}$$

Second, the existence of extra search costs for migrants, b, implies that

$$U_A^L > U_A^M \tag{4.32}$$

$$U_B^L > U_B^M \tag{4.33}$$

For speculative migration from A to B to be optimal, it would have to be true that $U_B^M - K > U_A^L$. However, this is impossible given equations (4.31) - (4.33).

$$U_A^L \geq U_B^L > U_B^M \;\; \Rightarrow \;\; U_B^M - K \not> U_A^L$$

Migration from B to A is optimal if $U_A^M - K > U_B^L$, which may hold when A is the larger region. □

4.4.3 Migration for a specific job

Migration for a specific job is the result of cross-regional hiring: a firm in one region hires a worker residing in the other region, so the worker has to move in order to take up the position. This means that migration for a specific job entails both a change of location and a change of employment status for the worker. The migration for a specific job is optimal when firms find cross-regional hiring profitable, and when workers are willing to accept such jobs. As in the case of speculative migration, migration for a specific job is not optimal in all equilibria, and only occurs in one direction, from the smaller to the larger region.

Intuitively, workers never find it optimal to move from region A for a specific job in region B. Region B is relatively smaller, so when such a cross-regional match

inevitably ends, the worker would find herself in a worse position than if she stayed unemployed in her home region A. Knowing this, the worker would only accept such a match if the wage was greater than in her local region A. However, paying more would make cross-regional hiring unprofitable for the employer, who could simply relocate to A and pay the local wage.

On the other hand, cross-regional hiring in the opposite direction, from region B to region A, can be profitable because the change in utility is positive. The worker is willing to bear the migration cost because by moving to A her value of unemployment increases, while the firm is willing to hire in B because the greater cost of hiring in a smaller, less efficient region is more than offset by the cheaper wages there.

Proposition 4. *Migration for a specific job is optimal if cross-regional vacancy creation is profitable for the firm, and if such a match maximises the worker's utility. Denote the weakly smaller region B, and the weakly larger region A. It is then never optimal for workers to move from region A for a specific job in region B. Specific migration from region B to A will be optimal if the following two conditions are satisfied.*

$$rU_A^L - rU_B^L < (r+\delta)(U_A^L - U_A^M - K) \tag{M2}$$

$$\beta S_A^M > U_B^L - U_A^M \tag{M3}$$

Proof. This proof proceeds in three steps. First, looking at the conditions for the profitability of posting a vacancy, I show that cross-regional hiring cannot happen in both directions. In the second step, I examine workers' utility under different types of cross-regional hiring to show that if migration for a specific job is optimal, it happens in one direction, from the smaller (B) to the larger (A) region, because workers' utility cannot be maximised by moving into the smaller region. In the third step, I demonstrate that local and cross-regional hiring in region B in general cannot coexist, and identify the point at which the latter becomes more profitable.

Step 1 In order for a given type of vacancy to be created, its vacancy creation condition (expressions (4.27) - (4.30)) must hold with equality. I am interested in equilibria with non-zero labour market activity, i.e. vacancy creation is always profitable in region A, and may be profitable in region B. This means that the vacancy creation condition for local hiring in A holds with equality:

$$c = \frac{M^A}{v^A}(1-\beta)\frac{y - rU_A^L}{r+\delta} \tag{4.27}$$

Cross-regional hiring in either region is then optimal if and only if this hiring can deliver the same marginal profit as local hiring in A.

Cross-regional hiring in B is optimal if vacancy creation condition (4.28) holds. Intuitively, jobs located in region A must generate the same expected profit regardless of where they are advertised.

$$c = \frac{M^A}{v^A}(1-\beta)\frac{y - rU_A^L}{r+\delta} = \frac{M^B}{v^B}(1-\beta)\frac{y - rU_A^M + \delta(U_A^L - U_A^M) - (r+\delta)K}{r+\delta} \tag{4.34}$$

$$\therefore\ (y - rU_A^L)\left[\frac{M^A}{v^A} - \frac{M^B}{v^B}\right] = \frac{M^B}{v^B}(r+\delta)(U_A^L - U_A^M - K) \tag{4.35}$$

The relative size of regions A and B implies that $\frac{M^A}{v^A} - \frac{M^B}{v^B}$ is weakly positive, i.e. vacancies are filled faster in the region with the larger labour market. For this equation to hold, the right-hand side of the equation also needs to be positive. The necessary condition for that is $U_A^L - U_A^M > K$.

The condition for cross-regional hiring in A can be derived analogously:

$$c = \frac{M^A}{v^A}(1-\beta)\frac{y - rU_A^L}{r+\delta} = \frac{M^A}{v^A}(1-\beta)\frac{y - rU_B^M + \delta(U_B^L - U_B^M) - (r+\delta)K}{r+\delta} \tag{4.36}$$

where the left hand side is the local vacancy creation condition in region A (4.27), and the right hand side is cross-regional hiring in A (4.30). Because the vacancy-filling probability in region A is the same for local and cross-regional postings, it implies that the surplus of both vacancies has to be the same:

$$y - rU_A^L = y - rU_B^M + \delta(U_B^L - U_B^M) - (r+\delta)K \tag{4.37}$$

$$\therefore\ r(U_A^L - U_B^L) = -(r+\delta)[U_B^L - U_B^M - K] \tag{4.38}$$

$U_A^L > U_B^L$, so the necessary condition for the optimality of cross-regional hiring is $U_B^L - U_B^M < K$. This is in contradiction with the necessary condition for cross-regional hiring in B, which requires $U_B^L - U_B^M > K$. As a consequence, cross-regional hiring cannot be optimal in both directions.

Intuitively, condition (4.35) states that cross-regional hiring in B is only optimal if the saving from avoiding the extra cost of search ($U_A^L - U_A^M$) makes up for the migration cost K and lower probability of filling a vacancy in region B, $\frac{M^A}{v^A} - \frac{M^B}{v^B}$. The opposite needs to happen for cross-regional hiring in A: migration costs have to exceed the extra cost of search, so that the match surpluses for local and cross-regional hiring in A are equalised. If $U_B^L - U_B^M - K > 0$, the surplus of cross-regional hiring in A would be too large compared to local hiring in A, leading to the perverse outcome where there is no local hiring in the region with the more efficient local labour market. For cross-regional hiring in A, then, $U_B^L - U_B^M - K$ must be negative.

Step 2 Migration for a specific job is optimal if it also maximises worker's utility. Usually, this condition would be satisfied automatically whenever the surplus of the match is positive. However, this is no longer the case for cross-regional hiring, because the worker's outside option changes when she accepts the offer. In the case of cross-regional hiring in region B, for example, the wage of a cross-regional match is based on U_A^M, the value of unemployment if the worker wanted to exit from the match voluntarily. However, her outside option at the time of deciding whether to accept the offer is U_B^L. As a result, a cross-regional job offer, less the obligatory migration cost, may not always exceed the value of staying unemployed.

For cross-regional hiring in B, this implies the following condition:

$$\beta S_A^M > U_B^L - U_A^M \tag{M3}$$

A cross-regional job offer will be accepted if the worker's share of the match surplus, βS_A^M, compensates for the difference in outside options[4]. A sufficient condition for this to hold is $U_B^L < U_A^M$. This is true when region A is sufficiently larger than

[4]Alternatively, this condition can be written as $W_A^M - K > U_B^L$, i.e. the value of employment in region A as a migrant, less the migration cost, has to exceed the local value of unemployment.

region B[5]. The utility gain from moving to the larger region then offsets the loss of utility from paying the migration cost and having to bargain as an unemployed migrant.

For the same reason, migration for a specific job into region B is never optimal. It would require the following inequality to hold:

$$\beta S_B^M > U_A^L - U_B^M \tag{4.39}$$

To see why this condition can never be satisfied, first plug in the expression for S_B^L and re-arrange:

$$\beta \frac{y - rU_B^M + \delta(U_B^L - U_B^M)}{r+\delta} - \beta K > U_A^L - U_B^M \tag{4.40}$$

$$\beta \frac{y - rU_B^L + (r+\delta)(U_B^L - U_B^M)}{r+\delta} - \beta K > U_A^L - U_B^M \tag{4.41}$$

$$\beta \frac{y - rU_B^L}{r+\delta} + \beta[U_B^L - U_B^M - K] > U_A^L - U_B^M \tag{4.42}$$

$$\beta S_B^L > (1-\beta)[U_B^L - U_B^M - K] + U_A^L - U_B^L + K \tag{4.43}$$

Given condition (4.12), the left hand side is smaller than K.

$$K > \beta S_B^L > (1-\beta)[U_B^L - U_B^M - K] + U_A^L - U_B^L + K \tag{4.44}$$

$$K > (1-\beta)[U_B^L - U_B^M - K] + U_A^L - U_B^L + K \tag{4.45}$$

$$-(U_A^L - U_B^L) > (1-\beta)[U_B^L - U_B^M - K] \tag{4.46}$$

The difference between value of unemployment in region A and B has to be smaller in magnitude than the extra search costs relative to migration cost. However, this is at odds with the condition for optimality of cross-regional hiring in A, equation (4.38). Plugging in the expression for $-(U_A^L - U_B^L)$ derived in this equation:

$$\therefore \quad -(U_A^L - U_B^L) = \frac{r+\delta}{r}[U_B^L - U_B^M - K] > (1-\beta)[U_B^L - U_B^M - K] \tag{4.47}$$

Because under (4.38) $U_B^L - U_B^M - K$ is negative, the above inequality would only hold if $-\beta > \frac{\delta}{r}$, which is impossible.

[5]However, condition (M3) may hold even when $U_B^L \geq U_A^M$, depending on the size of the match surplus.

Step 3 The vacancy creation conditions for hiring in region B, (4.28) and (4.29), make it clear that in general only one type of hiring will be optimal at any given time. The probability of filing a vacancy in region B is the same regardless of the type of vacancy (local or cross-regional), so unless the match surpluses are exactly equal to each other, there will be one type that delivers higher expected profits.

The two type of vacancies will be posted simultaneously if

$$\frac{M^B}{v^B}(1-\beta)\frac{y - rU_A^M + \delta(U_A^L - U_A^M) - (r+\delta)K}{r+\delta} = \frac{M^B}{v^B}(1-\beta)\frac{y - rU_B^L}{r+\delta} \quad (4.48)$$

$$rU_A^L - rU_B^L = (r+\delta)(U_A^L - U_A^M - K) \quad (4.49)$$

When $rU_A^L - rU_B^L > (r+\delta)(U_A^L - U_A^M - K)$, local hiring is more profitable. The condition for the profitablity of cross-regional hiring in region B is thus:

$$rU_A^L - rU_B^L < (r+\delta)(U_A^L - U_A^M - K) \quad \text{(M2)}$$

□

4.4.4 Steady-state equilibria with endogenous population

A regional labour market with endogenous population is described by the local labour market size and tightness, $m_A, m_B, \theta_A, \theta_B$, and regional population levels P_A, P_B. In an equilibrium, the unemployment and vacancy creation conditions (equations 4.23) – (4.30) must hold, alongside a set of conditions for the flow of employment and unemployment within and across regions, which will be given below.

I define a steady-state equilibrium as one in which regional population levels P_A, P_B, and hence the labour market variables $m_A, m_B, \theta_A, \theta_B$, are constant. Constant population levels are only possible when neither type of migration is optimal. This means that the steady-state equilibria when population is endogenous can be described by a set of equations similar to the autarkic regional equilibria in section 4.3. There is only local hiring in both regions, and no workers move speculatively.

The relevant values of unemployment and vacancy creation conditions thus are:

$$rU_A^L = \phi_A \mu(\theta_A) \beta \frac{y - rU_A^L}{r + \delta} \tag{4.23*}$$

$$rU_B^L = \phi_B \mu(\theta_B) \beta \frac{y - rU_B^L}{r + \delta} \tag{4.25*}$$

$$c = \frac{\phi_A \mu(\theta_A)}{\theta_A} (1 - \beta) \frac{y - rU_A^L}{r + \delta} \tag{4.27*}$$

$$c = \frac{\phi_B \mu(\theta_B)}{\theta_B} (1 - \beta) \frac{y - rU_B^L}{r + \delta} \tag{4.29*}$$

The description of equilibrium is completed by a set of equations which pin down the flow of employment and unemployment within regions, and specify that the migration flow across regions are zero. Equations (4.50) and (4.51) below capture how, within each region, the stock of workers whose matches are exogenously destroyed is offset by new local matches. The other two equations set population change over time in each region, $\dot{P}_A, \dot{P}_B$, to 0 because there is no cross-regional hiring ($\lambda_A^B = 0$) nor is it optimal to move speculatively ($U_A^M - U_B^L - K \leq 0$).

$$\frac{m_A}{\mu(\theta_A)} = P_A - \frac{\phi(m_A) m_A}{\delta} \tag{4.50}$$

$$\frac{m_B}{\mu(\theta_B)} = P_B - \frac{\phi(m_B) m_B}{\delta} \tag{4.51}$$

$$\dot{P}_A = \lambda_A^B \frac{m_B}{\mu(\theta_B)} + \mathbb{1}[U_A^M - U_B^L - K > 0] \frac{m_B}{\mu(\theta_B)} = 0 \tag{4.52}$$

$$\dot{P}_B = -\lambda_A^B \frac{m_B}{\mu(\theta_B)} - \mathbb{1}[U_A^M - U_B^L - K > 0] \frac{m_B}{\mu(\theta_B)} = 0 \tag{4.53}$$

A steady-state equilibrium exists if there is a combination of $m_A, m_B, \theta_A, \theta_B, P_A, P_B$ for which neither type of migration is optimal. There is no speculative migration if condition (M1) does not hold. There is no migration for a specific job if at least one of the conditions (M2), (M3) do not hold. These three conditions thus identify the set of $m_A, m_B, \theta_A, \theta_B, P_A, P_B$ at which steady-state equilibrium may exist. The condition for existence is summarised in Proposition 5.

Proposition 5. *Define a steady-state equilibrium as one where P_A, P_B are constant. Then a steady-state equilibrium exists if there is a set of endogenous variables*

$m_A, m_B, \theta_A, \theta_B, P_A, P_B$ such that the migration condition (M1) does not hold, and at least one of (M2), (M3) does not hold. The set of steady-state equilibrium variables is then:

$$\left[\frac{c\theta_B}{(1-\beta)\phi_B m_B} - \frac{c\theta_A}{(1-\beta)\phi_A m_A} < \frac{(1-\beta)b}{r+\phi_A\mu(\theta_A)} + \beta K\right) \cup$$
$$\left[\left[\frac{c\theta_B}{(1-\beta)\phi_B m_B} - \frac{c\theta_A}{(1-\beta)\phi_A m_A} > \frac{b}{r+\phi_A\mu(\theta_A)} - K\right] \cap\right.$$
$$\left.\left[\frac{c\theta_B}{(1-\beta)\phi_B m_B} - \frac{c\theta_A}{(1-\beta)\phi_A m_A} < \frac{rb}{r+\phi_A\mu(\theta_A))(r+\delta)} + \frac{rK}{r+\delta}\right]\right]$$

$m_A, m_B, \theta_A, \theta_B, P_A, P_B$ must satisfy the equilibrium values of unemployment (4.23) and (4.25*), vacancy creation conditions (4.27*) and (4.29*), and flow equations (4.50) – (4.53) .*

4.4.5 The adjustment process

When the conditions for a steady-state equilibrium described in Proposition 5 do not hold, the economy is going to be in a state of adjustment. Workers move from the smaller region B to the larger region A, adjusting regional population levels until the economy reaches another steady-state equilibrium. This may include a "black-hole" equilibrium in which all population and economic activity is concentrated in a single region (A).

There are two types of the adjustment process: via migration for a specific job, and via speculative migration. Which of these the economy finds itself at any given point in time depends on which of the three migration optimality conditions hold. Three binary conditions give eight possible combinations, including steady-state equilibria. They are summarised in Table 4.1.

If the adjustment happens via migration for a specific job, the economy of region A resembles that of an autarkic labour market, in that there is local hiring and no workers move out. The only difference is that there is also a steady inflow of workers from region B via the cross-regional hiring conducted there by firms located in A. As I show in Proposition 4, this is the only type of hiring in region B (apart from a knife-edge case in which local hiring occurs simultaneously).

Table 4.1: Migration conditions and types of equilibria

M1	**M2**	**M3**	**type**
does not hold	does not hold	does not hold	steady-state equilibrium
does not hold	holds	does not hold	steady-state equilibrium
does not hold	does not hold	holds	steady-state equilibrium
does not hold	holds	holds	adjustment via migration for a specific job
holds	does not hold	does not hold	adjustment via speculative migration
holds	holds	does not hold	adjustment via speculative migration
holds	does not hold	holds	adjustment via speculative migration
holds	holds	holds	adjustment via speculative migration

If the differences in the utility of the two regions are sufficiently large to make speculative migration optimal, all adjustment happens via this channel. The structure of regional economy A continues to resemble that of an autarkic region, but with a steady inflow of speculative migrants from region B who increase regional population level P_A. In region B, meanwhile, all search and hiring has ceased, because moving to A offers higher utility. All the workers who find themselves unemployed move away.

The adjustment described above continues until the relevant migration conditions change and the economy reaches a steady-state equilibrium. To understand the locations of the steady-state equilibria, and of the adjustment process, I look at how the migration conditions change with the endogenous variables. Conditions (M3) and (M1) are more likely to hold as regional divergence progresses: the increasing gap between regional utility, and the diminishing present value of the extra search costs together reinforce the optimality of either type of migration to the workers. The opposite is true for condition (M2): as regional divergence increases, it becomes less likely to hold. Cross-regional hiring becomes less profitable for firms over time because it requires the present value of the avoided extra search cost to exceed the gap between regional utilities. This intuition is summarised

in the three derivatives with respect to P_A below.

$$\frac{\partial M1}{\partial P_A} > 0 \tag{4.54}$$

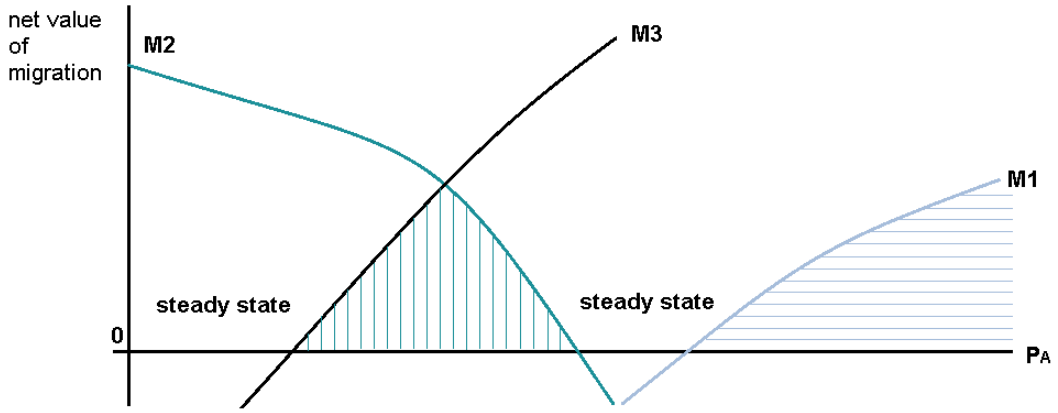
$$\frac{\partial M2}{\partial P_A} < 0 \tag{4.55}$$

$$\frac{\partial M3}{\partial P_A} > 0 \tag{4.56}$$

For the full derivation, see Appendices C.3.2 and C.3.3.

I combine these properties of the adjustment process in Figure 4.7. I sketch the three migration conditions as P_A increases; a positive value of the given migration condition means the condition holds. The combinations of the three curves create four intervals on the P_A line. When regional population levels are relatively similar (P_A is not much larger than a half of the total population), the economy is in a steady state equilibrium, because migration for a specific job (M3) is not (yet) optimal. The diagram shows that there may be multiple values of P_A for which this is the case, i.e. there may be multiple steady-state equilibria. When divergence increases, regional utility differences become sufficiently large to incentivise migration for a specific job. The economy adjusts in this way until all profits of cross-regional hiring are exhausted (M2 turns negative and ceases to hold). Because speculative migration is not attractive yet (M1 is negative), the economy enters another interval of steady states. The second adjustment, via speculative migration, comes when the degree of regional divergence is high. The extreme steady-state equilibrium after this adjustment is one where all workers have moved to the larger region A. The location of the different equilibria, and the implications of the dynamic process for regional divergence, are further explored in section 4.5.

Figure 4.7: Migration conditions and the location of steady-state equilibria



The diagram shows the net value of each migration condition as a function of the regional population in A. When a given line is above the x-axis, the condition holds. Pale blue: condition (M1) for speculative migration. Teal: condition (M2) for the profitability of cross-regional hiring. Black: condition (M3) for the optimality of migration for a specific job for the workers. Vertical shading: adjustment via migration for a specific job. Horizontal shading: adjustment via speculative migration.

4.5 Comparative statics: labour market frictions and regional disparities

In this section, I simulate the model of regional labour markets with endogenous population to analyse how regional divergence and migration depend on migration costs and labour market frictions. I also discuss what this model can tell us about the welfare implications of regional divergence and of large and persistent regional disparities.

4.5.1 Comparative statics: the role of search frictions and migration costs in regional divergence

The comparative statics presented below are the result of a simulation of the model with endogenous population (section 4.4). The parameter values used in the simulation are summarised in Table 4.2. I picked these values to loosely match monthly magnitudes in the real world, but they are not calibrated to correspond to a particular economy or a particular set of outcomes.

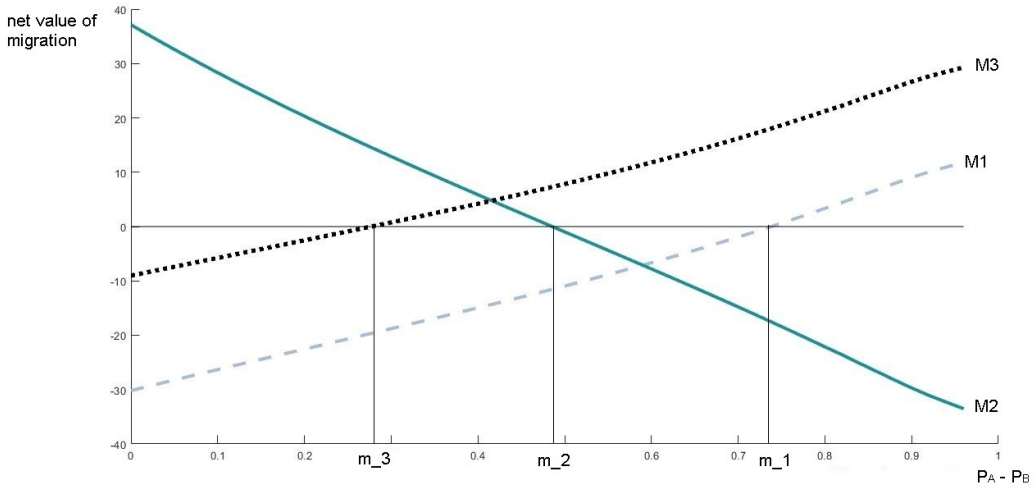
Table 4.2: Parameter values in simulations

parameter		value
y	output per period	1
r	discount rate	0.01
β	worker's bargaining power	0.5
δ	job-destruction probability	0.05
c	per-period cost of posting a vacancy	7
b	per-period cost of search in a new labour market	1
K	migration cost	12

An economy with this set of values can find itself in steady-state equilibria or different types of adjustment process, and thus with different levels of migration. To explore this, I plot the simulated migration conditions (M1) - (M3) as a function of regional population levels in Figure 4.8. It is a simulated version of the diagram in Figure 4.7. I denote m_1, m_2 and m_3 as the level of regional population differences at which the net value of the respective migration condition becomes positive or negative. These intersection points identify the location of steady-state equilibria and of the adjustment process, and determine the magnitude of the maximum potential migration of either type.

For the given parameter values, Figure 4.8 shows that when the difference in the regional populations is small, the regions are in a low-divergence steady state. Even though cross-regional hiring is profitable at low level of regional differences (M2), this type of matches only becomes optimal for workers when M3 turns positive, at $m_3 = 0.28$. When regional population differences fall between 0.28 and 0.49, workers move for specific jobs from region B to region A, and the economy is in a process of adjustment where regional divergence increases. The adjustment process ends when the economy enters another interval of steady-state equilibria. At $m_2 = 0.49$, cross-regional hiring ceases to be profitable for employers, because the lower wages in region B no longer outweigh the thinning labour market and the associated hiring costs. The economy gets "stuck" at a high-divergence steady state: despite the poor labour market conditions in region B, firms do not take advantage of the low wages, and workers do not emigrate. The economy is in a steady state because M1, the condition for the optimality of speculative migration,

Figure 4.8: Migration conditions and the location of equilibrium types



The diagram shows the net value of each migration condition as a function of the difference in regional populations. When a given line is above the x-axis, the condition holds. Pale blue, dashed: condition (M1) for speculative migration. Teal, solid: condition (M2) for the profitability of cross-regional hiring. Black, dotted: condition (M3) for the optimality of migration for a specific job for the workers. The total population of the economy is normalised to 1, so a difference of 1 corresponds to a case with all workers residing in region A.

does not turn positive until $m_1 = 0.74$. When it does turn positive, all unemployed workers in B move to A, worsening the labour market conditions in region B even further, and incentivising further speculative migration. The extreme outcome of this process is a steady state equilibrium with the entire population of the economy residing and working in region A.

This simulation shows that increasing returns to matching in local labour markets can drive regional divergence. First, the positive market size effects translate into a positive feedback cycle of migration and higher utility: the net value of moving from the workers' perspective, M1 and M3, increase with regional population differences. Second, cross-regional hiring can speed up this process by facilitating migration for a specific job when workers would not be willing to move speculatively. This can be seen in the adjustment process via migration for a specific job at relatively low levels of regional divergence.

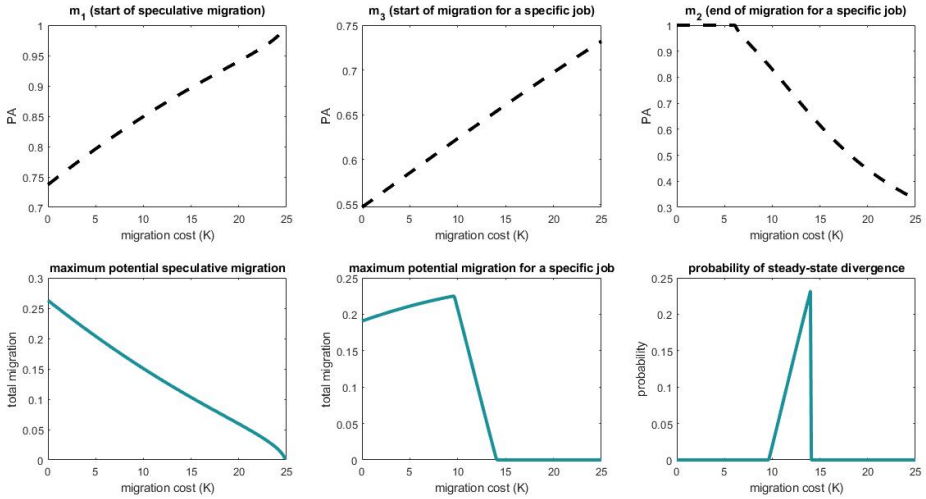
In the rest of this section, I analyse how the shape of regional divergence and the size and type of migration change with migration costs and labour market

frictions. In summary, I find a large variation in how costs and frictions affect the regional economy. Regarding migration costs, I show that the standard result holds: high migration costs reduce migration flows. At the same time, *medium* values of migration costs are associated with the highest probability of an economy getting stuck in a state of large and persistent regional differences. The impact of cross-regional search frictions depends on their type. Larger frictions in the form of the extra search cost in a new labour market (b) makes both cross-regional hiring and persistent regional differences more likely. On the other hand, the same effect can be achieved by *lowering* cross-regional search frictions by increasing matching spillovers between local labour markets, because it makes hiring in the smaller region relatively easier.

I start the analysis by looking at the impact of the one-off cost of migration K (Figure 4.9). Keeping the rest of the parameters constant, I recalculate the migration condition intersection points m_1, m_2, m_3 for different values of K; these are plotted in the top panels in Figure 4.9. m_1 and m_3 are increasing in K: as migration costs increase, the net value of either type of migration becomes positive at higher and higher level of regional differences. Moreover, as the plot for m_2 shows in the third panel, high migration cost K reduce the profitability of cross-regional hiring from the firm's perspective.

The bottom panels of Figure 4.9 summarise what these changes in m_1, m_2, m_3 imply for the size of migration flows. The first two plots show the *maximum potential migration* in a regional economy as K increases. For speculative migration, this equals the distance between 1 (black hole steady state) and m_1 (the point at which speculative migration becomes optimal). The maximum potential migration for a specific job is calculated as the difference between m_2 (at which cross-regional hiring stops being profitable) and m_3 (at which this type of migration becomes optimal for workers). Both types of migration decrease as K increases; in the case of migration for a specific job, when K approaches 15, m_3 overtakes m_2 and there is no level of regional divergence at which it is optimal.

Figure 4.9: Comparative statistics for the migration cost

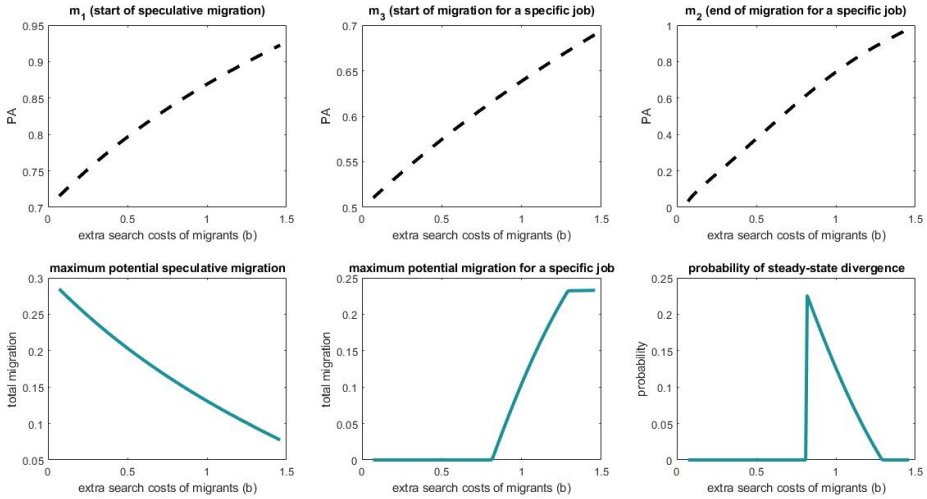


Comparative statistics for migration cost K. Top row: total speculative migration as a function of K, total migration for a specific job as a function of K, and the total length of the interval of regional population levels at which high-divergence steady states exist. Bottom row: the x-axis intersection points for each of the three migration conditions as a function of K. From left to right: M1, M3, M2.

Finally, the third bottom panel in Figure 4.9 looks at the probability of the economy entering a high-divergence steady-state equilibrium. These equilibria are located on the interval of regional population differences bounded by m_2 and m_1, so their probability is given by the fraction of the regional population differences line that corresponds to m_1 - m_2. The comparative statics simulation shows that this probability is higher for medium value of migration costs, where they are neither so large as to pose barriers to speculative migration, but not so small as to make the cost-sharing implicit in cross-regional hiring unattractive to workers.

I repeat the exercise for the extra cost of searching in a new labout market, b. I find that it also lowers the total speculative migration, but it has the opposite effect on migration for a specific job (Figure 4.10). In general, b reduces migration flows because it acts as a de facto another migration cost. However, unlike K, b must only be paid when moving speculatively. As a result, speculative migration declines, but specific–job migration increases in b; this can be also seen in the upward profile of m_2, the point in regional divergence at which cross-regional hiring stops being

Figure 4.10: Comparative statistics for the extra cost of search for migrants



Comparative statistics for the extra search cost b. Top row: total speculative migration as a function of b, total migration for a specific job as a function of b, and the total length of the interval of regional population levels at which high-divergence steady states exist. Bottom row: the x-axis intersection points for each of the three migration conditions as a function of b. From left to right: M1, M3, M2.

profitable. Intuitively, high b increases specific-job migration because it allows the worker and the firm in the match to avoid this cost, thus increasing the match surplus. Related to this is the interval of high-divergence steady states, which are only supported by medium values of b: it must be neither too low (no cross-regional hiring) nor too high (cross-regional hiring is optimal until speculative migration kicks in, so there is no steady state at high levels of divergence).

The comparative statics for b highlight that increasing search frictions do not necessarily translate into smaller or slower divergence. Extra cost of search may drive higher (specific-job) migration, which increases the differences between regions, and make high-divergence steady-state equilibria more likely, cementing these differences into persistent regional disparities.

Another type of cross-regional search friction is the existence of the regional matching functions themselves. If all vacancies advertised in the economy and all the economy's unemployed matched in a single matching function, the job-finding and vacancy-filling probabilities would be the same everywhere. There would be no

link between regional differences and the efficiency of matching in the given region.

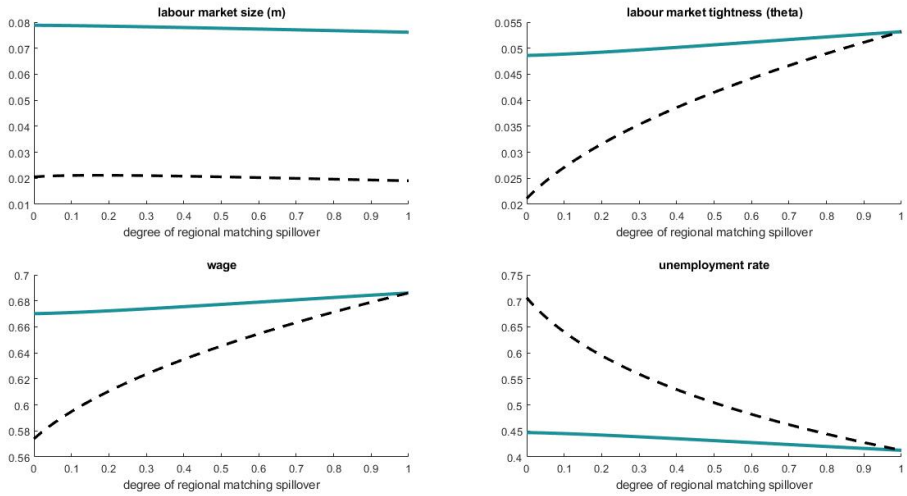
To understand how the "separateness" of regional matching functions affects the local labour markets, I extend the baseline model to encompass a range of matching functions with different degrees of regional matching spillovers, n (for the functional form of the general matching function, see Appendix C.4). n lies between 0 and 1. At 0, there is a separate local labour market matching technology for each region; at $n = 1$, all matching takes place in a single economy-wide technology. In between these two extremes lie labour markets with some degree of matching spillovers: the matching rates in a local labour market depend not only on the local search and vacancy posting, but also on the aggregate activity. Intuitively, some workers may be able to respond to vacancies posted in the other region, and the information about some vacancies may spread beyond the local labour market.

To illustrate the role of n, in Figure 4.11 I plot local labour market equilibria for different values of local matching spillovers n. The economy was simulated at a particular point in regional divergence, with region A relatively better off. However, the four panels in Figure 4.11 show that as the degree of separateness between the local matching functions decreases, the labour market conditions in the two regions converge, even as regional populations stay fixed. At $n = 1$, wages, labour market tightness and unemployment rates are the same across regions. This is despite the fact that the size of the labour market in region B does not increase: higher n means that the local m no longer determines the efficiency of matching in the local economy.

How does the degree of separateness in regional matching functions affect migration and regional divergence? To answer this question, I repeat the comparative statics exercise for n. The bottom panels in Figure 4.12 show that as these search frictions decline, so do both types of migration. In other words, even relatively small improvements in the inter-linkages between regional matching functions will reduce cross-regional hiring. As the third bottom panel shows, this translates into a sharp drop in the probability of a high-divergence steady state once n exceeds about 0.5.

The main driver of this result is the fact that large n reduces the utility gap between regions and hence the incentive to migrate. This utility effect outweighs

Figure 4.11: The effect of matching spillovers on regional disparities

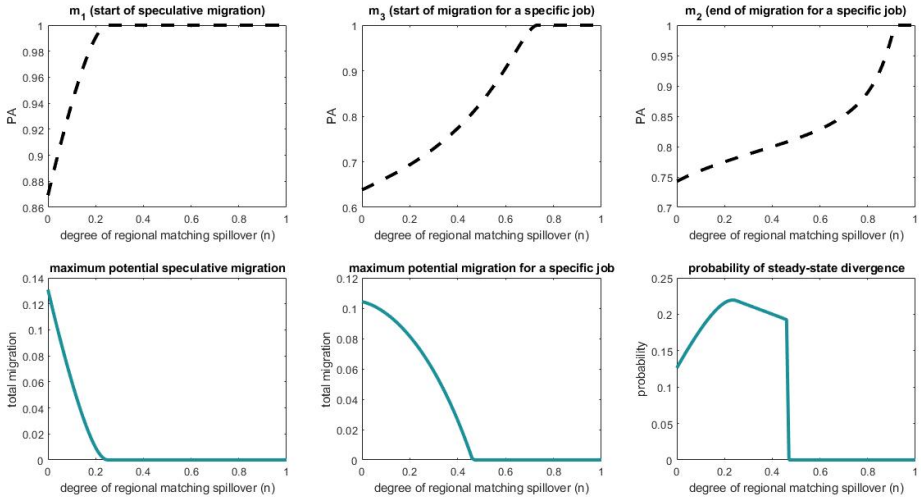


Comparative statics for n with fixed regional populations. The four panels plot equilibrium values of labour market size and tightness, and wages and unemployment rates as matching spillovers increase. Teal (solid): region A. Black (dashed): region B.

the direct impact of the reduction in cross-regional search frictions which would see cross-regional hiring increase as frictions decline. Further evidence of this can be seen in the top panels in Figure 4.12: as n increases, cross-regional hiring becomes more profitable for firms (third panel), but less attractive for workers (second panel). Because the latter reacts more strongly to increases in n, larger regional matching spillovers reduce total migration for a specific job, despite the fact that this type of hiring is becoming easier.

To analyse the direct role of cross-regional search frictions separately, I further extend the model by allowing for productivity differences across regions. In the baseline model in section 4.4, matches in both regions produce constant output y. I augment this to make output an increasing function of the regional population. This is in line with the standard assumption of agglomeration forces in spatial productivity, and ensures that as regional matching technologies become more inter-linked, the difference in regional utilities remains large. This is illustrated in a simulation of such an economy in Figure 4.13. In contrast to Figure 4.11, an increasing n causes convergence in unemployment rates but does not significantly diminish the gap in wages.

Figure 4.12: Comparative statistics for matching function spillovers



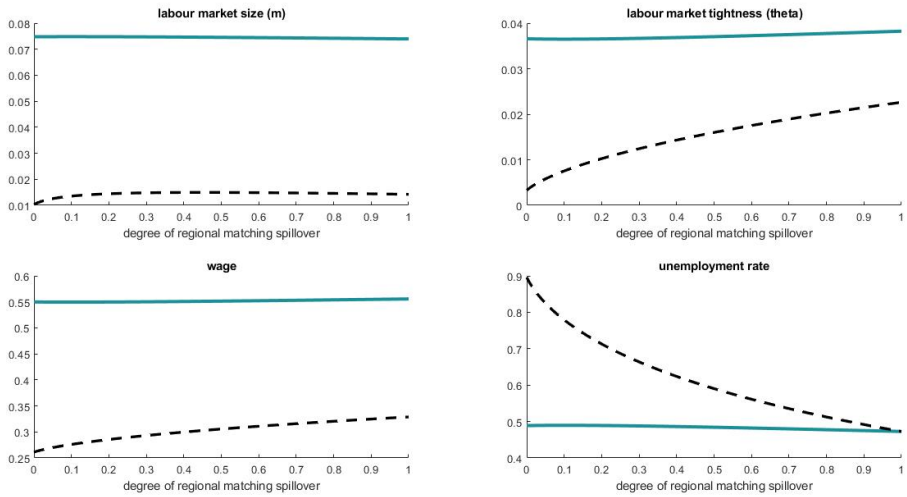
Comparative statistics for matching function spillovers n. Top row: total speculative migration as a function of n, total migration for a specific job as a function of n, and the total length of the interval of regional population levels at which high-divergence steady states exist. Bottom row: the x-axis intersection points for each of the three migration conditions as a function of n. From left to right: M1, M3, M2.

The comparative statics for n in this augmented model are presented in Figure 4.14. The simulations show that an improvement in the inter-linkages between local labour markets encourages migration for a specific job; in other words, reducing cross-regional search frictions does increase cross-regional hiring. It has the opposite impact on speculative migration. As a result, the reduction also speeds up regional divergence and the regions are hence more likely to get stuck in a high-divergence steady state. In general, the augmented model shows that the probability of large and persistent regional differences depends on the relative strength of agglomeration effects in the production and matching functions.

4.5.2 Welfare comparisons

In this model, individual utility is given by the discounted present value of a person's lifetime earnings. It does not depend on the region of residence beyond the local labour market conditions. This, together the with assumption about increasing returns to matching, implies that social welfare is maximised when all economic

Figure 4.13: The effect of matching spillovers on regional disparities when allowing for agglomeration in production



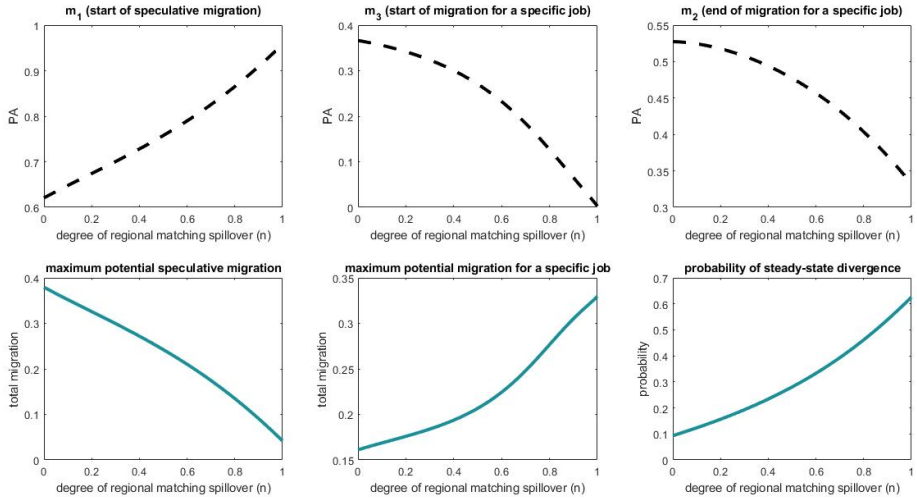
Comparative statics for n with fixed regional populations and with region-specific production functions. Regional match output is a positive and increasing function of regional population. The four panels plot equilibrium values of labour market size and tightness, and wages and unemployment rates as matching spillovers increase. Teal (solid): region A. Black (dashed): region B.

activity concentrates in a single region (a "black hole" equilibrium). This delivers the highest matching efficiency and match surplus, and thus incomes and social welfare.

The condition for this result is that the gap between regional utilities in this equilibrium exceeds the fixed migration cost: $\max[U_A^L - U_B^L] > K$. This assumption is also implicit throughout the chapter. When it holds, the outcome delivering the highest social welfare is one where all workers coordinate and move to region A at the same time. It would ensure that migration is optimal for everybody, and avoids the inefficiency of the process of regional divergence.

An implication of the fact that the "black hole" equilibrium is the preferred one is that regional divergence is desirable. The greater the regional disparity, the closer the economy is to the welfare-maximising equilibrium. A corollary to this statement is that migration is desirable too, because it speeds up regional divergence in the correct direction, from smaller to larger regions. This is in contrast to most standard models in the literature, where migration equilibrates regional differences. In this model, migration equilibrates regional differences in utility; the

Figure 4.14: Comparative statistics for the matching function spillovers when allowing for agglomeration in production



Comparative statistics for matching function spillovers n. Regional match output is a positive and increasing function of regional population. Top row: total speculative migration as a function of n, total migration for a specific job as a function of n, and the total length of the interval of regional population levels at which high-divergence steady states exist. Bottom row: the x-axis intersection points for each of the three migration conditions as a function of n. From left to right: M1, M3, M2.

difference is that this corresponds to a divergence, not convergence, in regional population levels and labour market conditions.

The existence of large and persistent regional differences is then mostly suboptimal because it means that the regional economies are not moving towards the social optimum. It is also undesirable because it increases inequality between the otherwise homogenous workers: those that were lucky and managed to move into the larger region have higher utility than those that did not have the opportunity.

In this light, a policymaker considering a reduction in cross-regional search frictions should weigh the benefits of faster divergence against the increased possibility of getting stuck in a high-divergence steady state (Figure 4.14).

4.6 Conclusion

The aim of this chapter was to establish the matching process within and between regions as one of the potential drivers of the large and persistent regional disparities

observed across most developed economies. I showed that the presence of increasing returns to scale in the local matching technology can generate large regional differences in wages, unemployment rates, and other labour market conditions. These differences then motivate workers to migrate from the smaller, less prosperous regions to the large, more prosperous ones, exacerbating these differences.

Cross-regional matching frictions can play a number of distinct roles in this process. First, the relative difficulty of searching in a new labour market is what makes cross-regional hiring profitable. Because the employers pay some of the migration costs, this type of migration may be optimal at relatively low levels of divergence, when speculative migration isn't. This means that this type of search frictions increases regional differences. Second, migration for a specific job also gives rise to a steady state when regional differences are large, precisely because it drives divergence when speculative migration is not yet optimal. When cross-regional hiring stops being profitable, but the gap in regional utilities is not yet large enough to encourage speculative migration, the economy will get stuck in a high-divergence steady state. The effect of increasing regional matching spillovers, i.e. making the local matching functions less local, is similar, even though it entails a reduction, not an increase, in cross-regional search frictions. A high degree of matching spillovers means that hiring efficiency is not that low in the smaller regions, which encourages cross-regional hiring. But, as in the case of extra search costs, it also makes the persistence of regional differences more likely, by first driving the divergence and then generating a set of steady-state equilibria.

The policy implications of this model depend on whether migration in the direction from the smaller to the larger regions (as is currently observed) constitutes a move towards or away from the social optimum. Under the assumptions of this model, the answer to that question is the former, so any policy that encourages migration of any type is welfare-improving insofar as it brings closer the "black hole" equilibrium.

A more realistic model would likely include some element of location-specific individual preferences and congestion externalities (such as housing), both of which would put the optimality of a "black hole" equilibrium in doubt. Under this

scenario, regional divergence is best avoided, and policy measures should focus on changing the direction of migration to equilibrate regional population levels. Encouraging cross-regional hiring by making local labour markets more intertwined, in combination with some form of migration subsidies, may be beneficial here. If workers are to move from a larger labour market to a smaller one, it would be easier to do so via migration for a specific job rather than speculatively, because the latter calls for the worker to pay the entire migration cost. In general, the principles that describe optimality conditions for both types of migration as well as the various types of equilibria would be still applicable. An extension of this model that merges local labour markets and cross-regional search frictions with congestion externalities and location preferences is a topic for further research.

5
Conclusion

This thesis demonstrates that the functioning of labour markets plays a significant role in within-country mobility. The ability to search for jobs remotely across regions makes migration easier, boosting mobility. The workers that face larger cross-regional search frictions thus move less, even if they may gain more from doing so.

The arguments presented here rely on a mixture of empirical and theoretical work. I use an existing dataset in a new way, which allows me to expand our understanding of what inter-regional migration actually looks like. Drawing on empirical IO literature, I am able to back out unobserved job offers and link their relative abundance to the observed migration patterns of the less and the more educated in the US. I build a structural model of job search in a multi-region environment where speculative migration is possible, to show how the empirically unobserved choice set heterogeneity shapes workers' employment and location. Estimating this model then lends further support to the argument that job search matters for migration, and that the less educated in the US move less partly because they lack cross-regional employment opportunities. Finally, I step back and analyse what the addition of cross-regional hiring tells us about the existence of large and persistent regional disparities characterising most developed economies. I show that, depending on the underlying agglomeration force, the interaction

between job search and migration can be a force for good, or it can serve to further deepen regional inequality.

This thesis enters the long-standing debate about whether the observed migration flows are the consequence of an equilibrium or disequilibrium process in the economy. The mainstream view in the literature is that most migration flows are an equilibrium process, with marginal utility equalised across space, as opposed to being driven by imbalances between regions. While I work within the same framework, my findings suggest that the existing migration flows are to an extent the result of search frictions between regional labour market. Chapter 4 shows that the existence of these frictions means that regional utility is no longer equalised. Were these frictions reduced or removed, workers would move more. This implies that the current flows are suboptimal.

Whether this assertion is correct depends on what these search frictions actually represent. The discrete choice model in Chapter 2 presents some indication of what they may be: the significance of the recruitment proxy variables means that at least some of the frictions originate in employer behaviour. However, whether this is the result of profit maximisation, or also some imperfections in the way the labour market operates – such as information asymmetry and thin labour market effects – is a question for further research.

In general, a better knowledge of the process of vacancy creation and hiring is essential to improve our understanding of local labour markets and migration for a specific job. This is an under-researched field of the literature, but the newly emerging data on online vacancies could potentially remedy that.

The role of communication technology and the internet in hiring and migration is another potentially fruitful area of research. The empirical evidence in this thesis is based on data from the second half of the 1990s when much of the online tools used routinely today were in their infancy. It is thus a priori unclear what the findings of this thesis would look like today. However, the large gap in the representativeness of online vacancy data between jobs for the less and the more educated suggests that much of the effects of frictions on migration and employment still apply.

Finally, this thesis creates the scope, and suggests some of the tools, to study the interaction between job search and migration, employment and location choices, within an economy. The focus on the differences in migration behaviour between education groups, while prevalent in the migration literature, is really just a proxy for a host of labour market characteristics that require further study. The process of cross-regional job search itself is much more nuanced than how it is represented in this thesis. The emerging literature on social networks and employment referrals (Witte (2019)), and the work on worker search behaviour (Belot et al. (2018)), are just some of the possible ways to study this further. The aim of this thesis was to show that further work on job search and migration could open potentially important channels in addressing some of the most pressing economic problems of today.

Appendices

A

Appendix to Chapter 2

A.1 Data

I use the Survey of Income and Program Participation (SIPP) because of its three key features: it is a (i) panel with (ii) monthly data that (iii) follows its respondents when they move. These three features are necessary to analyse the relationship between cross-regional job search and migration. I need panel data in order to observe the outcomes of the worker's job search, and I need information on his geographic movement to be able to study migration. Finally, the data needs to be of sufficiently high frequency to allow me to link the two, i.e. to be able to see the worker's labour market outcomes just after the move. SIPP is, to my best knowledge, the only major dataset that combines these three features (for a detailed comparison of the different data sources, see Hernández-Murillo et al. (2011)). High-level monthly data on employment and wages is often a snapshot of the economy, missing the panel dimension. On the other hand, the available panel data at a similarly high frequency (such as the Current Population Survey) does not track respondents when they move; or, when it does (such as the National Longitudinal Survey of Youth and the Panel Study of Income Dynamics), the data is collected at low frequency (annual, bi-annual, or every ten years in the case of census data), making it hard for me make reasonable inference about whether the

worker migrated with a job in hand or not. The SIPP combines all these features at a large enough sample size to allow me to conduct meaningful empirical analysis.

The data collection for SIPP occurs every four months, when the respondents provide monthly retrospective data. There are several consecutive SIPP panels, each lasting between three and four years, so that there is potentially up to four years of monthly data on each participating individual. The first SIPP panel started in 1984, but I focus on the data collected after the major re-design in 1996. There are five such panels: 1996, 2001, 2004, 2008, and 2014 (ongoing).

I use the 1996 panel, which covers data from December 1995 to February 2001, for two reasons. First, it covers a "normal" period with no large downturns or recessions. Second, it follows a large redesign of the survey aiming to improve the data quality and extend the length of the survey, and manages to avoid some of the budget cuts and complications that plagued the later surveys (a 2000 panel had to be cancelled after 8 months due to budget restrictions, while the 2004 panel ran for less than 3 years).

Each SIPP panel contains between 14,000 and 52,000 interviewed households, but like in all panel datasets, this number changes over time due to sample attrition. In the SIPP, individuals may exit the sample by becoming unresponsive, through entering the army or an institution (such as prison), or migrating overseas. At the same time, the panel sample can grow: because the SIPP is based on households, new individuals may enter the survey any time by becoming a member of the respondent household, e.g. via marriage. This results in a degree of flux in and out of the panel sample. Of the total of 115, 996 individuals surveyed in the 1996 panel, a fifth (21%) remained in the survey for a year or less.[1] Just over a half of the sample (52,558 individuals), remained in the dataset for the full 4 years. The descriptive statistics on the subsample used in my analysis are summarised in Table A.1. The dsitribution of people's choices is given in Table A.2.

Data on regional house prices, rents, and unemployment rates are drawn from supplementary datasets. The rent data is taken from the 5% sample of the 2000

[1]This may be due to inflow of new respondents as well as sample attrition.

Table A.1: Descriptive statistics for the sample of working men between the age of 25 and 50

	dropout	high school	college	total
age	38.60	38.37	39.26	38.65
	(7.427)	(7.221)	(7.475)	(7.327)
white	84.6%	86%	89.3 %	86.8 %
	(0.361)	(0.347)	(0.309)	(0.339)
employed	97.3%	98.4%	99.3%	98.5%
	(0.161)	(0.125)	(0.0833)	(0.120)
monthly wage ($)	1829.7	2677.1	4703.2	3156.8
	(1207.1)	(1987.5)	(4422.3)	(3019.0)
total household income ($)	3282.4	4550.4	7074.0	5124.1
	(2319.8)	(3060.3)	(5449.4)	(4043.1)
urban population	79.1%	77.8%	88.1%	80.9%
	(0.406)	(0.415)	(0.324)	(0.393)
migration propensity (state)	0.0499%	0.117%	0.286%	0.158%
	(0.0223)	(0.0343)	(0.0534)	(0.0397)
migration propensity (region)	0.0268%	0.0553%	0.156%	0.0806%
	(0.0164)	(0.0235)	(0.0394)	(0.0284)
total	2,093	10,431	4,196	16,720

Summary descriptive statistics for the sample used in the empirical sections of this paper. The values are averaged over the whole duration of the survey, 1996-1999. Migration propensities are calculated monthly.

census. I calculate median rent by state and then aggregate it up to the census regions used in the model. I also include data on quarterly data on median house price sales by region, for two reasons. First, rent and house price data correlate with both living costs and local amenities, which is why it makes sense to include two variables. Second, unlike the rent data, house prices vary over time. The third variable, monthly regional unemployment rate, is taken from the Bureau of Labor Statics.

The full list of variables and their sources is summarised in Table A.3 below.

Table A.2: Observed choices

option	selected	%
employment, Northeast	1948	16.36
unemployment, Northeast	139	1.17
employment, Midwest	2616	21.98
unemployment, Midwest	227	1.91
employment, South	3846	32.31
unemployment, South	339	2.85
employment, West	2492	20.93
unemployment, West	297	2.49
employment, home	10612	89
unemployment, home	968	8.1
employment, away	292	2.5
unemployment, away	53	0.5
total	11925	100

Summary of individual's choices. The upper panel presents distribution of choices from the 8 options in the logit model. The lower panel aggregates these options based on employment/unemployment and migration.

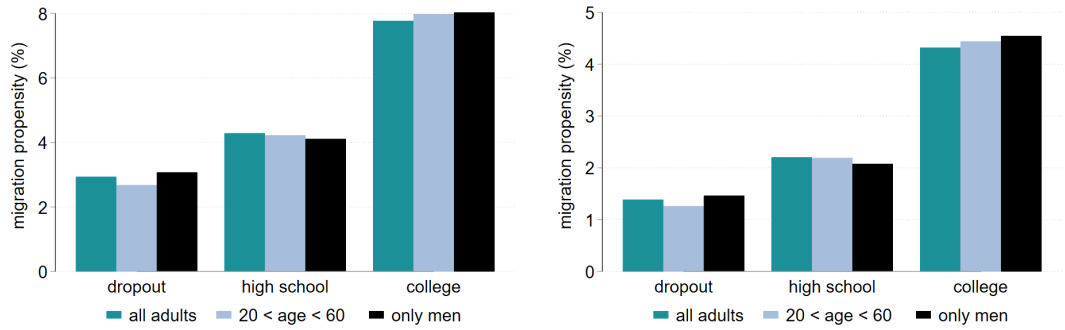
Table A.3: Definitions and sources of variables.

variable	definition	source
migration	dummy equal to 1 if person changes region of residence	SIPP
speculative migration	dummy equal to 1 if person migrates and is unemployed in the following month	SIPP
income	accepted wage offer: observed monthly wage rejected wage offer: imputed average monthly wage offer reservation wage: imputed average monthly lowest wage	SIPP
unemployment	dummy equal to 1 if person is unemployed in a given month	SIPP
unemployment rate, regional	monthly state-level unemployment rate, aggregated to regions, 1996-1999	Bureau of Labor Statistics
unemployment rate, individual	individual-specific probability that the worker is unemployed in a given month. Excludes months of and after migration.	SIPP
regional rent	median rent calculated from the 5% sample of 2000 Census	IPUMS
regional house price	quarterly median house sale price, by census region, 1996-2000	Federal Reserve Bank of St. Louis
migration cost	dummy equal to 1 if person changes region of residence	SIPP
switching cost	dummy equal to 1 if this month's employment or location differs from previous month	SIPP
young	dummy equal to 1 if individual is less than 35 years old	SIPP
kids	dummy equal to 1 if individual is in a family with at least 1 child	SIPP
married	dummy equal to 1 if individual currently married	SIPP
education	highest achieved education	SIPP
online vacancies	log odds ratio of the probability density mass of a given occupation in online postings and full universe	JOLTS, Burning Glass
spatial concentration	Gini index of distribution of sector-specific employment relative to regional population	County Business Patterns
employees per firm	average number of employees per firm, per sector, 1996-1999	County Business Patterns

A.2 Stylised facts about regional migration and education

Figure A.1: Propensity to migrate, US, 1996-1999

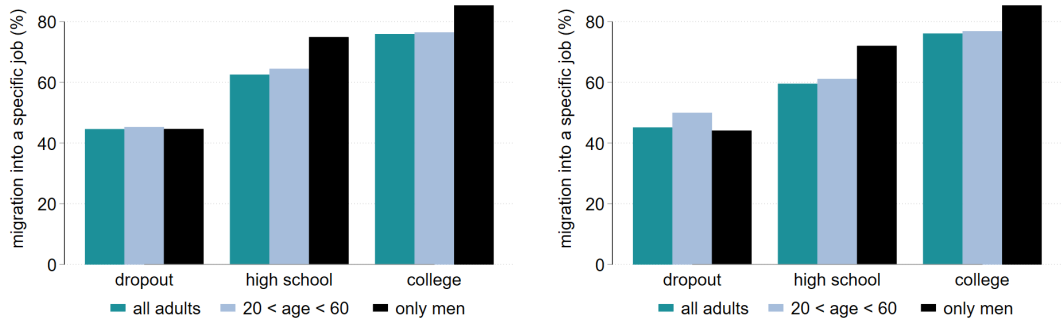
(a) Inter-state migration **(b)** Regional migration



Calculated from the Survey of Income and Program Participation, 1996-1999 panel. Adult is defined as anyone over the age of 18. States refer to the 50 US states. Regions refer to the 4 Census regions (see section 3.2 for definition).

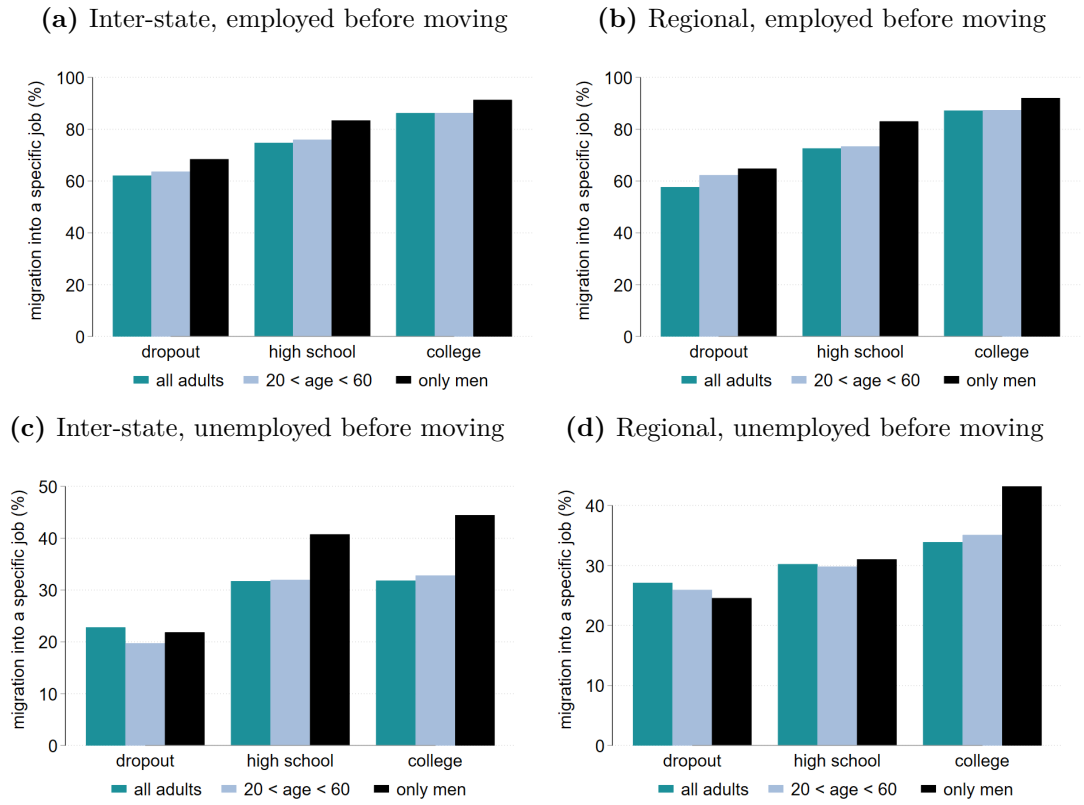
Figure A.2: Probability of moving into a specific job (conditional on moving), US, 1996-1999

(a) Inter-state migration **(b)** Regional migration



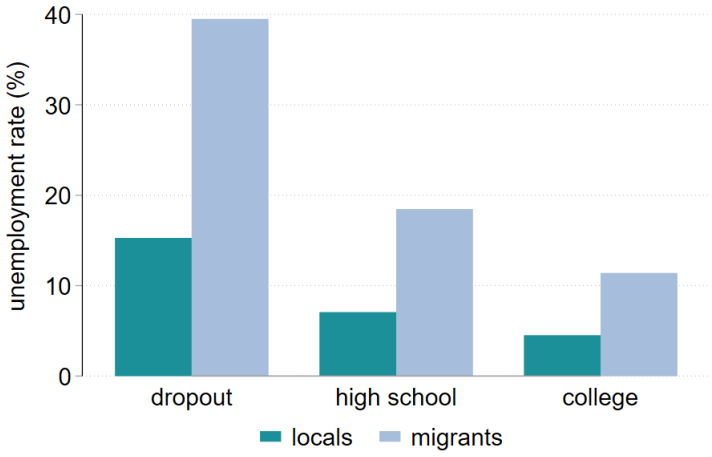
Calculated from the Survey of Income and Program Participation, 1996-1999 panel. Adult is defined as anyone over the age of 18. Migration for a specific job refers to migration followed by employment. States refer to the 50 US states. Regions refer to the 4 Census regions (see section 3.2 for definition).

Figure A.3: Probability of moving into a specific job, conditional on moving and employment state before migration, US, 1996-1999



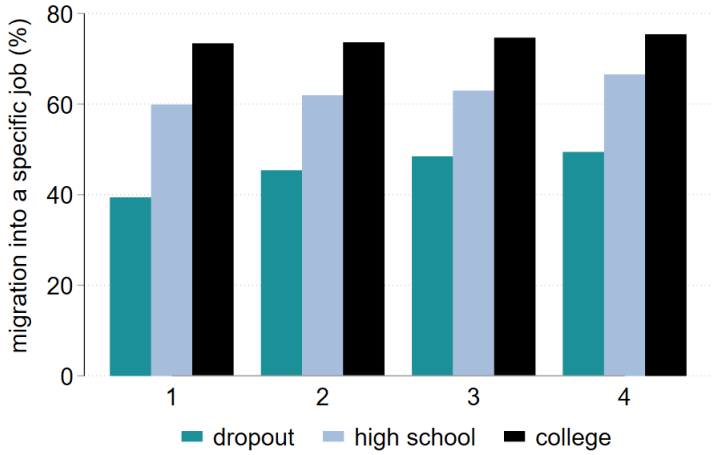
Calculated from the Survey of Income and Program Participation, 1996-1999 panel. Adult is defined as anyone over the age of 18. Migration for a specific job refers to migration followed by employment. States refer to the 50 US states. Regions refer to the 4 Census regions (see section 3.2 for definition).

Figure A.4: Unemployment rates for stayers and migrants, US, 1996-1999



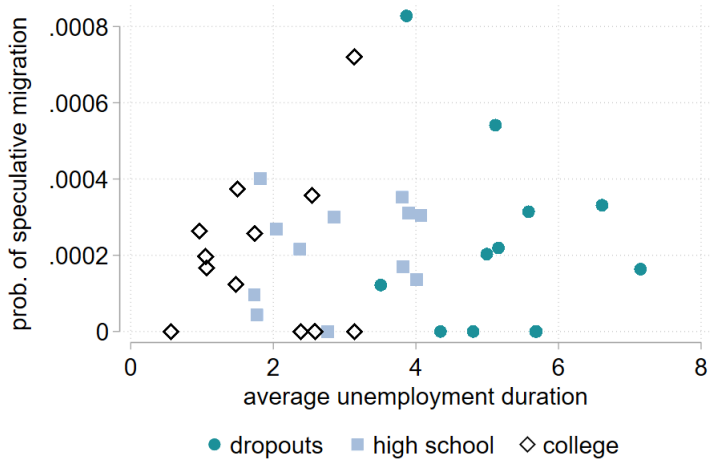
Calculated from the Survey of Income and Program Participation, 1996-1999 panel. Sample: adults (anyone over the age of 18). The bars represent average state-level unemployment rate for the given group.

Figure A.5: Migration for a specific job by education as a function of months after moving, US, 1996-1999



Calculated from the Survey of Income and Program Participation, 1996-1999 panel. Sample: adults (anyone over the age of 18). The bars represent average state-level unemployment rate for the given group.

Figure A.6: Speculative mirgation vs. unemployment duration, by education and sector, US, 1996-1999



Calculated from the Survey of Income and Program Participation, 1996-1999 panel. Sample: adults (anyone over the age of 18). The bars represent average state-level unemployment rate for the given group.

A.3 Post-migration outcomes

Table A.4: Impact of migration on wages, within-person comparison

	monthly wage premium		
	$ (average)	% (average)	fraction that is positive
dropouts	84	19%	73%
high school graduates	292	29%	67%
college graduates	1096	46%	78%

Calculated from the Survey of Income and Program Participation, 1996-1999 panel.
The sample are working males between the age of 25 and 50.

Table A.5: Impact of migration on unemployment

	non-empl. rate (%)		duration of unempl. spells		prob. of spell>6 months (%)	
	stayers	migrants (after move)	stayers	migrants (after move)	stayers	migrants (after move)
dropouts	18	26	5.5	4.2	29	20
high school	9	12	4.8	3.8	24	19
college	4	6	4.5	3.7	21	15

Calculated from the Survey of Income and Program Participation, 1996-1999 panel.
The sample are working males between the age of 25 and 50.

A.4 Imputing wage offers

The mean of the wage offer distribution is smaller than the mean of the observed wage because workers do not accept wages that are not large enough. This is a well-known problem in labour economics that is usually solved following a selection control approach introduced by Heckman (1979). It consists of first parametrically estimating the probability that a given worker would be employed, and then adding this selection correction term into a regression that explains observed wages as a function of worker characteristics. The coefficients from this second-stage regression then allow me to calculate average wage offer for a worker, conditional on his characteristics.

For the purposes of this model, however, I also have to control for the probability of observing wage from a particular region. Because the worker's decision is along

two dimensions (employment and location), I have to explicitly control for both workers choosing highest possible wages, and workers rejecting high wages in favour of living in anther region. One way to solve this problem is to add variables to the Heckman control function that determine workers' migration decision. The disadvantage of this approach is that, by relying on a rich set variables to capture both types of bias, the overall wage regression become difficult to estimate and convergence of results is challenging.

My preferred way of estimating wage offers follows Dahl (2002). In his estimates of state-level education premium, Dahl faces the same multinomial selection bias. His solution is to calculate, for each region, the probability that a local worker becomes employed there, and the probability that a worker moves into employment from elsewhere. He then uses a polynomial of these probabilities as a proxy for the correction function, adding it to the wage regression in the same way as the Heckman correction term.

The results of the Heckman wage regressions are presented in Table A.6; Dahl estimates are summarised in Table A.7. The wage regression, using the given correction function, is estimated separately for each region, using the wage data on workers working and living in the region at the time. I use race as exclusion restriction in the Heckman specification. Overall, the coefficients across the two regressions are quite similar, as is the (unreported) standard deviation of the error term, suggesting that both models can explain similar mount of variation in the observed wages. The Dahl specification is my preferred one because, thanks to the lower demands on estimation, it allows me to add more worker characteristics in the wage regression.

Table A.6: Heckman selection model for monthly regional wages.

	(1) Northeast	(2) Midwest	(3) South	(4) West
wage				
age	280.6687***	278.3296***	219.8698***	263.7825***
	(8.0228)	(5.6201)	(4.9114)	(6.7049)
education	1116.2263***	1021.3983***	1000.2749***	966.8457***
	(23.3704)	(16.8157)	(13.5038)	(18.2518)
_cons	2245.7523***	1741.0617***	2044.4947***	2462.5954***
	(60.8001)	(43.4558)	(36.1614)	(49.1740)
select				
age	0.0319***	0.0228***	0.0257***	0.0052
	(0.0056)	(0.0051)	(0.0041)	(0.0047)
education	0.3023***	0.3506***	0.1887***	0.2267***
	(0.0155)	(0.0138)	(0.0109)	(0.0121)
race	-0.1585***	-0.2641***	-0.1362***	-0.0898***
	(0.0105)	(0.0116)	(0.0089)	(0.0068)
_cons	1.5561***	1.5149***	1.7357***	1.4580***
	(0.0468)	(0.0405)	(0.0347)	(0.0371)
N	78530	108208	144165	95882
ll	-7.152e+05	-9.692e+05	-1.286e+06	-8.572e+05

Standard errors in parentheses

Both stages also include occupation groups and year dummy.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.7: Dahl selection model for monthly regional wages

	(1) Northeast	(2) Midwest	(3) South	(4) West
age	279.8684***	270.4605***	217.3884***	258.5599***
	(8.2174)	(5.7405)	(5.0433)	(6.8946)
education	1091.7691***	961.5562***	996.4935***	948.8347***
	(24.6079)	(17.5926)	(14.0792)	(19.1049)
year	73.9644***	82.4742***	88.6739***	113.9386***
	(12.0316)	(8.3644)	(7.3824)	(10.0632)
home_prob	-2421.5960*	-3933.6999***	-1485.0401*	262.5190
	(1005.0397)	(727.8249)	(668.3174)	(799.1905)
mig_prob	-4.065e+04***	-1.387e+04***	356.7142	-2.473e+04***
	(5571.2857)	(3179.5752)	(3265.3795)	(4907.0853)
home_prob2	10403.0351***	-901.3914	6360.9951***	-746.5879
	(1842.0758)	(1183.6315)	(1235.7193)	(1155.8361)
mig_prob2	38877.6788***	14905.4679***	-968.5881	23634.7809***
	(5710.3835)	(3239.1183)	(3312.2007)	(4903.5391)
cross_prob1	55462.4511***	14841.3148***	6372.0374	32798.8037***
	(6795.9589)	(4093.9859)	(4091.4949)	(6184.1172)
cross_prob2	-7915.8938***	4587.8734***	-4672.2692***	1460.6144
	(1774.5541)	(1099.2049)	(1186.0104)	(1083.5815)
cross_prob3	-7807.7315*	4440.2111	-4814.0601*	-4116.8270
	(3847.9920)	(2300.8156)	(2205.1413)	(3446.3211)
_cons	2231.6960***	2152.4964***	1894.9275***	1540.6217***
	(408.2322)	(292.2394)	(273.3837)	(323.9317)
N	72852	101648	134390	87900
ll	-6.856e+05	-9.368e+05	-1.241e+06	-8.204e+05

Standard errors in parentheses

The regression also includes occupation groups and year dummy.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

B

Appendix to Chapter 3

B.1 The role of initial population distribution

This is the proof of Proposition 1, which states that the equilibrium distribution of regional population does not depend on its initial distribution.

Proof. I will prove Proposition 1 for a 2-region case. Let α denote the population share in region A and $1-\alpha$ the population share in the other region, B. Let subscript E denote their equilibrium values, subscript of 0 the initial values, and t values in period t. As in the rest of the paper, I normalise the total population in the economy to 1.

I start by modeling the migration flows in and out of region A. In period 1, α_1 will be the initial population α_0 less any emigration (to region B), plus immigration (from region B). I use P_E to denote the probability that a worker residing in region A would emigrate to region B, while P_I is the immigration rate from B.

$$\begin{aligned} \alpha_1 &= \alpha_0 - P_E\alpha_0 + P_I(1-\alpha_0) \\ &= \alpha_0(1 - P_E - P_I) + P_I \end{aligned} \tag{B.1}$$

The above is a simplified, stylised version of the zero net migration condition. P_E and P_I depend on local unemployment rates, other population shares, and all model

parameters (through their impact on values of employment and unemployment and thus choice probabilities).

I can extend equation (B.1) to derive a recursive formula for α_t. Following the same logic and substituting in for α_1, the period 2 population share in A will be

$$\begin{aligned}\alpha_2 &= \alpha_1 - P_E\alpha_1 + P_I(1-\alpha_1) \\ &= \alpha_1(1 - P_E - P_I) + P_I \\ &= [\alpha_0(1 - P_E - P_I) + P_I](1 - P_E - P_I) + P_I \\ &= \alpha_0(1 - P_E - P_I)^2 + P_I(1 - P_E - P_I) + P_I\end{aligned}$$

In general, population in A at time t will be:

$$\alpha_t = \alpha_0(1 - P_E - P_I)^t + P_I\sum_{0}^{t-1}(1 - P_E - P_I) \tag{B.2}$$

To find equilibrium population share α_E, I allow migration flows to carry on infinitely, i.e $t \to \infty$. The first term of the expression in (B.2) will converge to 0. Because both P_E and P_I are probabilities, they will lie on the unit interval, and so will the term $(1 - P_E - P_I)$. As a result, the weight on the initial population share α_0 goes to 0 over time:

$$\alpha_0(1 - P_E - P_I)^t \to 0 \text{ as } t \to \infty$$

The second term in equation (B.2) converges to to a positive constant that depends on the relative size of emigration and immigration into the region:

$$P_I\sum_{0}^{t-1}(1 - P_E - P_I) \to \frac{P_I}{P_E + P_I} \text{ as } t \to \infty$$

Putting these two together, the overall expression for $\alpha_{t\to\infty} = \alpha_E$ is a constant independent of the initial population share.

$$\alpha_t \to \frac{P_I}{P_E + P_I} \text{ as } t \to \infty \tag{B.3}$$

I arrive at the same expression if I use an equilibrium condition for regional population share (equation (3.9)). Its stylised version states that the flows out of region A must be balanced by the immigration from region B:

$$\alpha_E P_E = (1 - \alpha_E)P_I$$

Solving for α_E yields the same expression as in equation (B.3).

$$\alpha_E = \frac{P_I}{P_E + P_I} \tag{B.4}$$

Proposition 1 holds if and only if some migration occurs, which directly follows from expressions (B.3) and (B.4). $\alpha_E = \frac{P_I}{P_E+P_I}$ is a real number only if at least one of P_E, P_I is strictly positive. If both of them were equal to 0 (i.e. no migration is optimal), α_t would instead converge to α_0. The zero net migration condition would return "anything goes" result, as there would be no migration for any set of regional population shares. □

B.2 Closed-form solutions for model moments

Employment-employment transition rate between regions

$$
\begin{aligned}
&T_{ee,mk} = \text{job-specific migration from employment in } m \text{ to } k \\
&= \alpha_m(1-u_m)\zeta_1\lambda_m f_k(z)^\tau \left[\begin{array}{c}\Pr(\mathbf{V_k} - K + \gamma_k > \\ \{\mathbf{V_m} + \gamma_m, U_m + \gamma_m, U_k + \gamma_k - K, ..., U_J + \gamma_J - K\})\end{array}\right] h_m(w) \\
&= \alpha_m(1-u_m)\zeta_1\lambda_m \\
&\quad \int\int \frac{exp(V_k(z)-K)}{exp(V_k(z)-K) + exp(V_m(w)) + exp(U_m) + \sum_{j\neq m}^{J} exp(U_j - K)} \mathrm{d}H_m(w)\mathrm{d}F_k(z)
\end{aligned}
$$

Employment-unemployment transition rate between regions

$$
\begin{aligned}
&T_{eu,mk} = \text{speculative migration from employment in } m \text{ to } k \\
&= \alpha_m(1-u_m)\lambda_m f_m(z)^\tau \left[\begin{array}{c}\Pr(U_k + \gamma_k - K > \\ \{\mathbf{V_m} + \gamma_m, \mathbf{V'_m} + \gamma_m, U_m + \gamma_m, ..., U_J + \gamma_J - K\})\end{array}\right] h_m(w) \\
&\quad + \alpha_m(1-u_m)\zeta_1\lambda_m \sum_{j\neq m}^{J} f_j(z)^\tau \left[\begin{array}{c}\Pr(U_k + \gamma_k - K > \\ \{\mathbf{V_m} + \gamma_m, \mathbf{V_j} + \gamma_j - K, U_m + \gamma_m, ..., U_J + \gamma_J - K\})\end{array}\right] h_m(w) \\
&\quad + \alpha_m(1-u_m)(1-\lambda_m - (J-1)\zeta_1\lambda_m) \left[\begin{array}{c}\Pr(U_k + \gamma_k - K > \\ \{\mathbf{V_m} + \gamma_m, U_m + \gamma_m, ..., U_J + \gamma_J - K\})\end{array}\right] h_m(w) \\
&= \alpha_m(1-u_m)\lambda_m \\
&\quad \int\int \frac{exp(U_k - K)}{exp(V_m(z)) + exp(V'_m(z)) + exp(U_m) + \sum_{j\neq m}^{J} exp(U_j - K)} \mathrm{d}H_m(w)\mathrm{d}F_m(z) \\
&\quad + \alpha_m(1-u_m)\zeta_1\lambda_m \sum_{j\neq m}^{J} \\
&\quad \int\int \frac{exp(U_k - K)}{exp(V_m(z)) + exp(V_j(z)-K) + exp(U_m) + \sum_{j\neq m}^{J} exp(U_j - K)} \mathrm{d}H_m(w)\mathrm{d}F_j(z) \\
&\quad + \alpha_m(1-u_m)(1-\lambda_m - (J-1)\zeta_1\lambda_m) \\
&\quad \int\int \frac{exp(U_k - K)}{exp(V_m(z)) + exp(U_m) + \sum_{j\neq m}^{J} exp(U_j - K)} \mathrm{d}H_m(w)
\end{aligned}
$$

Unemployment-employment transition rate between regions

$$
\begin{aligned}
&T_{ue,mk} = \text{job-specific migration from unemployment in } m \text{ to } k \\
&= \alpha_m u_m \zeta_2 \theta_m f_k(z)^\tau [\Pr(\mathbf{V_k} + \gamma_k - K > \{U_m + \gamma_m, ..., U_J + \gamma_J - K\})] \\
&= \alpha_m u_m \zeta_2 \theta_m \int \frac{exp(V_k(z)-K)}{exp(V_k(z)-K) + exp(U_m) + \sum_{j\neq m}^{J} exp(U_j - K)} \mathrm{d}F_k(z)
\end{aligned}
$$

Unemployment-unemployment transition rate between regions

$T_{uu,mk}$ = speculative migration from unemployment in m to k

$$
\begin{aligned}
&= \alpha_m u_m \theta_m f_m(z)^\tau [\Pr(U_k + \gamma_k - K > \{\mathbf{V_m} + \gamma_m, U_m + \gamma_m, ..., U_J + \gamma_J - K\})] \\
&\quad + \alpha_m u_m \zeta_2 \theta_m \sum_{j \neq m}^{J} f_j(z)^\tau [\Pr(U_k + \gamma_k - K > \{\mathbf{V_j} + \gamma_j - K, U_m + \gamma_m, ..., U_J + \gamma_J - K\})] \\
&\quad + \alpha_m u_m (1 - \theta_m - (J-1)\zeta_2 \theta_m) [\Pr(U_k + \gamma_k - K > U_m + \gamma_m, ..., U_J + \gamma_J - K)] \\
&= \alpha_m u_m \theta_m \int \frac{exp(U_k - K)}{exp(U_m) + exp(V_m(z)) + \sum_{j \neq m}^{J} exp(U_j - K)} \mathrm{d}F_m(z) \\
&\quad + \alpha_m u_m \zeta_2 \theta_m \sum_{j \neq m}^{J} \int \frac{exp(U_k - K)}{exp(U_m) + exp(V_j(z) - K) + \sum_{j \neq m}^{J} exp(U_j - K)} \mathrm{d}F_j(z) \\
&\quad + \alpha_m u_m (1 - \theta_m - (J-1)\zeta_2 \theta_m) \frac{exp(U_m)}{exp(U_k - K) + \sum_{j \neq m}^{J} exp(U_j - K)}
\end{aligned}
$$

Employment-employment transition rate within region

$T_{ee,mm}$ = employment-to-employment transitions within m

$$
\begin{aligned}
&= \alpha_m (1 - u_m) \lambda_m f_m(z)^\tau [\Pr(\mathbf{V_m} + \gamma_m \geq \{\mathbf{V'_m} + \gamma_m, U_m + \gamma_m, ..., U_J + \gamma_J - K\})] h_m(w) \\
&\quad + \alpha_m (1 - u_m) \lambda_m f_m(z)^\tau [\Pr(\mathbf{V'_m} + \gamma_m \geq \{\mathbf{V_m} + \gamma_m, U_m + \gamma_m, ..., U_J + \gamma_J - K\})] h_m(w) \\
&\quad + \alpha_m (1 - u_m) \zeta_1 \lambda_m \sum_{j \neq m}^{J} f_j(z)^\tau \left[\begin{array}{c} \Pr(\mathbf{V_m} + \gamma_m \geq \\ \{\mathbf{V_j} + \gamma_j - K, U_m + \gamma_m, ..., U_J + \gamma_J - K\}) \end{array} \right] h_m(w) \\
&\quad + \alpha_m (1 - u_m)(1 - \lambda_m - (J-1)\zeta_1 \lambda_m - \delta_m) \left[\begin{array}{c} \Pr(\mathbf{V_m} + \gamma_m \geq \\ \{U_m + \gamma_m, ..., U_J + \gamma_j - K\}) \end{array} \right] h_m(w) \\
&= \alpha_m (1 - u_m) \lambda_m \\
&\quad \int \int \frac{exp(V'_m(z))}{exp(V'_m(z)) + exp(V_m(w)) + exp(U_m) + \sum_{j \neq m}^{J} exp(U_j - K)} \mathrm{d}H_m(w) \mathrm{d}F_m(z) \\
&\quad + \alpha_m (1 - u_m) \lambda_m \\
&\quad \int \int \frac{exp(V_m(w))}{exp(V'_m(z)) + exp(V_m(w)) + exp(U_m) + \sum_{j \neq m}^{J} exp(U_j - K)} \mathrm{d}H_m(w) \mathrm{d}F_m(z) \\
&\quad + \alpha_m (1 - u_m) \zeta_1 \lambda_m \\
&\quad \sum_{j \neq m}^{J} \int \int \frac{exp(V_m(w))}{exp(V_j(z) - K) + exp(V_m(w)) + exp(U_m) + \sum_{j \neq m}^{J} exp(U_j - K)} \mathrm{d}H_m(w) \mathrm{d}F_j(z) \\
&\quad + \alpha_m (1 - u_m)(1 - \lambda_m - (J-1)\zeta_1 \lambda_m - \delta_m) \\
&\quad \int \frac{exp(V_m(w))}{exp(V_m(w)) + exp(U_m) + \sum_{j \neq m}^{J} exp(U_j - K)} \mathrm{d}H_m(w)
\end{aligned}
$$

Employment-unemployment transition rate within region

$$
\begin{aligned}
&T_{eu,mm} = \text{flows into unemployment within } m \\
&= \alpha_m(1-u_m)\delta_m \Pr(U_m + \gamma_m > \{U_j + \gamma_j - K, ..., U_J + \gamma_J - K\}) \\
&= \alpha_m(1-u_m)\delta_m \frac{exp(U_m)}{exp(U_m) + \sum_{j \neq m}^{J} exp(U_j - K)}
\end{aligned}
$$

Note that this moment does not include voluntary quits when the worker stays within her region. Given the set-up of this model (stationary), this would be equal to 0, since a worker would not accept a job that she did not prefer to unemployment.

Unemployment-employment transition rate within region

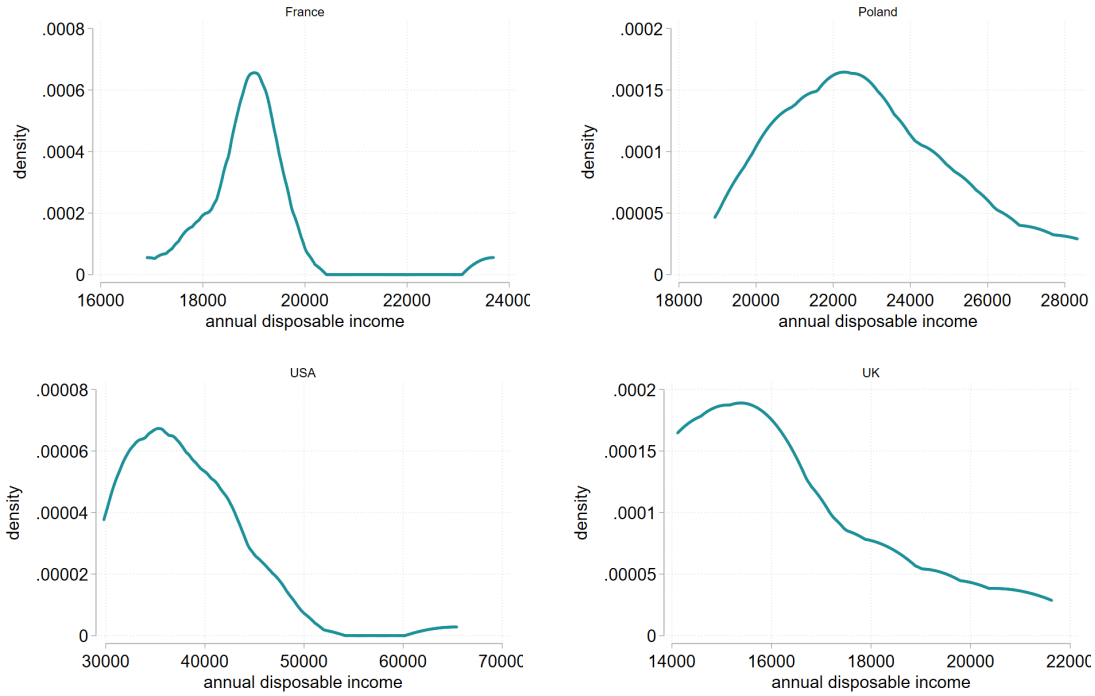
$$
\begin{aligned}
&T_{ue,mm} = \text{flows into employment within } m \\
&= \alpha_m u_m \theta_m f_m(z)^{\tau} [\Pr(\mathbf{V_m} + \gamma_m > \{U_m + \gamma_m, ..., U_J + \gamma_J - K\})] \\
&= \alpha_m u_m \theta_m \int \frac{exp(V_m(z))}{exp(V_m(z)) + exp(U_m) + \sum_{j \neq m}^{J} exp(U_j - K)} \mathrm{d}F_m(z)
\end{aligned}
$$

C
Appendix to Chapter 4

C.1 Motivation

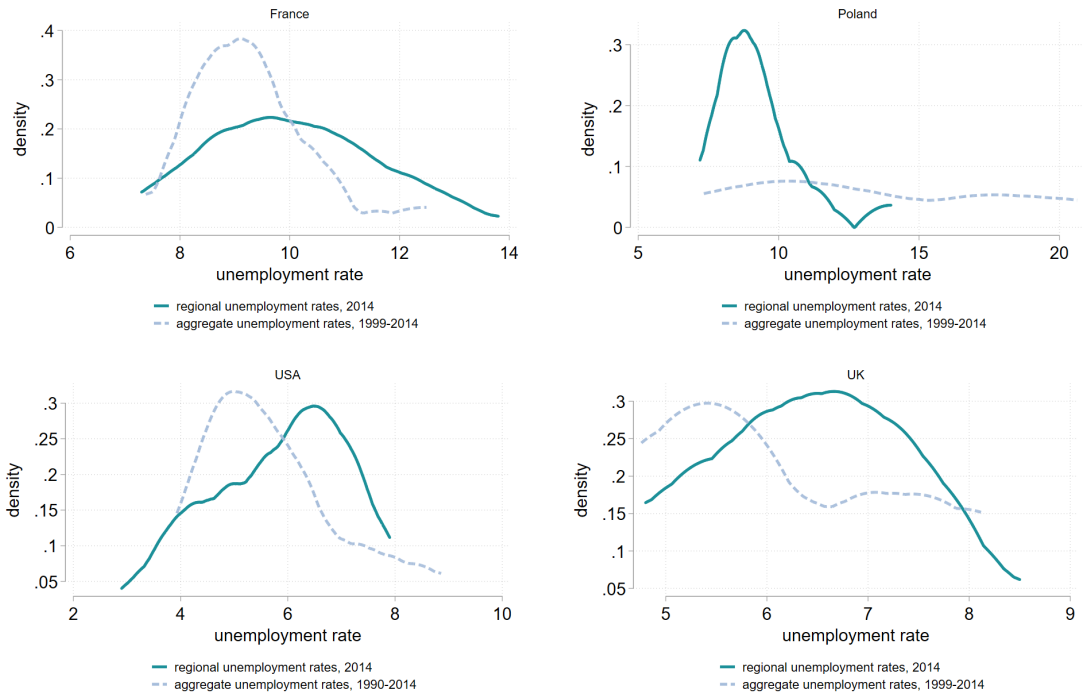
This paper is motivated by the five stylised facts about regional labour markets described in section 4.2. Below, I present evidence of these stylised facts for France, Poland, the UK, and the USA.

Figure C.1: Distribution of regional income



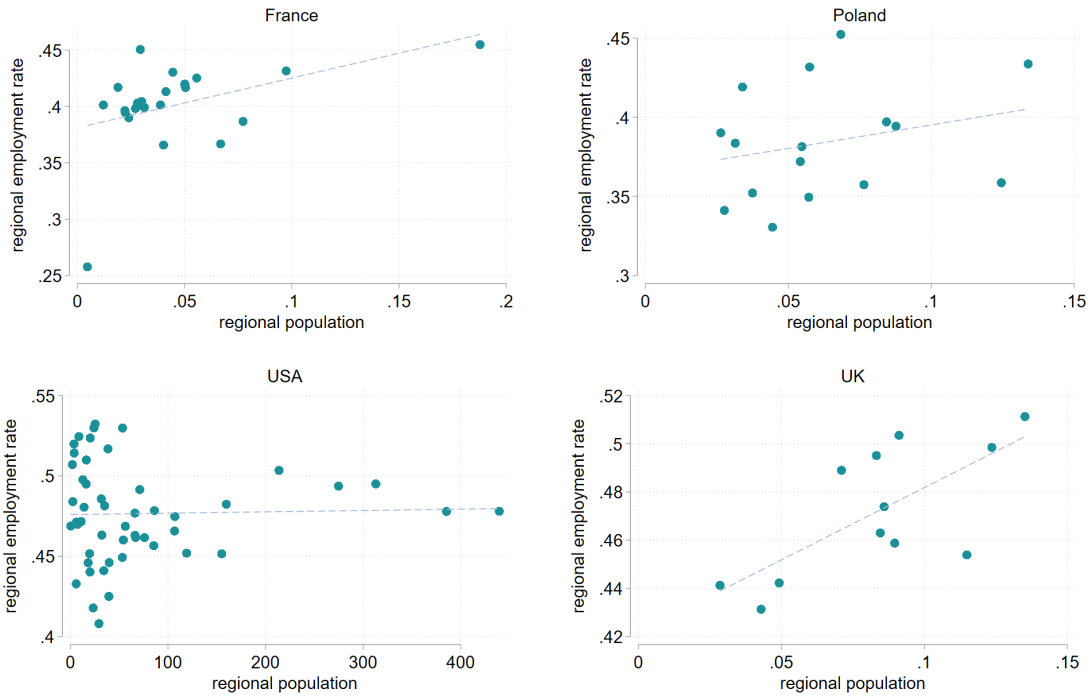
The graph plots kernel density of annual average regional disposable income in the year 2011. Data source: OECD regional database for TL2 – large regions.

Figure C.2: Distribution of regional unemployment rates relative to aggregate unemployment rate over time



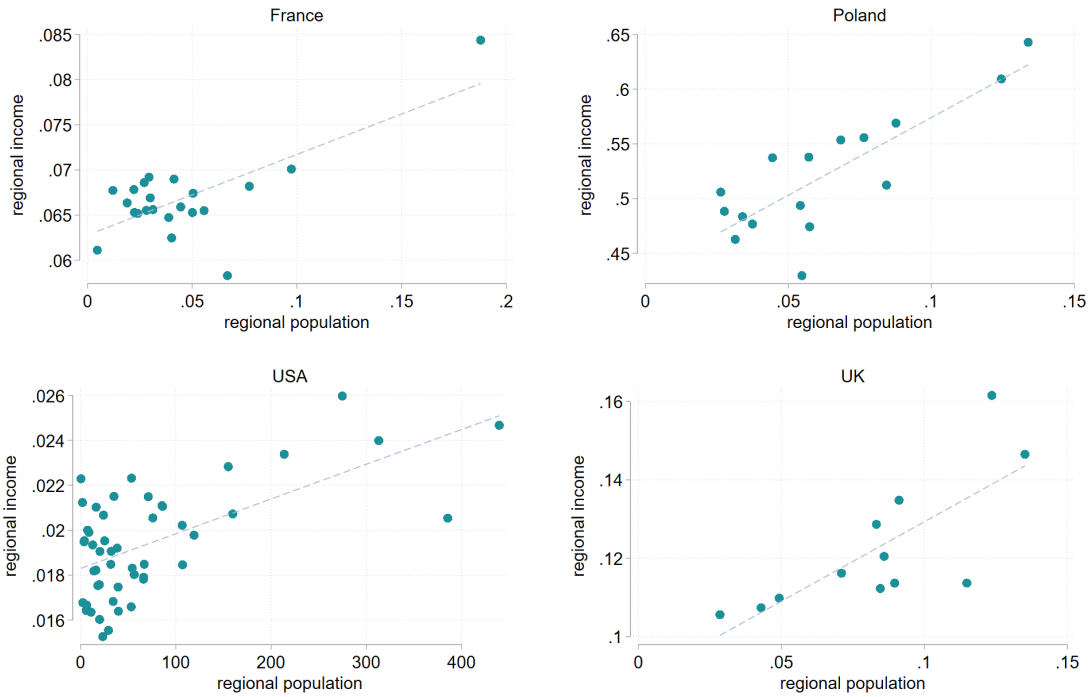
The graph plots kernel density of annual average regional unemployment rate in 2014, and the kernel density of the average aggregate unemployment rate during the years 1990-2014. Data source: OECD regional database for TL2 – large regions.

Figure C.3: Correlation between regional population size and employment rate



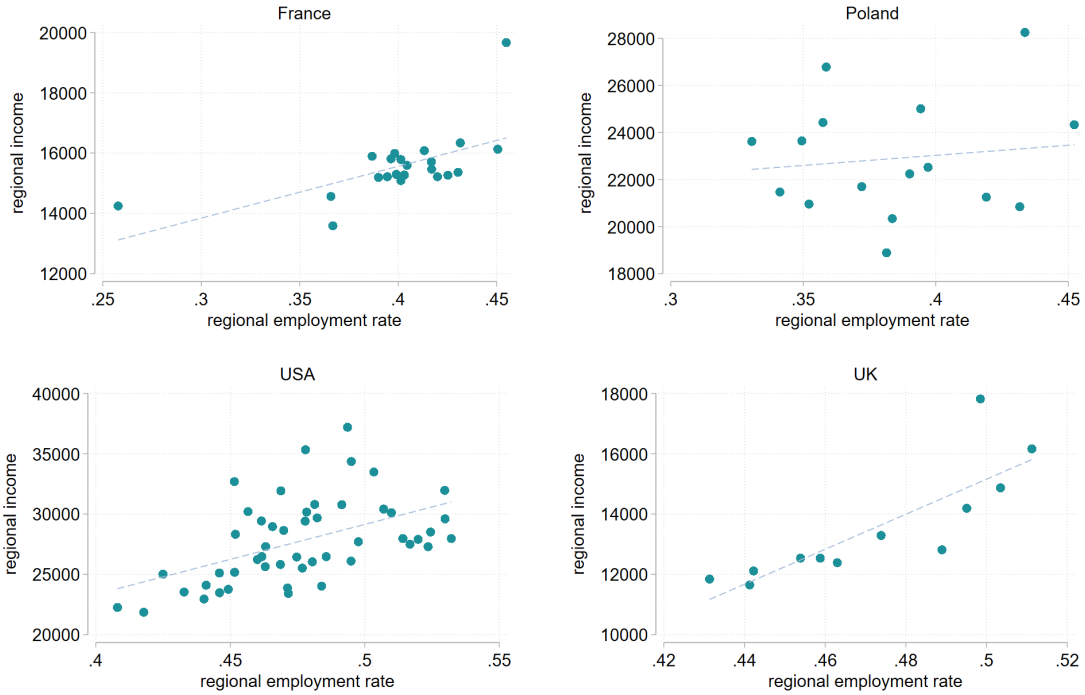
The graph plots relative population share of a region (regional population density for the US) against employment relative to total population. Data corresponds to averages over the years 1990-2014, based on availability. Data source: OECD regional database for TL2 – large regions.

Figure C.4: Correlation between regional population size and income



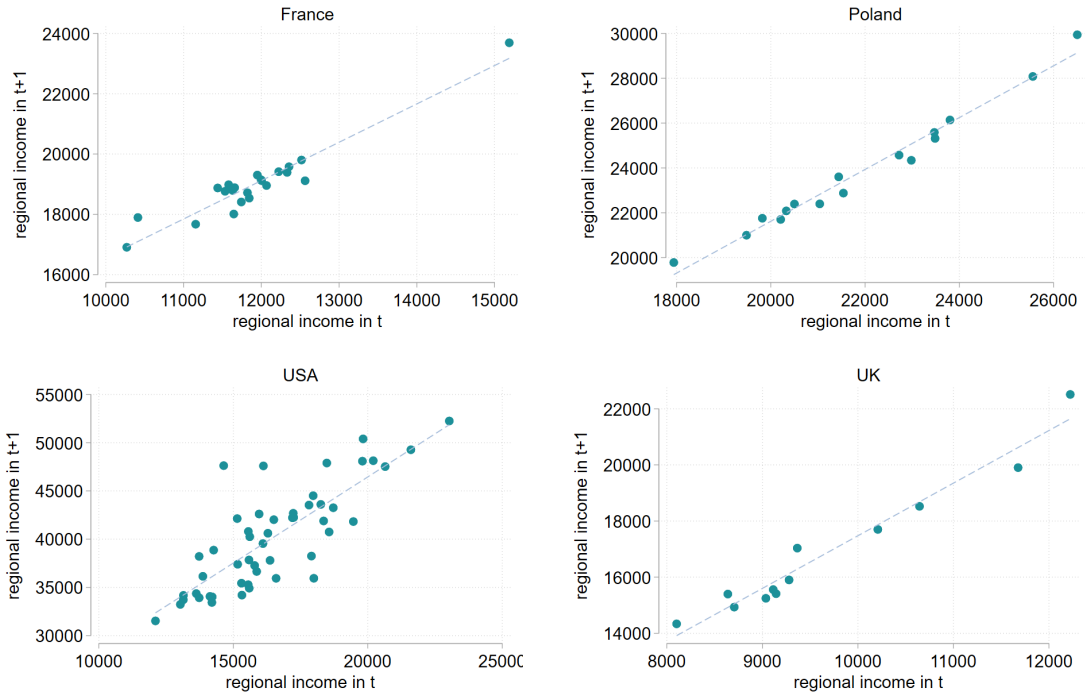
The graph plots relative population share of a region (regional population density for the US) against relative regional disposable income. Data corresponds to averages over the years 1990-2014, based on availability. Data source: OECD regional database for TL2 – large regions.

Figure C.5: Correlation between regional income and employment (wage curve)



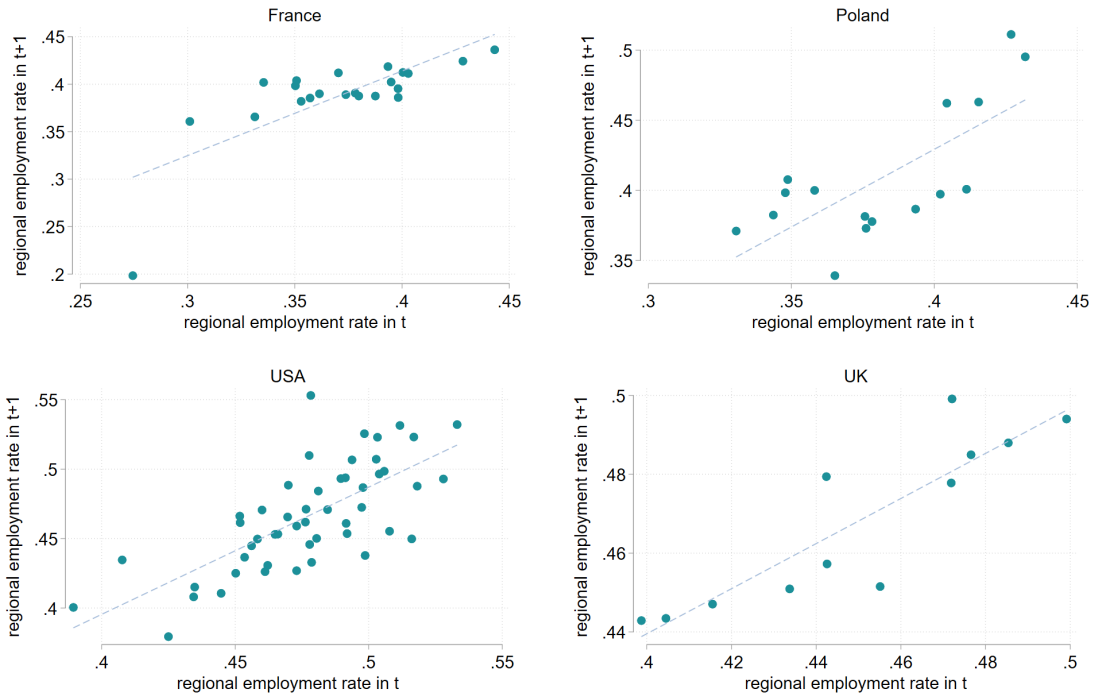
The graph plots employment rate (based on population) against regional disposable income. Data corresponds to averages over the years 1990-2014, based on availability. Data source: OECD regional database for TL2 – large regions.

Figure C.6: Correlation between regional income at the start and the end of the period



The graphs plot the correlation between regional income in the first (t) and the last year (t+1) of data availability: 1995 and 2011 (France); 2010 and 2012 (Poland); 1997 and 2013 (the UK); 1990 and 2014 (USA). Data source: OECD regional database for TL2 – large regions.

Figure C.7: Correlation between regional employment rate at the start and the end of the period



The graphs plot the correlation between population-based employment rate in the first (t) and the last year (t+1) of data availability: 1999 and 2014 (France); 1999 and 2014 (Poland); 1999 and 2014 (the UK); 1990 and 2014 (USA). Data source: OECD regional database for TL2 – large regions.

C.2 Local labour markets under autarky

C.2.1 The model

In this section, I derive the Bellman equations describing the behaviour of workers and firms in a model of local labour markets under autarky, as well as the Nash bargaining procedure that describes how the output of a match is shared between them.

There are two Bellman equations describing the behaviour of workers in each region. The per-period return to unemployment equals the probability that a worker finds a match times the gain of being employed (W) vs unemployed (U). The per-period return to employment consists of wage less the loss of being fired, times the probability of job destruction.

$$rU = \phi(m)\mu(\theta)[W - U] \tag{C.1}$$

$$rW = w + \delta[U - W] \tag{C.2}$$

Analogously, there are two Bellman equations describing the behaviour of employers, capturing their choice of local and cross-regional hiring in each region. The return to posting a vacancy equals the expected gain of filling it, less the per-period cost of advertisement on the labour market, c. The return to a filled job is the match output less wage paid to the worker, less the expected loss that materialises when the match is exogenously destroyed.

$$rV = \frac{\phi(m)\mu(\theta)}{\theta}[J - V] - c \tag{C.3}$$

$$rJ = y - w + \delta[J - V] \tag{C.4}$$

Workers and firms bargain over how to split the match output in a process of Nash bargaining; that is how the wage is set. The outcome depends on their relative bargaining power β and $1 - \beta$, respectively, and on their outside options, which are given by the Bellman equations above. Note that under the free entry condition,

the value of any vacancy can be at most 0, which sets the firm's outside option.

$$
\begin{aligned}
w &= \text{argmax}(W-U)^{\beta}(J)^{1-\beta} \\
\text{F.O.C: } W-U &= \beta(J+W-U) \\
\frac{w+\delta U}{r+\delta} - U &= \beta\left(\frac{y-w}{r+\delta} + \frac{w+\delta U}{r+\delta} - U\right) \\
w &= \beta y + (1-\beta) rU
\end{aligned} \tag{C.5}
$$

The equilibrium wage equals the worker's outside option rU plus her share of the match surplus, $\beta(y - rU)$.

To arrive at the values of unemployment outlined in the main model, we need to substitute the relevant wage expressions and expressions for the value of employment into the Bellman equations for the value of unemployment. Combining the Bellman equations for vacancies and filled jobs, setting the value of vacancy at 0, and subbing in the expressions for the relevant wages, gives us the vacancy creation conditions.

C.3 Local labour markets with migration

C.3.1 The model

In this section, I derive the equations describing the behaviour of workers and firms in a model with endogenous population from first principles.

There are eight Bellman equations describing the behaviour of workers: a value of employment and unemployment as a local or migrant in either of the two regions of the economy. The return to unemployment depends on the expected gain from being employed locally and cross-regionally, denoted by the first two terms. The last term captures the possibility of speculative move, if optimal, if the worker finds no matches. The difference between the value of unemployment for as a local and as a migrant lies is twofold: migrants have to pay an extra search cost b, and their outside option is different. The return to employment equals the relevant

wage less the expected loss of the match being destroyed.

$$rU_A^L = \lambda_A^A[W_A^L - U_A^L] + \lambda_B^A[W_B^M - K - U_A^L] + (1 - \lambda_A^A - \lambda_B^A)\max[0, U_B^M - K - U_A^L] \tag{C.6}$$

$$rU_B^L = \lambda_B^B[W_B^L - U_B^L] + \lambda_A^B[W_A^M - K - U_B^L] + (1 - \lambda_B^B - \lambda_A^B)\max[0, U_A^M - K - U_B^L] \tag{C.7}$$

$$rU_A^M = \lambda_A^A[W_A^L - U_A^M] + \lambda_B^A[W_B^M - K - U_A^M] + (1 - \lambda_A^A - \lambda_B^A)\max[0, U_B^M - K - U_A^M] - b \tag{C.8}$$

$$rU_B^M = \lambda_B^B[W_B^L - U_B^M] + \lambda_A^B[W_A^M - K - U_B^M] + (1 - \lambda_B^B - \lambda_A^B)\max[0, U_A^M - K - U_B^M] - b \tag{C.9}$$

$$rW_A^L = w_A^L + \delta[U_A^L - W_A^L] \tag{C.10}$$

$$rW_B^L = w_B^L + \delta[U_B^L - W_B^L] \tag{C.11}$$

$$rW_A^M = w_A^M + \delta[U_A^L - W_A^M] \tag{C.12}$$

$$rW_B^M = w_B^M + \delta[U_B^L - W_B^M] \tag{C.13}$$

Analogously, there are eight Bellman equations describing the behaviour of employers, capturing their choice of local and cross-regional hiring in either region. The return to posting a vacancy equals the expected gain of filling this vacancy, net of the per-period advertisement cost. The return to a filled job is the output of the match less the wage paid to the worker, less the expected loss of the match being destroyed.

$$rV_A^L = \frac{M^A}{v^A}[J_A^L - V_A^L] - c \tag{C.14}$$

$$rV_B^L = \frac{M^B}{v^B}[J_B^L - V_B^L] - c \tag{C.15}$$

$$rV_A^M = \frac{M^B}{v^B}[J_A^M - V_A^M] - c \tag{C.16}$$

$$rV_B^M = \frac{M^A}{v^A}[J_B^M - V_B^M] - c \tag{C.17}$$

$$rJ_A^L = y - w_A^L + \delta[J_A^L - V_A^L] \tag{C.18}$$

$$rJ_B^L = y - w_B^L + \delta[J_B^L - V_B^L] \tag{C.19}$$

$$rJ_A^M = y - w_A^M + \delta[J_A^M - V_A^M] \tag{C.20}$$

$$rJ_B^M = y - w_B^M + \delta[J_B^M - V_B^M] \tag{C.21}$$

The surplus of each match is split according to Nash bargaining. This yields the equilibrium wage for each of the four hiring scenarios. Note that under the

free entry condition, the value of any vacancy can be at most 0.

$$
\begin{aligned}
w_A^L &= \text{argmax}(W_A^L - U_A^L)^\beta (J_A^L)^{1-\beta} \\
\text{F.O.C: } W_A^L - U_A^L &= \beta(J_A^L + W_A^L - U_A^L) \\
\frac{w_A^L + \delta U_A^L}{r+\delta} - U_A^L &= \beta\left(\frac{y - w_A^L}{r+\delta} + \frac{w_A^L + \delta U_A^L}{r+\delta} - U_A^L\right) \\
w_A^L &= \beta y + (1-\beta) r U_A^L
\end{aligned} \tag{C.22}
$$

$$
\begin{aligned}
w_B^L &= \text{argmax}(W_B^L - U_B^L)^\beta (J_B^L)^{1-\beta} \\
\text{F.O.C: } W_B^L - U_B^L &= \beta(J_B^L + W_B^L - U_B^L) \\
\frac{w_B^L + \delta U_B^L}{r+\delta} - U_B^L &= \beta\left(\frac{y - w_B^L}{r+\delta} + \frac{w_B^L + \delta U_B^L}{r+\delta} - U_B^L\right) \\
w_B^L &= \beta y + (1-\beta) r U_B^L
\end{aligned} \tag{C.23}
$$

$$
\begin{aligned}
w_A^M &= \text{argmax}(W_A^M - U_A^M - K)^\beta (J_A^M)^{1-\beta} \\
\text{F.O.C: } W_A^M - U_A^M - K &= \beta(J_A^M + W_A^M - U_A^M - K) \\
\frac{w_A^M + \delta U_A^L}{r+\delta} - U_A^M - K &= \beta\left(\frac{y - w_A^M}{r+\delta} + \frac{w_A^M + \delta U_A^L}{r+\delta} - U_A^M - K\right) \\
w_A^M &= \beta y + (1-\beta) r U_A^M + (1-\beta)(r+\delta)K - (1-\beta)\delta(U_A^L - U_A^M)
\end{aligned} \tag{C.24}
$$

$$
\begin{aligned}
w_B^M &= \text{argmax}(W_B^M - U_B^M - K)^\beta (J_B^M)^{1-\beta} \\
\text{F.O.C: } W_B^M - U_B^M - K &= \beta(J_B^M + W_B^M - U_B^M - K) \\
\frac{w_B^M + \delta U_B^L}{r+\delta} - U_B^M - K &= \beta\left(\frac{y - w_B^M}{r+\delta} + \frac{w_B^M + \delta U_B^L}{r+\delta} - U_B^M - K\right) \\
w_B^M &= \beta y + (1-\beta) r U_B^M + (1-\beta)(r+\delta)K - (1-\beta)\delta(U_B^L - U_B^M)
\end{aligned} \tag{C.25}
$$

To arrive at the values of unemployment outlined in the main model, we need to substitute the relevant wage expressions and expressions for the value of employment into the Bellman equations for the value of unemployment. Combining the Bellman equations for vacancies and filled jobs, setting the value of vacancy at 0, and subbing in the expressions for the relevant wages, gives us the vacancy creation conditions.

C.3.2 Equilibria with endogenous population

Below, I derive the conditions for the optimality of migration, (M1), (M2) and (M3), in terms of the equilibrium endogenous variables $m_A, m_B, \theta_A, \theta_B, P_A, P_B$.

In all three cases, I use equation (4.24*) to substitute in for U_A^M, and then substitute in for U_A^L, U_B^L using the equilibrium vacancy creation conditions (4.27*) and (4.29*).

The condition for speculative migration, (M1), is:

$$U_A^M - K > U_B^L \tag{M1}$$

$$U_A^L - \frac{b}{r+\phi_A\mu(\theta_A)} - K > U_B^L \tag{C.26}$$

$$\frac{y}{r} - \frac{c(r+\delta)\theta_A}{r(1-\beta)\phi(m_A)\mu(\theta_A)} - \frac{b}{r+\phi_A\mu(\theta_A)} - K > \frac{y}{r} - \frac{c(r+\delta)\theta_B}{r(1-\beta)\phi(m_B)\mu(\theta_B)} \tag{C.27}$$

$$\frac{c(r+\delta)\theta_B}{r(1-\beta)\phi(m_B)\mu(\theta_B)} - \frac{c(r+\delta)\theta_A}{r(1-\beta)\phi(m_A)\mu(\theta_A)} > \frac{b}{r+\phi_A\mu(\theta_A)} + K \tag{C.28}$$

$$\frac{c\theta_B}{(1-\beta)\phi(m_B)\mu(\theta_B)} - \frac{c\theta_A}{(1-\beta)\phi(m_A)\mu(\theta_A)} > \frac{rb}{(r+\phi(m_A)\mu(\theta_A))(r+\delta)} + \frac{rK}{r+\delta} \tag{M1*}$$

The condition for the optimality of cross-regional hiring, (M2), is:

$$rU_A^L - rU_B^L < (r+\delta)(U_A^L - U_A^M - K) \tag{M2}$$

$$rU_A^L - rU_B^L < (r+\delta)\left(\frac{b}{r+\phi_A\mu(\theta_A)} - K\right) \tag{C.29}$$

$$-\frac{c(r+\delta)\theta_A}{(1-\beta)\phi(m_A)\mu(\theta_A)} + \frac{c(r+\delta)\theta_A}{(1-\beta)\phi(m_A)\mu(\theta_A)} < (r+\delta)\left(\frac{b}{r+\phi_A\mu(\theta_A)} - K\right) \tag{C.30}$$

$$\frac{c\theta_B}{(1-\beta)\phi(m_B)\mu(\theta_B)} - \frac{c\theta_A}{(1-\beta)\phi(m_A)\mu(\theta_A)} < \frac{b}{r+\phi(m_A)\mu(\theta_A)} - K \tag{M2*}$$

And the condition for migration for a specific job, (M3), is:

$$\beta S_A^M > U_B^L - U_A^M \tag{M3}$$

$$\beta\left(\frac{y - rU_A^M + \delta(U_A^L - U_A^M)}{r+\delta} - K\right) > U_B^L - U_A^M \tag{C.31}$$

$$\beta\left(\frac{y - rU_A^L}{r+\delta} + \frac{b}{r+\phi_A\mu(\theta_A)} - K\right) > U_B^L - U_A^L + \frac{b}{r+\phi_A\mu(\theta_A)} \tag{C.32}$$

$$\beta\left(\frac{y - rU_A^L}{r+\delta} - K\right) - \frac{(1-\beta)b}{r+\phi_A\mu(\theta_A)} > U_B^L - U_A^L \tag{C.33}$$

$$\frac{\beta c\theta_A}{(1-\beta)\phi(m_A)\mu(\theta_A)} - \beta K - \frac{(1-\beta)b}{r+\phi_A\mu(\theta_A)} > \frac{c(r+\delta)\theta_A}{r(1-\beta)\phi(m_A)\mu(\theta_A)} - \frac{c(r+\delta)\theta_B}{r(1-\beta)\phi(m_B)\mu(\theta_B)} \tag{C.34}$$

$$\frac{c(r+\delta)\theta_B}{r(1-\beta)\phi(m_B)\mu(\theta_B)} - \frac{c(r+\delta)\theta_A}{r(1-\beta)\phi(m_A)\mu(\theta_A)} > \frac{(1-\beta)b}{r+\phi_A\mu(\theta_A)} - \frac{\beta c\theta_A}{(1-\beta)\phi(m_A)\mu(\theta_A)} + \beta K \tag{C.35}$$

$$\frac{c\theta_B}{(1-\beta)\phi(m_B)\mu(\theta_B)} - \frac{c\theta_A}{(1-\beta)\phi(m_A)\mu(\theta_A)} > \frac{r}{r+\delta}\left[\frac{(1-\beta)b}{r+\phi(m_A)\mu(\theta_A)} + \beta K - \frac{\beta c\theta_A}{(1-\beta)\phi(m_A)\mu(\theta_A)}\right] \tag{M3*}$$

where I used equation (4.21) to substitute in for S_A^M.

C.3.3 Derivatives of the migration conditions with respect to regional population levels

As explained in the previous section of this Appendix, the optimality condition for speculative migration, (M1), can be expressed as

$$U_A^L - \frac{b}{r+\phi(m_A)\mu(\theta_A)} - K > U_B^L \tag{C.26}$$

Proposition 2 states that the labour market tightness and size increase in the population level of their region, so the local value of unemployment will increase too. When the local population level falls, so will the value of unemployment in that region. Finally, recall that $\phi(m), \mu(\theta)$ are both positive functions of m

and θ, respectively. As a consequence:

$$\frac{\mathrm{d}U_A^L}{\mathrm{d}P_A} > 0 \tag{C.36}$$

$$\frac{\mathrm{d}U_B^L}{\mathrm{d}P_A} < 0 \tag{C.37}$$

$$\frac{\mathrm{d}\phi(m_A)\mu(\theta_A)}{\mathrm{d}P_A} > 0 \tag{C.38}$$

$$\therefore \frac{\mathrm{d}M1}{\mathrm{d}P_A} > 0 \tag{4.54}$$

The result for the profitability of cross-regional hiring (M2) can be derived analogously. The condition can be expressed as:

$$rU_A^L - rU_B^L < (r+\delta)\left(\frac{b}{r+\phi_A\mu(\theta_A)} - K\right) \tag{C.29}$$

so as P_A increases, the left hand side expression increases, and the right hand side declines.

$$\frac{\mathrm{d}M2}{\mathrm{d}P_A} > 0 \tag{4.55}$$

Finally, the derivative of the optimality of migration for a specific job (M3) can be derived from the following version of equation (C.33):

$$U_A^L\left[1 - \frac{\beta r}{r+\delta}\right] + \beta\left(\frac{y}{r+\delta} - K\right) > U_B^L + \frac{(1-\beta)b}{r+\phi_A\mu(\theta_A)} \tag{C.33}$$

With increasing P_A, the left hand side increases because $1 - \frac{\beta r}{r+\delta} > 0$. At the same time, the right hand side decreases due to the same reasons as above. Overall:

$$\frac{\mathrm{d}M3}{\mathrm{d}P_A} > 0 \tag{4.56}$$

C.4 Modelling regional matching spillovers

To understand how the "separateness" of regional matching functions affects the local labour markets, I extend the baseline model to encompass an economy-wide matching function. In the extreme case where there is a single matching technology for both regions, the matching function is:

$$M^* = M(u_A + u_B, v^A + v^B) \tag{C.39}$$

The local job-finding and vacancy filling probabilities are equal to their economy-wide counterparts, and equivalent across regions:

$$\frac{M_A}{u_A} = \frac{M^*}{u_A}\frac{u_A}{u_A + u_B} = \frac{M^*}{u_A + u_B} = \frac{M_B}{u_B} \tag{C.40}$$

$$\frac{M_A}{v^A} = \frac{M^*}{v^A}\frac{v^A}{v^A + v^B} = \frac{M^*}{v^A + v^B} = \frac{M_B}{v^B} \tag{C.41}$$

The other extreme is the case with two separate regional matching functions, as presented in section 4. In between these two extremes lie labour markets with some degree of matching spillovers. The composite matching function is:

$$M_A = (1-n)M(u_A, v^A) + nM^*\frac{u_A}{u_A + u_B} \tag{C.42}$$

$$M_B = (1-n)M(u_B, v^B) + nM^*\frac{u_B}{u_A + u_B} \tag{C.43}$$

where n denotes the degree of regional matching spillovers, or the inverse of local matching function separateness.

Bibliography

Jason Abaluck and Abi Adams. What do consumers consider before they choose? identification from asymmetric demand responses. Working Paper 23566, National Bureau of Economic Research, June 2017.

Joseph G Altonji and Lewis M Segal. Small-Sample Bias in GMM Estimation of Covariance Structures. *Journal of Business & Economic Statistics*, 14(3):353–366, July 1996.

Michael Amior. Why are higher skilled workers more mobile geographically? the role of the job surplus. Discussion Paper 1338, Centre for Economic Performance, 2015.

Michael Amior and Alan Manning. The persistence of local joblessness. *American Economic Review*, 108(7):1942–70, July 2018.

Simon Baker, Seamus Hogan, and Christopher Ragan. Is There Compelling Evidence against Increasing Returns to Matching in the Labour Market? *Canadian Journal of Economics*, 29(4):976–93, November 1996.

Alison E. Barber, Michael J. Wesson, Quineita M. Roberson, and M. Susan Taylor. A tale of two job markets: Organizational size and its effects on hiring practices and job search behavior. *Personnel Psychology*, 52(4):841–868, 1999.

John Barron, John Bishop, and William C Dunkelberg. Employer search: The interviewing and hiring of new employees. *The Review of Economics and Statistics*, 67 (1):43–52, 1985.

Alexander Bartik. Moving costs and worker adjustment to changes in labor demand: Evidence from longitudinal census data. *Manuscript, University of Illinois at Urbana-Champaign*, 2018.

Emek Basker. Education, job search, and migration. *Journal of Regional Analysis & Policy*, 48:38–61, 11 2018.

Christian Bayer and Falko Juessen. On the Dynamics of Interstate Migration: Migration Costs and Self-Selection. *Review of Economic Dynamics*, 15(3):377–401, July 2012.

Paul Beaudry, David A. Green, and Benjamin M. Sand. Spatial equilibrium with unemployment and wage bargaining: Theory and estimation. *Journal of Urban Economics*, 79(C):2–19, 2014.

Michèle Belot, Philipp Kircher, and Paul Muller. Providing Advice to Jobseekers at Low Cost: An Experimental Study on Online Advice. 10 2018.

Olivier Jean Blanchard and Peter A. Diamond. The Aggregate Matching Function. NBER Working Papers 3175, National Bureau of Economic Research, Inc, November 1989.

Olivier Jean Blanchard and Lawrence F. Katz. Regional Evolutions. *Brookings Papers on Economic Activity*, 23(1):1–76, 1992.

David G. Blanchflower and Andrew J. Oswald. An Introduction to the Wage Curve. *Journal of Economic Perspectives*, 9(3):153–167, Summer 1995.

Hoyt Bleakley and Jeffrey Lin. Thick-market effects and churning in the labor market: Evidence from US cities. *Journal of Urban Economics*, 72(2):87–103, 2012.

George J. Borjas, Stephen G. Bronars, and Stephen J. Trejo. Self-selection and internal migration in the United States. *Journal of Urban Economics*, 32(2):159–185, September 1992.

John Bound and Harry J Holzer. Demand shifts, population adjustments, and labor market outcomes during the 1980s. Working Paper 5685, National Bureau of Economic Research, July 1996.

Raj Chetty and Nathaniel Hendren. The impacts of neighborhoods on intergenerational mobility ii: County-level estimates*. *The Quarterly Journal of Economics*, 133(3): 1163–1228, 2018.

Raj Chetty, Michael Stepner, Sarah Abraham, Shelby Lin, Benjamin Scuderi, Nicholas Turner, Augustin Bergeron, and David Cutler. The association between income and life expectancy in the united states, 2001-2014. *JAMA*, 315, 04 2016.

David E. Clark and William J. Hunter. The impact of economic opportunity, amenities and fiscal factors on age-specific migration rates. *Journal of Regional Science*, 32(3): 349, 1992.

Christopher T Conlon and Julie Holland Mortimer. Demand estimation under incomplete product availability. Working Paper 14315, National Bureau of Economic Research, September 2008.

Gordon B. Dahl. Mobility and the return to education: Testing a roy model with multiple markets. *Econometrica*, 70(6):2367–2420, 2002. ISSN 1468-0262.

Paul S. Davies, Michael J. Greenwood, and Haizheng Li. A conditional logit approach to u.s. state-to-state migration. *Journal of Regional Science*, 41(2):337–360, 2001.

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

Glenn Ellison and Edward L Glaeser. Geographic Concentration in U.S. Manufacturing Industries: A Dartboard Approach. *Journal of Political Economy*, 105(5):889–927, October 1997.

Martin Ellison, Godfrey Keller, Kevin Roberts, and Margaret Stevens. Unemployment and market size. *Economic Journal*, 124(575):119–148, 03 2014.

Paolo Epifani and Gino A. Gancia. Trade, migration and regional unemployment. *Regional Science and Urban Economics*, 35(6):625 – 644, 2005.

John Francis. Agglomeration, job flows and unemployment. *The Annals of Regional Science*, 43(1):181–198, 2009.

Simon Franklin. Location, search costs and youth unemployment: Experimental evidence from transport subsidies. *The Economic Journal*, 128(614):2353–2379, 2018.

P.A. Gautier and C.N. Teulings. Search and the city. *Regional Science and Urban Economics*, 39(3):251–265, May 2009.

Ahu Gemici. Family migration and labor market outcomes. *Manuscript, New York University*, 2011.

Edward L Glaeser and Joshua D Gottlieb. The economics of place-making policies. Working Paper 14373, National Bureau of Economic Research, October 2008.

Michelle Sovinsky Goeree. Limited information and advertising in the u.s. personal computer industry. *Econometrica*, 76(5):1017–1074, 2008. ISSN 1468-0262.

Michael J. Greenwood. Internal migration in developed countries. In M. R. Rosenzweig and O. Stark, editors, *Handbook of Population and Family Economics*, volume 1 of *Handbook of Population and Family Economics*, chapter 12, pages 647–720. Elsevier, 1997.

Paul Gregg, Stephen Machin, and Alan Manning. Mobility and joblessness. In *Seeking a Premier Economy: The Economic Effects of British Economic Reforms, 1980-2000*, pages 371–410. National Bureau of Economic Research, Inc, 2004.

John R. Harris and Michael P. Todaro. Migration, unemployment and development: A two-sector analysis. *The American Economic Review*, 60(1):126–142, 1970.

James J. Heckman. Sample selection bias as a specification error. *Econometrica*, 47(1): 153–161, 1979.

Rubén Hernández-Murillo, Lesli S. Ott, Michael T. Owyang, and Denise Whalen. Patterns of interstate migration in the United States from the Survey of Income and Program Participation. *Federal Reserve Bank of St. Louis Review*, 93(3):169–186, May 2011.

Brad Hershbein and Lisa B. Kahn. Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review*, 108(7):1737–72, July 2018.

Henry W. Herzog, Alan M. Schlottmann, and Thomas P. Boehm. Migration as spatial job–search: A survey of empirical findings. *Regional Studies*, 27(4):327–340, 1993.

Kate Ho, Joseph Hogan, and Fiona Scott Morton. The impact of consumer inattention on insurer pricing in the medicare part d program. *The RAND Journal of Economics*, 48(4):877–905, 2017.

Hilary Hoynes. The employment, earnings, and income of less skilled workers over the business cycle. Working Paper 7188, National Bureau of Economic Research, June 1999.

Michael P. Keane, Petra E. Todd, and Kenneth I. Wolpin. *The Structural Estimation of Behavioral Models: Discrete Choice Dynamic Programming Methods and Applications*, volume 4 of *Handbook of Labor Economics*, chapter 4, pages 331–461. Elsevier, 2011.

John Kennan and James R. Walker. The Effect of Expected Income on Individual Migration Decisions. *Econometrica*, 79(1):211–251, January 2011.

Patrick Kline and Enrico Moretti. Place Based Policies with Unemployment. *American Economic Review*, 103(3):238–43, May 2013a.

Patrick Kline and Enrico Moretti. Place Based Policies with Unemployment. *American Economic Review*, 103(3):238–43, May 2013b.

Paul Krugman. Increasing Returns and Economic Geography. *Journal of Political Economy*, 99(3):483–99, June 1991.

Vanessa Lutgen and Bruno Van der Linden. Regional equilibrium unemployment theory at the age of the internet. *Regional Science and Urban Economics*, 53:50 – 67, 2015.

Ofer Malamud and Abigail Wozniak. The impact of college on migration: Evidence from the vietnam generation. *The Journal of Human Resources*, 47(4):913–950, 2012.

J. J. McCall. Economics of information and job search. *The Quarterly Journal of Economics*, 84(1):113–126, 1970.

Daniel McFadden. A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica*, 57(5):995–1026, 1989.

Daniel McFadden et al. Conditional logit analysis of qualitative choice behavior. 1973.

Ian Molho. Spatial search, migration and regional unemployment. *Economica*, 68(270): 269–283, 2001.

Raven Molloy, Christopher L. Smith, and Abigail Wozniak. Internal migration in the united states. *Journal of Economic Perspectives*, 25(3):173–96, September 2011.

Enrico Moretti. Chapter 14 - local labor markets*. volume 4, Part B of *Handbook of Labor Economics*, pages 1237 – 1313. Elsevier, 2011.

Henry G. Overman and Diego Puga. Unemployment clusters across Europe's regions and countries. *Economic Policy*, 17(34):115–148, 04 2002.

Barbara Petrongolo and Christopher Pissarides. Scale Effects in Markets with Search. *Economic Journal*, 116(508):21–44, 01 2006.

Christopher Pissarides, Richard Layard, and Martin Hellwig. Unemployment and vacancies in britain. *Economic Policy*, 1(3):500–559, 1986.

Christopher A. Pissarides. *Equilibrium Unemployment Theory, 2nd Edition*, volume 1 of *MIT Press Books*. The MIT Press, March 2000.

Christopher A Pissarides and Jonathan Wadsworth. Unemployment and the Inter-regional Mobility of Labour. *Economic Journal*, 99(397):739–55, September 1989.

Tyler Ransom. Labor market frictions and moving costs of the employed and unemployed. Working Paper 12139., IZA, 2019.

Jordan Rappaport. Moving to nice weather. *Regional Science and Urban Economics*, 37 (3):375 – 398, 2007.

Jennifer Roback. Wages, rents, and the quality of life. *Journal of Political Economy*, 90 (6):1257–78, 1982.

Joan R. Rodgers and John L. Rodgers. The Effect of Geographic Mobility on Male Labor-Force Participants in the United States . *Journal of Labor Research*, 21(1): 117–132, January 2000.

G. Russo, P. Rietveld, P. Nijkamp, and C. Gorter. Spatial aspects of recruitment behaviour of firms: an empirical investigation. *Environment and Planning A*, 28(6): 1077–1093, 1996.

John Rust. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, 55(5):999–1033, 1987.

Benoît Schmutz and Modibo Sidibé. Frictional Labour Mobility. *The Review of Economic Studies*, 09 2018.

Arthur L. Silvers. Probabilistic income-maximizing behavior in regional migration. *International Regional Science Review*, 2(1):29–40, 1977.

Larry A. Sjaastad. The costs and returns of human migration. *Journal of Political Economy*, 70(5):80–93, 1962.

Katalin Tardos and Jacob Pedersen. Recruitment channels and strategies for employing the low-skilled workers in europe. *Szociologiai Szemle*, 21, 12 2011.

Ronald Jr. Warren. Returns to scale in a matching model of the labor market. *Economics Letters*, 50(1):135–142, January 1996.

Christopher H. Wheeler. Local market scale and the pattern of job changes among young men. *Regional Science and Urban Economics*, 38(2):101 – 118, 2008.

Ronald L. Whisler, B.S. Waldorf, G. F. Mulligan, and D. A. Plane. Quality of life and the migration of the college-educated: A life-course approach. *Growth and Change*, 39 (1):58–94, 2008.

Marc Witte. Job referrals and strategic network formation: Experimental evidence from urban neighbourhoods in ethiopia. *Manuscript, University of Oxford*, 2019.

Abigail Wozniak. Are college graduates more responsive to distant labor market opportunities? *The Journal of Human Resources*, 45(4):944–970, 2010.

Eran Yashiv. The Determinants of Equilibrium Unemployment. *American Economic Review*, 90(5):1297–1322, December 2000.