NUI MAYNOOTH

Ollscoil na hÉireann Má Nuad

Commuting flows & local labour markets:

Spatial interaction modelling of travel-to-work

Carson J. Q. Farmer, B.Sc., M.Sc.

Supervisor: Professor A. Stewart Fotheringham

Thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy (Geocomputation)

National Centre for Geocomputation, Faculty of Science
National University of Ireland, Maynooth
Maynooth, Ireland

September, 2011

Abstract

One of the most promising approaches to mitigating land-use and transportation problems is continued research on urban commuting. Commuting is essential to many individuals, allowing them to participate in the labour market and earn a living to meet their essential needs. As such, a better understanding of the determinants of commuting will ultimately lead to a better understanding of the complexities of employment, housing, and the many spatial processes underlying commuting. However, in order to understand the commuting process, it is important to examine the milieu within which commuting takes place: the local labour market (LLM). In this thesis, the interplay between commuting and LLMs is explored through the use of regionalisation techniques and spatial interaction models. It is shown that LLM characteristics play a significant role in intra-regional commuting patterns and that a failure to account for LLM conditions may seriously hinder the applicability of models of commuting. Specifically, it is found that there are many different LLMs across Ireland, and that these LLMs characterise the commuting patterns of population sub-groups. By incorporating these LLMs into models of commuting, this thesis shows that in addition to distance and working population size, the spatial structure of origins and destinations and a number of non-spatial attributes such as unemployment, housing density, and education, all significantly affect commuting flows. Furthermore, the distance decay component of these models appears to be capturing a combination of geographical distance

and regional differentiation due to LLM boundaries, leading to 'functional' distance decay. This concept of functional distance decay is a key finding of this thesis, and indicates that in addition to the configuration of origins and destinations, distance decay is also dependent on the spatial structure of LLMs, or more generally, the totality of surrounding conditions within which spatial interaction takes place.

Dedication

To Bob, Tia, Devon, MacKenzie, and Amanda. Thank you for all your love and support, I am forever grateful.

Acknowledgements

The work presented in this thesis could not have been completed without the help, guidance, advice, and encouragement of my colleagues, friends, and family. It is to these many people that I owe my deepest gratitude.

First and foremost, I would like to thank my supervisor Prof. A. Stewart Fotheringham for his many contributions of time, ideas, and humour; all of which has greatly contributed to my PhD experience. His knowledge and insight into anything and everything he encounters is inspiring and it has been an honour to work with him over these past three years.

I would also like to give a general thanks to all 3rd floor Iontas building occupants, especially those to whom I have gone for advice, help, guidance, and friendship. Everyone in the NCG and NIRSA, be it consciously or not, has greatly contributed to my personal and professional time at NUIM. More specifically, I acknowledge the post-docs for many helpful discussions and early career advice, the administrative staff for keeping me on track and organised, and the various senior academic staff members for providing general help with all things academic.

Thanks are also due to Prof. Chris Brunsdon for answering my many statistical questions, Dr. Alexei Pozdnoukhov for many interesting discussions, side-projects, and being my second supervisor. I would also like to extend my sincerest gratitude to my examination committee, Martin Charlton and Prof. Graham Clarke, whose insightful comments and suggestions have greatly improved the quality of this thesis.

I would also like to thank all my fellow PhD students, each of whom have made the long and sometimes frustrating, hours at the office all the more bearable. In particular, I'd like to thank Fergal, Cathal, Ishwari, and my fellow ISSP comrades for their helpful advice, insightful discussions, and general camaraderie.

Finally, I thank my family for providing a wonderful environment in which to grow, both as a person and as an academic, where curiosity and a willingness to learn were always encouraged. Without their encouragement, I would not be where I am today.

Funding

Research presented in this thesis was funded by the Irish Social Sciences Platform (ISSP), under the Programme for Research in Third Level Institutions (PRTLI), administered by the Higher Education Authority (HEA) and co-funded under the European Regional Development Fund (ERDF). Additional funding was provided by the Social Sciences and Humanities Research Council (SSHRC) of Canada, as part of the SSHRC Doctoral Fellowships Program. Partial funding by a Strategic Research Cluster (SRC) grant (07/SRC/I1168) by Science Foundation Ireland (SFI) under the National Development Plan (NDP) is also gratefully acknowledged.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 General Overview

"Of all the trips that are made daily within [an urban] area, none is more important than the journey-to-work. Nowhere is the rhythm of the city more evident than in its daily commuting patterns" (Taaffe et al., 1996, chap. 6, p. 184). The spatial interaction that takes place within a particular region is often a reflection of its spatial organisation, transportation network, land-use, and local labour market. According to Horner (2004), one of the most promising approaches to mitigating land-use and transportation problems is continued research on urban commuting. While commuting is only one aspect of travel in most countries, there is strong evidence to suggest that commuting and congestion are inextricably linked, leading to many social and environmental consequences. Furthermore, effective management of travel demand, through future needs prediction and 'what if' policy scenarios, requires detailed knowledge of the determinants of transport use (both public and private), the effect of urban structure on commuting patterns, and a myriad of other commuting-related factors (Commins & Nolan, 2010b). The link between home and work is also of vital importance for understanding individual issues such as stress, time away from family/home, and participation in the labour market, as well as more general social/environmental issues such as vehicle emissions, infrastructure pressures, traffic noise, and the *structure* of the local labour market.

Commuting has increased substantially over the past decades in most developed nations (e.g., Green et al., 1999; Sultana & Weber, 2007; Westin & Sandow, 2010), including in Ireland (e.g., Horner, 1999; Commins & Nolan, 2010a). Many factors have contributed to this development, including increased female participation in the labour force, higher education levels, increasing worker specialisation, improved infrastructure, and a lower propensity for internal migration (Westin & Sandow, 2010). In the US, travel-to-work trips account for approximately 20-25% of the total trips taken by individuals (Pisarski, 2002). Indeed, the most congested time of the day on the road network is typically when people are travelling to and from work (Redmond & Mokhtarian, 2001; Horner, 2004). In addition, people tend to plan other activities such as shopping, relative to their home and work locations (Redmond & Mokhtarian, 2001). Commuting is essential to many individuals; allowing them to participate in the labour market and earn a living to meet their essential needs (Hanson & Pratt, 1988; Horner, 2004).

There appears to be wide-spread agreement among geographers and transportation researchers interested in commuting that a better understanding of the determinants of commute times, distances, and residential and/or employment locations will ultimately lead to a better understanding of the complexities of employment, housing, the economy, society, and the many spatial processes underlying commuting (Horner, 2004). This is evidenced by the wide range of research topics surrounding commuting, including visualising commuter flows (e.g., Nielsen & Hovgesen, 2008), location demand (e.g., Anas & Chu, 1984), gender differentials in commuting and the workplace (e.g., Turner & Niemeier, 1997), rural/urban population dynamics (e.g., Renkow & Hoover, 2000), urban structure effects on commuting (e.g., Giuliano & Small, 1993; Giuliano & Narayan, 2003), travel-to-work mode of transport (e.g., Vega & Reynolds-Feighan, 2008; Commins & Nolan, 2010a,b), and most importantly for this thesis, modelling of commuting flows (e.g., Vermeulen, 2003; Gitlesen & Thorsen, 2000; McArthur et al., 2010). Additionally, commuting behaviour is widely cited as one of the most appropriate indicators of local labour market dimensions (e.g., Goodman, 1970; Ball, 1980; Gerard, 1958; Vance, 1960; Hunter, 1969), and as a result, is often used to delineate the boundaries of local labour markets. This linkage between the local labour market and commuting behaviour is an important observation and one which is central to the work carried out in this thesis.

In the recent literature, researchers have used spatial interaction models to try to understand commuting behaviour and the commuting process (e.g., Thorsen & Gitlesen, 1998, 2002; Gitlesen & Thorsen, 2000; Ubøe, 2004; Vermeulen, 2003; O'Kelly & Lee, 2005; O'Kelly & Niedzielski, 2008, 2009; McArthur et al., 2010, *inter alia*). These studies have provided valuable insights into the determinants of commuting and the interaction between residence and employment locations. Additionally, some of the above research has provided evidence that local labour market characteristics may play an important role in intra-regional commuting patterns. In particular, Thorsen & Gitlesen (1998) and Gitlesen & Thorsen (2000) find that special care should be taken regarding the benefits of residing and working in the same (employment) zone. They explicitly adjust their models of commuting to take into account these effects. In general however, there has been limited consideration of the effects of local labour market characteristics on commuting patterns and models.

The past several decades have seen increasing recognition that the standard administrative areas used by governments for policy making, resource allocation, and research may not be particularly meaningful in terms of representing areas of relative cohesiveness (Ball, 1980; Casado-Díaz, 2000). As such, there has been a move towards the identification and delineation of local labour markets (Goodman, 1970; Ball, 1980; Coombes & Openshaw, 1982; Casado-Díaz, 2000). Ball (1980) suggests that the use of a spatial definition of local labour market conditions is of particular relevance for several reasons, including as a means of presenting information and statistics on employment and socio-economic structures, as well as assessing the effectiveness of regional policy decisions and local government reorganisation. Thus, a failure to account for local labour market conditions may seriously hinder the applicability of models of commuting. In essence, failure to consider local labour market effects on commuting is a failure to consider the milieu within which commuting behaviour operates: the local labour market. In this thesis, we attempt to capture this local labour market effect by attempting to understand commuting behaviour from the perspective of the local labour market. The questions we are therefore asking are: "to what extent does commuting enforce the spatial boundaries of local labour markets?" and "how much do the boundaries of local labour markets influence commuting?"

1.2 Research objectives

These primary research questions are difficult to answer on their own, particularly because the effects of commuting on the local labour market and the effects of the local labour market on commuting are difficult to separate. Therefore, in order to address these research questions, we define a series of research goals and objectives which guide the remainder of this thesis. Our research goals can be separated into three distinct but related targets: 1) define and delineate local labour markets for Ireland in a logical and defensible manner, 2) develop a spatial interaction model of commuting which takes into account the problems and opportunities inherent in Irish commuting patterns, and 3) incorporate goals 1 and 2 into a single modelling framework to investigate the effects of local labour markets on models of commuting flows. Each of these three goals can be further refined to reveal individual research objectives, which represent the steps required to complete each goal. The following is a breakdown of the three research goals and their objectives:

1. Define and delineate local labour markets for Ireland in a logical and defensible manner[1]
    - Develop a useful definition of a local labour market in the context of the functional regionalisation literature,
    - Develop an efficient and effective regionalisation algorithm which is applicable in a range of travel-to-work contexts.
2. Develop spatial interaction model(s) of commuting
    - Develop an effective base-model for commuting in Ireland based on modern spatial interaction principals,
    - Develop more complex commuting models designed to extend spatial interaction theories to the unique context of commuting and, in particular, commuting in Ireland.
3. Local labour market effects on commuting behaviour
    - Explore the effects of the local labour market on parameters in the global commuting model by controlling for the effects of internal, external, and inter-regional commuting,

[1] While not a direct requirement of this objective, it is also beneficial to examine how different socio-economic and population sub-groups contribute to the aggregate local labour market.

- Implement local spatial interaction models at the level of the local labour market and explore parameter trends across Ireland.

1.3 Thesis structure

The three research objectives and their six research goals presented above form the basis for the remaining eight chapters of this thesis. The general trend of the chapters is from theoretical/critical to empirical/analytical. The specific structure and preliminary details of the remaining chapters is as follows:

Chapter 2: Local labour markets In this chapter we focus on the definition and delineation of local labour markets via functional regionalisation methods. Linkages between local labour markets and functional regions are also presented. We examine several competing alternatives in terms of both definitions and applications and conclude the chapter with an alternative definition and methodology that we feel addresses several shortcomings of existing procedures.

Chapter 3: Spatial interaction modelling Here we examine the evolution of spatial interaction theories and techniques, starting from the most basic gravity model and progressing to more modern forms of Poisson spatial interaction models. We also highlight the utility of spatial interaction modelling in the context of commuting research and examine the utility of local spatial interaction models, which 'sets the scene' for the remaining chapters of this thesis.

Chapter 4: Irish commuting data The empirical analyses carried out in this thesis require detailed travel-to-work data in order to delineate functional regions and model commuting flows. Here, we use data compiled for the entire

Republic of Ireland. This chapter examines recent trends in Irish commuting patterns and introduces the dataset used throughout the remainder of the thesis. We also examine the dataset in detail, providing some initial insights and highlighting specific features of the network of commuting flows.

Chapter 5: Generating functional regions In this chapter we focus on the application of the functional regionalisation procedure proposed in Chapter 2. We also implement several modifications to this earlier algorithm and introduce a generalisable technique for assessing the stability of the regionalisation(s). Results are presented and a discussion is undertaken. In addition, we highlight the utility of our regionalisation procedure by examining a range of socio-economic functional regions.

Chapter 6: Modelling commuting flows Here, we focus primarily on the development and empirical evaluation of a global spatial interaction model of commuting (based on Irish commuting data). Following the specification of a base-model, we present several model improvements and highlight their utility based on empirical evidence and results. Furthermore, we explore internal versus external commuting flows based on the functional regions (local labour markets) specified in the previous chapter. The statistical models and results presented in this chapter provide the catalyst required to extend our analysis to the *local* scale in the following chapter.

Chapter 7: Local labour market effects The goal of this chapter is to provide the evidence needed to tie the preceding chapters together and to present a unified exposition of local labour market effects on commuting. We start by calibrating an aggregate 'global' model of commuting flows *between* local labour markets, and finish off the chapter by calibrating a set of 'local' models of commuting *within* local labour markets. Each stage of the analysis in this chapter builds on previous results and ultimately leads to the primary conclusions drawn from the thesis.

Chapter 8: Choice set integration While not directly contingent on previous chapters, this chapter provides a ‘way forward’ in terms of integrating theories of spatial interaction and spatial choice. Here we focus on the concept of choice set generation, providing both theoretical and methodological definitions of choice sets for spatial interaction models. We test our theories on a subset of the full commuting dataset used in previous chapters and calibrate a set of spatial interaction models designed to take into account the sparse nature of commuting data. From this, we highlight the importance of considering travel-to-work as a two-part choice process.

Chapter 9: Conclusions In this final chapter, we present the main conclusions and recommendations derived from previous chapters. We summarise our main findings, discuss the implications of our results, and detail future directions for extending this research area. In particular, we provide linkages between the previous seven chapters and readdress some of the key objectives and goals presented in this Introduction. Finally, specific contributions of this research to the commuting literature, as well as to spatial interaction modelling in general, are discussed, followed by some final remarks.

1.4 Moving forward

Politicians and practitioners alike talk about the local labour market as if it were a well defined spatial unit. Statements such as ‘this process is operating at the level of the local labour market’, or ‘local labour market X has seen major changes in employment levels over the past three quarters’, are common; yet we lack explicit definitions of exactly what a local labour market is and how it structures, or is structured by, employment and the labour market. The definition and delineation of local labour market boundaries is therefore of great interest for many reasons, chief among them for informing policy. Furthermore, consideration of the effects of the local labour market on commuting and the various processes that drive commuting (i.e., residential choice, job search, mode choice, etc.) is of interest, particularly in an age where individuals are becoming more specialised, commuting longer distances, and competing at both local and global scales.

While the methods and analyses employed in this thesis are presented in the context of commuting in Ireland, they could be applied equally well to other forms of spatial interaction data in both the natural and human environments. In particular, the global models presented in Chapter 6 are applicable in any number of spatial interaction settings, including as inputs into urban and/or land-use models. Additionally, our unique treatment of local labour market effects in spatial interaction models in Chapter 7 presents many opportunities for extensions to other forms of 'local' regional effects, as well as extensions into ecology and physical geography. Furthermore, the work on choice set generation presented in Chapter 8 should provide an important step towards a more behavioural interpretation of spatial interaction models in general. Finally, the functional regionalisation procedure presented in chapters 2 and 5 has wide-ranging implications for policy in Ireland and abroad and may provide a 'catalyst' for the development of a generalisable framework for international comparison of local labour market dimensions.

Chapter 2

Local labour markets

2.1 Introduction

Journey-to-work behaviour is widely cited as one of the most appropriate indicators of local labour market dimensions (e.g., Goodman, 1970; Ball, 1980; Gerard, 1958; Vance, 1960; Hunter, 1969) and, as a result, is often used to delineate the boundaries of local labour markets. A local labour market is often defined as a geographical region in which a large majority of the local population seeks employment and the majority of local employers recruit their labour (Goodman, 1970; Ball, 1980; Coombes & Openshaw, 1982; Casado-Díaz, 2000). This type of definition is generally referred to as a homogeneous labour market and encompasses a region where aggregated supply and demand meet (van der Laan & Schalke, 2001). It is generally assumed that only a limited amount of inter-region travel-to-work occurs, such that individual local labour markets are relatively self-contained. Two forms of self-containment are usually considered: 1) demand-side self-containment, whereby the majority of jobs are filled by local residents and 2) supply-side self-containment, whereby the majority of residents who live in the region, also work in the region (Cörvers et al., 2009). In most cases, a homogeneous local labour market is delineated on the basis of the entire labour supply, rather than distinct sub-markets, which is particularly useful when used as the basis for spatial policies and reporting.

Alternative definitions of local labour markets are also available and may describe heterogeneous views on the labour market. These heterogeneous approaches tend to stress the existence of sub-markets (van der Laan & Schalke, 2001), or internal labour markets in which, even in a relatively small region, it is probable that employers who are spatially separated will have differing labour markets to some extent (Goodman, 1970). Indeed, by considering the labour market from different points of view, for example employers (e.g., Robinson, 1968), or workers (e.g., Goldner, 1955), it is again possible to arrive at entirely different definitions of local labour markets. In the case of employers for instance, a local labour market will generally be defined as a "geographical area containing those members of the labour force, or potential members of the labour force that a firm can induce to enter its employ under certain conditions, and those other employers which whom the firm is in competition for labour" (Robinson, 1968, p. 66). Conversely, the local labour market of an individual worker will likely be much more limited, consisting entirely of those jobs "about which he hears, preferably via a trusted grape-vine, and which meet his preconceptions of his ability to obtain and retain them, and some platform of satisfaction or expectation of improvement" (Goodman, 1970, p. 183). In general, these heterogeneous labour markets differ from their homogeneous counterparts in that they are not necessarily delineated on the grounds of the entire labour supplyand as such need not be related to the labour force in neighbouring regions. Goodman (1970) and van der Laan & Schalke (2001) cite practical examples of both heterogeneous and homogeneous local labour markets, though in general, a homogeneous local labour market is preferred, particularly in the context of spatial policies or research aimed at specific cities, locales, or movements of labour.

Regardless of the type of local labour market definition being used, delineation of local labour market boundaries has tended to be based on analyses of travel-to-work data, using commuting information (e.g., commuting distances) to determine potential trade-offs between housing and employment locations (e.g., Bhat & Guo, 2004; Rodriguez, 2004). Various theories of local labour market dimensions have been interpreted in the context of commutingand commuting research has often focused on a range of population groups and subgroups. For example, disaggregate data on worker gender has been used to distinguish between spatial choices made by male and female workers in terms of residential location (e.g., White, 1977; Singell & Lillydahl, 1986). Similarly, theories of spatial choice (e.g., Singell & Lillydahl, 1986), and interaction (e.g., Vermeulen, 2003) have been used to help understand the distances people are prepared to travel to bridge home and work, or the factors that influence an individual's choice of a place of residence or employment.

One can conclude from reviewing the research and definitions employed in local labour market delineation that there remain a number of competing alternatives and that agreement on the most appropriate form is unlikely and, indeed, seldom discussed (Hanushek, 1981). This likely stems from the fact that while the concept of a local labour market is generally well understood, any attempt to properly define it in terms of real world economic activity tends to come from *ad hoc* assessments of the situation or data at hand, rather than from generalisable concepts applicable to a range of applications and environments. However, while differences in local labour market definitions exist, some consensus does arise in practice, as the practical use of local labour markets invariably requires a homogeneous view on labour supply, primarily because it must be based on the integration of markets in a spatially confined region (van der Laan & Schalke, 2001).

2.2 Functional regions

A homogeneous local labour market is normally influenced by the direct and indirect interactions between places. These interactions can be conceived as functional interdependencies between local regions and, in the context of travel-to-work, relate to the behaviour of individual commuters. As such, local labour market boundaries are often based on functional regions, which represent the aggregate travel-to-work patterns of the local population. A number of procedures for operationalising various functional region definitions have been suggested (e.g., Masser & Brown, 1975; Slater, 1981; Coombes et al., 1986; Flórez-Revuelta et al., 2008), the most successful likely being that of Coombes et al. (1986), which is based on the work of Smart (1974), and to some extent Goodman (1970). This procedure has been positively evaluated by Eurostat (1992) against various other regionalisation procedures (Coombes & Casado-Díaz, 2005). The primary units of interest are Travel-to-Work Areas (TTWA), which are a classification of smaller areal units into larger regions based on commuter interactions. The aim of this procedure is to define as many TTWAs as possible, subject to certain statistical constraints which ensure that the regions remain statistically and operationally valid (Coombes & Casado-Díaz, 2005). In practise, a minimum level of supply- and demand-side self-containment of 75% is used, with a minimum number of internal workers set at approximately 3,500 working individuals. In cases of highly populated areas ($>$ 20,000 working individuals), the minimum level of self-containment can be reduced. Similar functional regionalisation procedures have been employed in the Netherlands (van der Laan, 1991) and other parts of Europe (Eurostat, 1992), and more recently in Spain (Casado-Díaz, 2000), and New Zealand (Papps & Newell, 2002).

While the Coombes et al. (1986) procedure has been extensively evaluated in many contexts, there remain several viable alternatives. These techniques can generally be broken down into three basic classes: 1) Numerical taxonomy techniques, 2) Multi-stage aggregation procedures, and 3) Central place aggregation methods. Functional regionalisation procedures based on numeric taxonomy principals were initially developed in the 1970s and 80s, and were introduced as alternatives to more *ad hoc* methods for defining functional regions. Such methods include the Markov chain analysis techniques of Brown & Horton (1970) and Brown & Holmes (1971), as well as the strategy of Masser & Brown (1975; 1980), which is based on refinements to Ward's (1963) hierarchical aggregation procedure. Alternatively, Slater (1975; 1976; 1981) developed techniques based on fitting a hierarchical structure to a (standardised) matrix of flows by sequentially identifying the strong components of the corresponding directed graph. More recent work by Cörvers et al. (2009) has also employed a similar procedure to that of Brown & Holmes (1971). Despite their wide-ranging applicability and objective nature, some have criticised these early procedures as being overly deterministic and based too heavily on statistical objectives (Coombes et al., 1986), leading to the development of techniques bounded in behavioural and economic theory, such as the various multi-stage aggregation procedures.

Multi-stage aggregation procedures include the original Coombes et al. (1986) algorithm and its variants, as well as several more recent techniques, including the procedures of van der Laan & Schalke (2001) and Konjar et al. (2010). This latter work actually presents a refinement of previous work on central place aggregation by Karlsson & Olsson (2006), in which three separate functional regionalisation procedures are presented based on the local labour market, commuting zones, and worker and employer accessibility. A similar approach based on the local labour market procedure of Karlsson & Olsson (2006) is employed in Drobne et al. (2010), where urban centres are chosen based on their national and international importance. While this requires central places to be specified *a priori*, it has the benefit of being based on classic urban economic modelling (van der Laan & Schalke, 2001). An additional example of this type of functional regionalisation is given by the Metropolitan Statistical Areas (MSA) in the U.S., which delineate a core urbanised region surrounded by regions showing a high degree of social and economic integration with the core as measured by commuting data (OMB, 2000). Building on the concept of aggregating regions to urban centres is the work of Hensen & Cörvers (2003), which uses a spatial interaction model based on Alonso's General Theory of Movement (GTM) to determine 'attractive' municipalities to which surrounding municipalities are assigned. Another example of model-based functional regionalisation is the interregional interaction model (IRIM) of Noronha & Goodchild (1992), which utilises the concept of functional distance to partition the U.S. into two functional regions based on student migration.

One thing that becomes immediately apparent when considering many of the regionalisation procedures developed to date, is the seemingly arbitrary choice of threshold values and parameters. These threshold values will largely determine the size and number of functional regions found in a given regionalisation and can greatly influence the results of any subsequent analyses. For example, earlier attempts to develop regionalisation algorithms for the generation of TTWAs (and by proxy, functional regions) made no attempt to theoretically defend their choice of arbitrary values, suggesting that the choice of a 75% level of self-containment "…could be defended as lying exactly half-way between perfect self-containment and a level of 50% which seems a reasonable minimum for thinking of an area in labour market terms at all" (Smart, 1974, as quoted in Coombes & Openshaw (1982)). While some authors have subsequently provided defensible arguments for the use of fixed parameter values (e.g., Coombes et al., 1982, 1986), and in some cases have reduced the required number of threshold values (e.g., Coombes & Bond, 2007), others have argued against the use of situation-dependant absolute threshold values (e.g., Cörvers et al., 2009), or have suggested alternatives, such as using relative instead of absolute values (e.g., van der Laan & Schalke, 2001). Despite these attempts, there remain few alternative methods in the geographical literature which do not rely to some degree on arbitrary criteria and even fewer still being used in practise. It is likely this reliance on absolute threshold values that has limited the application of disaggregate functional regionalisations in the past, as it is difficult to determine exactly which threshold values to use for a particular population sub-group.

An additional problem with many functional regionalisation procedures is that they provide little or no theoretical basis on which to validate a particular regionalisation. For example, many researchers have relied on subjective assessments of the configuration of functional regions, often based on the authors' perceptions of local environments and specific application contexts to determine the number and configuration of functional regions. As Noronha & Goodchild (1992) point out, it is difficult to maintain confidence in assessments of validity such as "these results agree *fairly well* with the regions used by the government for planning and statistics..." (Hollingsworth, 1971, emphasis added by Noronha & Goodchild (1992)), or "the functional regions produced by the approach outlined above . . . conformed *extremely closely* to an *intuitive knowledge* of the study area" (Brown & Holmes, 1971, emphasis added by Noronha & Goodchild (1992)). Furthermore, due to a lack of any theoretical foundation, it is impossible to "[. . .] know whether a given method of analysis has indeed recovered the regions that exist, or whether the results are artefacts of the spatial setting, noise, or the method of analysis itself" (Noronha & Goodchild, 1992).

One reason for the above problems is that often the objectives of the regionalisation procedure use to delineate local labour markets are not the same as the objectives of the original local labour market definition on which they are meant to be based. Indeed, more often then not, regionalisation methods have been based purely on the principals of numerical taxonomy and/or the statistical properties of the data, rather than considering the behavioural characteristics of the underlying process (i.e., commuting) (Noronha & Goodchild, 1992). As suggested by Coombes et al. (1986), this leads to methods which tend to be too deterministic and unable to cope with the stochastic nature of human behaviour. While the shortcomings of many of these procedures are easily pointed out, some exceptions can be found (e.g., Coombes et al., 1986; Casado-Díaz, 2000; van der Laan & Schalke, 2001; Watts, 2004), where some attempt to consider the direct and indirect patterns of commuters have been employed. However, as noted above, these regionalisation procedures have tended to rely on arbitrary economic or geographical criteria and the flows *between* regions, while failing to recognize a crucial geographic variable in interactions *within* regions: "the friction of distance" (Noronha & Goodchild, 1992).

2.3 An alternative

A functional region can be thought of as the result of the collective travel-to-work behaviour of all (or a sub-group of) individuals in a particular region. As such, the underlying process governing the development and maintenance of a functional region is fundamentally a question of spatial interaction: movements and communication over space that result from a series of decision processes (Fotheringham & O'Kelly, 1989). This spatial interaction can be thought of as a network of flows (interactions) between units of interest. In the language of networks then, each vertex in a network represents a region of interest (e.g., census area, town, county, etc.) and a pair of vertices can be connected by an edge which represents the level of interaction (flow) between them (e.g., number of commuters). Additional information on the nature of the interactions can be represented by the degree (number of neighbours) and strength (sum of connected flows) of individual vertices and the direction and weight of individual edges on the network. Treating travel-to-work patterns as a network of flows is not by any means a new concept in functional regionalisation procedures (e.g., Slater, 1981); however, it is one that is often overlooked in the geographical literature.

There is, however, a growing literature in Physics and related fields on methods designed specifically for finding groupings or clusters in network-based data, some of which are suited to the analysis of functional regions (e.g., Girvan & Newman, 2002; Newman, 2003; Palla et al., 2005). Many of these formulations, however, have objective goals or functions which are not compatible with our general understanding of functional regionalisation procedures (see for example the suggested 'best practices' for selecting a regionalisation procedure in Eurostat (1992)) and are therefore of no use to geographers in this context. Conversely, there have been a number of methodological advancements in the statistical physics (and related fields) literature following the work of Girvan & Newman (2002), focusing on the idea of finding 'communities', or groupings, in various types of networks. What these methods are designed to measure is the degree to which a network displays community structure: the (natural) division of network vertices into groups, where within group connections are dense and between group connections are more sparse (Newman & Girvan, 2004). Generally, community structure is determined based solely on the information encoded in the network itself, rather than additional exogenous attributes. Thus, the number and size of the communities are determined by the properties of the network itself. This concept of community structure is intuitively similar to our own definition of functional regions, where vertices represent regions and edges define the flows between regions. As such, this line of inquiry presents an excellent opportunity to import methods from outside the geographical literature to solve a fundamentally geographical problem. An example of a network with clear community structure is given in Figure 2.1.

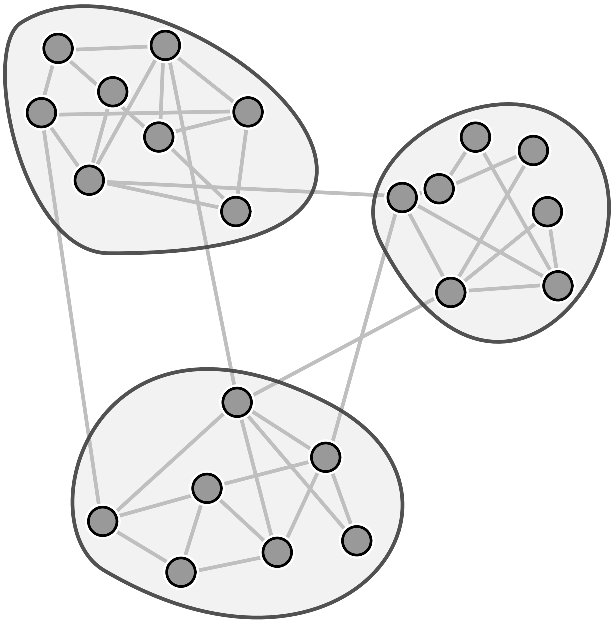

Figure 2.1: Example of a network with clear community structure (after Girvan & Newman (2002, p. 7821, fig. 1).

2.4 Modularity

One of the most widely used methods for quantifying the 'quality' of a particular partition of a network into communities is the modularity quality function, Q, of Newman & Girvan (2004). In the context of functional regionalisation, modularity is based on the notion that a good division of a network into functional regions is not necessarily one in which there are few flows between functional regions, but rather one in which there are fewer than expected flows between functional regions (Newman, 2006a). As such, the modularity function is designed to find statistically surprising spatial interaction by measuring the difference between the total fraction of flows that fall within functional regions and the expected fraction of such flows based on some null model (Newman, 2006a; Leicht & Newman, 2008; Porter et al., 2009). In general, the null model is based on the so-called configuration model, which is a random network conditioned on the count and weighting of the flows in the input network. Thus, given an observed trip matrix $\mathbf{T}$ (see Figure 3.1 for further information), representing the connections of a network of travel-to-work flows, modularity can be computed as

$$Q = \frac{1}{2T} \sum_{ij} \left[T_{ij} - \frac{t_i t_j}{2T} \right] \delta_{c_i, c_j}, \tag{2.1}$$

where T_{ij} is an element of the observed trip matrix $\mathbf{T}$, and the Kronecker delta function, δ, takes the value 1 if regions i and j are in the same functional region ($c_i = c_j$) and 0 otherwise. The expected flow magnitude between regions i and j is given by $t_i t_j / 2T$ where t_i and t_j are the magnitudes of the flows associated with regions i and j and $T = 1/2 \sum_{ij} T_{ij}$ is half the total flows for the entire network. This is akin to a χ^2 goodness-of-fit test and means that Q measures the normalised sum of deviations from a (conditioned) random network. In other words, the larger the value of Q, the more the observed network deviates from our null model. This formulation can be extended to include flow direction (common in travel-to-work data) by considering that $t_i^{in} t_j^{out} / T$ is the expected flow magnitude between j and i, where t_i^{in} and t_j^{out} are the separate incoming and outgoing flows of the regions (Leicht & Newman, 2008).

The goal is to find the configuration of functional regions that maximises Q (i.e., maximises deviation from our null model). In practice however, this problem has been shown to be computationally hard and therefore some form of heuristic is required to approximate modularity maximisation (Danon et al., 2005). A number of heuristics have been suggested in the literature, each differing in terms of the mathematical, statistical, and/or computational techniques employed (Porter et al., 2009). In essence, heuristics for maximising modularity are designed to balance the quality of the regionalisation with speed and computational costs. Examples include several 'greedy' algorithms which are designed to quickly maximise modularity in extremely large networks of millions of nodes (Newman, 2004b; Clauset et al., 2004), as well as more 'accurate' but slow algorithms such as, *inter alia*, spectral partitioning (e.g., White & Smyth, 2005; Newman, 2006a,b; Leicht & Newman, 2008), simulated annealing (e.g., Guimerà & Amaral, 2005; Guimerà et al., 2005; Reichardt & Bornholdt, 2006), and extremal optimisation (e.g., Duch & Arenas, 2005).

The spectral partitioning techniques mentioned above are particularly effective methods for heuristically maximising modularity and have the benefit of being computationally straight-forward. As given in Newman (2006a), modularity can reformulated in terms of a so-called modularity matrix $\mathbf{B}$. This makes it possible to take advantage of the spectral properties of the network's weighted adjacency matrix to heuristically optimise modularity via spectral partitioning. Spectral partitioning in its most basic form is the splitting of a network into two groups based on the properties of the network's Laplacian matrix L, with elements $L_{ij} = t_i\delta_{ij} - T_{ij}$ (where all variables are as previously defined, and $\delta_{ij} = 1$ if $i = j$ and 0 otherwise), and an index vector s, with $s_i = +1$ if the associated vertex belongs to group A, and $s_i = -1$ if the vertex belongs to group B. The elements of s are chosen such that the total flow between the two groups, $R = 1/4s^T Ls$, is minimised (see Porter et al., 2009). In the simplest case, the s_i are based on the signs of the elements of the leading eigenvector v of the network's adjacency matrix, such that $s_i = \mathrm{sgn}(v_i)$. Two-group partitioning is recursively applied to the smaller groups until some stopping rule is reached. For a comprehensive review of spectral partitioning methods, see Ng et al. (2002) and Spielman & Teng (2007). See also Newman & Girvan (2004) and Porter et al. (2009) for a presentation of spectral partitioning in the context of modularity maximisation.

When using the modularity matrix to optimise Q, the elements of the matrix are given by

$$B_{ij} = T_{ij} - \frac{t_i t_j}{2T}. \tag{2.2}$$

This, equation (2.1) can be reformulated using matrix notation as

$$Q = \frac{1}{4T} s^T \mathbf{B} s, \tag{2.3}$$

for a given partition of the dataset into two functional regions. Again, this can be extended to directed networks (which will be the focus for the remainder of this thesis) by symmetrising $\mathbf{B}$ using the procedure presented in Leicht & Newman (2008) to obtain

$$Q = \frac{1}{4T} s^T \left(\mathbf{B} - \mathbf{B}^T\right) s. \tag{2.4}$$

The matrix $\mathbf{B} - \mathbf{B}^T$ is now symmetric, though it is not the same as simply symmetrising the network itself and, as such, should yield more relevant results than simply ignoring flow direction. If the goal is to further subdivide the network into more than two functional regions (as is often the case in functional regionalisation exercises), the partitioning procedure can be applied recursively to each smaller functional region, keeping in mind that each functional region is actually part of a larger network of commuting flows (Leicht & Newman, 2008). This requires a generalisation of the previous modularity function so that we take into account the *change* in modularity of the overall network based on the division of a given functional region, g. Change in modularity is therefore given as

$$\Delta Q = \frac{1}{4T} s^T \left(\mathbf{B}^{(g)} - \mathbf{B}^{(g)T} \right) s, \tag{2.5}$$

where we take advantage of the fact that $s_i = 1$ for all i (s_i is either -1 or $+1$) and specify a generalised modularity matrix as

$$B_{ij}^{(g)} = B_{ij} - \frac{1}{2} \delta_{ij} \sum_{k \in g} (B_{ik} + B_{ki}), \tag{2.6}$$

such that $\mathbf{B}^{(g)}$ is the sub-matrix of $\mathbf{B}$ for the functional region g, with the average of the corresponding row and column sums subtracted from each diagonal element (Leicht & Newman, 2008). Note that modularity is in a sense a 'global' statistic because it measures the *overall* modularity of the network based on changes in an individual functional region. While it is possible that a globally optimal solution will lead to optimal local solutions, this is not necessarily a requirement. This is quite different from previous functional regionalisation procedures, where local self-containment and populations drive the shape and size of individual functional regions.

While the above formulation for maximising modularity produces relatively good divisions of the network, it is often beneficial to also include a 'fine-tuning' step (Newman, 2006b), where modularity is further maximised by using, for example, a simple KL-style procedure (see Kernighan & Lin, 1970)). Additional constraints on the algorithm may be employed if necessary. For example, it may be argued that similar self-containment restrictions to those found in the TTWA procedures may be necessary in a particular context, such as when issues of data collection and policy making require a minimal level of supply- or demand-side self-containment (e.g., Eurostat, 1992). In these cases, it is relatively straight-forward to add self-containment checks to the fine-tuning stage of the algorithm, at which point, if a certain level of self-containment is not reached, the split is rejected. While these additional constraints will invariably require additional parameters, in some cases the increased level of control will offset the disadvantages associated with additional parameters. Regardless of the fine-tuning procedure used in practice, the dual optimisation strategy has been shown to produce good results (Newman, 2006a) and is based on the fact that the spectral clustering approach provides a general specification of the network split, but the additional fine-tuning step allows us to reach the best possible modularity value. With this in mind, the general modularity maximisation procedure is as follows:

1. Construct the modularity matrix as in (2.2),
2. find the leading (most positive) eigenvalue of the symmetric modularity matrix, $\mathbf{B} - \mathbf{B}^T$,
3. divide the network into two communities based on the signs of the elements of the corresponding eigenvector,
4. fine-tune the results in order to further maximise modularity (note that additional fine-tuning conditions may also be applied here),
5. repeat the above for each community of the divided network (recursively), now using the generalised modularity matrix of (2.6),
6. continue dividing a given community until a proposed split makes a zero or negative contribution to the overall modularity,
7. once all communities can no longer be split without reducing the overall modularity, the algorithm ends.

From modularity, we now have a suitable definition of a functional region: a geographical region in which within-region interaction in terms of commuters' travel-to-work flows is maximised and between-region interaction is minimised. This definition and implementation benefits from several ideal properties, including: 1) it is based on statistical criteria: the structure of the observed network of flows is compared to a hypothetical random network, such that statistically surprising interactions are found and utilised to delineate the functional regions, 2) the procedure provides statistical justification for the computed arrangement of functional regions, halting further subdivisions when the interactions within a region are no longer significantly different from random, 3) the procedure is exhaustive, ensuring that all base units belong to one and, only one, functional region, 4) the procedure requires no tuning parameters (though these can be added if further control is required, see Section 5.2.1 and Section 5.2.2 for details), providing a consistent set of criteria on which to compare functional regions derived for different datasets and areas, and 5) because the procedure is not reliant on *a priori* information, it can be easily applied to the commuting patterns of minority groups, socio-economic sub-groups, and/or males and females separately, providing a range of disaggregate functional regionalisations for comparison.

2.5 Moving forward

The regionalisation procedure presented in this chapter provides a viable alternative to legacy functional regionalisation procedures which require the use of *ad hoc* tuning parameters and threshold values. The procedure benefits from a data-driven objective function, while still allowing for fine-grained control when needed. Furthermore, the large breadth of research surrounding modularity, community structure, and network theory means that our functional regionalisation procedure is built upon solid theoretical and empirical grounds. By tying in concepts from labour markets, network theory, and spatial interaction, our modularity-based functional regionalisation procedure can be an effective tool for exploring local labour market structure and delineating local labour market boundaries. In the remainder of this thesis, we utilise the theory and methodology behind modularity optimisation as the basis for a tangible definition of functional regions and, by proxy, local labour markets. In particular, we take advantage of the ideal properties of the functional regionalisation procedure initialised in this chapter and extend the procedure to take into account some of the issues and opportunities inherent in travel-to-work data. In the following chapter, we explore the concept of spatial interaction and chart the theoretical and methodological developments of this modelling regime in the context of commuting flows. Coupling what we learn in Chapter 3 with the local labour market theories touched on in this chapter, we move forward to the empirical stages of this thesis in Chapters 5 and 6 with a solid base upon which to develop our modelling framework for investigating the effects of the local labour market on commuting patterns.

Chapter 3

Spatial interaction modelling

3.1 Introduction

According to Fischer (2001, p. 195), "one of the major intellectual achievements and, at the same time, perhaps the most useful contribution of spatial analysis to social science literature is the development of spatial interaction models". Spatial interaction can be defined in general terms as the movement of individuals, commodities, capital, and information over (geographic) space resulting from a decision process (Batten & Boyce, 1986; Fotheringham & O'Kelly, 1989). Thus, spatial interaction in broad terms encompasses research into migration, shopping, recreation, commodity and capital flows, communication, transportation networks, and commuting (Haynes & Fotheringham, 1984). The fundamental principal underlying these types of interactions is that in each case, individuals trade off the benefit of the interaction (e.g., travel-to-work), with the costs associated with overcoming the spatial separation between the individual and their possible destination (Fischer, 2001). It is this core concept that has made theories of spatial interaction and, by proxy, spatial interaction modelling, important in terms of understanding spatial behaviour across a range of topics.

A model in its most general form is a simplified representation of an object or phenomenon of investigation for purposes of description, *explanation*, forecasting, or planning (Wegener, 2000). Similarly, a spatial model is a model extended to the spatial domain, such that a representation of both spatial and non-spatial (aspatial) attributes is developed. According to Wegener (2000), there are three categories of spatial models used in the literature, with each category differing in terms of their degree of formalisation. The most simplified form of spatial model, scale models, are a representation of real-world physical features such as the natural terrain or a transport network. Conversely, conceptual models are used to highlight the differences and abstract linkages between the components of the system or phenomenon under investigation. Thirdly, mathematical models, which tend to be the most complex form of spatial models, attempt to operationalise conceptual models through mathematical representations. In this sense, complex relationships and interactions can be represented and explained using mathematical constructs based on theories of spatial interaction and spatial choice.

3.2 Mathematical models

As suggested by Fotheringham & O'Kelly (1989), mathematical models are generally used for two primary purposes: explanation and prediction. In the context of commuting behaviour for example, we may be interested in *explaining* observed commuting patterns (i.e., understanding the determinants of commuting), *predicting* future commuting patterns (i.e., replicating commuting under difference scenarios), or both (i.e., using our understanding of commuting in a larger context, such as in urban models). In this current research, we are primarily interested in understanding the determinants of commuting and so we focus on explanation over prediction, though these goals are often not mutually exclusive.

Analysis of commuting and/or travel-to-work patterns has tended to focus on mathematical models of both individual and aggregate commuting behaviour and have emphasised the use of commuting data (i.e., commute distances) for determining trade-offs between housing and employment locations (e.g., Bhat & Guo, 2004; Rodriguez, 2004). In addition, work has involved the use of disaggregate data on worker gender to distinguish between spatial choices made by male and female workers in terms of residential location (e.g., White, 1977; Singell & Lillydahl, 1986). There is also a wealth of research on commuting patterns in large metropolitan areas of the US (e.g., Gordon et al., 1991; Giuliano & Small, 1993; Wachs et al., 1993; Taylor & Ong, 1995; Shen, 2000), and Europe (e.g., Warnes, 1972; Frost et al., 1998; Findlay et al., 2001; Giuliano & Narayan, 2003; Moss et al., 2004) which focus on variations in commute times among population groups in models of travel-to-work. In addition, a large number of analyses involve explicitly examining the impact of urban structure on commuting patterns (e.g., Giuliano & Small, 1993; Giuliano & Narayan, 2003), and more specifically, how issues in urban commuting can be resolved through examining the commuting patterns of minority groups (e.g., Gabriel & Rosenthal, 1996; Owen & Green, 2000).

Outside commuting research, the increasing use of mathematical models to explain spatial phenomenon such as diffusion of innovations (e.g., Rogers, 1993; Valente, 1996; Allaway et al., 2003), migration (e.g., Pellegrini & Fotheringham, 1999, 2002), and retail geography (O'Kelly, 1981; Clarke & Madden, 2001; O'Kelly, 2009, e.g.,) further attests to their utility as explanatory tools. Within the category of mathematical spatial models, there are a number of subcategories, many of which are separated on the basis of how a model deals with various aspects of the real world (i.e., Euclidean versus social or economic space), the underlying theoretical framework, or the specific modelling techniques applied. In order to properly understand the choices a commuter makes and the implications that this has for the local labour market, comprehensive models of travel-to-work behaviour based on a range of relevant variables can be developed. These models should attempt to take into account the economic, social, educational, and employment characteristics of individuals, as well as the spatial context within which individuals operate. Using this information, researchers can make predictions about which individuals are most likely to work in a particular region, industry, or both, why they are likely to seek employment in any particular industry or locale and, by proxy, how this leads to the underlying structure of the local labour market.

3.3 Spatial interaction modelling

Spatial interaction modelling as a separate research endeavour developed out of a need to mathematically model and understand the movements of individuals, information, and commodities in space. Proponents of spatial interaction modelling include economic geographers, regional scientists, and regional planners, as well as researchers studying topics in transportation, migration, diffusion, retailing, marketing, and commuting. Originally developed from theories of interacting particles and gravitational forces in physics, spatial interaction modelling has developed through a series of refinements in terms of functional form, conceptual representations of economic and spatial distances, as well as range of analytically rigorous technical improvements.

The phenomena of interest in most spatial interaction research are interactions of 'actors' between a set of 'origin' and 'destination' spatial units. Since our primary interest in this thesis is commuting, we are concerned with the flows of individuals travelling to/from home and work. As such, the relevant actors are individual commuters, and their homes and workplaces are their origins and destinations respectively. As in nearly all spatial interaction analyses, the relationship between commuter origins and destinations can be represented using some form of spatial interaction, or origin/destination (OD) matrix like the one shown in Figure 3.1. This type of interaction matrix, $\mathbf{T}$, usually represents flows of actors (commuters), T_{ij}, between n origins (home) and m destinations (work). By summing the flows across each row of the interaction matrix, we obtain the observed outflow from each origin, O_i, and similarly, by summing each column of the interaction matrix, we obtain the observed inflow into each destination, D_j. The sum of all flows in the interaction matrix, $T = \sum_{ij} T_{ij}$, therefore, represents the overall level of interaction in the OD matrix.

$$\mathbf{T} = \quad \begin{array}{c} \begin{bmatrix} T_{11} & \cdots & T_{1j} & \cdots & T_{1m} \\ \vdots & & \vdots & & \vdots \\ T_{i1} & \cdots & T_{ij} & \cdots & T_{in} \\ \vdots & & \vdots & & \vdots \\ T_{n1} & \cdots & T_{nj} & \cdots & T_{nm} \end{bmatrix} \begin{array}{c} O_1 \\ \vdots \\ O_i \\ \vdots \\ O_n \end{array} \\ \\ \begin{array}{ccccccc} D_1 & \cdots & D_j & \cdots & D_m & & T \end{array} \end{array}$$

Figure 3.1: Example spatial interaction (origin/destination) matrix. The row, column, and total sums (O_i, D_j, and T) represent the origin outflows, destination inflows, and overall level of interaction of the matrix respectively.

In the following sections, we briefly examine the methodological developments of spatial interaction modelling, drawing on insights from Haynes & Fotheringham (1984, chap. 4, p. 40-48), Fotheringham & O'Kelly (1989, chap. 2, p. 15-35), Fotheringham et al. (2000a, chap. 9, p. 213-234), and Roy (2004, chap. 1, p. 10-18). Through this structure, we trace the development of spatial interaction as a modelling framework, starting with a simple analogy to Newtonian gravitational theory from the physical sciences. We will then outline the progressive methodological developments of spatial interaction theory, highlighting the derivation of a general 'family of spatial interaction models' and some important variations. Following this, we will trace the development of a more theoretical basis for spatial interaction models based on behavioural and competing destinations theories. Finally, we finish off with some mention of spatial information processing, discrete choice, and the concept of spatial choice.

3.4 Spatial interaction theory

3.4.1 Gravity models & social physics

Perhaps the most well known early attempt to characterise observed regularities in spatial interaction data is that of Ravenstein (1885). In his early study, Ravenstein observed that large cities tended to draw migrants from other large cities, and that this effect appeared to drop off with distance (i.e., the number of individuals migrating from one city to another tended to be inversely proportional to the distance between the two cities). This basic observation lead to the development of a mathematical model for predicting migration flows between origins and destinations based on the Newtonian gravity model,

$$T_{ij} = k\frac{P_i P_j}{d_{ij}}, \tag{3.1}$$

where T_{ij} again denotes the flows between origin i and destination j, P_i and P_j represent the sizes of locations i and j, d_{ij} describes the distance between i and j (which is not required to be Euclidean distance), and k is a scaling parameter used to adjust the magnitude of T_{ij} relative to $P_i P_j / d_{ij}$. Subsequent refinements to this simple model allowed for the specification of P_i and P_j as the populations of i and j (e.g., Stewart, 1941). Realisation that the effect of distance and population may in fact vary depending on the context and type of flows being investigated (i.e., the 'distance decay' and population size effect will likely be different for migration than it would be for shopping or travel-to-work), lead some researchers to empirically estimate additional parameters to allow for these variations

$$T_{ij} = k\frac{P_i^{\alpha} P_j^{\lambda}}{d_{ij}^{\beta}}, \tag{3.2}$$

where β is a distance decay parameter which represents the degree to which estimated values of T_{ij} decrease with distance, and μ and α are parameters reflecting the relationship between T_{ij} and P_i and P_j respectively.

Various forms of (3.2) have been used and refined over the years (e.g., Ravenstein, 1885; Reilly, 1929; Stewart, 1941, 1942), and while the term 'gravity model' has remained popular, the model itself was widely criticised for its simplicity and overall lack of theoretical grounding. Indeed, the general model made no attempt to characterise the behaviour of the individuals it was designed to model. This lead some researchers to attempt to develop an acceptable theoretical framework for the gravity model in order to take advantage of its predictive power and simplicity (Fotheringham et al., 2000a). Some notable early attempts include Stouffer (1940), Zipf (1949), Dodd (1950), and Huff (1959). The sociological slant of Dodd (1950), in which the deterministic gravity model was given a more probabilistic base, and the principal of least effort put forward by Zipf (1949) were particularly relevant modifications to the standard gravity formulation.

According to Haynes & Fotheringham (1984), it was not until the work of Huff (1959, 1963, 1964) on consumer behaviour, and the concept of intervening opportunities of Stouffer (1940, 1960), that a truly behavioural interpretation of the gravity model was formulated. In particular, Huff's probabilistic retail model focused attention on the choice options of the *shopper* (Roy, 2004), rather than simple competition between *retailers*. Similarly, Stouffer attempted to explain the movements of migrants based on the number of opportunities at a given destination, keeping in mind that there may be intervening opportunities which might attract a migrant along the way. Additional advances in the form of a sequence of comments and replies (Niedercorn & Bechdolt Jr, 1969, 1970, 1972; Mathur, 1970; Allen, 1972) in the Journal of Regional Science surrounding a paper entitled "An economic derivation of the "Gravity Law" of spatial interaction" (Niedercorn & Bechdolt Jr, 1969), also provided a useful derivation of the gravity model, this time from economic principals and concepts of utility maximisation. This conceptual advance suffered from several important shortcomings however, including the fact that it was attempting to describe aggregate outcomes from individual level behaviour, a problem shared by the models of Huff (1959), and Stouffer (1940). Despite these shortcomings, it is somewhat surprising that the work of Niedercorn & Bechdolt Jr (1969), and perhaps more appropriately Niedercorn & Bechdolt Jr (1972), was not further pursued in the geographical literature given its economic/behavioural interpretation and applicability to the large body of literature on utility maximisation. A possible reason for this stems from the fact that the work of Wilson (1967, 1970, 1971) provided an entirely new way of looking at spatial interaction models from the perspective of statistical mechanics and addressed the issue of individual versus aggregate outcomes that previous modelling frameworks were unable to accommodate.

3.4.2 Entropy & spatial interaction

The pioneering work of Wilson (1967, 1975), produced a 'family of spatial interaction models' useful for a range of spatial interaction problems. Wilson's framework was particularly important because it attempted to provide a theoretical justification for spatial interaction models, albeit one not based on human behavioural properties, but rather on the statistical mechanics of the model and was a catalyst for a great deal of research on spatial interaction modelling in general (Fotheringham, 2001). Similar to early gravity models, Wilson's framework starts at the aggregate (macro) level of flows; however, Wilson then takes this a step further by working backwards to the individual (micro) level, in order to provide an explanation of the total observed interactions of the system (Haynes & Fotheringham, 1984; Roy & Thill, 2003; Roy, 2004). Using the language of statistical mechanics, Wilson considered a spatial interaction system in which the flows of individuals (e.g., commuter flows) between origin/destination pairs is simply a 'macrostate' of the overall system and that the individuals moving between origins and destinations (e.g., commuters) are individual 'microstates' which combine to produce the observed macrostate. Clearly, there are a number of different combinations of microstates that could lead to a given macrostate, such that

$$R = T!/\prod_{ij} T_{ij}! \tag{3.3}$$

is the number of microstates for a given macrostate, where T is the total flow of the spatial interaction system and T_{ij} is the distribution of flows (states) in the system. In other words, R is the number of ways that T actors can be allocated from origins i to destinations j. Based on the concept of entropy maximisation (or alternatively based on information minimisation, see Fotheringham & O'Kelly (1989)), the problem is to pick the distribution T_{ij} which maximises R (i.e., choose the macrostate which can be constructed from the largest number of microstates). This is based on the assumption that each possible microstate of a given macrostate is equi-probable (Roy, 2004) and, unless additional information on the system is available, this is the least biased estimate of the true distribution (Fotheringham & O'Kelly, 1989). It is also the outcome with the highest uncertainly and is therefore consistent with an information minimisation approach as suggested by Fotheringham & O'Kelly (1989, and references therein). For ease of computation, we can take the natural logarithm of R divided by T (this will not alter the optimisation)

$$\begin{aligned} H &= (1/T)\ln R \\ &= (1/T)(\ln T! - \sum_{ij} \ln T_{ij}!), \end{aligned} \tag{3.4}$$

and, when all T_{ij} are large, Stirling's approximation, $\ln x! = x \ln x - x$, can be used to further simplify the objective function to produce

$$H = (1/T)\left(T \ln T - T - \sum_{ij} T_{ij} \ln T_{ij} + T\right). \tag{3.5}$$

As in Fotheringham et al. (2000a), H can be rearranged to produce

$$H = -\sum_{ij} (T_{ij}/T) \ln(T_{ij}/T), \tag{3.6}$$

which, by denoting the proportion of all trips from origin i to destination j, T_{ij}/T, as p_{ij}, is equivalent to an alternative definition of entropy introduced by Shannon (1948), $H = -\sum_{ij} p_{ij} \ln p_{ij}$. This alternative concept of entropy can be interpreted as a measure of the uncertainty of the distribution T_{ij}/T and is explored further in Fotheringham & O'Kelly (1989), Roy & Thill (2003), Roy (2004) and references therein.

Clearly, the maximisation of (3.6) without additional constraints will always yield a solution where all T_{ij} values are as close to equal as possible (i.e., the solution with the maximum entropy/uncertainty). In a spatial interaction setting, this is obviously not ideal, as we will likely have some additional information about the spatial interaction system that we can use to derive more accurate estimates of T_{ij}. Based on the notation in Figure 3.1, the type of information that we have on a spatial interaction system might include:

1. all outgoing flows of the system O_i, or
2. all incoming flows of the system D_j, or
3. both O_i and D_j, or
4. neither O_i nor D_j and, usually also
5. the observed total (or average) trip length.

Wilson recognised that the above information could be entered into (3.6) as constraints on the maximisation procedure, in order to derive a range of models which could be used to solve various spatial interaction problems. Fotheringham & O'Kelly (1989, chap. 2, p. 2-3) provide examples of various situations where the following constraints on the general entropy maximisation spatial interaction model might be used. These constraints include an origin constraint

$$\sum_j T_{ij} = O_i, \tag{3.7}$$

a destination constraint

$$\sum_i T_{ij} = D_j, \tag{3.8}$$

and a trip length constraint

$$\sum_{ij} T_{ij} \ln d_{ij} = C, \tag{3.9}$$

where d_{ij} is the distance between i and j and C is to total trip length (cost) of the system. In deriving his family of models, Wilson allows for zero, one, or two of the above constraints to be imposed on the interactions. In addition, further control on the model may be gained by incorporating information on the attraction and/or propulsion of origins and destination via the inclusion of additional variables (attraction factors) in the model (Cordey-Hayes & Wilson, 1971; Fotheringham & O'Kelly, 1989). As such, given the type of constraint(s) used and the included variables, we end up with different members of the family of spatial interaction models. For example, the *unconstrained* model is the result of maximising (3.6) subject to (3.9) and produces the expression

$$T_{ij} = v_i^{\mu} w_j^{\alpha} d_{ij}^{\beta}, \tag{3.10}$$

where β is found so as to satisfy (3.9), v_i and w_j are origin and destination attraction (propulsion) variables, and μ and α are parameters reflecting the relationship between T_{ij} and v_i and w_j respectively, which may be obtained in the calibration of the model. In terms of practical applications, unconstrained models are perhaps less useful than some of the following constrained models, owing largely to the fact that there is no constraint on the outgoing and incoming flows of the system, often leading to relatively poor estimates of T_{ij}.

The *origin-* and *destination-constrained* models are similar in form, with the difference being that an origin-constrained model is used to reproduce the observed outflows from each origin and allocate them to the various destinations, whereas a destination-constrained model takes as given the inflows to each destination and allocates the interactions to the various origins. The origin-constrained model is given by maximising (3.6) subject to (3.7) and (3.9)

$$T_{ij} = \frac{O_i w_j^{\alpha} d_{ij}^{\beta}}{\sum_j w_j^{\alpha} d_{ij}^{\beta}}, \tag{3.11}$$

and similarly, the destination-constrained variant is given by maximising (3.6) subject to (3.8) and (3.9)

$$T_{ij} = \frac{D_j v_i^{\mu} d_{ij}^{\beta}}{\sum_i v_i^{\mu} d_{ij}^{\beta}}. \tag{3.12}$$

These are both 'share' models, in that they allocate shares of the number of individuals leaving an origin (arriving at a destination), amongst the various destinations (origins) according to the attributes of said destinations (origins). In the destination-constrained case therefore, $v_i^{\mu} d_{ij}^{\beta} / \sum_i v_i^{\mu} d_{ij}^{\beta}$ is simply j's share of T_{ij}. Origin- and destination-constrained models are appropriate for situations where, for example, we are interested in understanding the characteristics of a destination (origin-constrained) that make it attractive to migrants (e.g., Fotheringham et al., 2000b), or we want to model and predict net interregional commuting (destination-constrained) for use in a larger urban modelling framework (e.g., Vermeulen, 2003).

The final model in Wilson's family of spatial interaction models is the *doubly-constrained* spatial interaction model and is perhaps more appropriate in situations where prediction, rather than explanation, is the primary goal of the modelling exercise. This is because both outflows from origins and inflows to destinations are taken as exogenously defined and the model simply allocates these flows to the links between origins and destinations. In other words, it reproduces the observed interactions to a high accuracy, without providing any information on the attractiveness or propulsive of the origins and destinations. As a result, doubly-constrained spatial interaction models are often used in trip distribution problems, where estimates of inter-zonal flows are required and the number of trips originating and ending in each zone are known *a priori* (Fotheringham & O'Kelly, 1989). The form of the doubly-constrained model derives from maximising (3.6) subject to (3.7), (3.8), and (3.9), producing

$$T_{ij} = A_i B_j O_i D_j d_{ij}^{\beta}, \tag{3.13}$$

where

$$A_i = \left(\sum_j w_j^{\alpha} B_j D_j d_{ij}^{\beta} \right)^{-1}, \tag{3.14}$$

and

$$B_j = \left(\sum_j v_i^{\mu} A_i O_i d_{ij}^{\beta} \right)^{-1}. \tag{3.15}$$

The (interrelated) balancing factors A_i and B_j ensure that the origin and destination constraints are met. In practice, these balancing factors are iteratively adjusted during the calibration of the model (see Williams & Fotheringham, 1984). If additional information on the distribution of T_{ij} is known, then the balancing factors may contain appropriate attraction and/or propulsion factors (w_j^α in (3.14) and v_i^μ in (3.15)), whereas, if no additional information is available, these variables may be omitted or set equal to one.

Extensions to Wilson's entropy models are plentiful, both from a theoretical and practical perspective (Wilson, 1975). Notable theoretical improvements include the Alonso (1978) framework, which takes the family of spatial interaction models a step further by deriving a general 'theory of movement' which provides a much more general framework within which the more 'traditional' family of spatial interaction models can be derived as special cases (Fotheringham & Dignan, 1984; Fotheringham & O'Kelly, 1989; O'Kelly, 2004). The Alonso framework has, in turn, inspired additional work on spatial interaction modelling. This includes the work of Fotheringham & Dignan (1984), who show that by using the Alonso framework as a starting point, each of the four traditional spatial interaction models covered in this section are in fact extremal points on a continuum of models (Fotheringham et al., 2000a). According to Fotheringham & Dignan (1984), an infinite number of 'quasi-constrained' spatial interaction models can be generated based on how strictly various constraints are enforced (see Fotheringham & Dignan (1984) and Fotheringham & O'Kelly (1989) for a detailed treatment of this work). The Wilson framework has also been extensively applied to real-world geographical problems, though much of this work is conducted within the private sector and therefore remains difficult to find in academic journals (Birkin et al., 2010). Of particular relevance is the work on retail/business geography conducted largely within the university of Leeds (e.g., Clarke et al., 1998; Clarke & Clarke, 2001; Nakaya et al., 2007; Birkin et al., 2010, and references therein). Additional examples are also available (e.g., Wilson, 2000; Singleton et al., 2010, *inter alia*) and characterise the utility of this modelling framework as an effective decision support tool.

Several criticisms of the Wilson framework (and related work) have been raised in the literature, including the fact that only some of the models derived using this framework have a behavioural interpretation. Additionally, many have argued that analogies to physical processes used in both gravity models and entropy-type models are not always productive, or indeed relevant when the subject of enquiry is fundamentally a question of human spatial behaviour (O'Kelly, 2004). Furthermore, because of their reliance on aggregate spatial behaviour rather than the choices of individuals, these models were thought to lack explanatory power when it came to understanding fundamental spatial behaviour (Timmermans & Golledge, 1990). While some of these issues persist, the Wilson framework remains a popular means of formulating spatial interaction models, as evidenced by the recent celebration of Wilson's "Entropy in urban and regional modelling" in a special issue of Geographical Analysis (Geographical Analysis, 2010). Indeed, as we will show in Section 3.5.1, modern Poisson spatial interaction models are essentially equivalent to entropy maximising models when appropriate constraints are applied.

3.4.3 Local models & distance decay

According to Fischer (2001, p. 195), since the pioneering work of Wilson, there has been very little innovation in terms of spatial interaction modelling in quantitative geography. One principal exception to this claim is the 'competing destinations' (Fotheringham, 1983a, 1986, 1987) framework for spatial interaction and spatial choice. The competing destinations model was initially presented theoretically by Fotheringham (1983a,b, 1984, 1986) and has subsequently been empirically verified in many settings (eg., Fotheringham & Trew, 1993; Pellegrini et al., 1997; Thorsen & Gitlesen, 1998; Guldmann, 1999; Gitlesen & Thorsen, 2000; Fotheringham et al., 2001, *inter alia*). In this section, we will outline some of the developments that lead to the specification of a competing destinations model and the implications this has for the interpretation of spatial interaction models.

Spatial interaction models of the type presented in Section 3.4.2 have enjoyed a long and productive history in quantitative geography. They have been used in a large number of theoretical and applied settings and have proven to be useful techniques for exploring and predicting human spatial behaviour. However, these models have traditionally been applied at a 'global' level, meaning that one set of model parameters is generated for the entire study area and, these parameters are assumed to apply equally well across the entire study region. This is akin to assuming that there is a single explanation for human spatial behaviour that applies equally well across, for example, the entire Republic of Ireland! If the spatial interaction system under investigation has spatially varying relationships (i.e., spatial non-stationarity), then these assumptions will likely be invalid, leading to poor predictions and locally misleading explanations. One method for 'discovering' spatial non-stationarity is to examine the spatial distribution of model residuals, with the assumption that any patterns in the residuals are the result of 'spatial effects'. This is common in the spatial econometric literature and is usually accompanied by some measure of spatial autocorrelation in the residuals.

There have been several calls in the quantitative geography literature for a more 'spatially local' outlook on spatial data analysis in general (e.g., Openshaw et al., 1987; Getis & Ord, 1992; Anselin & Getis, 1992; Fotheringham & Rogerson, 1993), and a change in focus from *similarities* between places to *differences* across space (Fotheringham, 1997; Lloyd, 2011). These calls have since spawned the development of a range of local spatial analysis methods [e.g., local point pattern analysis (e.g., Openshaw et al., 1987, and variants), local spatial autocorrelation (e.g., Getis & Ord, 1992; Ord & Getis, 1995; Anselin, 1995), and local regression models such as GWR (Brunsdon et al., 1996, 1998; Fotheringham et al., 1996, 2002)], particularly in the wake of advances in geographic information systems (GIS). While this trend towards local methods in quantitative geography is relatively recent, the idea of calibrating local spatial interaction models has long been recognised as an important means of revealing hidden spatial relationships concealed by the 'averaging' effects of global spatial interaction models (Fotheringham et al., 2000a).

In a global spatial interaction model, the system under consideration is usually made up of a number of origins and destinations similar to the situation depicted in Figure 3.2a. However, in order to explore the dimensions of spatial non-stationarity and heterogeneity in parameter estimates, local spatial interaction models, where the flows *from* only one origin (Figure 3.2b), or *to* only one destination (Figure 3.2c) can be used. These local spatial interaction models are generally termed origin- and destination-specific models and can be used to compare model performance, parameter estimates, and behaviour across and/or between origins and destinations (Haynes & Fotheringham, 1984). In certain situations, it will make more sense to use one form of local model over another, depending on the type of constraint(s) being used, the nature of the flows being analysed, and the specific questions being asked.

The general form for the (origin-specific) local spatial interaction model (unconstrained) is

$$T_{ij} = w_j^{\alpha_i} d_{ij}^{\beta_i}, \tag{3.16}$$

where the variables are defined as before (see Section 3.4.2), and the parameters α_i and β_i are each specific to origin i (Haynes & Fotheringham, 1984). Notice that the term v_i^{μ} in the original model is no longer needed, as there is now only one origin (thus, origin-specific variables are redundant). Similar models can be formulated for other members of the family of spatial interaction models.

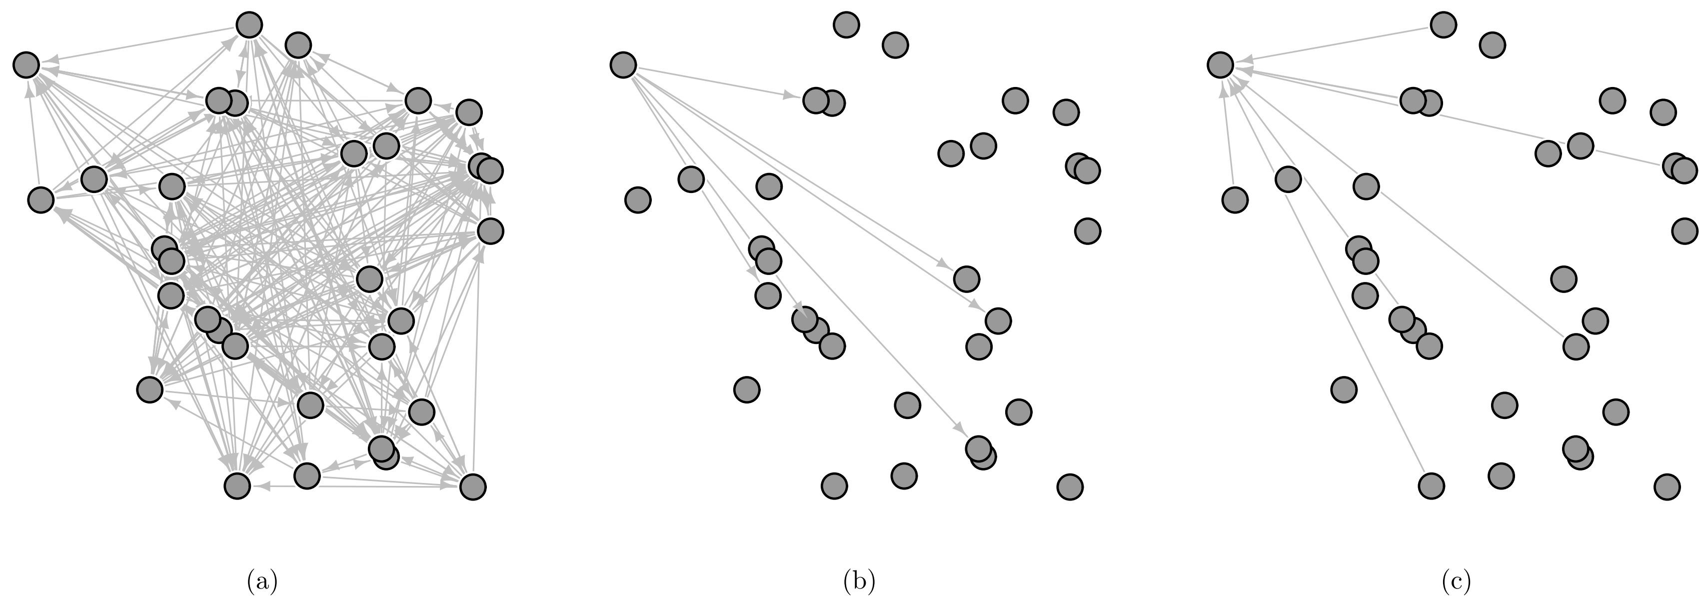

Figure 3.2: Example of traditional (a), origin-specific (b), and destination-specific (c) spatial interaction systems (after Haynes & Fotheringham (1984, chap. 4, fig. 4.1)).

A destination-constrained variant of the origin-specific local model would of course be of little use, as it would simply reproduce the observed outflows to each destination. Conversely, a destination-specific version of the destination-constrained local model is possible, and is given by

$$T_{ij} = \frac{D_j v_i^{\mu_j} d_{ij}^{\beta_j}}{\sum_i v_i^{\mu_j} d_{ij}^{\beta_j}}. \tag{3.17}$$

Origin- and destination-specific local spatial interaction models have been used in the past, with some early[1] work (e.g., Curry, 1972; Johnston, 1973; Leinbach, 1973; Stillwell, 1978; Griffith & Jones, 1980, *inter alia*) leading to intense debate over spatial structure and the interpretation of distance decay parameters (Fotheringham, 2001; O'Kelly, 2004). This debate centred around the observation that in almost all studies using origin-specific spatial interaction models, centrally located origins tended to display *less negative* distance decay parameters than their peripheral counterparts, leading to differing opinions on the true meaning of the distance decay parameter.

Prior to this debate, it was generally assumed that differences in distance decay parameters were primarily the result of differences in spatial behaviour and/or preference structures driven by differences in socio-economic conditions (Roy, 2004). Based on these assumptions, we would interpret the observed differences in distance decay between accessible and inaccessible regions to mean that, *ceteris paribus*, individuals in more accessible regions perceive distances differently to those in less accessible regions! Clearly this explanation is not particularly convincing and it was shown by Griffith & Jones (1980) and Fotheringham & Webber (1980), *inter alia*, that distance decay and the spatial structure of origins and destinations (i.e., the spatial arrangement of origins and destinations in the spatial interaction system) are in fact interde-

[1]Although Stewart (1942) used a fairly simplistic gravity style model for his work on undergraduate enrolments, his treatment of origin specific distance decay parameters was relatively advanced and pre-dated most work on origin/destination specific models.

pendent (Roy, 2004). Indeed, it has been suggested that the traditional spatial interaction model may in fact be a poor specification of reality (particularly in the presence of agglomeration or competition forces), owing to the fact that a relationship between interaction patterns and spatial structure exists, but is not addressed in the model.

Fotheringham (1980, 1981) took this one step further by outlining specifically what it was about spatial structure that was not captured by existing spatial interaction models - competition. In a series of papers surrounding 'a new set of spatial-interaction models', Fotheringham (1983a,b, 1984, 1986) suggested that the 'missing link' in traditional spatial interaction models was the competition for interactions that each destination faces from all other destinations in the spatial interaction system[2]. Building on this notion, he developed what has been termed the 'competing destinations' spatial interaction model. This model takes as given the fact that there is competition between potential destinations and that such competition is often the result of a multi-stage decision-making process whereby individuals first select a particular milieu (macro-destination) and then subsequently select from specific destinations (micro-destinations) within this milieu (Roy, 2004). As such, the inclusion of a relevant measure of destination competition within the spatial interaction framework will alleviate the aforementioned model misspecification. One way of capturing this destination competition is via a Hansen-type (Hansen, 1959) measure of accessibility

$$c_j = \sum_{j \neq k} w_k^{\alpha} d_{jk}^{\beta}, \tag{3.18}$$

where c_j represents the accessibility of destination j to all other destinations k and all other variables and parameters are as defined before. This measure of accessibility is easily integrated into the traditional spatial interaction mod-

[2]Justification for this assertion is given fully in Haynes & Fotheringham (1984), Fotheringham & O'Kelly (1989), and Fotheringham et al. (2000a).

elling framework to produce, for example, an origin-constrained competing destinations model

$$T_{ij} = \frac{O_i w_j^{\alpha} c_j^{\gamma} d_{ij}^{\beta}}{\sum_j w_j^{\alpha} c_j^{\gamma} d_{ij}^{\beta}}, \tag{3.19}$$

where all variables are as defined before. The corresponding unconstrained, destination-constrained and doubly-constrained models are similar to (3.19), with appropriate constraints and/or variables added and/or removed. The parameter γ in the above model represents the agglomeration $(+\gamma)$ or competition $(-\gamma)$ forces acting in the current spatial interaction system and will likely vary in sign given the types of flows being examined.

In addition to providing theoretical and practical explanation to the general spatial interaction model, the competing destinations framework has advantages over some modern spatial models which use econometric techniques to 'control' for spatial structure via, for example, autoregressive terms or univariate measures of spatial autocorrelation. Indeed, rather than considering only the spatial relationships in the *data themselves*, the competing destinations framework explicitly places the focus of spatial structure effects on the *actual process* (i.e., competition for interactions) generating the observed data, providing a theoretically and empirically superior interpretation of the results.

3.4.4 Discrete choice & competing destinations

As suggested by Fotheringham & O'Kelly (1989, chap. p. 67), "Most spatial interaction results from some sort of spatial choice (...) and consequently, there is a strong relationship between modelling spatial interaction and modelling spatial choice." In other words, if spatial interaction refers to the collective spatial choices of many individuals, then spatial choice attempts to explain the choices of a single individual. With this in mind, the competing destinations framework can be further explored in the context of discrete choice, spatial cognition, and utility maximisation to arrive at a fully behavioural interpretation of the competing destinations models[3].

At about the same time that researchers working on spatial interaction models of aggregate flows were attempting to redefine spatial interaction in terms of human spatial behaviour, others were attempting to borrow methods and theories from other disciplines, including economics, retailing, transportation, psychology, and marketing. Perhaps the most successful of these lines of inquiry resulted from the realisation that the spatial interaction models such as (3.11) and (3.12) were similar in form to the discrete choice models being developed in travel demand and marketing research (e.g., McFadden, 1980; Ben-Akiva & Lerman, 1985; Train, 1986). This afforded the opportunity to consider the choice process of a single individual, rather than the aggregate choices of many individuals and provided a more behavioural framework for understanding spatial choice and spatial interaction.

Developments in discrete choice models for spatial choice situations initially emerged from travel demand studies (e.g., Horowitz, 1980; Ben-Akiva & Lerman, 1985) and later were applied to topics such as migration (e.g., Pellegrini & Fotheringham, 1999, 2002), residential choice (e.g., Onaka & Clark, 1983; Longley, 1984; Clark & Onaka, 1985; Bhat & Guo, 2004), consumer store choice (e.g., Lerman & Liu, 1984; Timmermans, 1984; Fotheringham & Trew, 1993; Pellegrini et al., 1997), and recreational choice (eg., Peterson et al., 1983; Louviere & Timmermans, 1990).

Despite being relatively popular models for examining spatial behaviour, early discrete choice models were subject to several relatively impractical assumptions (Timmermans & Golledge, 1990), including the assumption that the utility of a destination is independent from the attributes of other des-

[3] An in-depth coverage of discrete choice models and spatial choice is not the goal here and, as such, interested readers are refereed to Fotheringham & O'Kelly (1989) and Fotheringham et al. (2000a) for further information on these topics.

tinations in the set of possible destinations. This property is known as the Independence from Irrelevant Alternatives (IIA) property and it implies that 1) the models are not able to account for spatial competition and agglomeration effects (Timmermans & Golledge, 1990) and 2) the models assume that all destinations in the choice set are evaluated and compared equally (Sobel, 1980; Wrigley, 1985; Fotheringham, 1986). These two very important limitations are not conducive to behavioural research and are ultimately due to the fact that early spatial choice models derive from economically-based aspatial travel demand studies (Fotheringham, 1984). Several separate model specifications have been formulated to account for some aspects of the IIA property (e.g., Daganzo, 1979; Borgers & Timmermans, 1986); however, these models are still unable to account for misspecification due to an inability to capture the effect of the spatial configuration of destinations (Timmermans & Golledge, 1990). In addition, given the potential number of possible destinations in any given situation (i.e., shopping, housing, employment), a simultaneous evaluation of all alternatives is not likely (Fotheringham, 1986; Gitlesen & Thorsen, 2000; Elhorst & Oosterhaven, 2006).

Additional limitations to the discrete choice framework as a basis for spatial choice modelling stem from research on spatial cognition and perception, which suggest that individuals perceive destinations as macro-destinations, often with 'fuzzy' (i.e., not discrete) boundaries between the larger macro regions. This makes it difficult to apply standard discrete choice models, even when macro-destinations are explicitly considered, because it requires an *a priori* definition of macro-regions. Further theoretical differences stem from the fact that in spatial choice situations, the differences between destinations are inherently spatial. In other words, both the location of each destination (be they macro or micro) and the location of the individual making the choice are relevant. This is not the situation in aspatial discrete choice models where the location of alternatives is not fixed. It is for these reasons (as well as several more technical reasons which are covered to various extents in Fotheringham (1983a, 1986), Fotheringham & O'Kelly (1989), and Fotheringham et al. (2000a)) that spatial choice models which explicitly consider space have been developed, including spatial choice variants of the competing destinations framework.

In essence, the competing destinations approach to spatial choice modelling employs the idea that individuals process spatial information hierarchically and that they will evaluate (i.e., make spatial choices) about macro-destinations first. This reduces the problem of *a priori* destination choice to one of probability: an individual is more likely to select a particular destination in a small macro-destination than a similar destination in a larger macro-destination, *ceteris paribus*. This assertion is based on the 'psychophysical law', which states that individuals tend to mentally underestimate the size of large objects (Stevens, 1957). Thus, the competing destinations framework addresses model misspecification in two ways: firstly, if a relationship between spatial structure and distance decay (and by proxy, the flows) does exist (see Section 3.4.3), then the introduction of a competition variable solves this problem. Conversely, if this relationship cannot be shown to exist, but we assume individuals process information hierarchically (a fair assumption given the range of empirical work to support it), then the competing destinations framework is still superior to alternative modelling frameworks because it addresses the IIA assumption via implicit hierarchical choice processes (Fotheringham, 1986). In order to illustrate the concept of hierarchical spatial information processing, we turn to the simple example provided in Figure 3.3.

Suppose an individual i, has a choice of n destinations, each with equal opportunities for interaction ($w_j = w \ \forall j \in n$) and each located equidistant from i ($d_{ij} = d_i \forall j \in n$) as in Figure 3.3a. Alternatively, suppose that not all n destinations are evenly spaced around i, such that some clusters (macro-

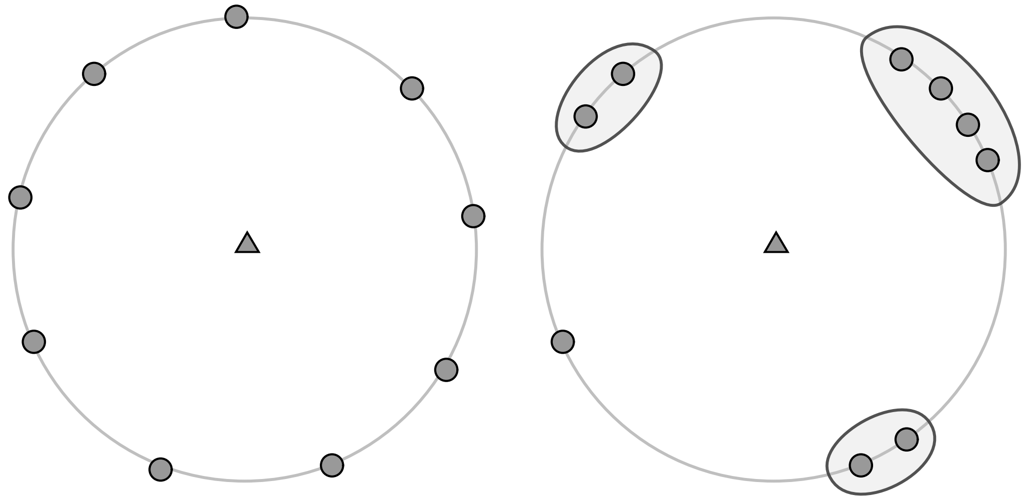

(a) (b)

Figure 3.3: Example of two spatial arrangements of 9 destinations based on non-higherchical (a) and higherarchical (b) spatial information processing. Note that the grey shaded regions denote macro-destinations and the origin (triangle) is equidistant from all destinations in both figures (after Fotheringham (1983b, p. 1121, fig. 1) and Fotheringham (1986, p. 405, fig. 1)).

destinations) may be formed, so that macro-destination k contains n_k potential destinations and has $n_k w$ opportunities. An example of this type of spatial interaction system is given in Figure 3.3b, with a single central origin and 9 destinations separated into 4 macro-destinations. If we assume that spatial structure bears no effect on the interactions of the system, then Figure 3.3a and Figure 3.3b would produce the same results. However, it is clear that this is not always the case, as in situations where agglomeration or competition forces are present. Indeed, Figure 3.4 shows graphically what is likely to happen to our perception of the attractiveness of a destination given the number of alternative destinations that are in close proximity (in the same macro-destination) to it. Note that the linear relationship in Figure 3.4 depicts the relationship assumed by a traditional origin-constrained spatial interaction model. Conversely, a more realistic relationship implies that evaluation of alternative destinations occurs in a hierarchical manner, such that macro-destinations are evaluated first, then the set of destinations (both macro and micro) within said macro-destination are evaluated in turn. Fotheringham (1986) shows mathematically that, where agglomeration and/or competition forces are present, a hierarchical processing strategy is likely and, when included in spatial interaction models via a competing destinations formulation, will lead to a more accurate representation of reality. It should be noted however, that the competing destinations framework may not apply in all situations and, in fact, may actually be misspecified when the requisite underlying assumptions are not met (Ubøe et al., 2008).

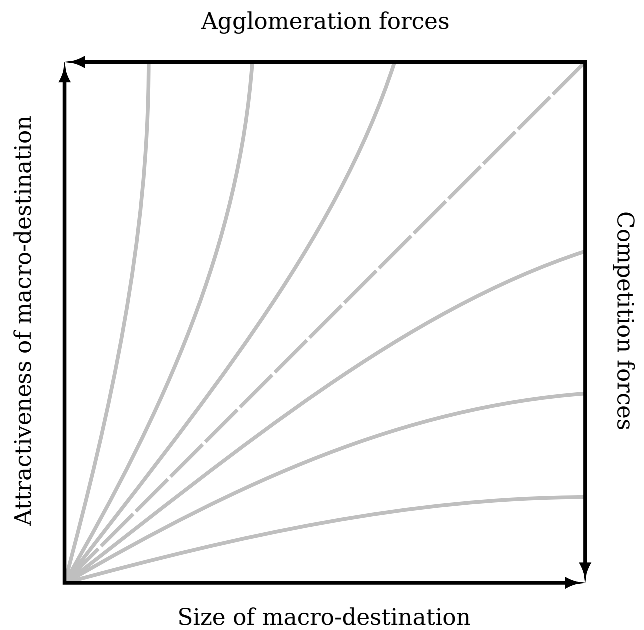

Figure 3.4: Theoretical relationships between the perceived attractiveness of a macro-destination and the number of destinations within said macro-destination. Note that as perceived attractiveness of a destination increases, competition forces decrease and similarly, as the size of a macro-destination increases, agglomeration forces decrease (after Fotheringham (1983b, p. 1122, fig. 2)).

In short, the spatial choice foundation of the competing destinations model implies that individual spatial behaviour is constrained by an individual's ability to process spatial information and is ultimately a function of the size and composition of their potential macro- and micro-destinations. The competing destinations framework is now supported by a great deal of empirical work and ultimately represents an advance in the application of behavioural theories to spatial interaction modelling. The integration of concepts from spatial interaction and spatial choice theories into a single unified framework is highly conducive to the transference of ideas and concepts between geography, psychology, and economics and, provides a powerful framework within which the remainder of this thesis operates.

3.5 Current trends

In the past, it has been common to estimate the parameters of spatial interaction models via regression by linearising the equations in terms of their parameters (Fotheringham & O'Kelly, 1989; Tiefelsdorf & Boots, 1995). Frequently, equations (3.10), (3.11), and (3.12), or preferably, their competing destinations variants, are linearised by taking the logarithms of both sides of the equation, so that, for example, (3.10) becomes

$$\ln T_{ij} = \ln k + \boldsymbol{\mu} \ln \mathbf{v}_i + \boldsymbol{\alpha} \ln \mathbf{w}_j + \beta \ln d_{ij}, \tag{3.20}$$

where $\mathbf{v}_i$ and $\mathbf{w}_j$ may now be interpreted as vectors containing selected origin and destination attributes and $\boldsymbol{\mu}$ and $\boldsymbol{\alpha}$ are the corresponding vectors of parameters. It is then usual to use $\ln T_{ij}$ to estimate the values of $\ln k$, μ, α, and β in an ordinary least squares (OLS) regression analysis (Flowerdew & Aitkin, 1982) via

$$\ln T_{ij} = \ln k + \boldsymbol{\mu} \ln \mathbf{v}_i + \boldsymbol{\alpha} \ln \mathbf{w}_j + \beta \ln d_{ij} + \varepsilon_{ij}, \tag{3.21}$$

where the random error term ε_{ij}, is assumed to be independent and identically distributed (i.i.d). This is the well-known log-normal spatial interaction model and, while potentially convenient to apply in practice, it suffers from several significant shortcomings, many of which are discussed in Fotheringham & O'Kelly (1989) and Flowerdew & Aitkin (1982), *inter alia*. These issues include a) the unrealistic assumption that flows are (log) normally distributed, b) the bias created by the logarithmic transformation (i.e., back-transformation bias), c) the failure of the assumption of homoscedastic error terms, and d) the problem of zero-valued interactions (i.e., the logarithm of zero is undefined). Tiefelsdorf & Boots (1995) also point out the difficulty in applying constrained interaction models in this context.

3.5.1 Poisson spatial interaction models

Flowerdew & Aitkin (1982) have argued that many of the above problems stem from an incorrect specification of the model itself and suggest that, rather than regression based on the normal distribution (as in OLS), a more appropriate distribution for modelling spatial interaction counts would be the Poisson distribution. Indeed, recognising that the number of individuals commuting from i to j must be a non-negative integer and assuming that there is a constant probability of an individual in i commuting to j, that the population of i is relatively large, and that the movement of commuters is independent, the number of individuals travelling from i to j can be considered to be the outcome of a Poisson process with mean λ_{ij}. Hence, the probability that t_{ij} individuals recorded as travelling from i to j is equal to T_{ij} is

$$\Pr(T_{ij} = t_{ij}) = \frac{e^{-\lambda_{ij}} \lambda_{ij}^{t_{ij}}}{t_{ij}!}. \tag{3.22}$$

Since λ_{ij} is unknown, it can be estimated via a spatial interaction model, where it is now assumed that λ_{ij} is logarithmically linked to a linear combination of the logged independent variables in (3.21), as in

$$\lambda_{ij} = \exp(I + \boldsymbol{\mu} \ln \mathbf{v}_i + \boldsymbol{\alpha} \ln \mathbf{w}_j + \beta \ln d_{ij}), \tag{3.23}$$

or more generally,

$$\lambda_{ij} = \exp(\mathbf{x}_{ij}^{'} \boldsymbol{\beta}), \tag{3.24}$$

where $\mathbf{x}_{ij}$ is a vector of explanatory variables associated with the origins and destinations and $\boldsymbol{\beta}$ is the corresponding vector of unknown parameters. This general formulation has been shown to be equivalent to other forms of spatial interaction models, including the entropy maximising spatial interaction models from Section 3.4.2 (Baxter, 1982, 1984; Griffith, 2010). In fact, all four models from the family of spatial interaction models given in Section 3.4.2 can be operationalised by particular specifications of the explanatory variable vector $\mathbf{x}_{ij}$ in (3.24). For example, the origin-constrained variant of a simple Poisson spatial interaction model is given by

$$\lambda_{ij} = \exp(I + \varphi_i + \boldsymbol{\alpha} \ln \mathbf{w}_j + \beta \ln d_{ij}). \tag{3.25}$$

This is essentially equivalent to an unconstrained model with an origin-specific constant term φ_i, which in the econometric literature is termed an origin-specific fixed-effect. In essence, φ_j represents an origin dummy covariant or set of indicator variables[4] for the relevant factor levels of i, $\forall i \in n$. Similarly, the relevant competing-destinations variant of this model is given by

$$\lambda_{ij} = \exp(I + \varphi_i + \boldsymbol{\alpha} \ln \mathbf{w}_j + \gamma \ln c_j + \beta \ln d_{ij}). \tag{3.26}$$

[4] For small- to medium-sized spatial interaction problems, the addition of n (or m for destination-constrained) dummy variables is entirely manageable; however, for large n or m, the size of the model matrix quickly becomes unwieldy! In Appendix A, we briefly highlight this problem and mention a software solution based on sparse matrix representations which is able to circumvent large matrix issues to some extent. These sparse model matrix techniques are used in all subsequent modelling exercises in this thesis.

The equivalence of the Poisson spatial interaction model to other spatial interaction model variants is significant because it provides a simple and comparative framework for fitting a range of spatial interactions models. Furthermore, since Poisson regression models belong to the family of generalised linear models (GLMs) (Nelder & Wedderburn, 1972), Poisson spatial interaction models benefit from several key advantages, including standard diagnostics for model specification, tests for assessing the significance of explanatory variables, efficient algorithms for obtaining maximum likelihood (ML) parameter estimates, and the potential for extending the model to include multiple additional explanatory variables beyond the standard size and distance variables (Guy, 1987; Davies & Guy, 1987).

3.5.2 Generalised linear models

The GLM framework emerged in the statistical literature in the early 1970s as a way of combining various statistical models in a single unifying framework. A GLM in its most basic form is a generalisation of OLS regression which relates a dependent variable T_{ij}, to a set of explanatory variables $\boldsymbol{x}_{ij}$. We assume that the observed values of T_{ij} come from a distribution in the exponential family, which includes the normal, binomial, and Poisson distributions, *inter alia.* The conditional mean of the distribution λ_{ij}, depends on the explanatory variables through

$$E[T_{ij}] = \lambda_{ij} = g^{-1}(\boldsymbol{x}'_{ij}\boldsymbol{\beta}) \tag{3.27}$$

where $E[T_{ij}]$ is the expected value of T_{ij}, $\boldsymbol{x}'_{ij}\boldsymbol{\beta}$ is a linear function called the linear predictor, $\boldsymbol{\beta}$ is a vector of unknown regression coefficients which are typically estimated by ML, and $g(\cdot)$ is a known link function which relates the linear predictor to the mean. Typically, the variance of T_{ij} is a function of the mean $V[T_{ij}] = V(\lambda_{ij})$. The probability density function for the exponential family is given by

$$\Pr(T_{ij} = t_{ij}) = \exp\left(\frac{t_{ij}\theta - b(\theta)}{a(\phi)} + c(t_{ij}, \phi)\right), \tag{3.28}$$

where θ is the canonical (or location) parameter which depends on the linear predictor and ϕ is a dispersion parameter which is often known *a priori*. The functions a and b are also known *a priori* and determine which member of the exponential family is used, whereas the function c is a normalising constant. Recalling the probability density function for the Poisson distribution given in (3.22), we see that Poisson regression is a special case of the GLM framework when $b(\theta) = \exp(\theta)$, $a(\phi) = 1$, and $c(t_{ij}, \phi) = \ln t_{ij}!$. The link function for the Poisson model with mean as in (3.24) is the log link function, because, as above $\lambda_{ij} = g^{-1}(\boldsymbol{x}'_{ij}\boldsymbol{\beta})$, implying that $\boldsymbol{x}'_{ij}\boldsymbol{\beta} = g(\lambda_{ij})$ and, since the mean from (3.24) is $\lambda_{ij} = \exp(\boldsymbol{x}'_{ij}\boldsymbol{\beta})$, then $\boldsymbol{x}'_{ij}\boldsymbol{\beta} = \ln(\lambda_{ij})$ means $g(\cdot) = \ln$, is the log link function. This is also the canonical link function for the Poisson, such that $g(\cdot) = \theta = \ln$ (Cameron & Trivedi, 1998) and leads to a log-linear relationship between the mean and the linear predictor. Since the variance in the Poisson model is equal to the mean, the dispersion parameter is fixed at $\phi = 1$ and the variance function becomes $V(\lambda_{ij}) = \lambda_{ij}$.

It is enough, for the purposes of this thesis, to recognise that Poisson regression may be easily applied in the GLM framework (Poisson GLM; hereafter refereed to simply as Poisson) and that maximum likelihood estimation (MLE) of the parameters in, for example, equation (3.26), is possible via the iterative weighted least squares (IWLS) method (Nelder & Wedderburn, 1972), which has been proven to converge at the maximum likelihood solution (McCullagh & Nelder, 1989). In fact, it has been shown that the MLE of the parameters of a Poisson spatial interaction model are identical to those obtained from MLE of the corresponding entropy-based model (Fotheringham & Dignan, 1984).

For additional information, a detailed account of the derivation of the Poisson model in the GLM framework is given in Cameron & Trivedi (1998), and the theory and application of MLE for GLMs are given in McCullagh & Nelder (1989) and Venables & Ripley (2002).

3.5.3 Over-dispersion

The advantages of using a Poisson specification for spatial interaction models have been clearly stated in the past (e.g., Flowerdew & Aitkin, 1982; Fotheringham & Williams, 1983; Davies & Guy, 1987; Flowerdew & Lovett, 1988; Lovett & Flowerdew, 1989; Burger et al., 2009, *inter alia*) and the improved model specification afforded by a formal statistical modelling approach clearly has advantages when testing the validity of a particular modelling exercise. Furthermore, the above treatment of Poisson regression as a GLM means that non-specialist software and methods can be used to derive MLE of model parameters in a flexible and comparable framework. However, as stated by Davies & Guy (1987, p. 301), "as in any statistical analysis, inferential rigour is contingent on assumptions. More specifically, the validity of the inferential procedures used in statistical modelling is theoretically dependent upon the underlying probability model [...] being correctly specified." While the Poisson model is often used for modelling count data, it may in fact be a gross misspecification of reality. As such, the Poisson distribution may not always be appropriate for modelling counts, such as commuting flows.

An additional potential issue with using Poisson spatial interaction models stems from the restrictive assumption of equidispersion. In practice, count data often display *over-dispersion* (Cameron & Trivedi, 1998), where the conditional variance is larger than the conditional mean. Generally, the magnitude of over-dispersion is clear from comparing the sample mean and variance of the dependent variable T_{ij}. However, it is important to consider that the addition of explanatory variables in a Poisson regression will decrease the conditional variance of the dependent variable to some degree, while leaving the conditional mean unchanged. This has the effect of reducing over-dispersion in some cases, though if the sample variance is over twice the sample mean, then it is likely that the data will remain overdispersed even after the inclusion of regressors (explanatory variables) (Cameron & Trivedi, 1998). If the (remaining) over-dispersion is not taken into account by the model, its effects on on standard errors and t-statistics can be quite significant, with effects similar to the failure of the assumption of homoskedasticity in a linear regression model. Indeed, it has been shown that the while the parameter estimates from a Poisson model are actually quite robust to invalid assumptions (Davies & Guy, 1987; Cameron & Trivedi, 1998), when over-dispersion is present, the standard errors can be biased downward to a great degree.

There are several reasons why over-dispersion in models of count data may arise in practice, including omitted explanatory variables[5], unobserved heterogeneity not taken into account by the model, an incorrectly specified functional form, and/or violation of the assumption of independence within counts. Furthermore, over-dispersion in spatial interaction models may be due to a violation of the assumption of independent flows (i.e., there may be [spatial] dependence between commuters), or an over-abundance (excess) of zero flows. Due to the widespread use of Poisson models for count-based regression analysis, the issue of over-dispersion has been examined a great deal (e.g., McCullagh & Nelder, 1989; Dean & Lawless, 1989; Dean, 1992; Breslow, 1990; Greene, 1994; Gardner et al., 1995; Cameron & Trivedi, 1990, 1998), including in the context of spatial interaction models (e.g., Baxter, 1985; Davies & Guy, 1987; Congdon, 1993; Burger et al., 2009). Several methods for detecting and controlling for over-dispersion in Poisson models have been suggested, includ-

[5]We have already seen misspecification due to omitted spatial structure effects in spatial interaction models in Section 3.4.3.

ing quasi- and pseudo-likelihood methods as well as more formal parametric modelling approaches using, for example, the negative binomial model (Baxter, 1985; Davies & Guy, 1987; Cameron & Trivedi, 1990; Venables & Ripley, 2002, and references therein).

Quasi-likelihood

Conventional Poisson spatial interaction models require us to specify the full probability model for the dependent variable and, while the parameters of the model may be estimated empirically, it is assumed that the conditional distribution of the dependent variable is Poisson. If, for instance, the mean/variance ratio of the dependent variable is not equal to one (i.e., the Poisson assumption of the mean equal to the variance is violated), it is possible to relax this Poisson assumption by only requiring parametric specification of the first and second moments of the conditional distribution. This is akin to specifying a relationship between the mean and variance, such that the variance is a function of the mean. This *quasi-likelihood* approach is a robust modelling approach originally proposed by Wedderburn (1974) and may be used to control for over-dispersion by allowing the dispersion parameter ϕ, in (3.28) to be estimated from the data. Thus, the variance is now a function of the mean, up to the scaling parameter ϕ (Davies & Guy, 1987; Zeileis et al., 2008). Parameter estimates of the quasi-likelihood Poisson model are identical to those of the standard Poisson model (and therefore, so are any predictions), however, inference based on standard errors requires an adjustment due to the modelled over-dispersion.

Pseudo-likelihood

As alluded to above, the Poisson model, while a convenient model for modelling discrete data such as counts, may not always be appropriate for modelling commuting flows. In fact, in cases where the Poisson probability model is *incorrectly* assumed, all model-based tests will be rendered liberal. If the covariance structure were known, it could be taken into account in the model, but more often than not, the form of misspecification is unknown. In such cases, model parameters can typically still be estimated consistently using MLE (consistent estimation of the model parameters only requires that the first moment of the conditional distribution be correctly specified), but for valid inference (i.e., standard errors) in such models, a consistent covariance matrix estimate is required. This can be obtained via *pseudo-likelihood* (Huber, 1967; White, 1982), or the *sandwich variance estimator* (Kauermann & Carroll, 2001, and references therein), which provides a consistent estimate of the covariance matrix without any distributional assumptions, even if the underlying (Poisson) model is misspecified. Since neither the quasi- nor pseudo-likelihood methods require the full distribution of the dependent variable to be specified, they do not correspond to models with fully specified likelihoods. This limits comparison of these models with other, more formally specified models (i.e., we cannot compare the full likelihoods); however, due to their ability to 'relax' the restrictive Poisson assumptions and accommodate misspecification, they may provide a much better fit than the standard Poisson model.

Negative binomial

A third way to correct for over-dispersion in the Poisson model and one which is frequently employed in econometric analysis of count data, is to assume a negative binomial distribution for the dependent variable T_{ij} (Lawless, 1987; Burger et al., 2009). The negative binomial model can be regarded as a generalisation of the Poisson distribution with an additional dispersion parameter allowing the conditional variance to exceed the conditional mean. It can be formulated as a mixture of Poisson distributions, such that an unobserved gamma-distributed random variable conditions the distribution of T_{ij} to be Poisson with mean μ, and variance $\mu + \mu^2/\theta$ (Venables & Ripley, 2002). Thus, the marginal distribution of T_{ij} is then negative binomial, with probability function

$$\Pr(T_{ij} = t_{ij}) = \frac{\Gamma(\theta + t_{ij})}{\Gamma(\theta) t_{ij}!} \frac{\mu^{t_{ij}} \theta^{\theta}}{(\mu + \theta)^{\theta + t_{ij}}}, \tag{3.29}$$

where $\mu = \lambda_{ij} = \boldsymbol{x}'_{ij}\boldsymbol{\beta}$ as before, θ is a dispersion parameter, and Γ is the gamma function. Note that for known θ, this is of type (3.28), and is therefore another special case of the GLM framework (with $\phi = 1$) with all the benefits of a properly specified GLM (Zeileis et al., 2008). Furthermore, the negative binomial model is a more general model than the standard Poisson model, such that the Poisson model is actually a special case of the negative binomial model when θ^{-1} is approximately zero (see Figure 3.5).

The negative binomial model allows for the possibility that there is unobserved heterogeneity in the flows due to, for example, omitted variables. This is often theoretically justifiable, in particular because it is usually impossible to collect all the explanatory variables relevant to explaining the variation in T_{ij}. Furthermore, the negative binomial model can be interpreted theoretically as a regular Poisson spatial interaction model, whereby the data are Poisson, but there is gamma-distributed unobserved region-to-region heterogeneity, reflecting the fact that the true mean (and therefore choice process) is not known. A likelihood ratio test can be employed in order to test whether the negative binomial distribution is a significant improvement over the Poisson regression (Cameron & Trivedi, 1998). A small p-value in this case would indicate that the negative binomial model is a significantly better fit than the Poisson model.

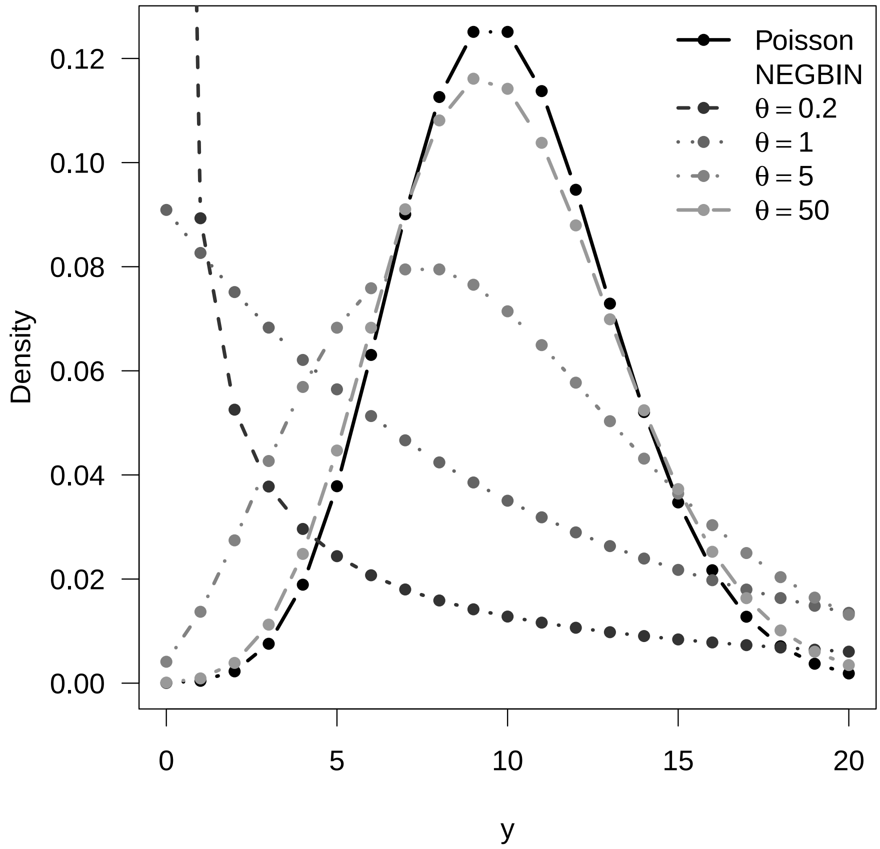

Figure 3.5: Negative binomial distribution compared with Poisson distribution with $\mu = 10$. Various values of θ are provided for comparison.

3.6 Moving forward

The simplicity and generalisability of the Poisson model (and its variants) make it a useful and viable alternative to more restrictive spatial interaction models. Furthermore, where problems may arise due to omitted variables and/or other forms of (unknown) misspecification, corrections may be employed which have been shown to provide consistent estimates of the standard errors both in an *ad hoc* fashion (i.e., pseudo- and quasi-likelihood methods) and parametrically via the negative binomial model. In the remainder of this thesis, we employ Poisson, negative binomial, and variants of these models (see chapters 6, 7, and 8) in the context of commuting flows to investigate and explain variations in model parameters and outputs due to, among other factors, local labour market characteristics. In the following chapter, we empirically define and delineate multiple sets of functional regions (i.e., the spatial manifestation of local labour markets) to be used in chapters 6 and 7, for our investigation of both global and local spatial interaction models of commuting.

Chapter 4

Irish commuting data

4.1 Introduction

Recent trends in commuting in Ireland indicate an increase in the number of individuals who regularly commute, as well as a significant increase in the number of long-distance commutes (Morgenroth, 2002; Commins & Nolan, 2010a), particularly in rural communities. According to the most recent (2006)[1] census of the population (CSO, 2007), approximately 82% of the Irish working population commute to their place of work, with an increase in the proportion of workers commuting to work of 24.78% since 1986 (see Figure 4.2b). Figure 4.2a provides some additional context in terms of the increase in the number of commuters over the last two decades. A more recent comparison with estimates from 2002 indicates that there has been an increase of 19.11% in the total number of individuals commuting to work. Overall increases in commuting levels throughout Ireland may be explained by the rapid economic and demographic changes seen over the last several decades throughout the country. Indeed, there has been a significant increase in the proportion of

[1]In fact, the most recent census is now the 2011 Census of Population; however this 2011 census has only recently been completed (April 10, 2011), and only preliminary travel-to-work data has been released to date.

employed individuals over the last few decades in Ireland, owing largely to increases in the rate of female participation in the labour force (Commins & Nolan, 2010a,b, and Figure 4.2a).

Figure 4.1: Map of the Island of Ireland with Irish cities having a population greater than ten thousand. Place-names are scaled by population size, and Northern Ireland is provided for reference.

An equally important aspect of commuting trends in Ireland is the distribution of distances travelled to and from work. Trends in the last two decades indicate significant increases in the distances people are willing to travel to

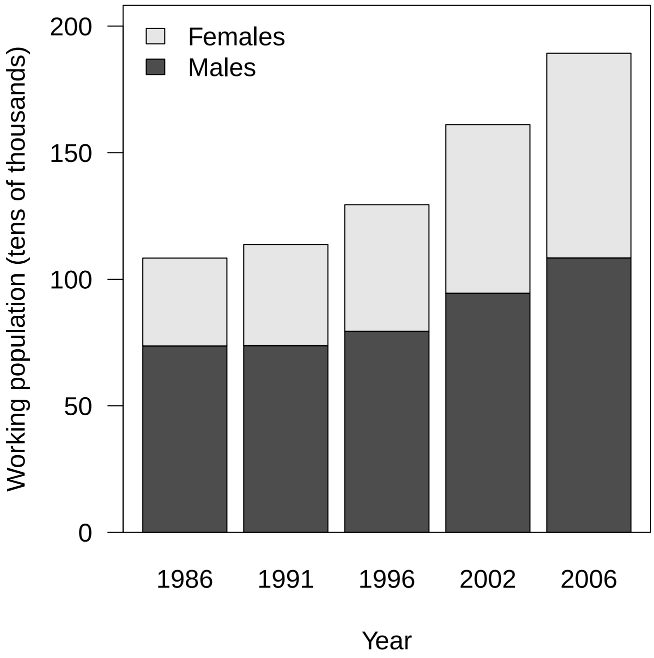

(a)

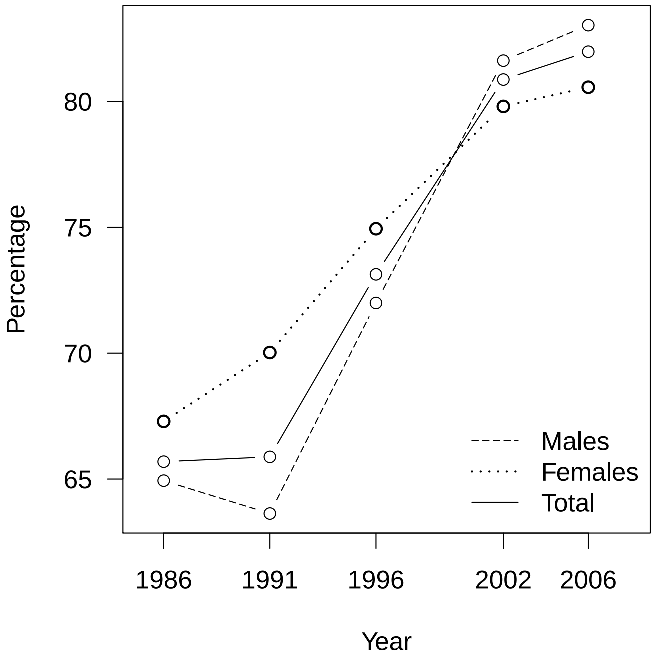

(b)

Figure 4.2: Change in the volume (a) and proportion (b) of workers who regularly commute in Ireland. Males, females, and the total commuting population are provided for comparison (source: CSO (2007, Table 1)).

bridge home and work. While the mean distance travelled has decreased to some degree in recent years (see change from 2002 to 2006 in Figure 4.3), this figure is likely to increase again given the recent high levels of unemployment (14.6% nationally in April 2011) and job competition (CSO, 2011). This increase in average commuting distances is most certainly a result of the increase over the last decade in the proportion of long-distance commutes ($d > 25$ kilometres), which has increased from 10.9% in 1996 to 12.4% in 2006 (Morgenroth, 2002). This has implications for both the size and configuration of local labour markets in Ireland, as well as general spatial interaction patterns of travel-to-work. It also has the potential to greatly affect estimated distance decay parameters in spatial interaction models of commuting and will likely cause reduced estimated distance decay when global models are calibrated for large areas of the country (i.e., the whole of Ireland).

Ireland is a particularly relevant country for examining travel-to-work patterns for several reasons. From a social and policy perspective, Ireland's rapid economic changes in recent years, including an extreme 'boom' period during the 1990s known as the Celtic Tiger (Breathnach, 1998; Bartley & Kitchin, 2007; Walsh, 2007) and the more recent economic decline since late 2007[2] (Kitchin et al., 2010; Commins & Nolan, 2010a), provide a unique perspective on the spatial implications of economic activity. Increased economic growth will undoubtedly have implications for the spatial patterns of commuting and similarly, rapid decreases in job availability and a rise in unemployment rates will lead to changes in the number of individuals travelling to work and their spatial patterns of commuting.

[2]See Subsection 4.2.2, 'Caveat lector' in Section 4.2

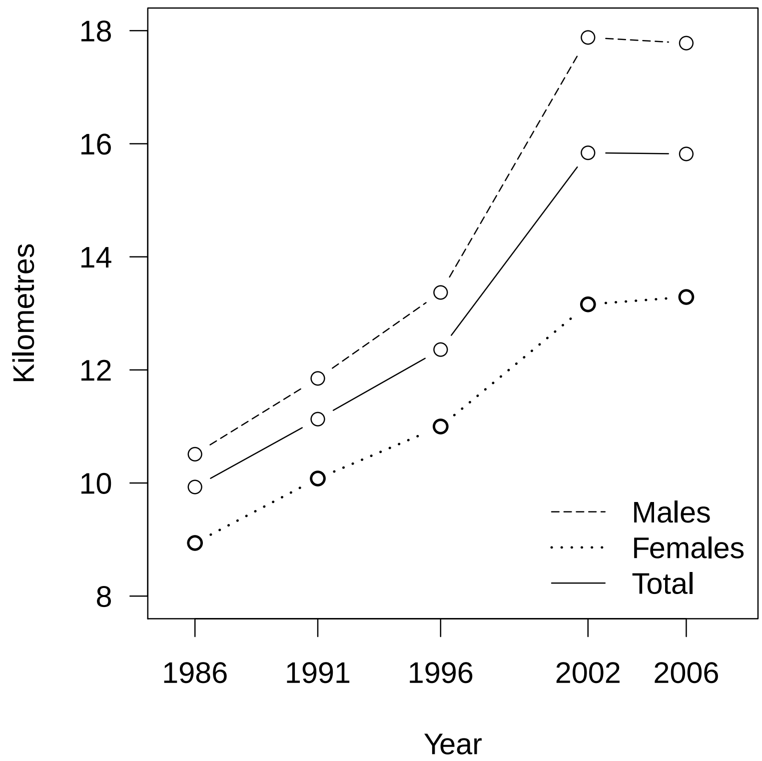

Figure 4.3: Change in the mean distance travelled to work (sources: CSO (2007, Table 5) and CSO (2006a)).

4.2 Primary dataset

The empirical analyses carried out in this thesis require travel-to-work data of the form described in Figure 3.1, where the rows and columns of the matrix represent the origins and destinations of workers and the entries in the matrix represent commuting flows. In order to derive such a dataset, comprehensive data on the travel-to-work movements of a population are required. In Ireland (see Figure 4.1), this type of data is available as part of the Place of Work Census of Anonymised Records (POWCAR) from the Census of Population of Ireland 2006 (CSO, 2006a). This dataset contains anonymised, geo-coded journey-to-work details (origin and destination) of all employed individuals in the Republic of Ireland who regularly commute, as well as a range of demographic and socio-economic characteristics, including sex, age group, socio-economic group, employment type, and means of travel.

As part of the Census 2006 processing programme in the Republic of Ireland, the place of work details of all employed individuals who undertook a journey to work were geo-coded and made available as part of the POWCAR dataset. These anonymised records, including information on the demographic and socio-economic characteristics of each individual, were made available to researchers for the purposes of analysis and research on commuting patterns (CSO, 2006a). Anonymity was ensured by filtering out all identifiable information (i.e., household number, person number within household, etc.) and by re-coding sensitive variables (i.e., those variables where the number of categories could lead to the identification of an individual when combined with other information on the record). According to the POWCAR user guide (CSO, 2006a), the POWCAR records *only* cover individuals who, at the time of the census, were a) enumerated in a private household, b) 15 years of age or older, c) enumerated at home, and d) indicated that their present principal status was working for payment or profit. In 2002 a place of work *sample* of anonymised records was released covering a 15% random sample of individuals satisfying the above criteria. In 2006, however, all individuals falling within the above scope were coded to place of work.

Table 4.1: Summary of POWCAR address returns (after (CSO, 2006a, p. 5)).

	Persons	Percentage
Total working population	1834472	100
Address was matched exactly	1097896	60
Address was matched to the same street or town	282953	15
Address was blank or un-codable	136853	7
Address was uncodable	1020	0
Works mainly at or from home	107202	6
No fixed place of work	209548	11

The location of the place of work was coded for each person in the sample on the basis of the reply to two questions from the census forms: 1) "What is (was) the full name of the Organisation you work(ed) for in the main job?", and 2) "What is (was) the full address at which you actually work(ed)?" (where "Work mainly at or from home" and "No fixed place of work" were also options). Where possible, the employer name and address was matched against addresses in the An Post GeoDirectory (GeoDirectory, 2011). Where the coder could not find an exact match, they coded to a near match if they could find a GeoDirectory address on the same street or in the same town as the address stated on the form (see Table 4.1).

4.2.1 Issues

The use of census-based data is not without its problems and there are several issues that are inherent in the use of this type of data for any analyses of social phenomena. Perhaps the most intuitive argument against the use of census data is based on the inherent temporal lag in data collection. Since most national censuses are only completed once every 5 to 10 years, some researchers have argued that the temporal lag between datasets makes comprehensive analysis over time difficult. Furthermore, others have suggested that the aggregation of census data due to privacy issues leads to the well-known issue of the modifiable areal unit problem (MAUP) (Gehlke & Biehl, 1934; Openshaw, 1977; Openshaw & Taylor, 1979) and, that results based on aggregate data are suspect. These comments are certainly valid, which is why most researchers interested in journey to work patterns avoid the use of standard statistical datasets, preferring to use data derived from national transportation or activity surveys or censuses which include disaggregate travel-to-work data, such as the POWCAR dataset. Clifton & Handy (2003) have indicated that there has been a recent resurgence in qualitative methods for travel behaviour research. Comprehensive reviews of the use of various qualitative techniques in travel behaviour research are given in Grosvenor (2000) and Clifton & Handy (2003). These reviews touch on a number of different qualitative methods, including attitudinal surveys (e.g., Kuppam et al., 1999), in-depth interviews (e.g., Jones et al., 1985), and focus groups (e.g., Polena & Glazer, 1991). Additional work using travel diaries has been used to help understand household commuting patterns (e.g., Stopher, 1992), as well as to determine the impact of telecommuting on household travel behaviour (e.g., Pendyala et al., 1991). Alternately, Carr (2008), has used an employer-based survey to assess interest in public transportation for work commute.

Additional complaints regarding quantitative data for travel-to-work studies includes the fact that, because quantitative data are only able to measure observable behavioural patterns, we are unable to take into account the social context within which the behaviour occurs (Røe, 2000). Furthermore, proponents of qualitative methods in travel-to-work research have tended to regard theoretical modelling of spatial behaviour as inadequate for an understanding of social relations (Werlen, 1993), leading some to suggest that mathematical models are unable to provide causal explanations and are limited in their scope to simple representations of spatial behaviour (Sayer, 1992). Many of these arguments are based on (dated) misconceptions rather than concrete evidence (Fotheringham, 2006) and it is clear that modern spatial interaction methods, theories, and practices are now much more tuned to the behaviour of the individuals they are modelling and there is a significant body of literature dealing with the collection, handling, analysis, visualisation, and theoretical underpinnings of such methods.

4.2.2 Caveat lector

Prior to the commencement of the research documented in this thesis, Ireland began to experience major economic decline, resulting in the country falling into recession in September of 2008. This recession appears to have occurred in the aftermath of the global financial crisis, which began in late 2007-2008. In addition to rising unemployment and a collapse of the banking industry, Ireland has experience a range of social and economic changes in recent years which will undoubtedly affect the commuting patterns and behaviour of the Irish workforce. It is not the place of this thesis to speculate on the causes and consequences of this major economic change, however, the results from this thesis must be considered within the correct context: the data utilised in this thesis reflect a time of economic and social prosperity for Ireland and any results and/or conclusions based upon them will reflect the tone of that time. It will therefore be of interest to re-examine the empirical evidence reported in this thesis once new census information is available in order to examine the effects of major economic change on commuting and the local labour market.

4.3 Network of flows

For the purposes of this thesis, the origin and destination of each working individual in the Republic of Ireland were geo-coded to their corresponding ED and from this we generated a generalised network of flows between all of the EDs. This allowed us to create a range of spatial interaction networks/matrices (see Figure 3.1) for various population sub-groups and/or combinations of disaggregations of the POWCAR data in order to provide more fine-grained control over our subsequent analyses. The full network of all commuters contains 3409 origins and destinations, with 222,484 non-zero flows, which form a single (weakly) connected network. In general, EDs are only directly connected to a limited number of relatively close neighbours, leading to a network with only 2% of the 11,621,281 possible connections in the spatial interaction matrix with flows > 0.

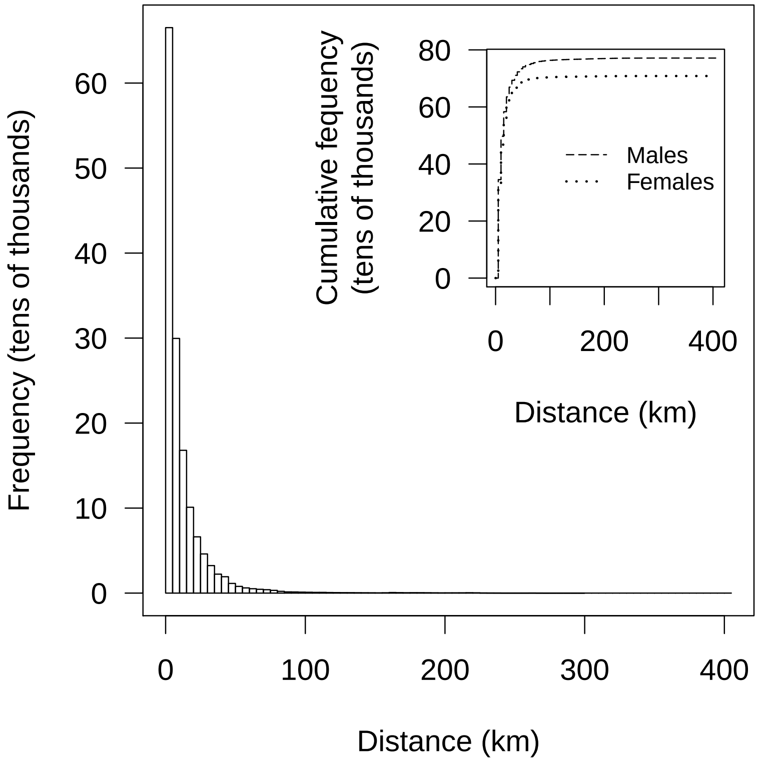

Figure 4.4: Frequency distribution of distances within the commuting network. Males, females, and the total employed population are provided for comparison.

The distribution of travel-to-work distances within the commuting network is heavily skewed towards shorter distances, with the frequency of commuters at shorter distances higher for females than males and *vice versa* for longer commute distances (see Figure 4.4). Table 4.2 provides a summary of commuting distances for various sub-groups of the population to help us understand the influence of population sub-groups on the the aggregate whole. In general, the distance summaries are in agreement with the literature on commute times and distances for these various socio-economic groups. For example, maximum commuting distance for males is higher than for females and white-collar workers tend to travel greater distances to bridge home and work than skilled and unskilled labourers. Similarly, younger workers appear to travel slightly longer distances than their more senior counterparts. A comprehensive evaluation of these trends in the context of functional regionalisation is given in Section 5.4.3.

Table 4.2: Summary statistics for commuting distances for several population sub-groups, as well as the total Irish work-force.

	Median	Mean	Std. Dev.	Max.
Males	5.97	12.35	21.21	401.40
Females	5.81	10.92	17.30	349.90
White-collar	7.14	13.45	20.96	375.80
Labour	4.86	10.14	18.10	401.40
Young	6.48	12.59	19.85	401.40
Experienced	5.26	10.50	18.86	386.20
Private	8.12	13.88	20.20	386.20
Public	7.66	12.96	18.77	298.80
Third-level	7.17	13.59	21.18	375.80
Secondary	5.17	10.25	17.93	401.40
All workers	5.89	11.66	19.45	401.40

An initial presentation of the commuting network is given in Figure 4.5, where linkages between each ED are represented by a semi-transparent line. This spatial representation of the commuting network leads to highly dense linear features emanating between the larger cities and towns throughout the Republic of Ireland. The regions in and around Dublin become completely obstructed by the number and density of interactions, which hints at the monocentric nature of employment in the Greater Dublin Area (GDA) (Convery et al., 2006; Vega & Reynolds-Feighan, 2008). The clear, linear features culminating at many of the larger cities throughout Ireland are particularly interesting in the context of this thesis, as these 'important' cities will largely dictate the shape, configuration, and number of local labour markets throughout the country, and influence (and be influenced by) the commuting patterns of the Irish workforce.

An additional visualisation of the Irish commuting data is given in figures

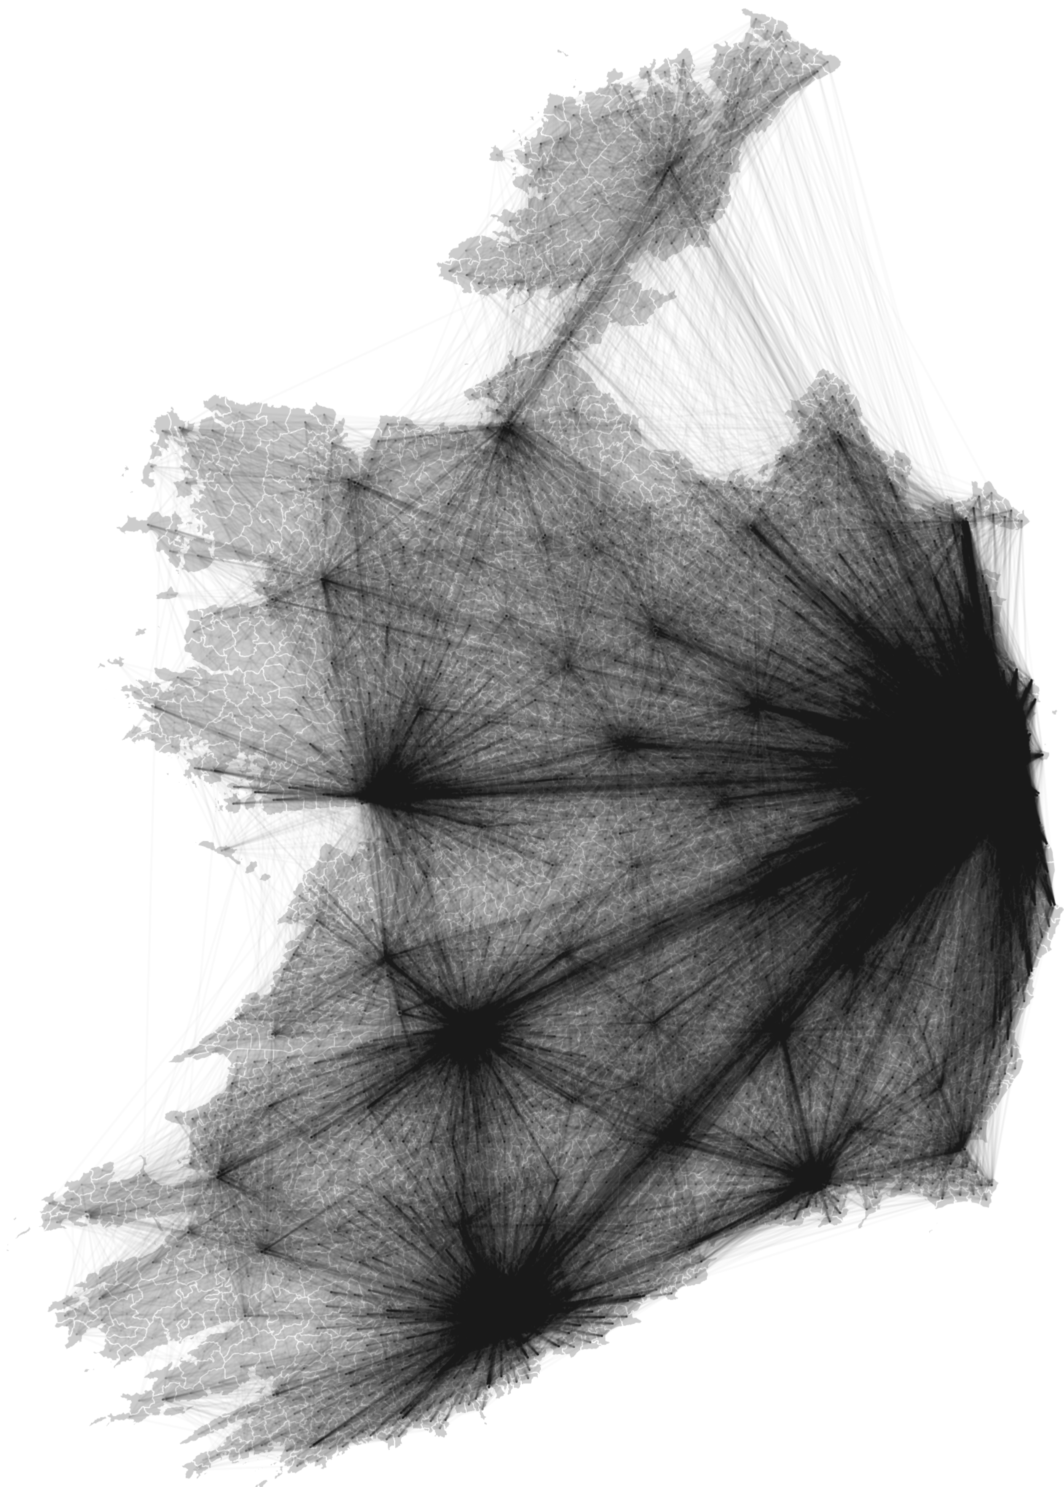

Figure 4.5: Map of all commuting flows within Ireland. Flows emanate from the centroid of each ED.

4.6 and 4.7. In these figures, the mean outgoing flow direction from each origin ED is plotted as an arrow. Figure 4.6 is the un-weighted mean direction and highlights the regional trends in commuting flows. There are approximately seven primary commuting 'sinks' throughout the country (e.g., Letterkenny, Sligo, Galway, Dublin, Limerick, Cork, and Waterford), with several smaller regional catchment areas spread throughout. Clearly, Dublin has the largest catchment region, and in some cases appears to 'poach' workers from nearby regions such as Waterford and some regional catchment areas such as Athlone and Navan. Conversely, Figure 4.7 represents the weighted mean direction of outgoing flows from each origin ED. Put another way, the mean direction in Figure 4.7 is weighted by the number of individuals travelling in each direction, such that the directions presented in the figure are more representative of the direction that the majority of workers travel from each ED.

This alternative representation of mean direction reveals some interesting patterns of commuting and is likely more representative of the 'true' commuting patterns of the Irish workforce. Smaller, regional catchment areas become much clearer and rather than a limited number of large regional catchments areas, we now see a large number of smaller local labour sinks. This provides a useful perspective on the potential structure of the local labour market and will be further examined in Chapter 5. For now, it is clear that the number and configuration of regional catchment areas is different to those in Figure 4.6, with many more commuting sinks being revealed and a significant reduction in the size of the Dublin catchment region. The extent to which Dublin's pull extends up and down the east-coast of the country is now much more visible, though its pull on regions south of Dublin towards Waterford and Wexford is now much less pronounced. These regional variations in mean commuting direction will become more evident in Chapter 5, where the flows of commuters will be used to form discrete functional regions. Furthermore, Figure 4.6 and

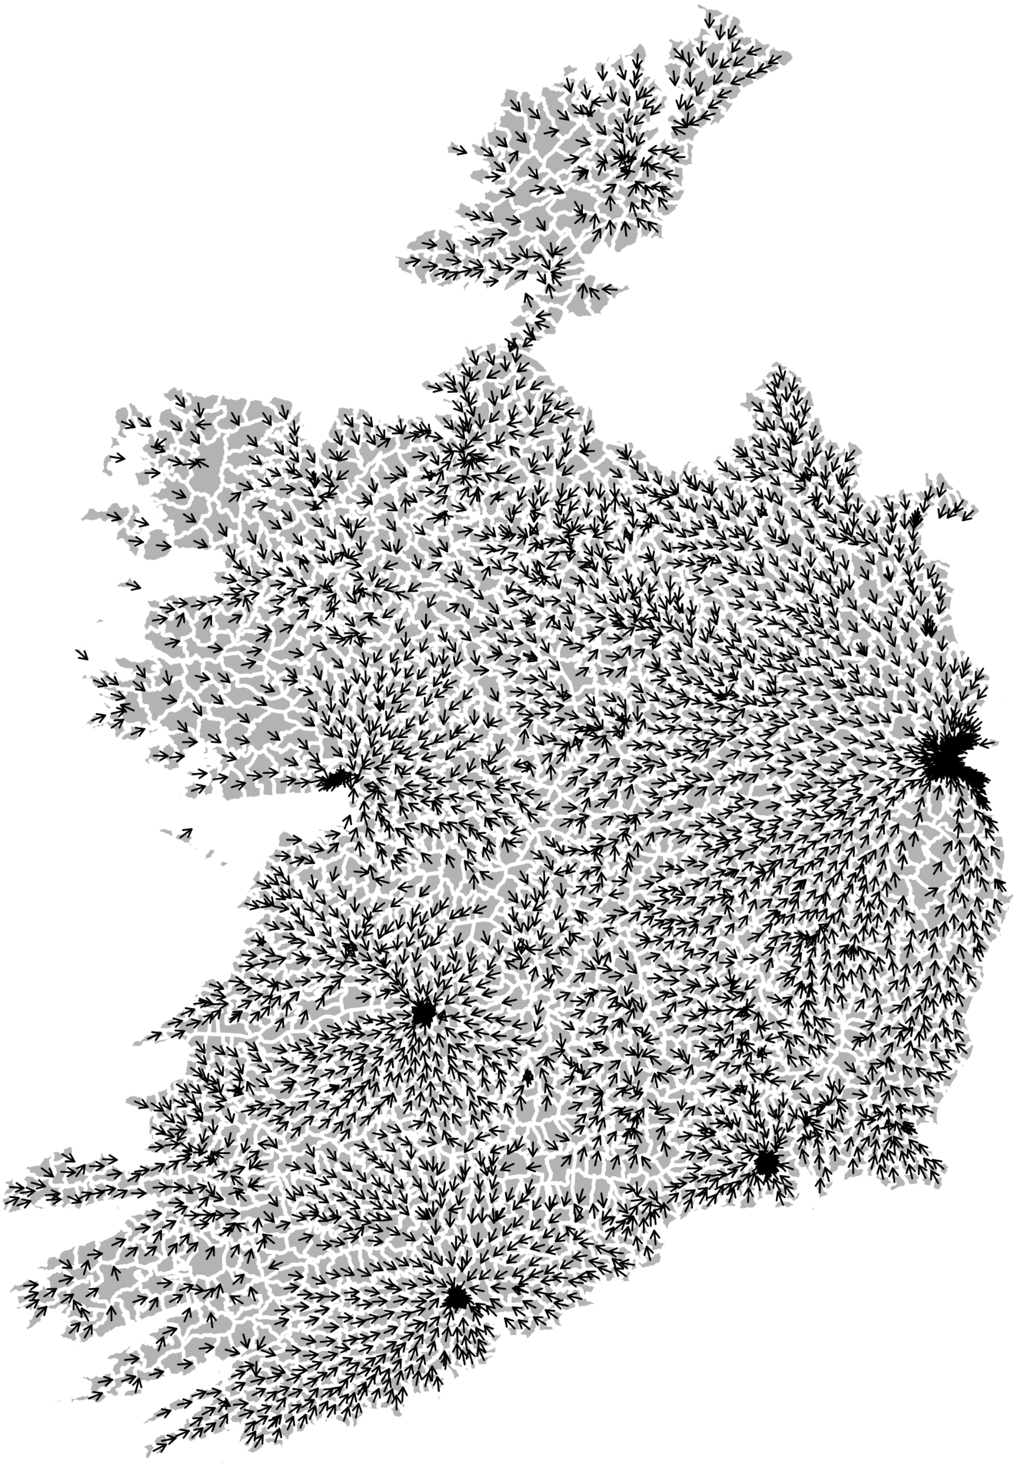

Figure 4.6: Directional flow diagram for all EDs in Ireland. Directions represent the mean direction travelled from each origin ED.

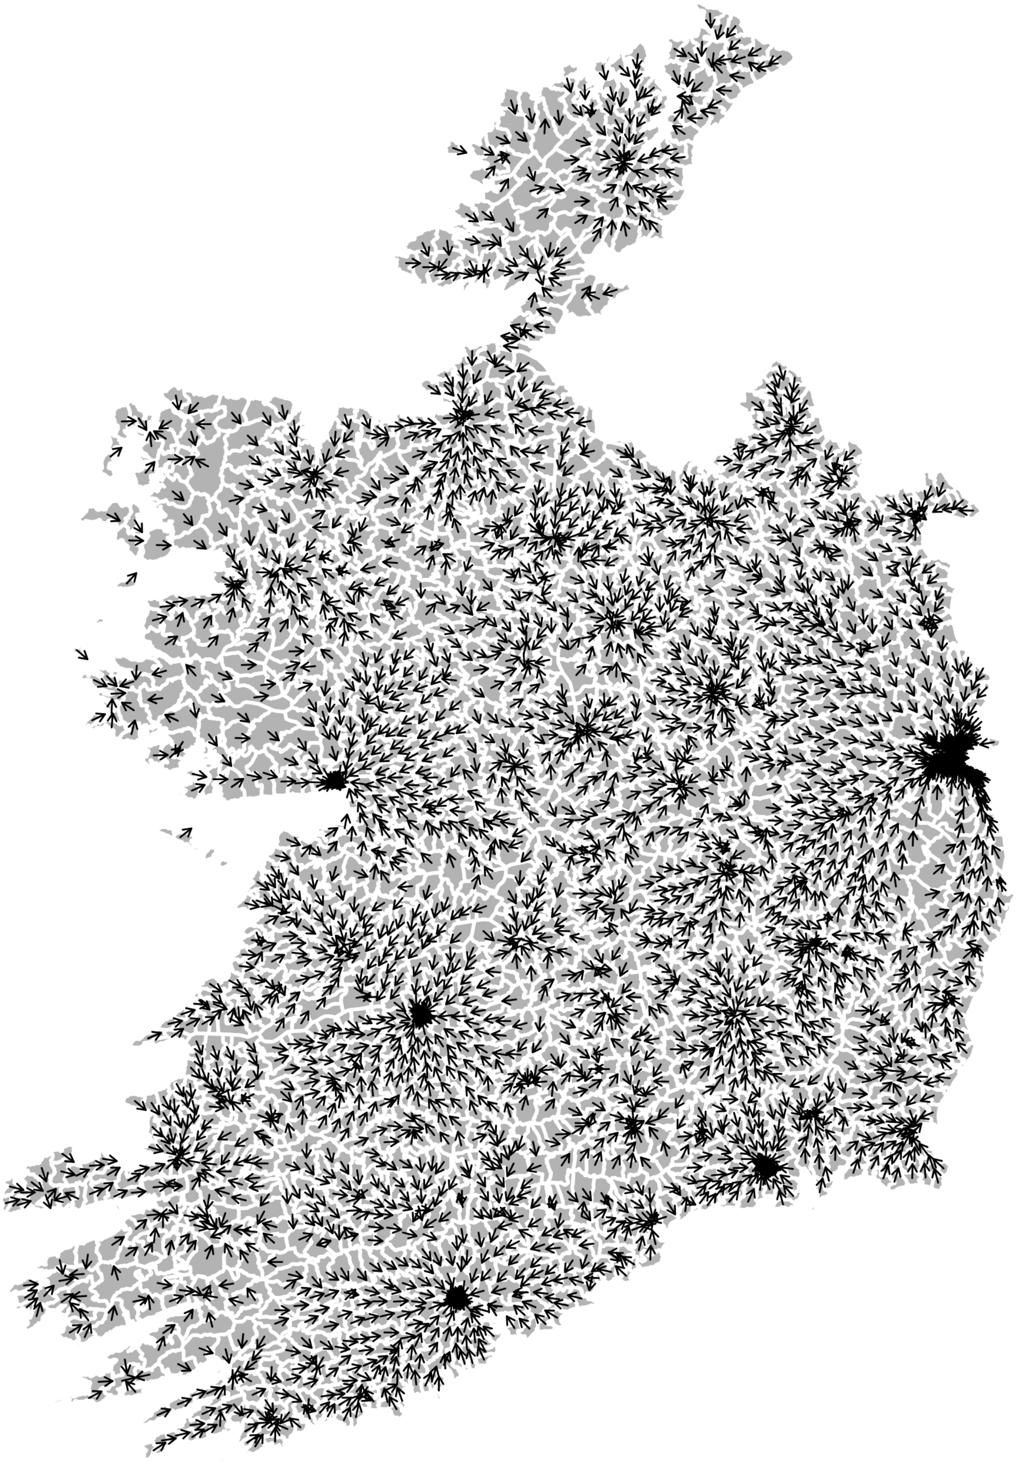

Figure 4.7: Directional flow diagram for all EDs in Ireland. Directions represent the weighted mean direction travelled from each origin ED.

Figure 4.7 will become particularly relevant in comparison with the regionalisations produced in Figure 5.3 from Chapter 5.

Many of the patterns of commuting represented here can be explained by examining the distribution of population, or perhaps more appropriately, workers. A comparison of total outgoing flows from a region with the resident working population (Figure 4.8) reveals a near perfect correlation ($r = 0.9816$) between these two values. Indeed, when looking at the ratio of working population to outgoing flows, we find that only 107 of the total 3409 EDs have a ratio less than 0.5, with the majority of EDs having a ratio between 0.77 and 0.87. This is certainly to be expected, as it is clear that regions with a high number of workers will also likely have a large number of outgoing flows. The extent to which this is the case will undoubtedly become an important factor in our spatial interaction models of commuting flows in Chapter 6. Several additional variables also play a key role in determining the magnitude and ultimate destination of outgoing flows, including education level, employment rate, population, proximity, and others. These additional variables will be explored further in Chapter 6 and will be primarily based on variables obtained or derived from the 2006 Census Small Area Population Statistics (SAPS) from the Central Statistics Office Ireland (CSO, 2006b).

4.4 Moving forward

The comprehensiveness of the available travel-to-work data for Ireland makes Ireland an ideal study area for modelling commuting patterns and behaviour. In most countries, travel-to-work data are sparse, selective (i.e., only represent a limited spatial area, employment sector, or population group), or are based on relatively small samples, which can lead to significant bias. Conversely, the travel-to-work data for Ireland is a full census of the entire working population

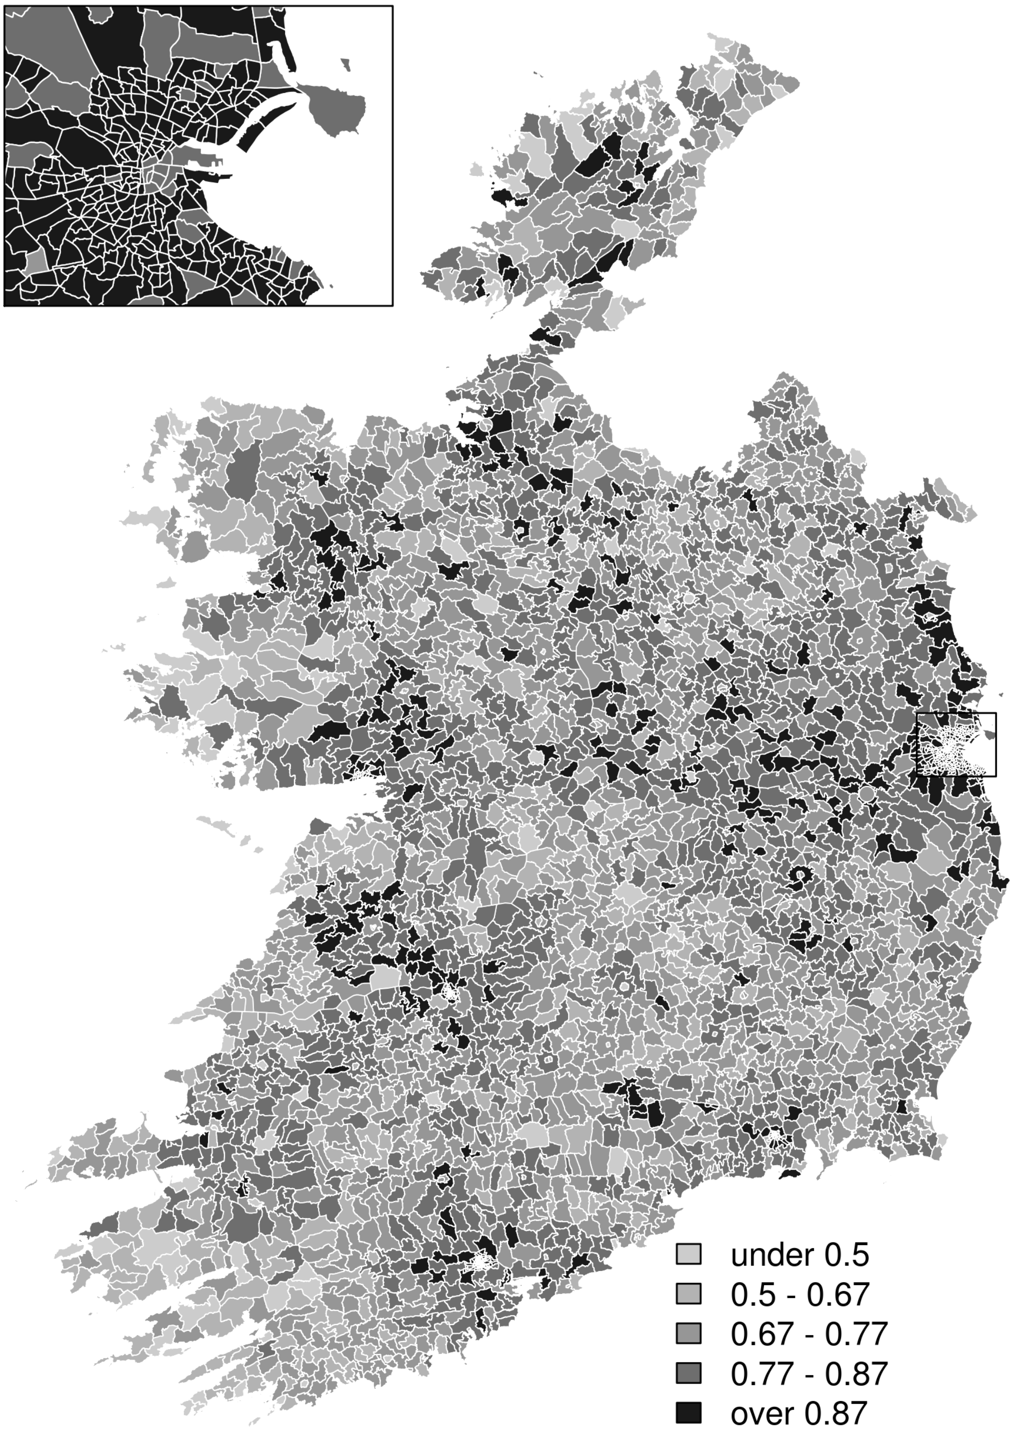

Figure 4.8: Map of the ratio of outgoing commuting flows to the number of resident workers. Darker values represent regions where the number of resident workers and the number of outgoing flows are similar.

and, as such, contains detailed micro-data on commuting behaviour across all employment sectors, population sub-groups, and regions. This comprehensive dataset provides a unique opportunity to examine the spatial patterns of commuting across a range of scales and disaggregations and, at a spatial resolution rarely feasible in other datasets. In particular, it allows us to explore commuting behaviour at the level of the local labour market, providing the means to ask and, answer, questions surrounding local labour market effects on commuting behaviour. The next step is to use the data described and explored in this chapter to begin to address our three research goals and their respective objectives. In the following chapter, we utilise the generalised commuting network and its disaggregate variants to delineate functional regions: the spatial manifestation of local labour markets. We then utilise the census data and commuting network presented here to to develop an effective base-model for commuting in Ireland. Finally, we merge the results from Chapter 5, with the base-model from Chapter 6 to arrive at a more nuanced spatial interaction model designed to extend spatial interaction theories to the unique context of commuting and, in particular, commuting in Ireland.

Chapter 5

Generating functional regions

5.1 Introduction

In Section 2.4 we introduced the concept of modularity as a useful definition of a functional region: a geographical region in which within-region interaction in terms of commuters' travel-to-work flows is maximised and between-region interaction is minimised. Furthermore, we presented an efficient heuristic designed to delineate functional regions by maximising the modularity objective function. This heuristic was shown to have several key features (Section 2.5) that make it a theoretically useful method for delimiting local labour market boundaries. In essence, the modularity optimisation routine presented in Section 2.4 treats travel-to-work patterns as a network of flows in order to delineate functional regions useful for representing and analysing local labour markets. This method differs from previous functional regionalisation procedures in that it requires few or no arbitrary threshold values or fine-tuning parameters in its formulation. Indeed, because the procedure uses the modularity quality function to determine the 'best' regionalisation for the given network of commuting flows, it requires no *a priori* specification of functional region size or count and is not reliant on the underlying population values.

This approach does however, lack two features which are beneficial (though not necessarily common) in a functional regionalisation procedure: 1) the procedure fails to consider the distances between origins and destinations, which can lead to non-contiguous functional regions, or functional regions which are overly large or complex, and 2) the procedure is entirely deterministic and as such may be sensitive to data uncertainty and/or the stochastic nature of human behaviour (Coombes et al., 1986).

We now consider solutions to both these problems in the following two sections (Section 5.2.1 and Section 5.2.2), followed by a discussion of disaggregate functional regions in Section 5.3. The remaining sections focus on empirical results (Section 5.4), including the application of the proposed regionalisation procedure to a simulated dataset in Section 5.4.1 and a network of commuting flows for the entire Republic of Ireland (described in Chapter 4) in Section 5.4.2. Furthermore, the procedure is applied to the commuting patterns of population sub-groups in Section 5.4.3. Brief concluding remarks and a link to the following chapter is provided Section 5.5.

5.2 Algorithm adjustments

In order to ensure that the derived functional regions are suitable for practical use (i.e., contiguous regions usable for statistical reporting and/or policy analysis), some additional constraints on the algorithm presented in Section 2.3 may be imposed to ensure that the functional regions display certain desirable characteristics (e.g., a large number of spatially contiguous regions, minimum self-containment levels, etc.).

5.2.1 Geographical weighting

The literature on spatial interaction/spatial choice, including commuting research using travel-to-work data, suggests that there are limits on the distances individuals are willing to travel to bridge home and work (Singell & Lillydahl, 1986; Vermeulen, 2003). In the context of functional regions, it is therefore possible to factor into our equations the fact that longer commutes are less likely, providing a more nuanced approach to functional regionalisation. As such, W_{ij} in Equation (2.2) can be replaced with an adjusted weighted adjacency matrix A_{ij} using a Gaussian-type inverse distance weighting scheme

$$A_{ij} = W_{ij} \exp\left(-d_{ij}^2/h^2\right), \tag{5.1}$$

where d_{ij} is the distance between region i and j, and h is a parameter used to control the bandwidth of the Gaussian operator. A small value of h results in very rapid distance decay and more compact spatial units, whereas a larger value of h will result in a smoother weighting scheme, with potentially non-contiguous functional regions. In practice, this bandwidth may be defined manually, or preferably by some form of automatic bandwidth selection procedure such as that of Sheather & Jones (1991)[1]. Note that other means of determining the optimal bandwidth may be employed, including fitting a spatial interaction model to obtain a distance decay parameter to be used in place of h. It should also be pointed out that *distance* in this case need not be geographical distance, but may instead be time distance, or any other form of economic distance.

[1] For an explanation of this method, see Section 5.4.2

5.2.2 Assessing stability

The validity of a particular regionalisation is undoubtedly related to the stability of the detected functional regions. A functional region may be thought of as stable if, for example, it remains relatively invariant to random- or sampling-error and/or noise. In this sense, we are interested in distinguishing between functional regions which reflect the true travel-to-work patterns observed in the dataset and those generated as a result of random effects, data uncertainties, or measurement error (Nemec & Brinkhurst, 1988). An effective method for assessing the stability of functional regions, or more generally, clusters, is via resampling methods (Hennig, 2007). Indeed, there are a number of studies which utilise 'bootstrap resampling' (Efron & Tibshirani, 1993) to assess the stability of detected clusters, as well as to determine the 'true' number of clusters in a dataset (e.g., Smith & Dubes, 1980; Felsenstein, 1985; Milligan & Cooper, 1985; Pillar, 1999; Kerr & Churchill, 2001; McKenna, 2003; Hennig, 2007; Farmer et al., 2010). This type of analysis is particularly useful in a regionalisation context, as it explicitly considers the variability within the observed interactions and does not necessarily rely on any predefined distribution to assess the significance of a particular functional regionalisation.

In the context of our current regionalisation algorithm, bootstrap resampling involves generating a large number b, of random 'bootstrap samples' from the original network of flows, applying the functional regionalisation algorithm to the original interaction network and creating bootstrap replicates of the original dendrogram by (re)applying the functional regionalisation algorithm to each bootstrap sample. These dendrogram replicates are each compared with the original dendrogram and significance is based on the proportion of dendrogram replicates that are similar to the original dendrogram. Additionally, this framework can be extended to assess the stability of individual functional regions (e.g., Hennig, 2007), ensuring that for a given regionalisation, only meaningful functional regions are retained, leaving spurious regions to be combined with the (geographically) nearest stable cluster. This maintains spatial contiguity of the functional regions, as well as reduces the effects of spurious functional regions on the final regionalisation. An additional benefit of the 'cluster-wise' assessment of stability is that relative to other bootstrap approaches, determining the stability of a functional region requires fewer bootstrap replications to produce useful results (Hennig, 2007). The procedure is implemented as follows:

1. regionalisation of the original dataset is performed and a dendrogram is obtained,
2. a number of bootstrap replicates of the original dendrogram are created as above and the regionalisation is applied to each,
3. for each functional region in the original dendrogram, the most similar functional region in each bootstrap regionalisation is found using the Jaccard coefficient γ as a measure of similarity (Jaccard, 1901)[2],
4. the level of similarity is recorded and stability for each functional region is assessed based on the mean similarity $\bar{\gamma}$ over all resampled datasets.

It is beneficial to examine multiple different resampling strategies when attempting to determine the stability of the detected functional regions. Hennig (2007) provides variations on four alternative procedures for resampling or modifying a dataset, some of which are applicable here. However, in the context of functional regionalisation, it is important to consider the special properties of spatial data, keeping in mind that there is often important information contained in the locations of the origins/destinations. As such, regular non-parametric bootstrapping techniques are not suitable for the analysis of spatial clusters because they assume no spatial structure in the network connections. Since it is this spatial structure that we are most interested in,

[2]The Jaccard coefficient is defined as the size of the intersection divided by the size of the union of the sample sets and is used here to measure the similarity between functional region membership.

removing it via the non-parametric bootstrap would be counter-productive. Viable alternative remsampling strategies include: a) replacing (a proportion of) flows with noise, b) adding a small amount of noise to (a proportion of) the flows, or c) using only a subset of the original flows (i.e., generating a sub-graph of the original network). Other resampling schemes may be viable, including several semi-parametric bootstrap techniques. The first two strategies above are based on the notion that the observed interaction matrix may include measurement error and, as such, it is useful to test whether the functional regionalisation remains the same under additional measurement error. Note that for particularly sparse networks the proposed bootstrapping algorithm may be more susceptible to errors in the observed flows (as well as outliers) than a full network. This is because the spatial structure may be lost if excessive resampling is performed, even when trying to maintain the spatial structure in the dataset via a model or the addition of noise.

5.2.3 Additional adjustments

Additional adjustments to the original algorithm may also be employed if necessary. Indeed, it may be argued that similar self-containment restrictions to those found in the TTWA procedures (see Section 2.2) may be necessary in a particular context, such as when issues of data collection and policy making require a minimal level of supply- or demand-side self-containment (Eurostat, 1992). In these cases, it is relatively straight-forward to add self-containment checks to the fine-tuning stage of the algorithm, at which point, if a certain level of self-containment is not reached, the split is rejected. While these additional adjustments will invariably require additional parameters, in some cases the increased level of control will offset the disadvantages associated with additional parameters.

5.3 Disaggregate functional regions

While some authors have explicitly incorporated disaggregate travel-to-work into their analyses of functional regions (e.g., Green et al., 1986; Coombes et al., 1988; Casado-Díaz, 2000; Watts, 2004), to date, very few applications of functional regions have focused on disaggregate patterns of commuting, such as commuting patterns of different minority groups, socio-economic groups, and/or the differences in commuting patterns of males and females. This lack of disaggregate functional regionalisations is likely due to the fact that methods which rely upon the specification of population threshold values may be sensitive to changes in population sizes induced by disaggregating the population into sub-groups. However, when aggregate data are used to determine the boundaries of functional regions, the patterns of commuting for the various constituent sub-groups of an area often become averaged. Thus, it is of interest to examine the spatial expression of different population sub-groups in terms of functional region boundaries in order to properly understand the structure of the local labour market (Green et al., 1986). In particular, the differences between male and female specific functional regions is of interest and there is a substantial body of literature to suggest that differences do indeed exist (e.g., White, 1977; Singell & Lillydahl, 1986; McLafferty & Preston, 1991; Rosenbloom & Burns, 1993; McLafferty, 1997; Prashker et al., 2008). There is also a wealth of research on commuting patterns of population sub-groups which focus on variations in commute times and distances (e.g., Warnes, 1972; Gordon et al., 1991; Giuliano & Small, 1993; Wachs et al., 1993; Taylor & Ong, 1995; Frost et al., 1998; Shen, 2000; Findlay et al., 2001; Giuliano & Narayan, 2003; Moss et al., 2004). In the context of local labour markets and functional regions, it is also important to examine the structure of functional regions for those more susceptible to unemployment, such as unskilled workers, and younger age groups. Following a more general analysis of aggregate functional regions in Section 5.4.2 and Section 5.4.2, Section 5.4.3 will examine and compare the structure of functional regions for males and females, as well as for select socio-economic groups, two age-group aggregations ($15 - 30$ years old and ≥ 30 years old), and private versus public modes of travel-to-work.

5.4 Regionalisation results

5.4.1 Simulated dataset

We now present some empirical results using our suggested functional regionalisation strategy. Firstly, we examine a simulated network similar to those presented in Newman (2004a) and Leicht & Newman (2008), with additional spatial structure. The computer-generated network (Figure 5.1a) is made up of 32 vertices (regions) with randomly assigned x and y coordinates on the unit square. Structure is imposed on the network by separating the vertices into three predefined 'groups', where between- and within-group flow directions and magnitudes are differentially applied. Specifically, the network contains only two groups if we ignore the magnitude of the flows and none if we ignore both magnitude and direction. This provides a useful test-case for the proposed algorithm: a geographical network with known structure. If the proposed functional regionalisation algorithm performs as expected, it should be able to distinguish between all three groups when considering both flow direction and magnitude and only between two groups when considering flow direction alone. In addition, because the groups are spatially defined, the geographical weighting should ensure that the groups remain spatially contiguous.

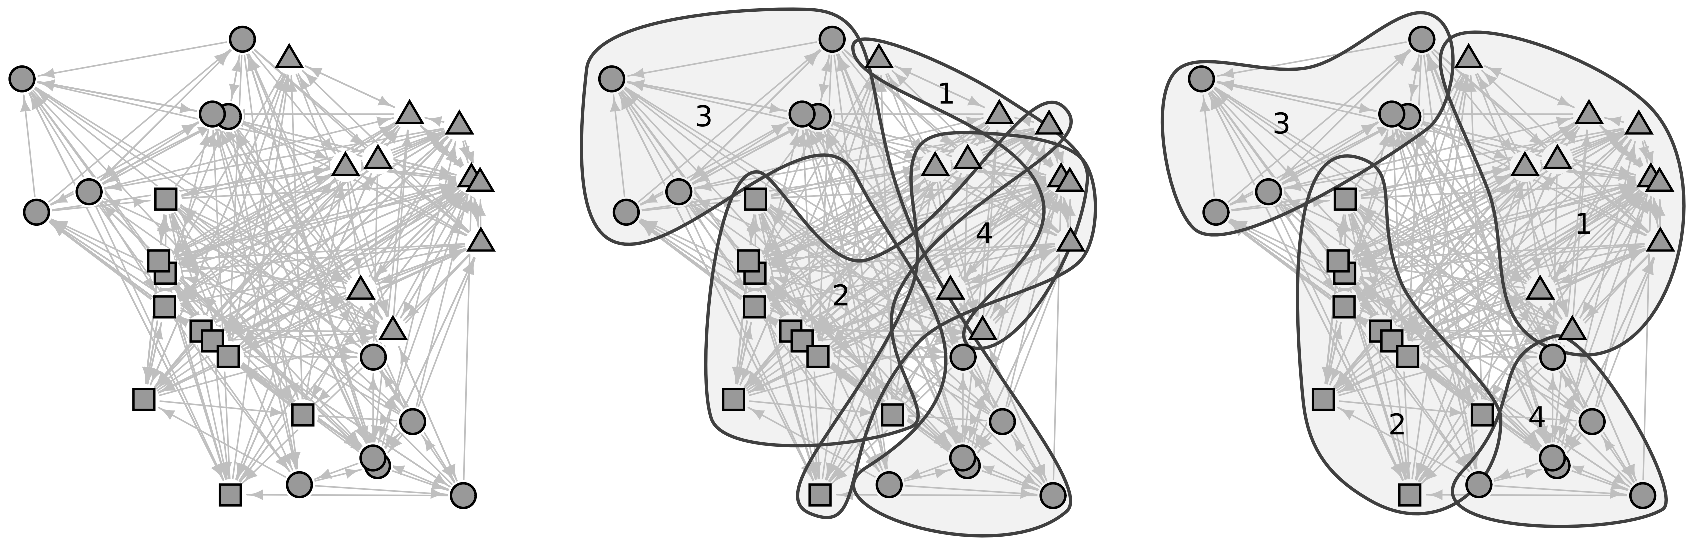

(a) (b) (c)

Figure 5.1: Community detection algorithm applied to the simulated network (a) described in the text by the original (non-weighted) version (b), and the geo-weighted variant (c). The shaded regions denote the detected functional regions (numerical labels), and the vertex shape denotes the known groupings (Group A: square, Group B: triangle, and Group C: circle).

Figure 5.1 shows the results of applying the proposed functional regionalisation algorithm to the simulated network dataset. In this example, the simulated dataset contains two spatially contiguous groups (Groups A and B) and one non-contiguous group (Group C) which occupies both the North-West and South-East corners of the plot (Figure 5.1a). When the original (non-weighted) version of the algorithm is applied to the simulated dataset (Figure 5.1b), it finds and separates Group C (Region 3) from the rest of the vertices; however, it is unable to properly distinguish between Groups A and B and as a result leaves 'floating' regions (Regions 1, 2, and 4), or non-contiguous functional regions. This forces Group C to stretch across the study region, which is undesirable for a practical functional region. Conversely, the geo-weighted version of the algorithm (Figure 5.1c) is not only able to highlight all three groups, but also splits Group C into two separate functional regions (Regions 3 and 4) to enforce spatial contiguity.

To assess the stability of this regionalisation, we employ two of the resampling strategies described in Section 5.2.2. Firstly, k proportion of flows are replaced with random noise with $r \in \{0, 1, 2, 3\}$[3] (Figure 5.2a) and, secondly, random noise with $r \in \{-3, -2, \ldots, 3\}$ is added to k proportion of the network flows (Figure 5.2b). Negative flows are not allowed and are set to zero if and when, they arise. The effect of increasing levels of uncertainty is observed by applying the above resampling techniques to increasing values of $k \in \{0.05, 0.10, \ldots, 1.00\}$ for both resampling schemes. As expected, both plots in Figure 5.2 show a clear decrease in the level of stability as the proportion of 'noisy' flows is increased. Region 4, which was generated as a result of the geographical weighting, is consistently less stable than the other detected functional regions under both resampling schemes. This is due to the fact that this region is generated as a result of the partitioning procedure,

[3] A maximum flow value of 3 is used here, as it lies squarely within the range of values for both within- and between-group flows.

rather than purely from the observed interactions in the data. This being the case, stability of the overall regionalisation (i.e., the entire study region) is influenced more heavily by the spatial configuration of the vertices than by the amount of noise introduced via the bootstrap replications. As such, regionalisations which are less impacted by geographical weighting will tend to be more stable than those which have been disproportionately affected by the weighting scheme. It may also be noted that despite having 100% of the flows replaced with noise, we still observe $\bar{\gamma} \geq 0.5$ (see Figure 5.2a)[4]. This again is due largely to the geographical weighting of the flows, although some level of similarity ($\bar{\gamma} \leq 0.3$) would be expected even without geographical weighting due to random variations in the added noise.

The results of applying the proposed algorithm to the above simulated dataset indicate that the algorithm is behaving as expected, and that given a network with known structure, the algorithm is able to detect and delineate the underlying structure. Furthermore, the algorithm is able to produce a spatially contiguous regionalisation of the network groups which is sufficiently stable for use with real-world datasets. These results provide confidence when moving forward to a real-world dataset where the underlying structure of the network is not known *a priori*.

5.4.2 Real-world dataset

Regionalisation

We now apply the modified regionalisation procedure to the Irish travel-to-work network presented in Section 4.2. Figure 5.3 shows two regionalisations of the network using the proposed functional regionalisation algorithm; the first *without* geographical weighting (Figure 5.3a), and the second *with* geo-

[4]Note that $\bar{\gamma}$ values less than 0.5 are generally indicative of an unstable cluster (Hennig, 2007).

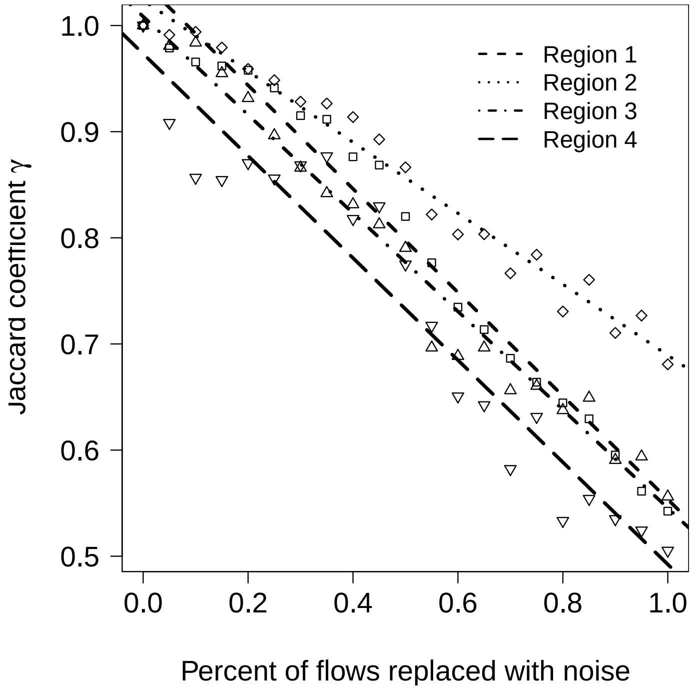

(a)

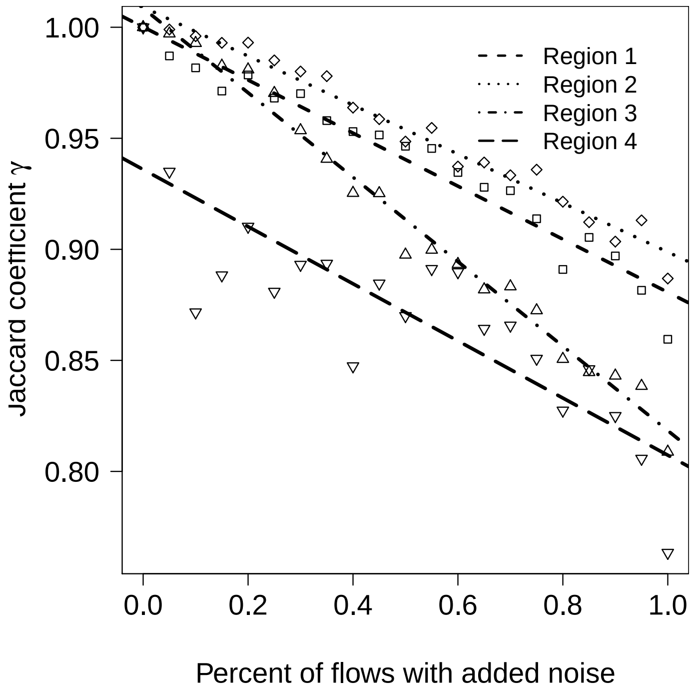

(b)

Figure 5.2: Stability of detected functional regions under varying levels of simulated measurement error. The plotted series represent the mean Jaccard coefficient over b bootstrap samples for each value of k by replacing edge weights with noise (a), and adding noise to the flows (b).

graphical weighting (Figure 5.3b). In the first regionalisation, there are 15 unique functional regions derived from the 3409 initial EDs. The average area of the functional regions is 4686.07km^2, with a range of 679.92 to 9911.89km^2. Conversely, the geo-weighted variant has a total of 65 functional regions, with a much smaller average size of 1070.07km^2, and a range of 57.44 to 5401.44km^2. For the geo-weighted variant, the size of the bandwidth (h) for the distance weighting function (Equation (5.1)), for the present analysis was computed using the bandwidth selection procedure of Sheather & Jones (1991). This is an optimal bandwidth selection procedure for kernel density estimation (KDE) based on choosing a value for h that minimises a kernel-based estimate of mean integrated squared error (MISE). In practice, different appropriate values of h are 'plugged in' to a KDE function in order to minimise some objective function based on the asymptotic MISE. To adjust for spatial heterogeneity in ED size and composition, we employed an adaptive bandwidth selection procedure whereby a value of h is computed separately for each origin in the interaction matrix based on the distances to its corresponding destinations. This produces a range of bandwidths with $\bar{h} = 7.31$km and $\sigma = 1.22$km, and ensures that the computed functional regions are not overly dependant on the size of the EDs, but rather reflect the interactions between them. For a review of bandwidth selection procedures and the technicalities of the Sheather & Jones method, see (Jones et al., 1996).

In order to assess the utility of these two different regionalisations, we turn again to the concept of self-containment and, more specifically, supply- and demand-side self-containment. The indexes of self-containment used here are taken from van der Laan & Schalke (2001) and are similar to the measure of self-containment from Green & Owen (1990). The indexes describe the ratio of individuals both living and working in a region, to the total number of individuals working in a region (demand-side) and the total number of em-

	Count	$\langle$Area(km^2)$\rangle$	Modularity	Supply	Demand
Non-weighted	15	4686.07	0.67	83.70 %	77.16 %
Geo-weighted	65	1070.07	0.53	59.01 % (87.46%)	53.94 % (88.44%)
Counties	34	2057.96	0.50	62.70 % (83.96%)	58.44 % (85.08%)

Table 5.1: Summary of measures for comparing the proposed functional regionalisation procedure with the non-weighted variant, and Irish Counties. Adjusted self-containment measures for geo-weighted and county regions are provided in brackets.

ployed individuals living in a region (supply-side). The higher the index, the more self-contained, or closed, a region is in terms of supply and/or demand of workers. In this case, the self-containment levels of the non-weighted functional regions are significantly higher than those of the geo-weighted variant (Table 5.1). Modularity provides an additional metric by which to measure the 'fitness' of our regionalisations. Again, the non-weighted variant has a larger modularity value than the geo-weighted variant. However, for operational purposes, functional regions of the size obtained using the non-weighted algorithm are not useful and a finer-scale regionalisation is required. The geo-weighted variant is designed to facilitate smaller, more compact functional regions by emphasising spatially local connections; it is therefore of no surprise that it generates a regionalisation with a higher number of functional regions.

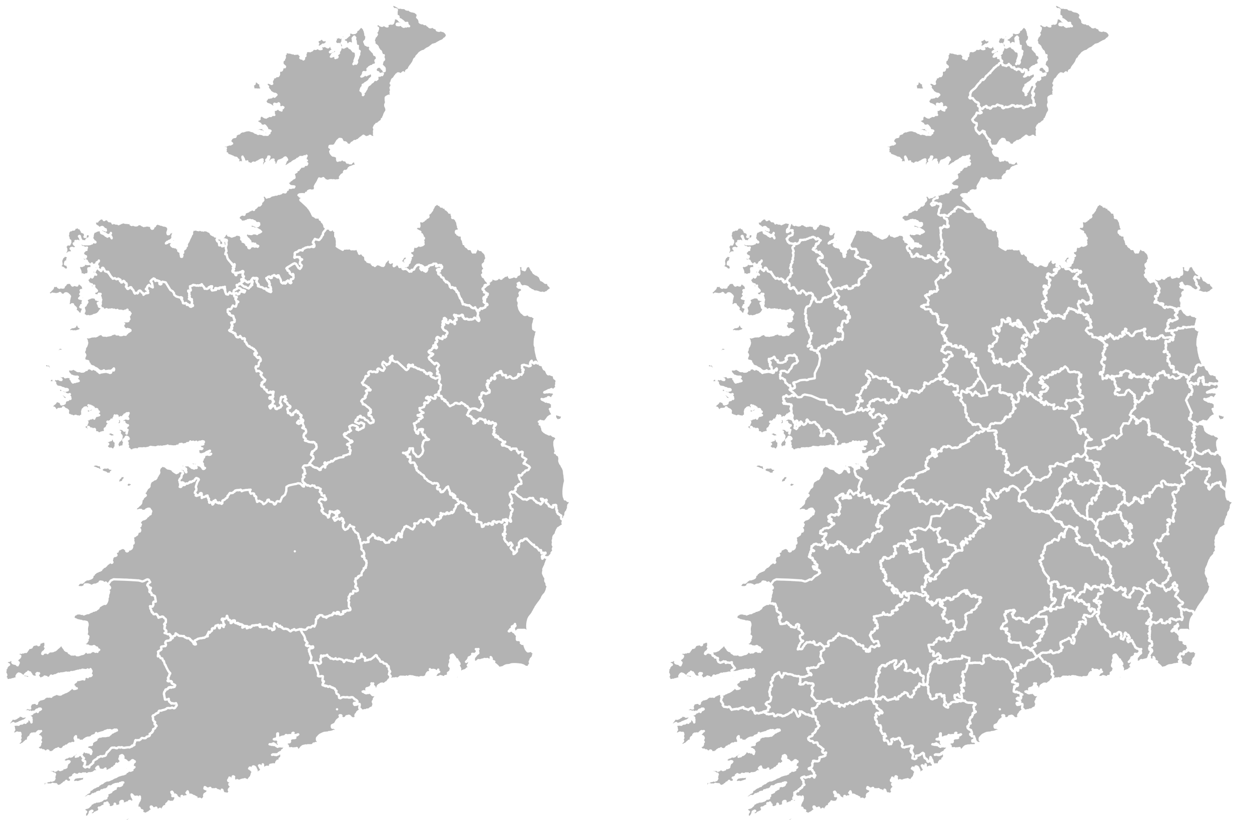

(a) non-weighted

(b) geo-weighted

Figure 5.3: Detected functional regions using both the non-weighted, and geo-weighted functional regionalisation algorithm introduced in Section 5.2.1.

When comparing the two functional regions with the nearest (in terms of size and numbers) available official regionalisation for the Republic of Ireland (Irish Counties), we find that despite there being approximately twice as many functional regions (geo-weighted) as there are Counties, the levels of self-containment and modularity are relatively close (while modularity for the functional regions is higher, self-containment is lower). This suggests that the functional regions provide a more relevant set of regions for measuring employment characteristics, even at the finer scale. Indeed, since self-containment is undoubtedly a function of region size, if the self-containment ratios SC, are adjusted to take into account the area of each region $(1 - SC/Area)$, we find that the geo-weighted functional regions (Supply=87.46 %, Demand=88.44 %) perform better than the Counties (Supply=83.96 %, Demand=85.08 %) in terms of self-containment.

It is also interesting to evaluate these regionalisations in the context of the directional flow diagrams presented in Section 4.3. Figures 5.4 and 5.5 show the above non-weighted and weighted functional regionalisations overlaid with the mean directional flows from Section 4.3. It is clear that in most cases, the direction of incoming flows corresponds to the boundaries of the two regionalisations; with non-weighted flows matching to the boundaries of the non-weighted functional regions and weighted-flows corresponding to the geographically-weighted functional regions. While the weighting styles of figures 5.4 and 5.5 differ from figures 4.6 and 4.7, the similarities are indisputable and provide additional confidence in our regionalisation results.

Hierarchical Structure

The radial dendrogram in Figure 5.6a represents the hierarchical structure of the regionalisation given in Figure 5.3b. Each subsequent stage divides the network of commuting flows into increasingly smaller functional regions

Figure 5.4: Comparison of non-weighted functional regionalisations from Figure 5.3 with the corresponding directional flow diagram from Section 4.3.

Figure 5.5: Comparison of geo-weighted functional regionalisations from Figure 5.3 with the corresponding directional flow diagram from Section 4.3.

and provides some clues as to the effectiveness of the overall procedure. For instance, the initial split at stage 1 of the procedure separates Dublin and surrounding regions from the rest of Ireland. This is to be expected, as the regions surrounding Dublin have been shown to be highly mono-centric in terms of employment (Convery et al., 2006; Vega & Reynolds-Feighan, 2008), feeding mostly into the greater Dublin area. The next split in the dendrogram (stage 2) separates the remaining EDs into two distinct groups of functional regions. This split, while derived from the commuting flows of the workers in the region, reflects differences in both the area and the underlying population size of the functional regions in the different branches. As a result, the functional regions can be separated into hub[5] and periphery regions based on this stage 2 split.

In terms of area, the hub functional regions have a cumulative distribution function (CDF) which is significantly different from the periphery regions (two-sample KS-test: $D = 0.7515$ with p-value < 0.01). Furthermore, while the distribution of population values for the hub and periphery regions are significantly different ($D = 0.3758$ with p-value < 0.01), neither are significantly different from the global population distribution. This suggests that the combination of hub and periphery regions display a similar population structure to the overall population, explaining the variation in functional region areas and providing a consistent base by which to compare the different regions. It is interesting to note that contrary to many regionalisation procedures, the one proposed here does not produce functional regions with a uniform population distribution (i.e., all functional regions have similar population values), but rather one that maintains the population structure of the underlying spatial units. This has the benefit of limiting the effects of the modifiable areal unit

[5]Note that the Ministry for the Environment, Community and Local Government has defined a series of 'hub' and 'gateway' regions for Ireland as part of the National Spatial Strategy (MECLG, 2002). However, the hub and periphery regions defined here are not part of this framework and should only be considered within the context presented in this thesis.

problem (MAUP) in terms of population, while also providing a more relevant set of spatial units with which to work. An additional aspect of the separation between hub and periphery regions becomes apparent when we consider the size and configuration of towns and cities located within the functional regions of these two categories. For the most part, the hub regions include one to two larger cities or towns, whereas the periphery regions tend to encompass several smaller towns which all feed into the larger functional region. Additional splits in the dendrogram beyond these major separations correspond to further subdivisions of these general categories into more specific local regions.

Stability

The stability of the above regionalisation was measured using a similar bootstrap procedure to the simulated data from Section 5.4.1, with a resampling scheme based on replacing k proportion of the flows with random noise drawn from a Poisson distribution with λ equal to the initial observed flow. For our analysis, k was fixed at 0.15 to simulate possible errors (e.g., sampling errors) in the POWCAR dataset. The distribution of stability values for the regionalisation is given in Figure 5.6b. Under the given resampling scheme, only seven functional regions have a $\bar{\gamma}$ less than 0.6 (with only one less than 0.5), suggesting that despite the added error in the dataset, the regionalisation procedure is able to consistently extract the initial set of functional regions.

The least stable functional regions in the regionalisation are primarily associated with the periphery functional regions, where strong centralised commuting flows are less common. Indeed, 72.72% of the functional regions with $\bar{\gamma} < 0.80$ are contained within a single subgroup of the periphery regions (after stage 53) and are labelled in Figure 5.6a as Unstable. This result is not unexpected, as the regionalisation algorithm is a divisive hierarchical method and, as the stages progress, the individual subregions are grouped only with

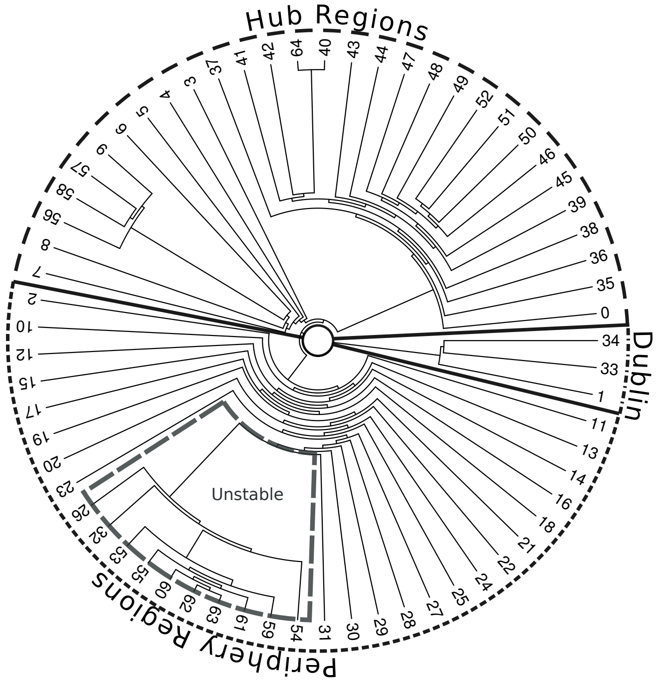

(a)

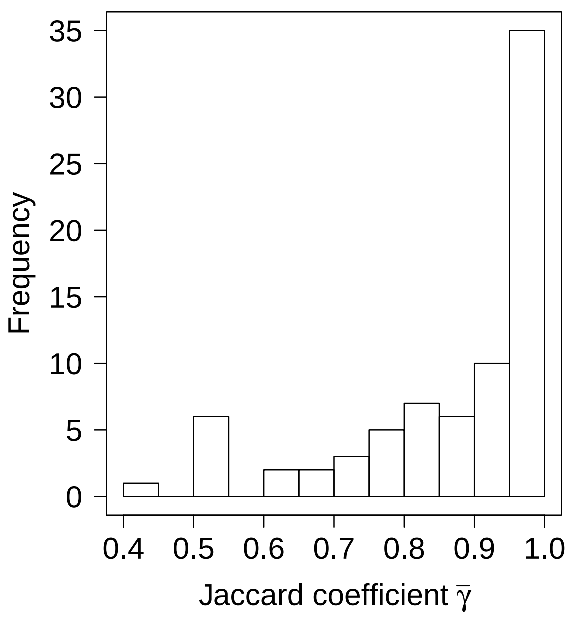

(b)

Figure 5.6: Hierarchical structure of the weighted regionalisation (a) and the distribution of stability values (b), for the regionalisation given in Figure 5.3b. Note that stability values below 0.5 are considered unstable.

the remaining subregions, leading to potentially less stable regions towards the end of the regionalisation procedure.

5.4.3 Population sub-groups

In this section, we examine the differences between various disaggregations of the POWCAR data into population sub-groups. The various groups in this case are not mutually exclusive, but rather represent components of the population that are potentially relevant to policy and resource allocation issues. The various sub-groups considered here include 1) Males - All male individuals who regularly commute, 2) Females - All female individuals who regularly commute, 3) White-collar - All individuals in the 'Employers and managers', and 'Higher' and 'Lower professional' socio-economic groups, 4) Labour - All individuals in the 'Manual skilled, 'Semi-skilled', and 'Unskilled' socio-economic group, 5) Third level - All individuals who possess either a Bachelor degree, professional qualification, both a degree and a professional qualification, postgraduate certificate or diploma, or a postgraduate degree and/or a Doctorate, 6) Secondary - All individuals who have no formal education, or education up to primary education, lower secondary, upper secondary, upper technical or vocational qualification, both upper secondary and technical, or vocational qualification, 7) Young - All individuals from 15 to 39 years of age who regularly commute, and 8) Experienced - All individuals from 40 years of age and greater who regularly commute. In addition to these 8 socio-economic groups, we also consider private and public means of transport separately, which for the present study includes all individuals who regularly commute using private transportation, such as a motor-cycle or scooter, a car (as driver or passenger), or a lorry or van, and all individuals who regularly commute using public transportation, such as via bus, minibus or coach, as well as train or other railcar. We compare the results of these various sub-groups with the renationalisations developed in Section 5.4.2. Note that an individual's socio-economic group is determined by their occupation and employment status (see CSO (2006a) for further information).

	Count	$\langle$Area(km^2)$\rangle$	Modularity
Males	63	1119.49	0.5115
Females	65	1085.04	0.5560
White-collar	55	1278.90	0.4708
Labour	81	868.06	0.5597
Third level	56	1255.62	0.4676
Secondary	59	1191.38	0.5929
Young	55	1278.02	0.5143
Experienced	64	1103.72	0.5601
Private	69	1023.37	0.5316
Public	18	4008.24	0.1992

Table 5.2: Summary of count, mean area, and modularity of disaggregate functional regions.

The various regionalisations for the population subgroups are given in Figure 5.7. The wide ranging configurations of functional regions speaks to the importance of considering both aggregate and disaggregate commuting patterns when performing functional regionalisations. In order to provide a comprehensive comparison of the various sub-group functional regionalisations with the aggregate functional regionalisation presented in Section 5.4.2, hub and periphery regions have been determined in the same manner presented in Section 5.4.2. Based on this assessment and the results presented in Figure 5.7, we find that many of the relationships that we would expect from the literature are observed when considering regionalisations produced from the sub-group populations. In the remainder of this section, these relationships are presented and the corresponding sub-groups are compared in order to provide a clearer understanding of how these various sub-groups combine to produce the observed aggregate functional regionalisation.

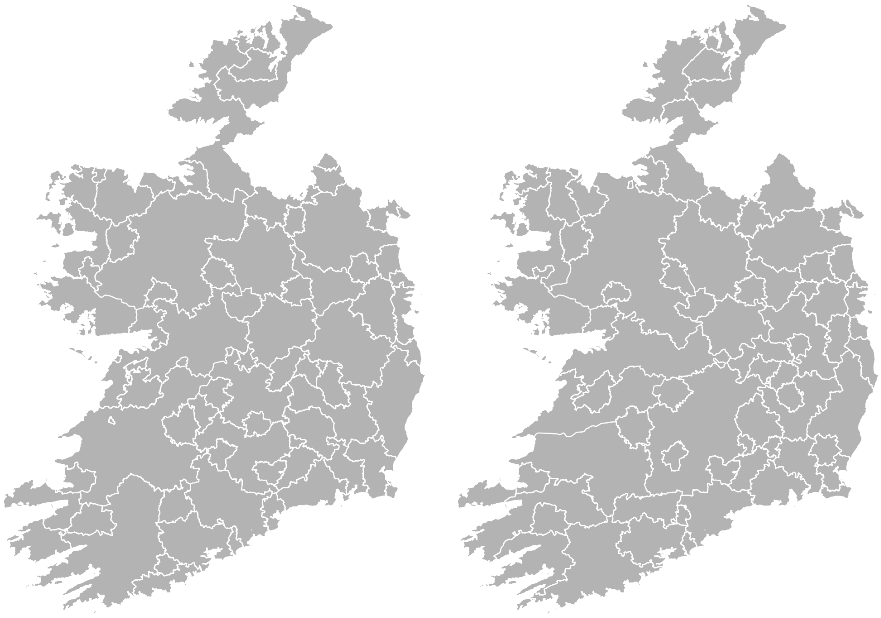

(a) Males (b) Females

Figure 5.7: Detected functional regions for various disaggregations of the Irish commuting data (see also Table 5.2) [continued on pages 107 and 108].

Males vs. Females While both male and female functional regionalisations are quite similar to the aggregate functional regions in terms of the *number* of functional regions, their size and spatial distributions suggest that there are important variations that bear consideration. For example, there are more female than male functional regions in general across Ireland, which is to be expected given previous work on gender-specific functional regionalisations (e.g., Green et al., 1986). Furthermore, female specific functional regions appear to be smaller and more numerous around Dublin, the largest urban centre in the country. Similarly, in regions where the male and female functional regions are similar (surround the same central place), the female hub regions tend to be more compact than their male counterparts, whereas the periphery regions tend to be less compact.

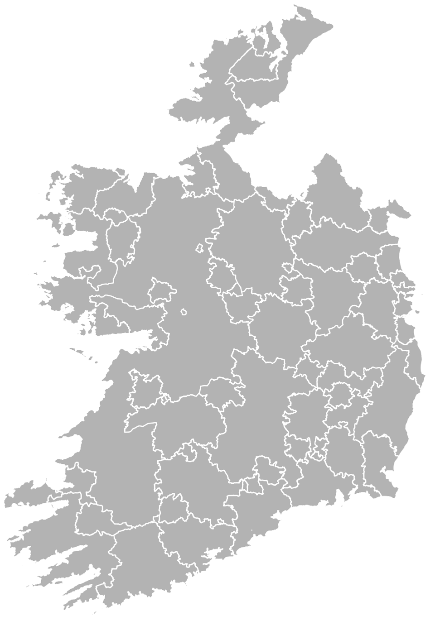

(c) White-collar

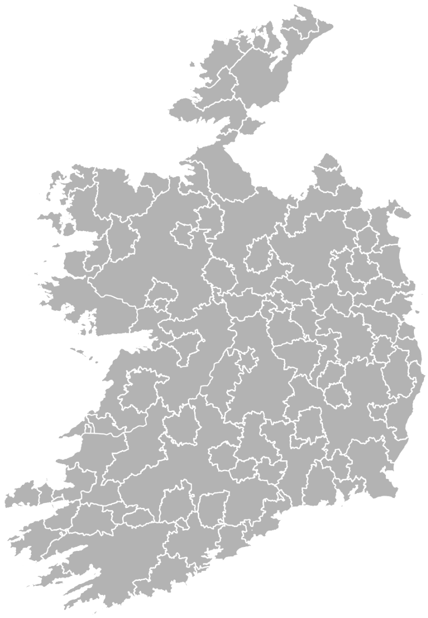

(d) Labour

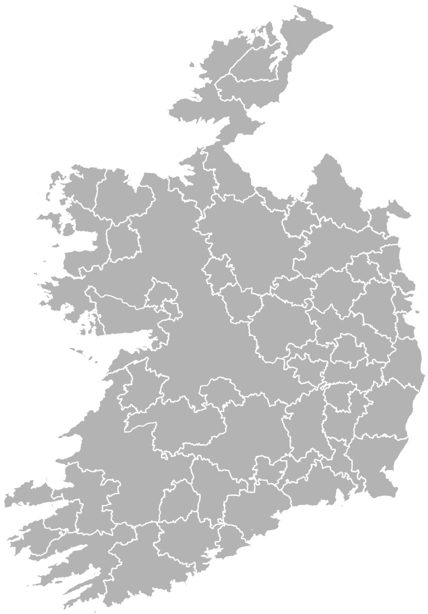

(e) Young

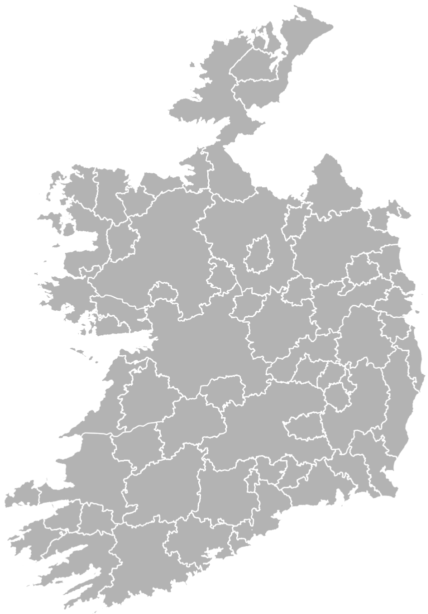

(f) Experienced

Figure 5.7: [continued from page 106]

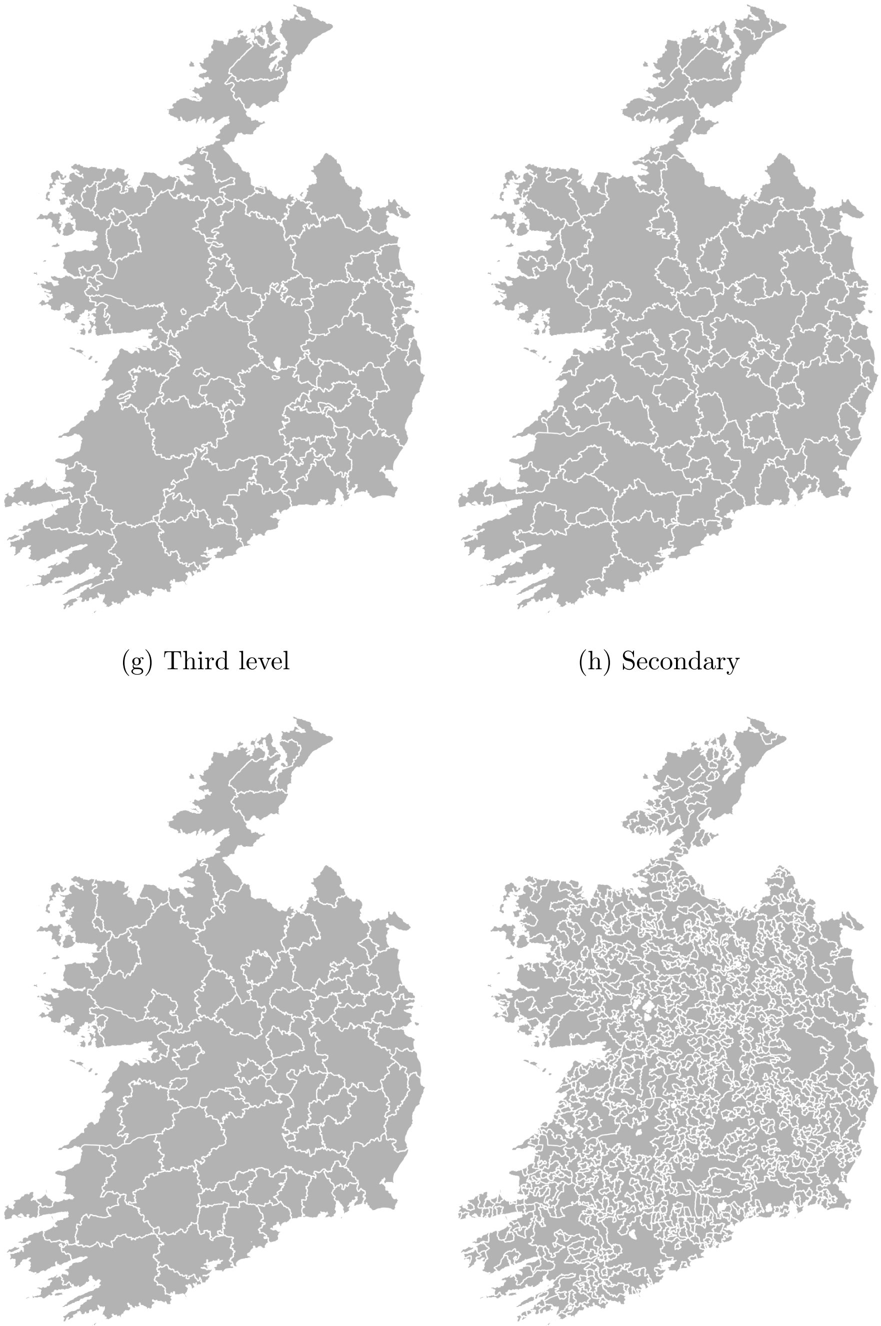

(g) Third level

(h) Secondary

(i) Private

(j) Public

Figure 5.7: [continued from page 107]

Labourers Labourers account for approximately 31.56% of the total working population; however, unskilled labourers in Ireland are comparably few, accounting for only 32.95% of labourers, or 10.40% of the total working population. Despite the relatively small population of unskilled labourers, the functional regionalisation for all labourers provides some important insights into the travel-to-work patterns of this subgroup. For instance, as expected from the literature on sub-group commute times, this group has the largest number of functional regions (81 compared with 65 for the aggregate regionalisation). This is significantly higher than any of the other subgroups and suggests that, similar to results from Casado-Díaz (2000), there appears to be a relationship between occupational status and the number of functional regions, such that lower paid labourers have a higher number of functional regions than their higher-paid counterparts. This relationship is particularly prominent in urban regions, where we see a strong tendency towards 'localised' functional regions.

White-collar Contrasting with the labourers, we see that white-collar workers have the fewest number of functional regions of the three occupational groups considered here, with much fewer functional regions than the aggregate regionalisation. Again, these results are in agreement with previous literature on occupational status and work trips (Wheeler, 1967, 1969; Cubukgil & Miller, 1982), providing concrete justification for the larger functional regions of white-collar workers in both the hub and periphery regions. Despite there being far fewer functional regions for white-collar workers than for labourers, there is no evidence to suggest 'nesting' of functional regions between these two subgroups. In other words, the spatial configuration of functional regions between these two groups is quite different, likely reflecting the differences in housing structure of these two groups. It is important to note that, besides public transportation users, the functional regionalisation for white-collar workers has the lowest modularity value (0.4708). While the low modularity value for public transportation users could be attributed to a sparse and therefore unstable matrix, the lower modularity value for white-collar workers is more likely due to the non-modular nature of commuting within this subgroup, as this group accounts for a large proportion of the total working population (39.21%). This low modularity value is important, in that it further emphasises the non-modular, or non-contiguous nature of commuting patterns for white-collar workers.

Education Differences in the functional regions generated as a result of disaggregating by level of education suggest a relationship between level of education and commuting distance, such that workers with a higher level of education (43.08% of the aggregate population), also commute the furthest. This relation has been observed in other contexts (e.g., Rouwendal & Rietveld, 1994; Vermeulen, 2003) and is most likely due to the fact that the higher-educated work-force tend to be more specialised and, as a result, can be expected to have more specific job preferences, thus increasing their overall job search regions. Additionally, the income levels of those with higher levels of education tend also to be higher and, as a result, these individuals will make different trade-offs between residential quality and commuting costs (Vermeulen, 2003). All these factors combine to produce the observed differences in functional region size and count and are especially prevalent in the rural regions of the country.

Young vs. Experienced The relationship between young and experienced workers shows that there are fewer functional regions for younger workers (55) than for more experienced workers (64). There are several possible reasons for this, including the fact that more experienced users may have had more time to optimise their housing/employment locations, increased job security leading to smaller job search areas and increased experience leading to a more selective job search region. Further to this, there is increasing evidence to show that the average level of education has risen over the past decades, such that younger people are generally higher educated (Vermeulen, 2003). Upon closer inspection, we find that in fact, 51.41 % of younger commuters (excluding those who are not old enough to have completed secondary or lower) have completed third level education or higher, whereas only 33.03 % of experienced workers have completed third level education or higher, providing further explanation for the observed larger functional region size of the younger commuters.

Public vs. Private The most pronounced differences between sub-groups is undoubtedly the comparison between private and public modes of transportation, both in terms of configuration and the number of detected functional regions. There are many socio-economic explanations for the differences observed here, though more than likely, the strongest impact is due to differences in the availability of public transportation across the country. There is a very clear separation between periphery and hub regions for the public transportation functional regions, with much larger functional regions around cities such as Dublin, Cork, Limerick and Galway, and smaller, non-contiguous functional regions found throughout the rest of the country. Indeed, the 18 functional regions derived for public transportation are actually made up of approximately 8 'local', or hub regions and 10 non-contiguous periphery regions. This is contrasted with the 69 private transportation functional regions, which to a large degree mimic the number and configuration of the aggregate functional regions. This is likely the result of several factors, including the extremely sparse nature of the public transport commuting matrix (only 9.44% of all commutes are via public transportation). Furthermore, there has been a major shift towards the use of private transportation for journeys-to-work in Ireland in the past decade, with the proportion of individuals commuting via private car increasing from 46.3% in 1996 to 57.1% in 2006 (Commins & Nolan, 2010a).

General trends While it is clear that there are major variations in sub-group commuting patterns and, as a result, count and configuration of functional regions, there are also several features that remain relatively stable throughout the various regionalisations; several of which are also prominent in the aggregate functional regionalisation. For instance, throughout the regionalisations there remain several consistent 'periphery zones' which tend to be much larger in size than their surrounding urban (hub) functional regions. Areas where these periphery functional regions are most common include the north- and south-west corners of central Ireland, including portions of counties Mayo, Cork, Galway, and Limerick. While this general spatial configuration of periphery functional regions remains relatively stable throughout the various sub-group regionalisations, the *number* of functional regions in these areas varies between sub-groups.

5.5 Moving forward

To date, aggregate patterns of commuting have been the primary means of establishing the boundaries of functional regions; however, the analysis presented here shows clear differences in functional region characteristics between population sub-groups. While several of these differences are predictable in nature (i.e., are strongly associated with periphery or hub regions in all sub-groups), others are less obvious at the aggregate level. It is therefore important to consider how the structure of the aggregate functional regions reflects the intricacies of sub-group commuting behaviour. The work presented here provides an efficient means of evaluating disaggregate functional regionalisations in order to examine the structure of the aggregate functional regions and, by association, the local labour market. With no need to specify *a priori* threshold values or self-containment criteria, modularity maximisation is a useful tool for both exploring the local labour market structure, as well as providing a point of origin for defining aggregate functional regions.

By weighting the commuting flows to take into account the geographical distances between regions, the procedure presented here has been shown to find stable, spatially-constrained functional regions in both a simulated and real-world geographical network. In addition, stability of the regionalisations produced via the proposed functional regionalisation procedure was tested using bootstrap resampling techniques designed to measure the effects of noise and/or random error on the algorithm's performance. Furthermore, the regionalisation procedure presented satisfies several criteria which are desirable when implementing a general regionalisation framework, including limiting the need for tuning parameters or threshold values. While no single functional regionalisation, whether it be based on aggregate or disaggregate data, can capture the true structure of complex commuting patterns (Green et al., 1986), it is clear that structure does indeed exist and that modularity provides an intuitive means of describing and evaluating said structure. In the following chapter, we utilise the aggregate functional regionalisation(s) derived in this chapter as a means of representing the boundaries of local labour markets. These local labour market boundaries are then used to explore the effects of local labour markets on commuting patterns, as well as investigate the localised travel-to-work patterns of Irish commuters.

Chapter 6

Modelling commuting flows

6.1 Introduction

In Chapter 5, we highlighted the modular nature of commuting in Ireland, and in particular, demonstrated that regional variations in commuting patterns do indeed exist. Furthermore, these regional variations can be seen to produce functional regions where, for the most part, workers both live and work. These functional regions are used here to define the boundaries of local labour markets, which are an integral component of the function of local, regional, and national labour markets. In this chapter, we further explore the concept of local labour market effects on commuting by examining variations in commuting behaviour at local, regional, and national levels. In particular, we are interested in understanding the interplay between origin attributes and destination choice, and how this leads to the patterns of commuting observed across the landscape. We examine this behaviour through the use of spatial interaction models of commuting, building on the concepts and theories introduced in Chapter 3. Specifically, we focus on the development of a Poisson spatial interaction model for commuting based on data for the entire Republic of Ireland (see Chapter 4). We further develop this model by integrating and exploring some of the model variants introduced in Section 3.5.3. The primary goal of this chapter is therefore to determine if, and how, commuting behaviour influences and is influenced by the local labour market.

The remainder of this chapter focuses on both the theoretical interpretation and empirical evaluation of commuting model outputs. To this end, we first develop an initial Poisson spatial interaction model of commuting in Section 6.2 and present an overview of the variables used in our empirical model in Section 6.3. Model validation and extensions are also explored, and lead to improvements based on a negative binomial spatial interaction model. This leads to the derivation of several models which are designed to take into account any latent local labour market effects, as well as parametrically control for over-dispersion in the commuting flow data.

6.2 Initial model

The primary objective of the modelling efforts in this chapter is to model commuting behaviour (i.e., commuting flows) with a range of relevant explanatory variables which describe the characteristics of a set of origins/destinations. To this end, we start with the general Poisson spatial interaction model presented in Section 3.5.1

$$\lambda_{ij} = \exp(\mathbf{x}_{ij}^{'}\boldsymbol{\beta}).$$

In this case, our dependent variable λ_{ij}, which is a matrix of commuting flows (flows) between each ED in the Republic of Ireland, will be regressed against several independent variables obtained or derived from the 2006 Census Small Area Population Statistics (SAPS) from the Central Statistics Office Ireland (CSO, 2006b). These variables are summarised in Table 6.1, and include the Euclidean distance between each ED, accessibility of each origin (see Section 3.4.4 and the following Section [6.3]), as well as a range of variables describing the attributes of each origin, including a) the number of available workers, b) the level of unemployment, c) the proportion of individuals with a third-level degree or higher, and d) the number of individuals per household. Relevant attributes of *destinations* (e.g., the number of available jobs), are not currently available from the POWCAR or SAPS datasets, which relegates us to the use of a destination-constrained spatial interaction model. This is not necessarily a problem, as we are ultimately interested in modelling the flow of commuters arriving at each ED, and as such, can take these values as given. In this case, we simply constrain our model to reproduce the observed in-flows via a destination constraint (fixed-effect) similar to the model presented in Equation 3.25. Our general Poisson spatial interaction model then becomes

$$\begin{aligned} \lambda_{ij} &= \exp(I + \psi_j + \boldsymbol{\mu} \ln \mathbf{v}_i + \gamma \ln c_i + \beta \ln d_{ij}) \\ &= \exp(I + \psi_j + \chi \mathrm{w}_i + \omega \mathrm{u}_i + \delta \mathrm{e}_i + \eta \mathrm{h}_i + \gamma \mathrm{c}_i + \beta \mathrm{d}_{ij}), \end{aligned} \tag{6.1}$$

where I is an intercept term, ψ_j is a destination-specific fixed-effect, c_i is the logged accessibility of origin i, d_{ij} is the logged distance between origin i and destination j, and w_i, u_i, e_i, and h_i are the logged origin attributes workers, unemployment, education, and housing respectively. The parameters associated with these variables are given by γ, β, χ, ω, δ, and η respectively. Complete definitions of these variables, as well as their expected relationship with the dependent variable are given in the following section. The above model (6.1), is a destination-constrained competing-destinations[1] Poisson spatial interaction model, and is the base model upon which all subsequent models in this thesis are built.

[1] In actual fact, this model is a destination-constrained *origin-centric* competing-destinations model, which is different from a standard competing-destinations model where competition is destination-centric. We elaborate on this concept, and provide justification for this form of accessibility measure towards the end of Section 6.3.

Table 6.1: Summary statistics for the model variables.

	Min.	Median	Mean	Std. Dev.	Max.
flows*	0.00	0.00	0.13	4.43	3085.00
distance (km)	0.11	137.97	143.76	73.65	461.46
accessibility	153.82	1025.71	10328.39	37917.78	1327926.32
workers[1]	27.00	264.00	610.19	1038.33	18080.00
unemployment[2]	0.01	0.05	0.06	0.04	0.44
education[3]	0.01	0.15	0.16	0.06	0.52
housing[4]	1.58	2.91	2.88	0.28	4.00

*dependent variable (commuters); 1 working individuals; 2 unemployment rate
3 proportion; 4 persons/household

6.3 Variables

An initial evaluation of the dependent variable is given in Figure 6.1, which shows a histogram of the observed commuting flow frequencies. Figure 6.1a shows the frequency of all commuting flows, including all zero flows, whereas Figure 6.1b is given to better highlight the range of non-zero flows which would otherwise be obscured by the large number of zeros. Note that the log transformed flows[2] are shown to highlight the full distribution of commuting flows. Clearly the distribution of commuting flows exhibits both substantial variation (i.e., over-dispersion) and a large number of zeros. Indeed, the variance-mean ratio for flows is extremely high (109.424), and while this is a strong indication that the Poisson model may be inappropriate in this case, it remains to be seen whether the inclusion of our independent variables will reduce this effect. In the following two sections, these issues will be further explored and addressed.

The independent variables used here were chosen for our analysis based on a number of previous studies of commuting in Ireland and internationally. They represent a range of factors are have been deemed relevant to commuting behaviour/patterns in Ireland and were selected based on their relevance and

[2]As there are zero flows, we add 0.5 to each flow when computing the logarithm (for visualisation purposes only).

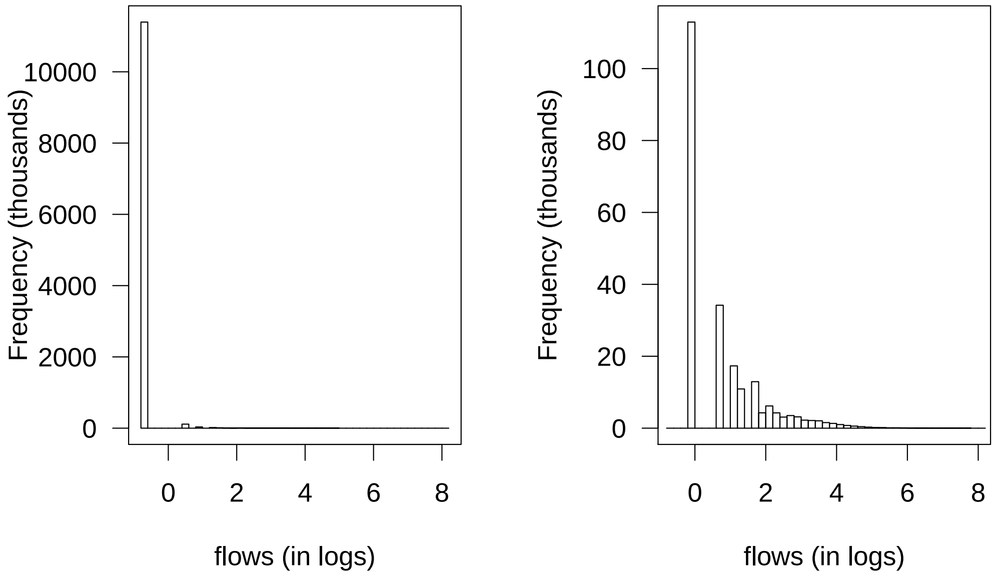

Figure 6.1: Frequency distribution of commuting flows (in logs) for the Dublin area, including all flows (left) and all non-zero flows (right).

availability[3]. Prior to the final variable selection, an evaluation of collinearity was performed, leading to the rejection of several previously chosen variables, including the proportions of unmarried (single) individuals, and urban land-use in each ED.

The set of independent variables chosen for our model(s) can be separated into attributes of *origins*, and attributes of *space*. In terms of attributes of origins, these include workers, unemployment, education, and housing. The variable workers refers to the number of workers residing in each origin ED and is derived directly from the CSO SAPS database (CSO, 2006b). This estimate of the labour force includes all individuals (aged 15 years and older) who are currently working for payment or profit, or unemployed, and are based on principal economic status as measured in the Irish census. The estimates of working population used here are designed to capture all individuals who are

[3]For this analysis to have utility to others interested in commuting patterns in Ireland (and elsewhere), we have selected variables which are readily available as part of the regular Irish (and other countries') census. This helps to ensure that our results remain both spatially and temporal relevant/comparable.

potentially inside the labour force and excludes individuals who are a) currently looking for their first regular job, b) a student or pupil, c) looking after home/family, d) retired from employment, and e) unable to work due to permanent sickness or disability. We would expect a positive relationship between workers and flows because a larger available working population will a) have a larger number of potential commuters (i.e., by sheer numbers), and b) have an increased level of internal competition, thus forcing many workers to seek employment outside the ED. A similar relationship between unemployment and flows is be expected. While this assertion is, at first glance, counter-intuitive, it is supported by a significant amount of empirical evidence surrounding the spatial mismatch problem/hypothesis (Kain, 1968, 1992). Originally developed to explain black unemployment and poverty in the U.S., the spatial mismatch hypothesis applies equally well to the case of out-commuting in high unemployment zones: there is a significant (spatial) mismatch between housing and jobs, such that unemployment is concentrated in towns and cities with a lack of jobs, leading local *working* residents to commute to external regions with a higher concentration of jobs (Zenou, 2000; Gobillon & Selod, 2007).

The education variable was similarly derived from the 2006 census, based on data from SAPS Theme 10 (Education). It describes the proportion of the working population with a third level degree or higher, which includes all individuals who possess either a Bachelor degree, professional qualification, both a degree and a professional qualification, postgraduate certificate or diploma, or a postgraduate degree and/or a Doctorate. It has already been shown that education levels have a strong influence on commuting patterns (e.g., Rouwendal & Rietveld, 1994; Vermeulen, 2003, and results from Section 5.4.3), and while there are sub-regional variations in levels of education, education is used here to represent the overall level of education of an ED. Based on results from Section 5.4.3, as well as the extensive literature on education, income levels, residential choice, we would expect there to be a weakly positive relationship between education and flows, with origin EDs having a higher proportion of well educated workers producing a higher number of commuters. Similarly, we expect the housing variable, which represents the number of individuals per household for each ED, to also have a positive effect on commuting flows. This variable is derived from the SAPS Theme 6 (Housing) data tables, and is used here as a proxy for many socio-demographic factors, such as housing density, family structure, and 'urbanness'.

A simple indication of the relationship between flows and all four origin attributes, along with the frequency distributions for the origin attributes, is given in Figure 6.2. Because our dependent variable is a count, a simple scatter-plot produces many ties, obscuring a large number of points; therefore, box-plots of grouped values are used here to emphasise the partial relationships between flows and the origin attributes (Zeileis et al., 2008). The use of a log-transformed dependent variable[4] is appropriate given that our Poisson model uses a log link function and similarly we use the log-transformed independent variables because it is the log-transformed variables that are used in the Poisson model. Indeed, unless otherwise indicated, when referring to the independent variables in the remainder of this thesis, we are in fact referring to the log-transformed variables.

The attributes of space include both the distance, and accessibility variables. These variables are part of any standard spatial interaction model[5], and are used to measure the effect of the spatial separation/arrangement of origins and destinations. In particular, the distance variable is an important aspect of any spatial interaction model, as it controls the magnitude of the distance decay parameter, which reflects the relationship between travel-to-work distance,

[4]All records with zero-valued flows have been omitted for visualisation purposes.

[5]The inclusion of an accessibility term is indicative of a competing destinations model, though in this case, it is an origin-centric measure of accessibility.

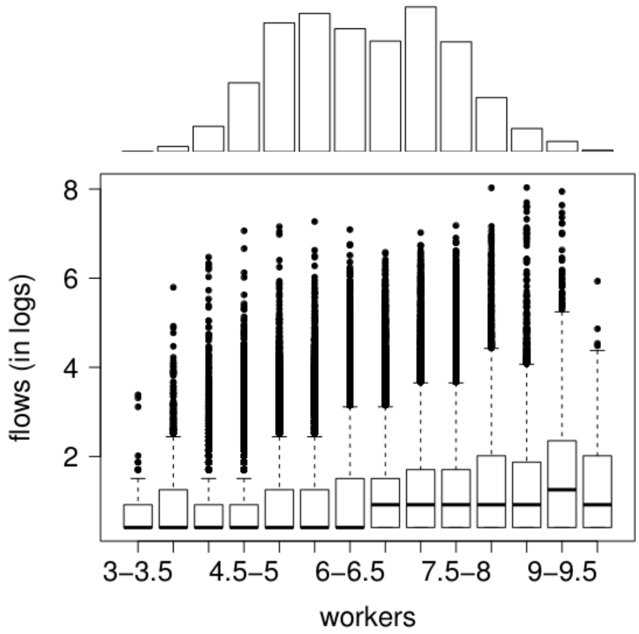

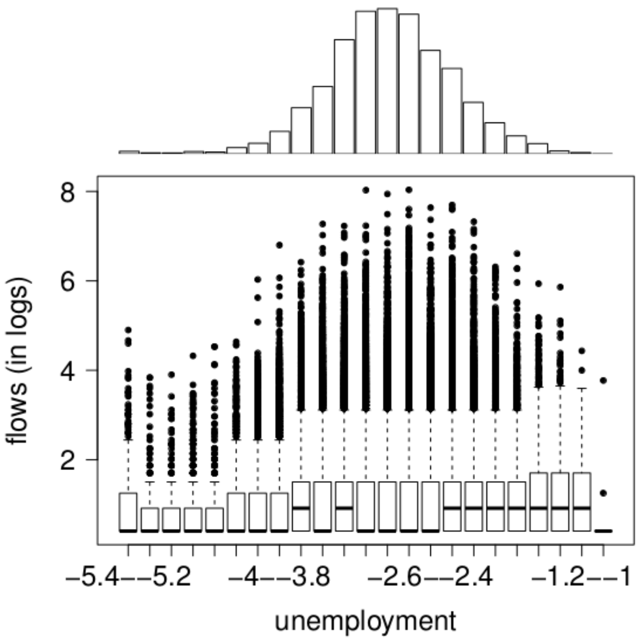

Figure 6.2: Plot of commuting flows (in logs) with the attributes of origins. The frequency distribution for all four variables are provided in the upper margin of their respective plots for reference [continued on following page].

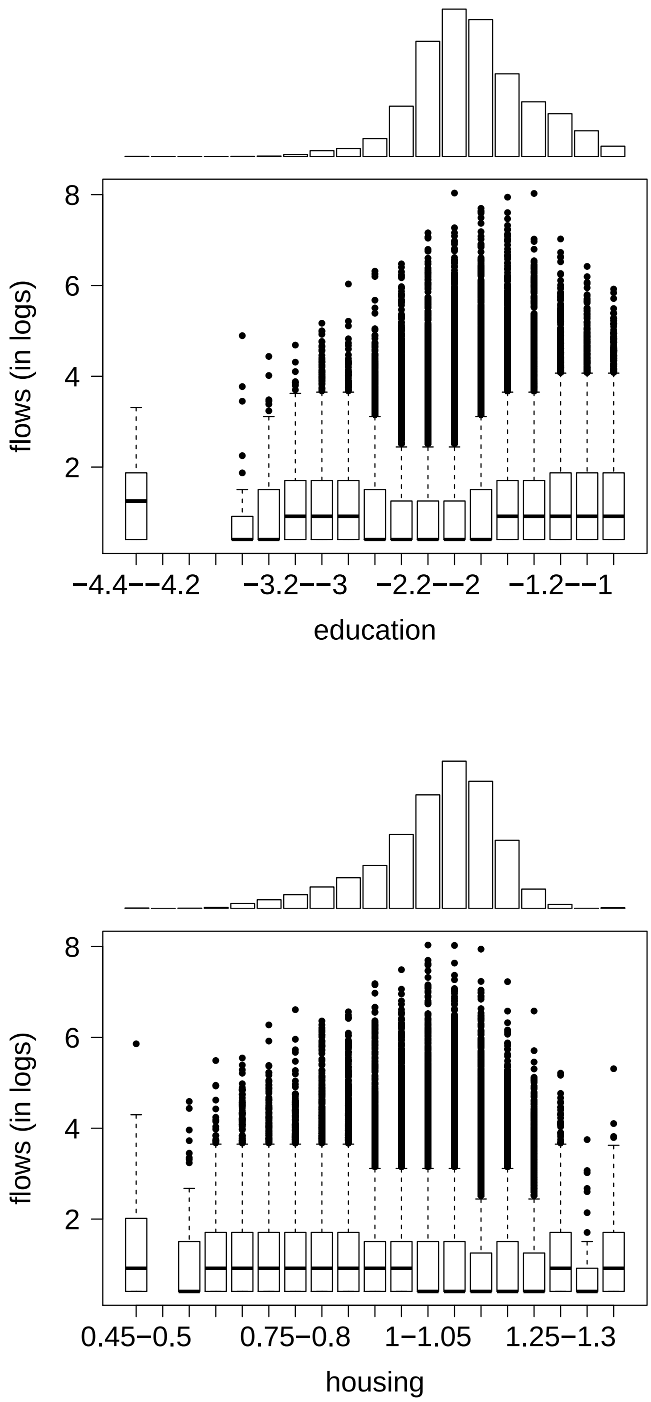

Figure 6.2: [continued from page 121]

and the commuting flows of the system. For the purposes of our analysis, we use straight-line distance[6] between the centroid of each origin and destination ED. As inter-ED distances are zero, and not of particular relevance to this work, we remove inter-ED interactions from the final model(s).

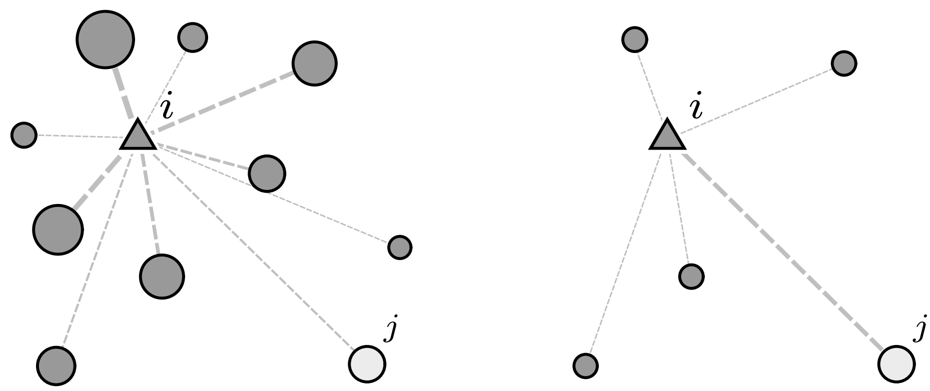

(a) (b)

Figure 6.3: Comparison of origin-centric accessibility to a destination where the origin i is surrounded by many large destinations (a), and where the origin is surrounded by only a few smaller destinations (b). Destination j is denoted by a lighter shade (lower left corner), and its relative accessibility to i (triangle) is represented by the dashed line thickness.

The accessibility term used here is based on the formulation presented in Equation 3.18, from Section 3.4.3; with some important modifications to suite the modelling task at hand. Firstly, because our focus is a destination-constrained model, the accessibility of destinations is no longer relevant. It is now much more important to represent the accessibility of origins, which is conceptually quite different. With an origin-centric accessibility measure, we are measuring the competition for interaction between origins and destinations, relative to an origin, rather than relative to a destination. Figure 6.3 provides

[6]Road-network-based distances would normally be preferred, however, previous modelling exercises (not shown) revealed no significant differences straight-line or road-network distances. As such, the more computationally simple straight-line distance is used throughout this thesis.

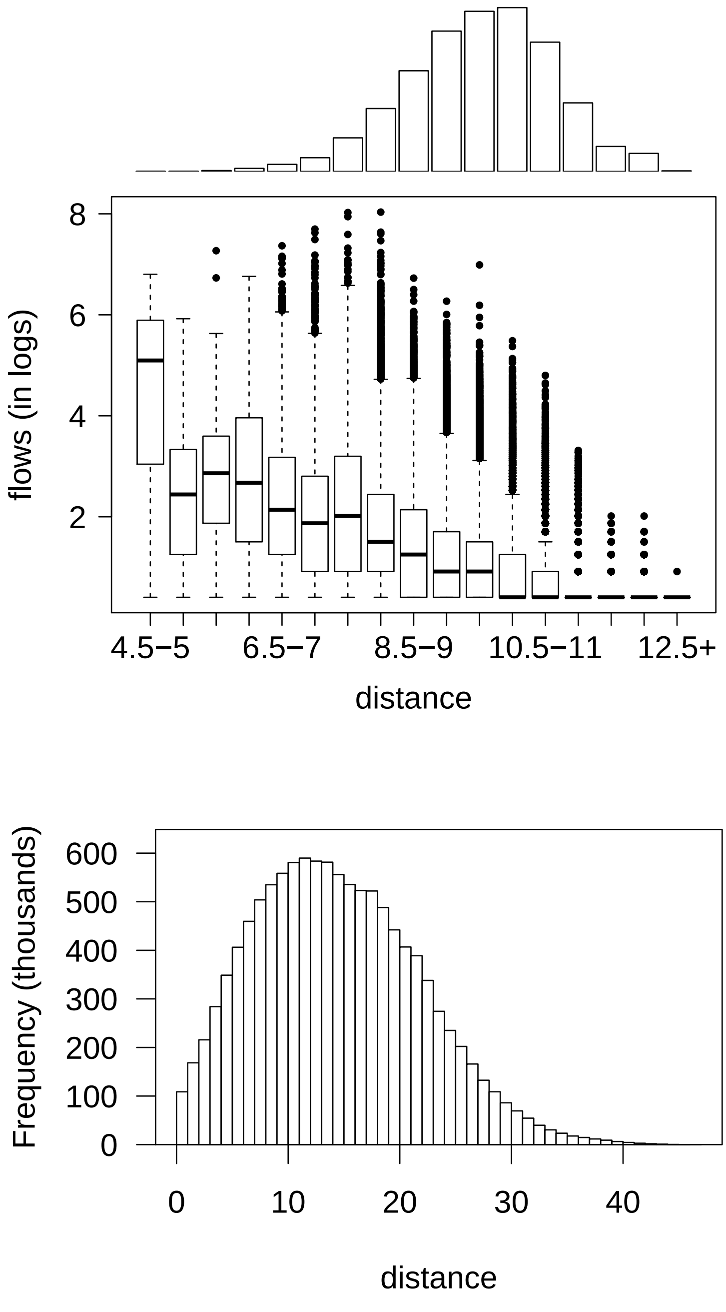

Figure 6.4: Plot of commuting flows (in logs) with both distance [this page], and accessibility [following page]. The frequency distribution for both distance and accessibility are provided in the upper margin of their respective plots for reference. Since distance and accessibility can be difficult to interpret in their log forms, we also include the 'raw' distance and accessibility histograms for reference.

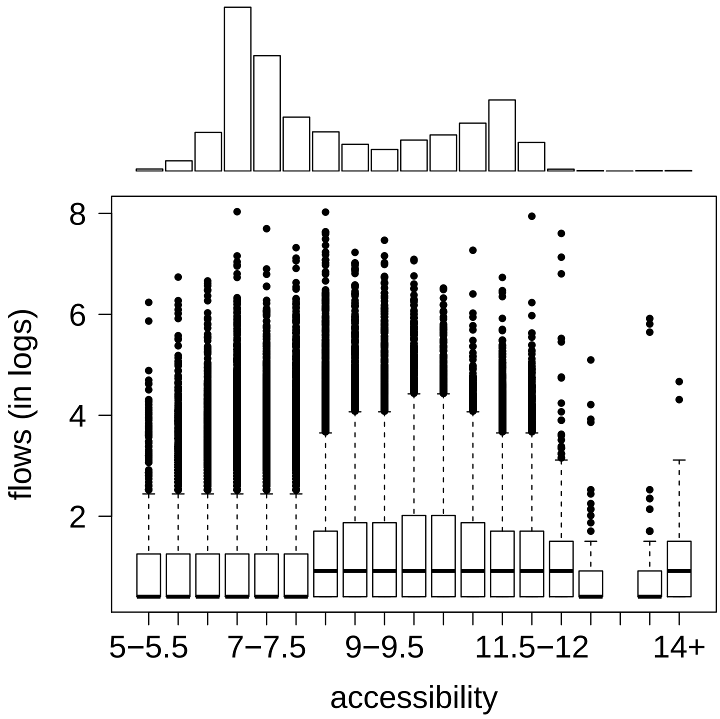

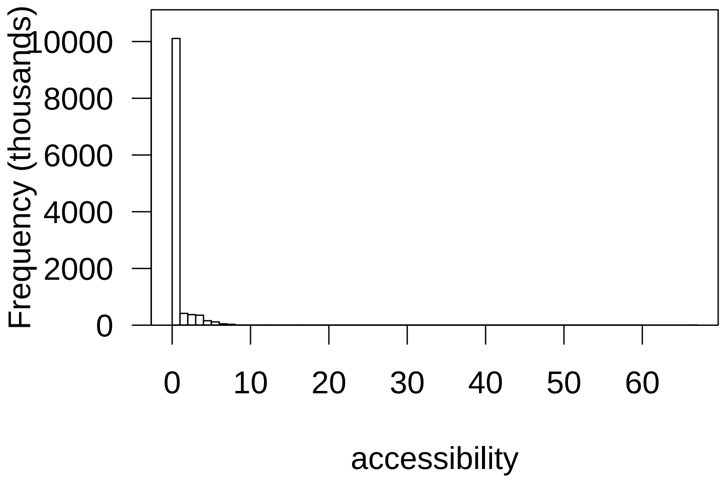

Figure 6.4: [continued from page 124]

a graphical representation of this relationship. This is conceptually similar to the concept of intervening opportunities that emerged from the spatial interaction literature in the past (Stouffer, 1960), such that we are measuring the level of 'intervening opportunities' that each individual faces when selecting a destination. From Figure 6.3, it can clearly be seen that an origin i, that is surrounded by many large potential destinations will be less likely to evaluate a particular destination j, from the set of possible m destinations, than a similar origin with a limited number of nearby destinations. In other words, *ceteris paribus*, a centrally located (highly accessible) origin will be less likely to evaluate a given destination than an isolated one.

Similar to Figure 6.2, Figure 6.4 depicts the relationship between flows and the attributes of space; along with the frequency distributions for both these variables. Additionally, because distance and accessibility can be difficult to interpret in their log forms, we also include the 'raw' histograms for reference. In particular, the distribution of commuting distances is in line with the large majority of empirical work on commuting in Europe and worldwide. This is quite an intuitive relationship and indicates that the friction of distance is indeed an important factor in the commuting process. The measure of accessibility presented here is slightly more difficult to interpret. As expected, there are a large number of 'inaccessible' EDs, with fewer moderate, to highly accessible ones. This distribution is reflected in the spatial (and aspatial) distribution of cities and towns throughout Ireland (see Figure 4.1 for a simplistic example of this). In terms of functional relationships, the inverse relationship between flows and distance is relatively clear, whereas the non-linear relationship between flows and accessibility likely reflects the bimodal distribution of accessibility, depicted in both the logged and 'raw' histograms.

6.4 Initial model results

We start by fitting the basic Poisson model (6.1) presented above. The second column in Table 6.2 (labelled POIS) contains the parameter estimates and standard errors for this initial model. All parameters are significant, with p-value< 0.01 (according to the regular Wald test), and their signs are in agreement with our expectations and the wider commuting literature. The distance decay parameter (-1.8766) is well within the expected range for commuting patterns in a country the size of Ireland, and indicates that for every percent *increase* in distance (in metres) between origin i, and destination j, we expect a 1.87% *decrease* in the number of commuters. Most of the remaining independent variables elicit similar (but weaker) responses in the dependent variable, whereas housing actually has the largest impact on flows of all 6 independent variables.

These results appear useful. However, based on our earlier exploratory analysis, it is possible that our model is misspecified due to over-dispersion (i.e., the Poisson assumption of equi-dispersion does not hold). As alluded to in Section 3.5.3, we can attempt to control for over-dispersion via pseudo- and quasi-likelihood methods. Columns 3 and 4 (labelled PPOIS and QPOIS respectively) of Table 6.2 provide modified standard errors from the and pseudo- and quasi-likelihood Poisson models, and inference is now based on these new standard errors. In both cases, all parameters remain significant, however, the standard errors are slightly more appropriate, particularly in the case of PPOIS. From the quasi-Poisson model, we also compute a dispersion parameter, which in this case is $\phi = 0.2260$. Alternatively, we can compute a dispersion estimate from the Poisson model such that $\mathrm{VAR}[y] = \phi\mu$, which yields $\phi = 1.4836$. Since neither of these values are equal to one, we can infer that our commuting flows display 'non-Poisson' properties. We can test this observation more formally by performing an auxiliary OLS regression (with-

Table 6.2: Model outputs based on initial evaluation of models.

	POIS	PPOIS	QPOIS	NEGBIN
(Intercept)	15.482	15.482	15.482	21.9374
	(0.0723)**	(0.2513)**	(0.0343)**	(0.169)**
distance	-1.8766	-1.8766	-1.8766	-2.5364
	(8e-04)**	(0.0082)**	(4e-04)**	(0.0026)**
accessibility	-0.6842	-0.6842	-0.6842	-0.5034
	(0.001)**	(0.0131)**	(5e-04)**	(0.0025)**
workers	0.6622	0.6622	0.6622	0.7506
	(0.001)**	(0.0124)**	(5e-04)**	(0.0029)**
unemployment	-0.1796	-0.1796	-0.1796	-0.0287
	(0.0022)**	(0.0166)**	(0.001)**	(0.0055)**
education	0.1295	0.1295	0.1295	0.1063
	(0.0033)**	(0.0265)**	(0.0016)**	(0.0084)**
housing	2.0211	2.0211	2.0211	0.6467
	(0.0099)**	(0.0989)**	(0.0047)**	(0.0272)**
-log Likelihood	1623421			858152
AIC	3253670			1723133
BIC	3302380			1771842
pseudo-R2	0.1945			0.5286

*$p < 0.05$; **$p < 0.01$

out intercept) between the observed and fitted values of the Poisson model (Cameron & Trivedi, 1990, 1998, p. 78 eq. 3.39). In this case, we reject the null hypothesis of Poisson (p-value < 0.01), indicating once again the presence of over-dispersion.

While controlling for over-dispersion by adjusting the standard errors in the regular Poisson regression (as we have done with the quasi- and pseudo-Poisson models) is an effective means of overcoming overdisperion and/or heterogeneity, they are in fact *ad hoc* adjustments, and it is therefore beneficial to explore more formal, parametric alternatives such as the negative binomial model. The negative binomial regression results are given in column 5 of Table 6.2 under the heading 'NEGBIN'. The range of parameters and their associated standard errors are similar to those from the earlier Poisson model and variants, save for some noticeable changes in magnitude for unemployment, housing, and distance. In both cases, the magnitude of the parameter has decreased, with a significant decrease in the magnitude of unemployment. This is not entirely surprising given unemployment's relationship with flows (Figure 6.2). Additionally, the distance variable has significantly increased in magnitude. This change is of particular interest, and the reasons for this large shift will become clear once we inspect the predicted probabilities of the Poisson and negative binomial models. For now, the shift to a stronger distance decay likely results from the negative binomial model's ability to capture over-dispersion resulting from the large number of zero flows.

Appreciable improvements in both the log likelihood, and information criteria (AIC and BIC) are observable between the Poisson and negative binomial models. Both statistics strongly support the use of the negative binomial form of the model. Furthermore, the pseudo-R^2 measures indicate that the negative binomial model is once again an improvement[7]. We also compare the negative binomial model with the Poisson model more formally via a likelihood ratio test (Lawless, 1987; Burger et al., 2009). Based on this result (p-value < 0.01), we once again reject the null hypothesis that the restriction implicit in the Poisson model is true. This provides additional support for implication the that over-dispersion is a problem here, and that the negative binomial model is a viable alternative to the Poisson model.

Before completely rejecting the Poisson model however, it may be beneficial to compare the expected probabilities generated by the Poisson and negative binomial models[8] to the observed probabilities. The most intuitive way to do this is to compare the observed and predicted *counts* for zero flows, and large

[7]Here we use McFadden's pseudo-R^2 measure (Agresti, 1990). It is important to note however, that pseudo-R^2 values must be treated with caution, as they are not equivalent to the OLS-based R^2 measures, and therefore should not be interpreted in the same manner. We use them here purely as a descriptive measure, and any/all inferences based on these measures are tempered by additional information and diagnostic tools.

[8]Since the quasi- and pseudo-Poisson models are not associated with a formal likelihood, predicted probabilities cannot be computed for these models.

flows (flows > 500), to get an idea of the model fit(s) at both tails of the distribution of commuting flows. This comparison highlights some interesting features of these two models: 1) the Poisson model grossly under-predicts the number of zero flows (11,206,613, with a difference of -191,322), contrarily, 2) the negative binomial model does quite well at predicting zero flow counts (11,400,328, with a difference of only +4,940), especially given the large number of zero flows in the dataset (11,395,388). For larger flows, 3) the negative binomial model does not fare as well, with a severe *over*-prediction of large flows (1,859, with a difference of +1720), similarly, 4) the Poisson model over-predicts the number of large flows (189-139=+50). While the fit at the upper end of the distribution may be problematic for prediction purposes, the large number of zeros in the dataset points to the increased importance of fitting this end of the distribution. As alluded to earlier, the stronger distance decay in the negative binomial model likely results from a better fit to zero and low value flows, which for the present dataset is of primary importance.

Table 6.3: Comparison of predicted counts for zero and large flows (Poisson and negative binomial models).

	POIS	NEGBIN	OBSERVED
Zero flows	11204066	11400328	11395388
Large flows†	189	1859	139

† flows > 500

Another common means of evaluating model fits is to assess them based on their residuals. In the context of GLMs, it is often of interest to examine the deviance residuals, against the fitted values. Deviance residuals are simply the square root of the contribution of an observation to the deviance, with the same sign as the raw residual, or, more formally, $d_i = \text{sign}(t_{ij} - \hat{\mu}_i)\sqrt{2(l(t_{ij}) - l(\hat{\mu}_{ij}))}$ where $l(\hat{\mu}_{ij})$ is the log-density of t_{ij} at $\mu = \hat{\mu}$ and $l(t_{ij})$ is the log-density of t_{ij} evaluated at $\mu = t_{ij}$. The deviance residuals are often considered the most useful for diagnostic purposes (Venables & Ripley, 2002), in particular because the sum of squares of these residuals is equal to the deviance, or residual deviance, which is essentially a GLM generalisation of the residual sum of squares from OLS. If a model fits the data well, the deviance residuals should form a band of points along the x-axis, within a range of about ± 3. Figure 6.5 displays the deviance residuals plotted against the fitted values for both the Poisson and negative binomial models. The negative binomial deviance residuals fall mostly within a range of -3.5 to 4, with extremely low deviance residuals at the upper end of the distribution. This pattern does not suggest a poor fit for any particular observation or subset of observations, except for perhaps $t_{ij} \approx 0$, which appears to display slightly higher positive residuals than we would like. This relatively good fit for the negative binomial model is in contrast to the Poisson model, whose residuals range from approximately -88 to 108, with a pattern of positive residuals at small values of t_{ij} and negative residuals for larger predicted commuting flows. Clearly then, the negative binomial model fits better than the Poisson model based on a visual inspection of the residuals.

An additional measure of goodness-of-fit is based on a χ^2 test on the residual deviance, which was mentioned in the preceding paragraph. An alternative representation of the residual deviance is simply 2 times the difference between the log-likelihoods of the current model M, and the saturated model S[9], or $D = 2(\mathcal{L}(M) - \mathcal{L}(S))$. Since the residual deviance is essentially the difference between the deviance of the given model and the deviance of the saturated model, a small residual deviance indicates a good fit to the data. This will result in the χ^2 test not being significant. In this case, we do not reject the null hypothesis for either model (p-value$>$ 0.05), indicating that they both fit the data reasonably well. If however, we would like to know which model is closest to the 'true' model, we can test this using the Vuong closeness test

[9]A saturated model is one in which there are n parameters (i.e., one parameter for each observation), leading the predicted values to be exactly equal to the observed values.

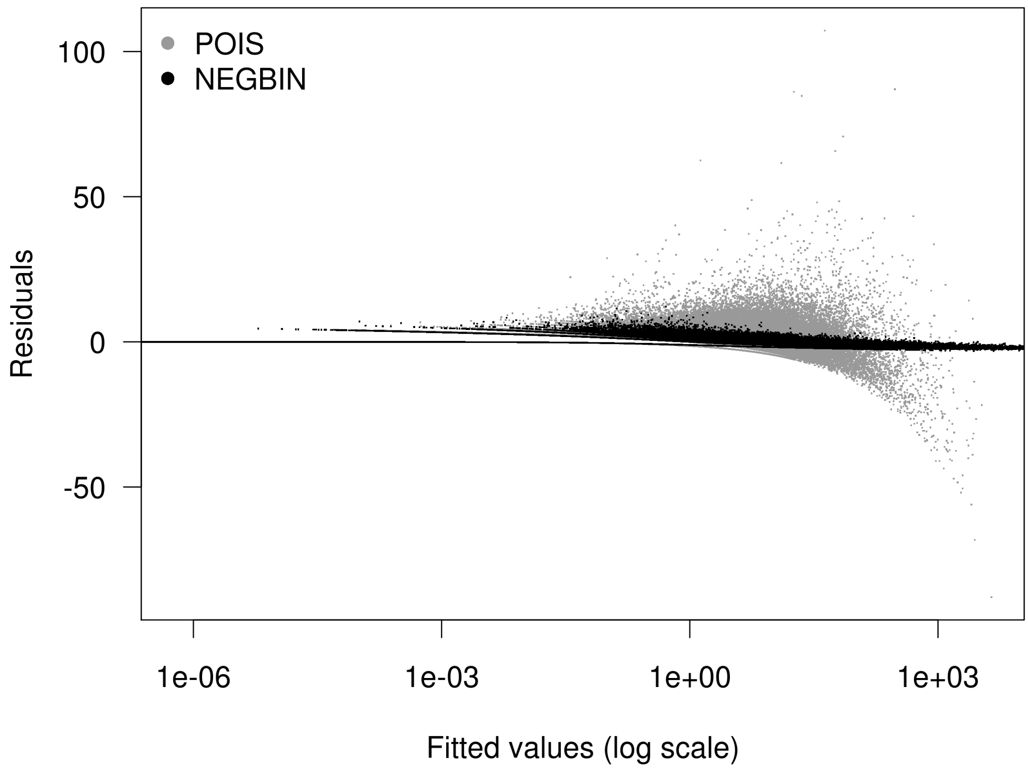

Figure 6.5: Deviance residuals for the initial Poisson and negative binomial models plotted against their corresponding fitted values. The fitted values are plotted on a log-scale. Note that the range of predicted values for the two models differ, with the negative binomial model predicting much larger maximum values than the Poisson model.

(Vuong, 1989). The Vuong test can be used to test the null hypothesis that two (non-nested) models are equally close to the true model. It is based on a comparison of the predicted probabilities of the two models, and under the null hypothesis that the two models are indistinguishable, it is normally distributed. Generally, large, positive values indicate that model one is preferred over model two, and *vice versa* for large, negative values. Intermediate values indicate that neither model is preferred. A Vuong test between the Poisson and negative binomial models indicates that the negative binomial model is preferred over the Poisson model (p-value < 0.01).

One explanation for the superior fit of the negative binomial model is that, in addition to observed zeros due to distance decay and other relevant parameters, there may be structural[10] zeros in the dataset which are induced by restrictions on commuters' movements due to the influence/conditions of the local labour market within which they live and work. In other words, it is possible that there are latent labour market effects that are causing overdispersion, excess zeros, and general non-Poissoness in the commuting flows [11]. We can test this assertion more formally by checking to see if the internal flows come from the same, or different distribution to the overall commuting flows. In this case, we base internal versus external flows on the functional regions developed in 5. The χ^2 test for homogeneity is then used to evaluate whether the internal flows come from the same distribution as the overall flows. This leads us to reject the null hypothesis (p-value < 0.01), suggesting that local commuting flows appear to come from a different distribution (and by proxy, are generated by a different process). This observation lend credence to our earlier statements about local labour market effects, and provides the motivation needed to explore these effects further in the following Section.

6.5 Internal & external commuting

Now that we have a suitable model which appears to fit relatively well (negative binomial), we can explore the effects of the spatial structure of local labour markets on our model and discuss some of these variations in the context of national commuting patterns. We start by specifying an indicator variable (internal) that characterises whether a commuting flow is within, or between

[10] In fact, these zeros are not true structural zeros, but are zeros induced by the spatial structure of local labour markets.

[11] We provide an alternative interpretation for some of these problems (in particular, excess zeros) in Chapter 8 based on choice set limitations. This alternative interpretation appears to have strong merit, however, it requires further evaluation before it can be a complete alternative to the models already presented here.

local labour markets. Here we utilise the spatial boundaries of the functional regionalisation presented in Chapter 5 (Figure 5.3b) to represent the boundaries of the local labour markets. In this case, we have 65 spatially contiguous local labour markets, each encompassing a given region of the overall study area. Summary statistics for these local labour markets are given in Table 6.4. The internal variable indicates that a large majority of the interactions in the commuting dataset occur *between* local labour markets. This is to be expected, given that the largest local labour market (in terms of the number of underlying EDs) is a periphery region (see Section 5.4.2 for an explanation of this term) with 174 internal EDs. Despite the smaller number of internal *connections* (2.72% of the total connections), over one quarter (25.81%) of the positive flows occur within local labour markets, which accounts for 67.61% of the total interaction (i.e., sum of flows) in the dataset! Clearly then the boundaries of local labour markets are a relevant descriptor of commuting trends, and in the following paragraphs, we will show that there are significant differences between model fits when we take into commuting within/across these boundaries.

Table 6.4: Summary statistics for local labour market regions. Counts refer to the number of underlying EDs contained within a given local labour market (functional region), and coverage is the percentage of the total commuting flows that begin in a given local labour market.

	Min.	Median	Mean	Std. Dev.	Max.
Areas (km^2)	57.45	585.05	1081.39	1081.46	5027.75
Counts	3	31	52	46	174
Population	2192	31919	65228	91098	517269
Coverage (%)	0.09	0.91	1.54	1.36	5.10

In order to explore the differences between internal and external commuting, we fit two additional models, each designed to address local labour market effects in slightly different ways. Firstly, we fit a negative binomial model, again of the form given in Equation 3.29, with μ similar to (6.1), with one additional indicator variable distinguishing between internal (1) and external (0) commuting. As such, λ_{ij} becomes

$$\lambda_{ij} = \exp(I + \varphi \mathrm{i}_{ij} + \psi_j + \chi \mathrm{w}_i + \omega \mathrm{u}_i + \delta \mathrm{e}_i + \eta \mathrm{h}_i + \gamma \mathrm{c}_i + \beta \mathrm{d}_{ij}), \tag{6.2}$$

where all variables and parameters are as defined before and the additional variable i_{ij} and associated parameter φ are used to measure the extent to which internal commuting differs from external commuting. Our second modified commuting model is designed to capture the differences between internal and external commuting by estimating a different set of parameters for each of these two commuting actions. This is equivalent to specifying a categorical interaction term with each of the 6 parameters in the model, such that

$$\lambda_{kij} = \exp(I_k + \psi_j + \chi_k \mathrm{w}_{ki} + \omega_k \mathrm{u}_{ki} + \delta_k \mathrm{e}_{ki} + \eta_k \mathrm{h}_{ki} + \gamma_k \mathrm{c}_{ki} + \beta_k \mathrm{d}_{kij}), \tag{6.3}$$

where $k = 1$ for external commuting, and $k = 2$ for internal commuting. This is akin to allowing a different slope (and intercept) for external versus internal commuting while keeping the destination constraints (fixed-effect) constant. The practical implication of this model is that we arrive at a different set of parameters for external versus internal commuting in a single model:

$$\lambda_{ij} = \begin{cases} \exp(I + \psi_j + \chi_1 \mathrm{w}_i + \omega_1 \mathrm{u}_i + \delta_1 \mathrm{e}_i + \eta_1 \mathrm{h}_i + \gamma_1 \mathrm{c}_i + \beta_1 \mathrm{d}_{ij}) & k = 1 \\ \exp(I + \psi_j + \chi_2 \mathrm{w}_i + \omega_2 \mathrm{u}_i + \delta_2 \mathrm{e}_i + \eta_2 \mathrm{h}_i + \gamma_2 \mathrm{c}_i + \beta_2 \mathrm{d}_{ij}) & k = 2 \end{cases}$$

This allows us to take advantage of a *pooled* estimate of variance, which is more discriminatory than simply fitting two separate models (Venables & Ripley, 2002). In fact, this model provides similar parameter estimates to two separate models (not shown), where models for internal and external commuting are each fitted separately (while holding destination constraints constant).

Model parameters and diagnostics for both these models (as well as the previous negative binomial model for reference) are given in Table 6.5. Once again, nearly all parameters are significant (p-value< 0.01), with expected sign and magnitude. One exception is the external-specific education variable. The results from Chapter 5 provide a possible explanation for this observation: well educated individuals tend to commute further afield, and are therefore more likely to cross local labour market boundaries. The direct consequence of this is that since commuters within a local labour market will likely encompass individuals from a range of education backgrounds (people in general will prefer to live/work in relatively close proximity, regardless of education level), we would not expect there to be any significant variation in commuting patterns due to education levels within a given functional region. Conversely, for commutes which *cross* functional region boundaries, because highly educated individuals are more likely to commute long distances than their less educated colleagues, we would expect a significant (albeit weak) positive relationship between commuting volume and education levels for external commuting. This observation is in direct correspondence with our findings from Chapter 5, Section 5.4.3.

Additional differences between the models arise as a result of their alternative formulations. For example, while the model based on (6.2) (denoted INTERNAL in Table 6.5) has nearly identical parameters to the earlier negative binomial model (denoted NEGBIN), the addition of the internal indicator variable has increased its likelihood (both AIC and BIC also indicate that the INTERNAL model is an improvement over the initial NEGBIN model). Furthermore, the internal parameter is statistically significant, indicating that the difference between internal and external commuting is important. Indeed, by computing $\exp(\varphi)$ (where φ represents the parameter for internal) to arrive at the direct effect of internal on commuting flows, we find that the volume of internal commuting is approximately 1.06 times higher than that for external

Table 6.5: Model outputs from internal/external models.

	NEGBIN	INTERNAL	DUAL	
			External	Internal
(Intercept)	21.9374	21.7618	22.2907	19.9392
	(0.169)**	(0.1702)**	(0.1723)**	(0.183)**
distance	-2.5364	-2.5209	-2.6078	-2.143
	(0.0026)**	(0.0032)**	(0.0037)**	(0.0056)**
accessibility	-0.5034	-0.5018	-0.4525	-0.6379
	(0.0025)**	(0.0025)**	(0.0028)**	(0.004)**
workers	0.7506	0.7508	0.7924	0.5846
	(0.0029)**	(0.0029)**	(0.0034)**	(0.0053)**
unemployment	-0.0287	-0.0292	-0.0016	-0.084
	(0.0055)**	(0.0055)**	(0.0066)	(0.0093)**
education	0.1063	0.1076	0.0709	0.0797
	(0.0084)**	(0.0084)**	(0.0103)**	(0.0138)**
housing	0.6467	0.6458	0.4139	0.9796
	(0.0272)**	(0.0273)**	(0.0328)**	(0.0454)**
Internal dummy		0.059		
		(0.007)**		
-log Likelihood	858152	858118	848759	
AIC	1723133	1723067	1704363	
BIC	1771842	1771790	1753171	
pseudo-R^2	0.5286	0.5286	0.5228	

$^{*}p < 0.05$; $^{**}p < 0.01$

commuting *ceteris paribus*.

Returning now to the internal/external-specific commuting model (denoted DUAL in Table 6.5), it is interesting to note that the distance decay for internal commuting ('Internal' sub-heading) is weaker than that for external commuting. This is likely a result of latent local labour market effects which may be altering the nature of distance decay; interactions *across* local labour market boundaries are reduced due to latent costs in working outside ones local labour market (i.e., differences in local labour market characteristics, etc.). In other words, distance decay is essentially capturing a combination of geographical distance (or other forms of distance) and regional differentiation due to labour market boundaries, leading to 'functional' distance decay (Noronha & Goodchild, 1992). This functional distance decay is then an additive factor of distance and external friction, hence the slightly reduced level of distance decay for internal commuting. In our DUAL model, the difference between distance decay parameters may be interpreted as the 'cost' associated with crossing a local labour market boundary. We will return to this concept of 'functional distance decay' in the following chapter.

Comparing the internal models in a more formal sense to the previous negative binomial model suggests that the DUAL model is a significant improvement over the other negative binomial-based models. This can be confirmed by examining the information criteria, and likelihoods of the models. Additional diagnostic tests, such as the Vuong test confirm this observation. For instance, Vuong tests between all three models indicates that INTERNAL is preferred over NEGBIN, and DUAL is preferred over both NEGBIN and INTERNAL (p-values < 0.01). While the deviance residuals of the three models (not shown) are quite similar, differences do arise when we examine their predicted probabilities (see Table 6.6), as in the previous section. For example, a comparison of predicted numbers of zero flows reveals that the INTERNAL model predicts nearly the same number of zeros as the NEGBIN model, whereas the DUAL model predicts over 3000 fewer zeros, bringing it significantly closer to the observed count (11397225-11395388=+1837). Similarly for large flows (flows > 500), the DUAL model provides an estimate that is approximately 38.30% less than the negative binomial model, and significantly closer to the true number of large flows (1147-139=+1008). The similarity of the NEGBIN and INTERNAL models is not surprising given the small difference in these two models, however, the significant improvements in fit induced by the DUAL model strongly points to local labour market effects in action.

As we have alluded to previously, treating internal and external flows separately improves the fit of our model, and increases the likelihood of our model substantially. It also provides a means of examining the differences in com-

Table 6.6: Comparison of predicted counts for zero and large flows (negative binomial models).

	NEGBIN	INTERNAL	DUAL	OBSERVED
Zero flows	11400328	11400361	11397225	11395388
Large flows†	1859	1865	1147	139

† flows > 500

muting when travelling within/between local labour markets. In addition to the differences in distance decay parameters due to 'functional distance decay' mentioned previously, several other important observations can be made. For example, external commuting is more dependent on the number of available workers than internal commuting. This is an intuitive result, because we would expect that regardless of the number of workers in a particular region, a larger majority of them will commute locally than externally *ceteris paribus*, which is certainly what we are observing here. In terms of education levels however, there is very little difference between internal and external commuting (i.e., small positive parameter in both cases). This presents a reduction in education effects from the NEGBIN and INTERNAL models, and indicates that education plays a similar role in commuting behaviour regardless of the spatial structure of local labour markets.

The largest difference in magnitude between internal and external specific parameters is for the housing variable. We find that housing has a stronger impact on *local* commuting than it does on *external* commuting. Since housing can be interpreted as a measure of how 'urban' a particular region is, this means that highly 'urban' areas will generate more local flows than they will external flows *ceteris paribus*. Unemployment is not significant for external flows (i.e., flows from an origin to a destination not within the same local labour market), whereas for local flows, the effect of unemployment has actually increased from the NEGBIN model. This is not surprising given that those most at risk for unemployment (i.e., lower education, income, skills), are also those less likely to commute longer distances, and by proxy, less likely to commute outside their local labour market (see Section 5.4.3). Origin-centric accessibility has a similarly negative effect on commuting flows and appears to have a slightly stronger (negative) impact *within* local labour markets than *between* them. Again, this is an intuitive result, and indicates that more centrally located origins will produce fewer local commuters than it will external commuters, but that *ceteris paribus*, in both cases more isolated origins will display higher levels of commuting in general.

Each of the above model diagnostics and comparisons provide increasing evidence that there are important differences between commuting at the national and local levels, and in particular, that there are important local labour market effects that must be taken into account when considering travel-to-work behaviour. In summary, we have found strong evidence that commuting flows decrease rapidly with distance, and that the centrally located origins (i.e., higher accessibility) will tend to export fewer commuters than their isolated neighbours *ceteris paribus*. Furthermore, we see that when over-dispersion is accounted for via the negative binomial model, the effects of unemployment and housing density on commuting flows decreases, indicating that these variables may be associated with heterogeneity in commuting behaviour. When accounting for local labour market effects, we find that there is a significant positive increase in internal commuting flows, and that considering internal and external commuting separately significantly increases the fit of our model. These observation lend credence to our earlier statements about local labour market effects, and provides the motivation needed to explore these effects further in the following chapter.

6.6 Moving forward

In this chapter, a range of spatial interaction models have been empirically tested and evaluated in the context of travel-to-work patterns. We have found that, in addition to distance and working population size, the spatial structure of origins and destinations, as well as a number of aspatial attributes such as unemployment, housing density, and education, all significantly affect commuting patterns. Furthermore, we have shown that variations in commuting processes appear to exist between (and within) local labour markets as defined by the functional regions developed in Chapter 5. Additionally, we have provided evidence for the superiority of negative binomial-based spatial interaction models over standard Poisson spatial interaction models in the context of sparse commuting flows. All of these assertions are supported by concrete statistical results and spatial interaction theory, which provides the evidence necessary to move forward with our selected empirical model.

What we have yet to explore is the extent to which commuting behaviour and patterns vary across local labour markets. This is a fundamentally different perspective on local labour market effects and one that has important implications for understanding travel-to-work behaviour. Questions such as 'how do the commuting patterns within a particular region influence the local labour market'? and perhaps more importantly, 'how does the spatial structure of local labour markets influence commuting patterns and behaviour?' are corollaries of this type of perspective. The 'global' perspective on local labour market effects that we have pursued in this chapter hints towards answers to these questions. However, a more explicit treatment of 'local' versus 'global' commuting patterns will ultimately lead to a more thorough understanding of variations in commuting behaviour/patterns. In the following chapter, we utilise the insights and models developed in this chapter, coupled with the local labour market boundaries derived in Chapter 5, to arrive at a single ‘inter-regional’, as well as a set of ‘region-specific’ commuting models. Both types of models are designed to explore the spatial relationship(s) between travel-to-work patterns/behaviour and local labour markets.

Chapter 7

Local labour market effects

7.1 Introduction

In the preceding chapter we developed a destination-constrained competing-destinations Poisson spatial interaction model for commuting based on a number of attributes of origins and space. We subsequently showed that such a model can be improved if we take into account over-dispersion due to unobserved heterogeneity in commuting flows using the negative binomial model. This latter model was evaluated by exploring the effects of local labour market structure on the parameters of the model by controlling for the effects of internal versus external commuting. To this end, the functional regions defined in Chapter 2 and implemented in Chapter 5 were used to represent the spatial manifestation of local labour markets and to characterise internal and external commuting. This has lead to strong evidence in favour of the idea that local labour market effects are acting on commuting flows in the Irish commuting data and provides the impetus needed to continue our analysis at a finer scale.

In this chapter, we focus primarily on the empirical evaluation and interpretation of two complementary modelling frameworks. The goal is to provide the final evidence needed to tie the preceding chapters together to provide a unified and coherent exposition of local labour market effects on commuting. With this in mind, in the following section we calibrate an aggregate 'global' model of commuting flows *between* local labour markets in order to get a sense of the overall level of inter-regional commuting in Ireland. Based on these results, we make several comparisons with previous results at the national level and highlight some of the more novel revelations of the inter-regional model. Subsequent to this, we calibrate 'local' commuting models at the level of the individual local labour markets in order to characterise the spatial variations in intra-regional commuting across Ireland. Additional insights and comparisons between the local and global models are provided and a segue into the concluding chapter (Chapter 9) is provided in Section 7.4.

7.2 Inter-regional commuting

Our initial modelling exercise in this section is concerned with inter-regional commuting patterns and outcomes. Here, we have aggregated commuting flows and the various attributes used in our models to the boundaries of the local labour markets derived in Chapter 5. New values for each of the attributes of origins and attributes of space were computed based on the new origin populations and spatial structure. Based on the summaries provided in Table 7.1, we can assume that, aside from changes due to aggregation, the newly derived variables are within expected ranges and distributions.

An evaluation of the dependent variable is provided in Figure 7.1, which shows two histograms of aggregate commuting flows; one using raw commuting flows (left), and the other in logs (right). Clearly the distribution of inter-regional commuting flows exhibits both substantial variation (i.e., over-dispersion) and a large number of zeros. Indeed, the variance-mean ratio for flows is extremely high (33074.58) and strongly indicates that the Poisson

Table 7.1: Summary statistics for model variables.

	Min.	Median	Mean	Std. Dev.	Max.
flows*	0.00	1.00	276.84	3025.97	92997.00
distance (km)	4.07	140.95	146.18	72.51	398.45
accessibility	344.13	1603.47	6961.99	13986.61	82099.62
workers[1]	1060.00	14818.00	32001.94	47107.78	280452.00
unemployment[2]	0.04	0.07	0.07	0.02	0.13
education[3]	0.13	0.17	0.17	0.03	0.27
housing[4]	2.59	2.82	2.82	0.10	3.17

*dependent variable (commuters); 1 working individuals; 2 unemployment rate
3 proportion; 4 persons/household

model is still an inappropriate distribution for our model. As this is the case and we have previous evidence indicating that the negative binomial model provides superior performance for modelling commuting flows, we will continue using our negative binomial-based model in the remainder of this chapter[1].

[1]As mentioned previously, a potentially much more power modelling framework (particularly in the presence of excess zeros) is presented in Chapter 8 and may provide a viable alternative models in future analyses.

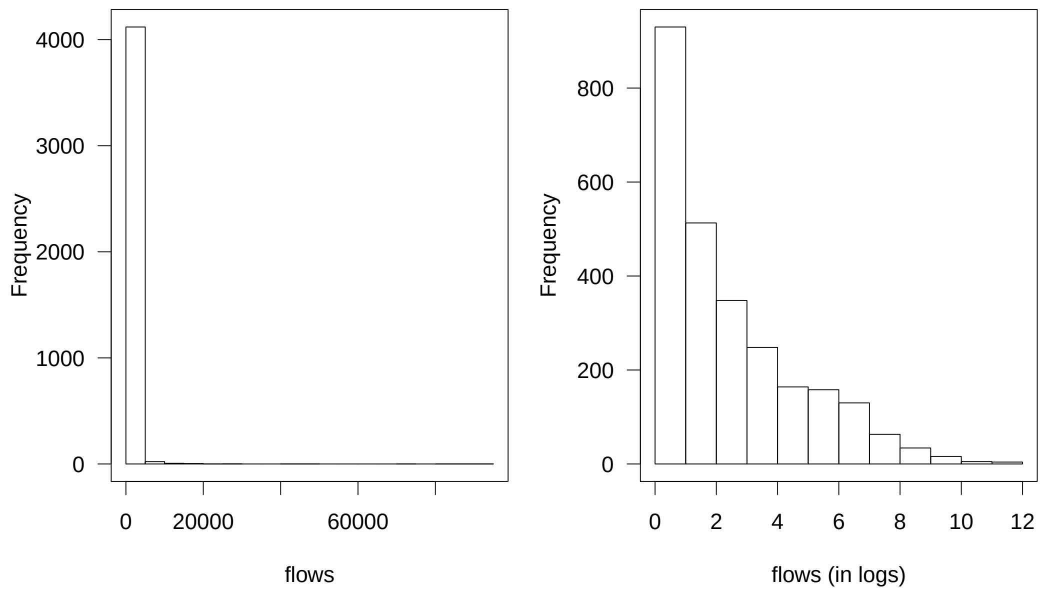

Figure 7.1: Frequency distribution of inter-regional commuting flows. Flows are given as raw counts (left) and in logs (right).

We start by calibrating the basic negative binomial model based on Equation 6.1 presented earlier. As in previous modelling exercises, Table 7.2 contains parameter estimates, standard errors, significance levels, and several model diagnostics for the inter-regional model (INTER-REGION). According to a partial Wald test, the accessibility and unemployment parameters are not significant when considering inter-regional flows. The remaining parameters are significant and display some important differences when compared with previous model parameters. Refitting the model with unemployment and/or accessibility removed (not shown) does not appear to provide any significant improvement in fit, which is confirmed by a likelihood ratio χ^2 test in which the null hypothesis is not rejected (p-value> 0.05). Furthermore, the magnitude and sign of the remaining parameters does not appear to be affected by the inclusion/removal of the non-significant parameters. The general suitability of this initial model can be assessed using a similar χ^2 test (based on a comparison with a null model) and suggests that the model fits the data reasonably well (p-value< 0.01). This observation is confirmed by the pseudo-R^2 (0.7196) measure provided in Table 7.2.

Further evaluation of the fit of our inter-regional model is provided by an examination of the deviance residuals as in the previous chapter. Figure 7.2 shows the deviance residuals from this model with the predicted commuting flows in logs to better visualise the deviance residuals for the large flows between local labour markets. As in previous models, the bulk of our deviance residuals fall between approximately -3 and +3, with a few outliers at medium predicted flows. It is interesting to note the patterns in the deviance residuals for smaller predicted flows, denoted in a lighter shade of grey in Figure 7.2. In fact, these five linear features are associated with observed flows of less than or equal to 5 and appear to suggest that for relatively small flows, the INTER-REGION model trends towards decreasing prediction magnitude. These de-

Table 7.2: Model outputs based on flows between local labour markets.

	INTER-REGION
(Intercept)	41.2426
	(0.7806)**
distance	-3.7585
	(0.0329)**
accessibility	-0.0604
	(0.0306)*
workers	1.1508
	(0.0217)**
unemployment	0.1923
	(0.0993)
education	1.0003
	(0.1238)**
housing	-3.5466
	(0.6682)**
-log Likelihood	10355
AIC	20854
BIC	21310
pseudo-R2	0.7196

*$p < 0.05$; **$p < 0.01$

viance patterns were not present/obvious in the previous negative binomial models and suggests that the INTER-REGION model may in fact be deficient for predicting smaller flows. This observation may be explained in part by the significantly fewer number of smaller flows resulting from the aggregate nature of inter-region commuting. Besides the potential problems associated with these residual patterns, the magnitude of the deviance residuals indicate that this model behaves relatively well for medium to large flows.

Examining the parameters more closely, we find that there are some major differences is sign and magnitude when compared with the models fitted to the full commuting dataset. Certainly we would expect differences in the parameters given that the models are fit to different data and are modelling different processes; however, these differences are important and bear further investigation. In all cases[2], the magnitudes of the parameters are larger than

[2]Here we are referring only to those parameters that were found to be significant.

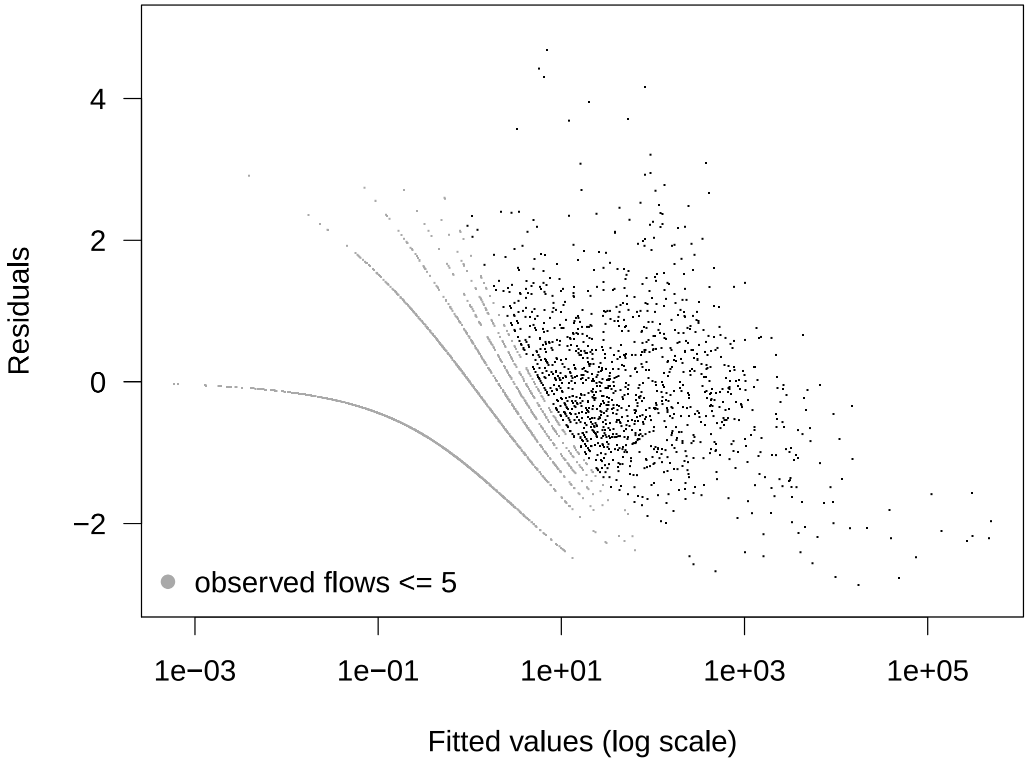

Figure 7.2: Deviance residuals from modelling commuting flows *between* local labour markets (Table 7.2), plotted against their corresponding fitted values. Note that the fitted values are plotted on a log-scale to better visualise the large flows between local labour markets.

expected from previous models, indicating that the individual variables have a larger influence on commuting flows than in previous models[3]. This being the case, it is important to consider the parameters of *this* model in the correct context; they are based on modelling a different process than previous models (i.e., inter-regional commuting only).

Looking first at the housing variable, we see that the parameter is strongly negative, indicating that, contrary to previous models, regions with a high number of occupants per household will tend to produce fewer outgoing com-

[3]Note that in many modelling contexts, it does not make sense to compare parameters from different models across different (or modified) datasets, however, since our attributes are in logs, we are comparing changes in *percentages* rather than raw *values*, which allows us to compare parameters between the various model forms.

muters than local labour markets with a large number of single-person households. This is an interesting finding and, one that, despite being superficially contradictory, actually corroborates previous findings. For example, in our DUAL model from Section 6.5, we found that internal commuting had a larger housing parameter than external commuting, indicating that, for a given level of housing density, more commuters would commute locally than externally *ceteris paribus*. Now that we are essentially modelling external commuting exclusively, we would expect fewer out-commuters in general and, for a particularly 'urban' local labour market (i.e., one with many individuals per household), we would expect fewer still; hence the strong negative parameter. In fact, according to our inter-regional model, we can expect over 3% fewer commuters from region i to j for every 1% increase in the number of individuals per household in i, *ceteris paribus*.

The strong negative housing parameter is partly offset by the relatively strong positive parameter for workers, which is twice as high as in previous models. In the context of inter-regional commuting, it indicates that regions with a large number of resident workers will export a much larger number of out-commuters to a given destination. Indeed, for a 1% increase in the number of resident workers in a local labour market, we should observe a 1.2% increase in the number of commuters *ceteris paribus*. In a similar vein, whereas education has had a relatively small (but significant) impact on commuting flows in previous models, a parameter of 0.8018 in our inter-regional commuting model suggests a much stronger relationship between the proportion of highly educated individuals and commuting flows. This likely stems from the fact that well-educated individuals are more likely to commute longer distances than their less educated counterparts (i.e., they will have better access to jobs, higher income levels, and a larger job search area), leading to a higher volume of inter-regional commuting for regions with a higher proportion of well educated individuals.

Differences in distance decay between inter-regional commuting and commuting in general reveal further insights into the relationship between commuting and the structure of local labour markets. It is useful to compare inter-regional distance decay with internal, external, and global distance decay parameters from our previous models. The inter-regional distance decay indicates that for every percent *increase* in distance between origin i and destination j, we can expect a 4.01% *decrease* in the number of commuters. In comparison with previous distance decay parameters, this is quite high, and points again to the concept of functional distance decay. For instance, global distance decay from our NEGBIN model in the previous chapter was -2.5364, which is over 1.5 times smaller than inter-regional distance decay presented here. This provides further evidence that distance decay is influenced by local labour market effects, suggesting that it may be the result of regular distance decay, plus an additive regional differentiation factor. For inter-regional commuting, this effectively produces a 'functional distance decay' as depicted in Figure 7.3, and discussed in Noronha & Goodchild (1992) and Chapter 6 of this thesis.

We can visualise this difference empirically by examining the changes in predicted commuting flows for varying distances based on the NEGBIN and INTER-REGION models (Figure 7.4). The range of predicted values is given by 1.5 times the upper and lower interquartile range and the central measure is given by the median. To properly compare predictions for various distances, all variables (except distance) were fixed at respective their means. The predicted values are in logs to represent the log-linear relationship between flows and distance. Ignoring for a moment the obvious difference in flow magnitude between the two curves (which can be explained by the difference in sheer volume of commutes for the two models [i.e., flows from a single ED versus

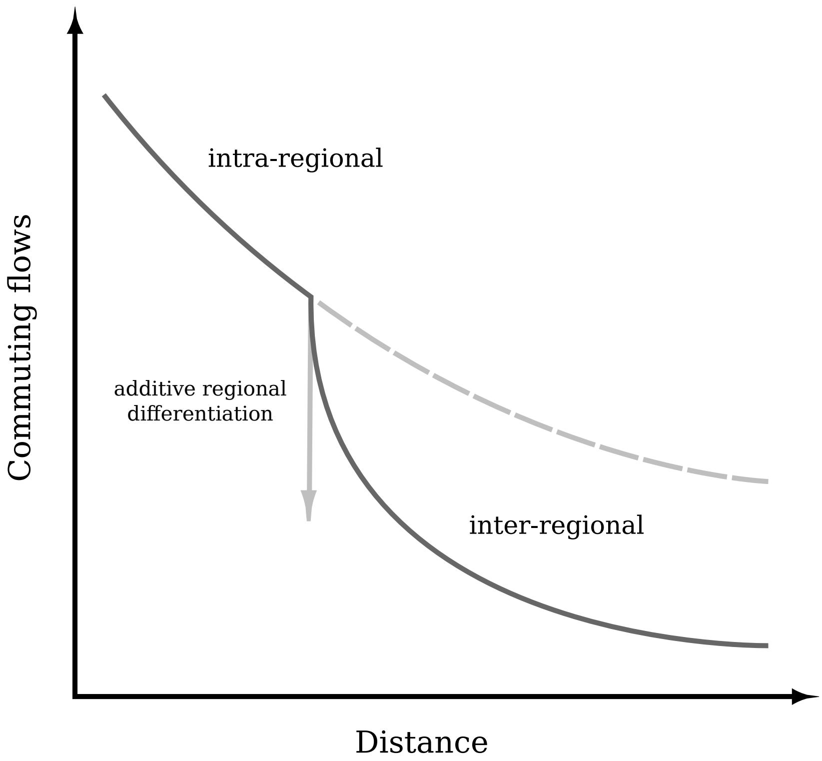

Figure 7.3: Conceptual representation of functional distance showing commuting flows as a function of distance. The additive regional differentiation induces an increase in distance decay (or similarly, a decrease in commuting flows) for inter-regional commuting. In theory, we would expect fewer commuters between origin i and destination j if they are not in the same local labour market, *ceteris paribus*.

those from an entire local labour market]), we see that at relatively short distances, a much greater decrease in the predicted number of flows for a given increase in distance is observed for inter-regional commuting than commuting in general. As a corollary to this, differences between the two curves at larger distances are negligible. As expected, the difference in distance decay is even more pronounced if we examine internal commutes only using the parameters from the DUAL model (not shown).

What we have shown in this section is that there are major differences in commuting patterns when we look at commuting from the perspective of inter-

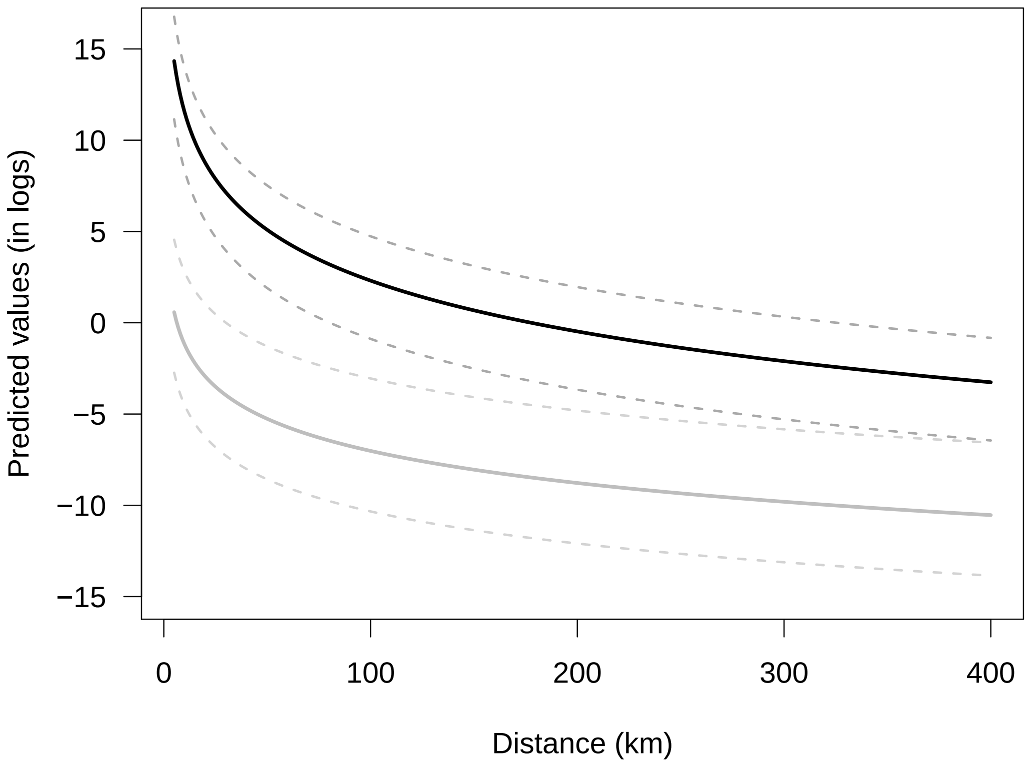

Figure 7.4: Predicted commuting flows for varying distances based on the NEGBIN (grey) and INTER-REGION (black) models. The upper and lower dashed lines for each curve represent 1.5 times the upper and lower interquartile range of predicted values and the solid dark line presents the median predicted value. Note that the minimum distance plotted is 5km, as distances smaller than this are not realistic for inter-regional flows.

local labour market commuting flows. This again highlights the importance of considering the configuration of local labour markets across the study region when attempting to accurately model commuting flows. Furthermore, the concept of functional distance decay presented in the previous chapter and explored further in the preceding section, provides a viable explanation for the significant change in distance decay between global, internal, external, and inter-regional commuting flows. Our exposition of local labour market effects on commuting flows is nearly complete; we now have sufficient evidence to argue that the spatial structure of local labour markets significantly influences commuting flows. However, we have yet to explore the *extent* to which local labour markets influence commuting flows. In the remainder of this chapter, we therefore take our comparison of local labour market effects on commuting patterns one step further, by evaluating *local* spatial interaction at the level of the local labour market. From this, we explore parameter trends/differences across Ireland and expose the varying extent to which local labour markets have an influence on travel-to-work.

7.3 Intra-regional commuting

Using the local labour market regions from previous chapters, we now calibrate a series of local spatial interaction models similar in vein to those introduced in Section 3.4.3. However, the local models utilised here differ from standard origin– or destination-specific models in that, rather than calibrating a separate model for each origin/destination, we calibrate a separate model for each *local labour market*. This provides a different set of parameters for each local labour market and gives an indication of the behavioural and structural differences between them. The strength of this approach lies in the fact that we can now focus on a) how local variations in parameter estimates lead to the observed global estimates, b) how local estimates of distance decay vary and how these relate to the underlying population structure of the labour market, and c) how variations in the strength of different parameters over space explain variations in commuting processes.

Local spatial interaction models at the level of the local labour market are simply local labour market-specific models similar in form to Equation 6.1, with the added condition that a different set of parameters is calibrated for each local labour market, such that

$$\lambda_{kij} = \exp(I_k+\psi_{kj}+\chi_k \mathrm{w}_{ki}+\omega_k \mathrm{u}_{ki}+\delta_k \mathrm{e}_{ki}+\eta_k \mathrm{h}_{ki}+\gamma_k \mathrm{c}_{ki}+\beta_k \mathrm{d}_{kij}) \quad \forall k \in L, \tag{7.1}$$

where k now represents a given labour market from the set of all local labour markets L, in the study region. Rather than fit one extremely large and unstable model, we calibrate $|L|$ different spatial interaction models[4], each with its own intercept, parameters, and destination constraints. For the current application, which is based on the local labour markets used in previous chapters, we have 65 individual spatial interaction models. Summaries of the 6 parameter estimates from our local models are given in Table 7.3.

Table 7.3: Summary statistics for significant local labour market-specific parameter estimates.

	Min.	Median	Mean	Std. Dev.	Max.
distance	-3.1566	-2.4012	-2.3286	0.4710	-1.2800
accessibility	-1.3805	-0.5504	-0.2991	0.9945	4.7853
workers	-0.2652	0.6944	0.6158	0.3179	1.3648
unemployment	-0.6935	-0.0713	0.2230	1.1512	4.3852
education	-1.0498	-0.3536	-0.1648	0.9858	3.7723
housing	-5.1413	-0.6825	-0.0987	3.4765	10.8027

Based on the above summaries, we firstly examine the relationship between the global parameter estimates and the distribution of local parameter estimates. It is interesting to note that in all cases, the global parameter falls well within the distribution of local parameters, and is relatively close to the mean[5] for each distribution. Furthermore, we find that, except in the case of distance, the range of estimates for each parameter extends from positive to negative values. This indicates that at the level of the local labour market, travel-to-work may in fact be much more complicated than perceived at a global level. The above observations become even more obvious if we visually examine the distribution of local parameter estimates as in Figure 7.5, where

[4] $|L|$ represents the cardinality of the set L, which is simply a measure of the number of elements in L.

[5] In all cases, the global parameter is well within one standard deviation of the mean of the local parameters.

the corresponding global parameter value is also provided for comparison.

What we immediately notice from Figure 7.5 is the apparent outlier visible in the frequency distributions for accessibility, unemployment, education, and housing. This outlier is associated with a local labour market just north of Dublin which encompasses Skerries and portions of Malahide (see Figure 4.1 for these locations on a map of Ireland). While it *is* possible that this Dublin satellite region is characterised by a significantly different commuting process, it is more likely that the relatively small size of the local labour market is leading to unstable results. In fact, since this local labour market is made up of only 6 EDs, leading to $6^2 - 6 = 30$ observations, the ratio of observations to parameters (note there are $n + 6 + 1$ parameters) is relatively small, leading to very few degrees of freedom. As such, the parameters estimated for this region (and by the same token, parameters for three additional local labour markets which have already been discounted) should be used with caution.

Another observation that can be gleaned from Figure 7.5 is the non-uniform distribution of parameter estimates. This emphasises the variability of parameters across the study area, and provides an indication that our earlier assertions regarding spatially varying parameters were correct. This is corroborated by the spatial distribution of parameter values given in Figure 7.6. These maps of local parameters support our hypothesis that commuting processes are not necessarily spatially invariant. Furthermore, the significance of the various parameter estimates[6] (see figures 7.6 and 7.7) also appears to be spatially varying. Indeed, save for distance (and for the most part, workers), the parameters for each individual variable are only significant in specific local labour markets. In particular, unemployment is only significant in a relatively small subset of local labour markets with a cumulative distribution function

[6] Testing for significance in this case may suffer from the multiple testing problem and should therefore be treated with caution. While corrections for multiple comparisons are available (e.g., Benjamini & Yekutieli, 2001, and references therein), these were not performed here and are left as an exercise for the reader.

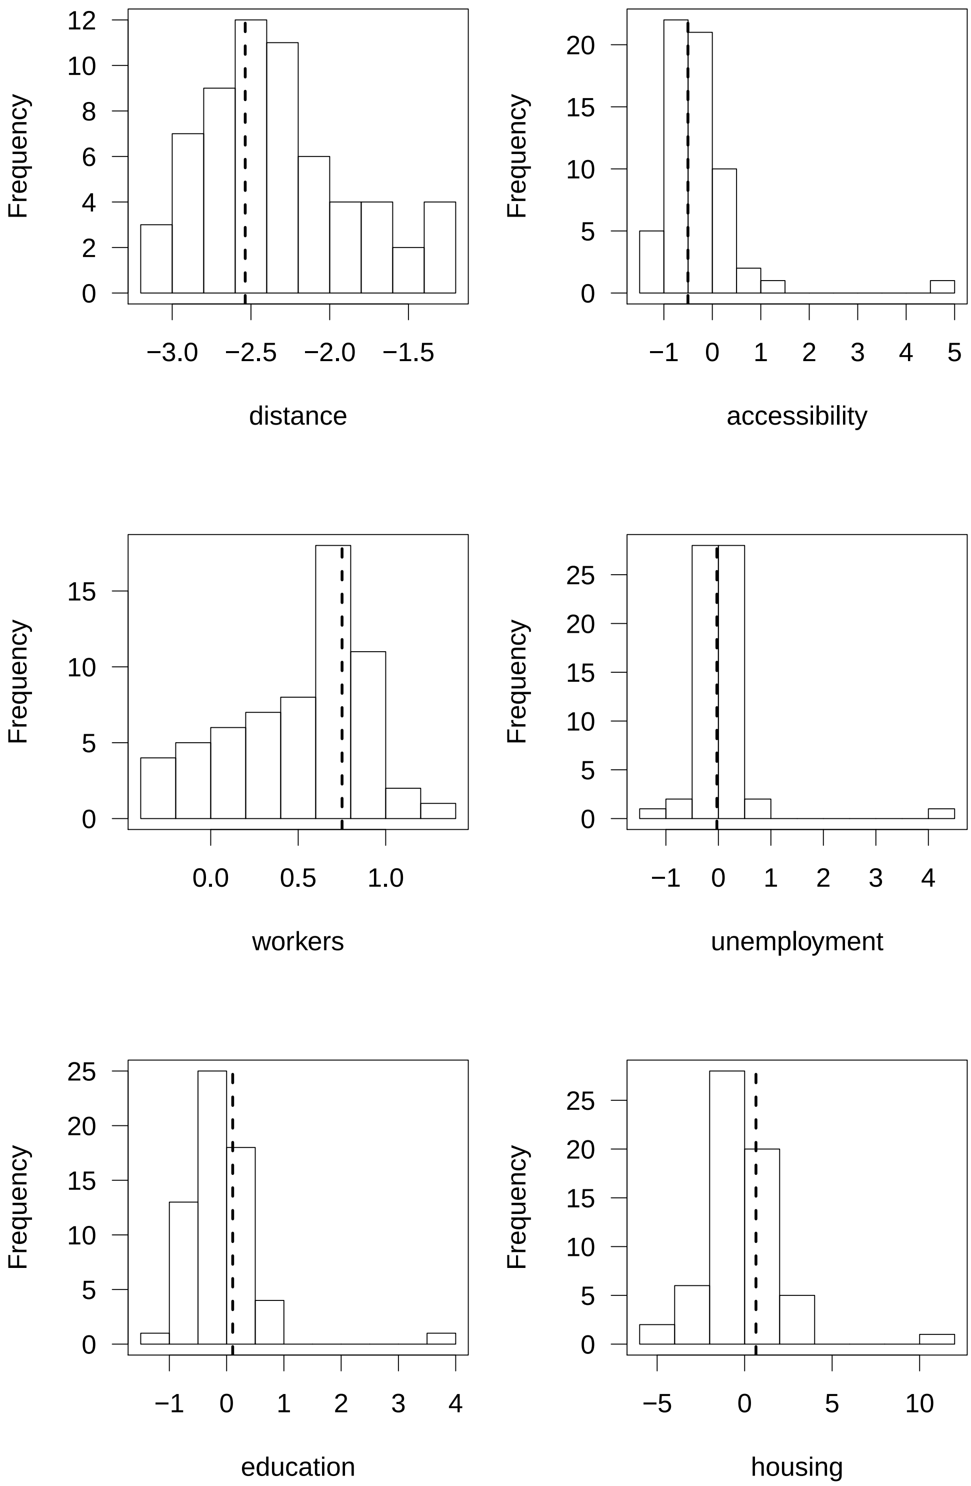

Figure 7.5: Frequency distribution of local parameter values for distance, accessibility, at the four attributes of origins. The corresponding global parameter value (based on the NEGBIN model) is provided for reference/comparison.

(CDF) for unemployment that lies below the global unemployment CDF (i.e., regions where unemployment is relatively low compared to the global distribution). This assertion is valided using a two-sample Kolmogorov-Smirnov test with p-value< 0.01. Similar observations can be made for education and housing, though their relationship with their underlying attribute values is less clear.

In terms of actual parameter estimates, the number of available workers appears to be important (i.e., relatively large parameters) where there are many workers available to begin with and less important in regions where there are comparatively fewer workers. Similarly, periphery local labour markets, which tend to have a mix of high and low numbers of resident workers, appear to be associated with mid-ranged parameter estimates. In general, the workers parameter appears to act as a 'balancing factor' in most regions, only emphasising the impact of resident workers in regions where there are many. One exception to this general rule is the greater Dublin area (GDA), where the number of resident workers appears to be less important than expected given its size. A simple explanation for this is that high levels of commuting in and around Dublin will occur regardless of the number of resident working individuals due to other factors such as unemployment rates (which are relatively high in Dublin), education, and housing occupancy levels. It is also worth noting that low-to-negative parameter estimates for workers were generally not significant.

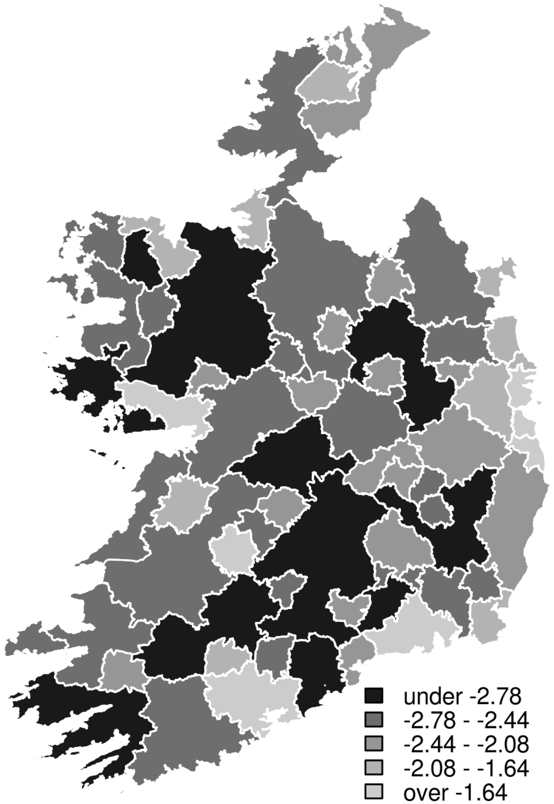

(a) distance

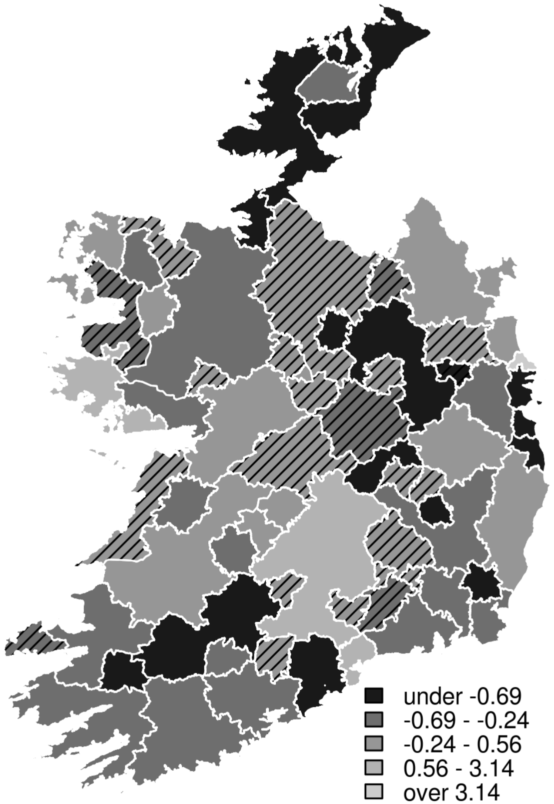

(b) accessibility

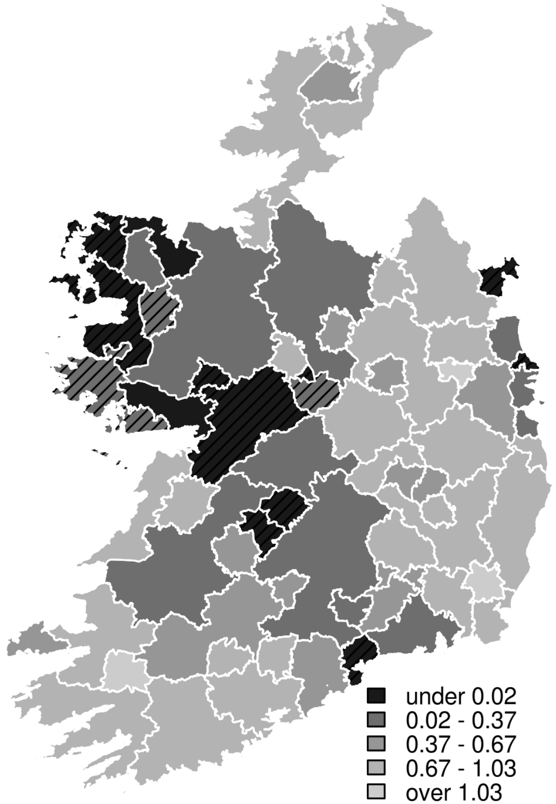

(c) workers

Figure 7.6: Local labour market-specific parameter estimates for the 6 variables in our model(s). Note that non-significant estimates (p-value > 0.05) are denoted by diagonal cross-hatching [continued on following page].

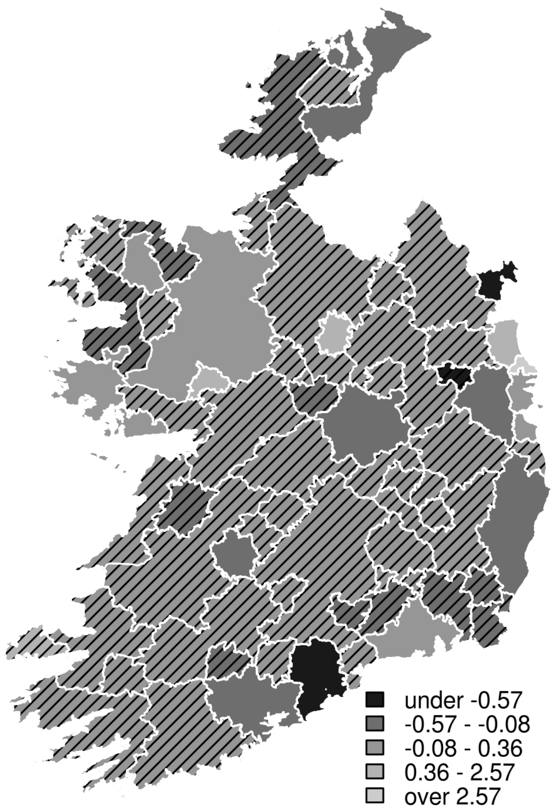

(d) unemployment

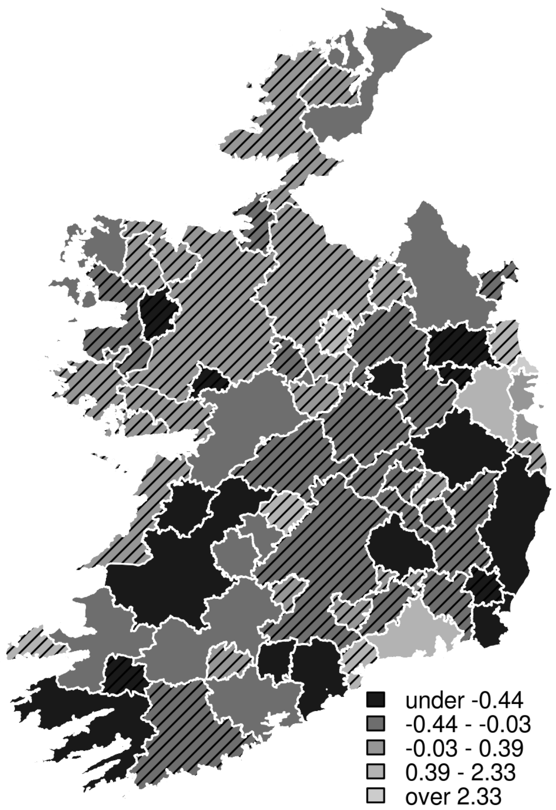

(e) education

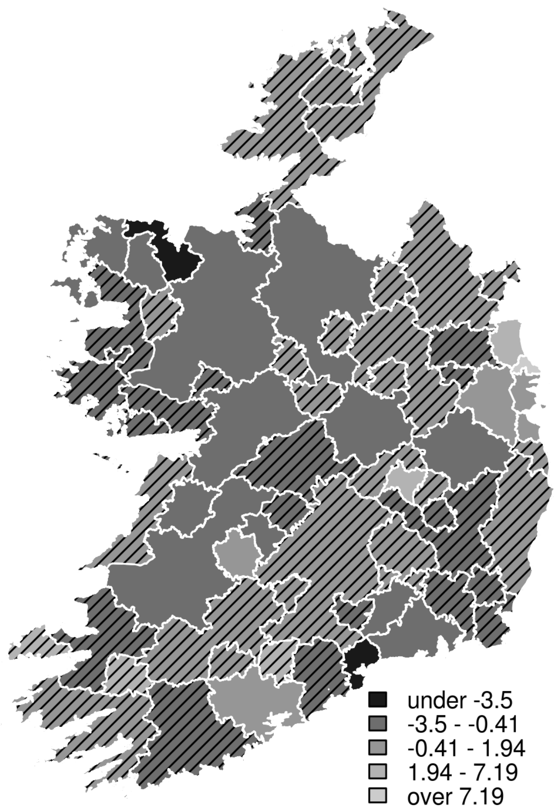

(f) housing

Figure 7.6: [continued from page previous page]

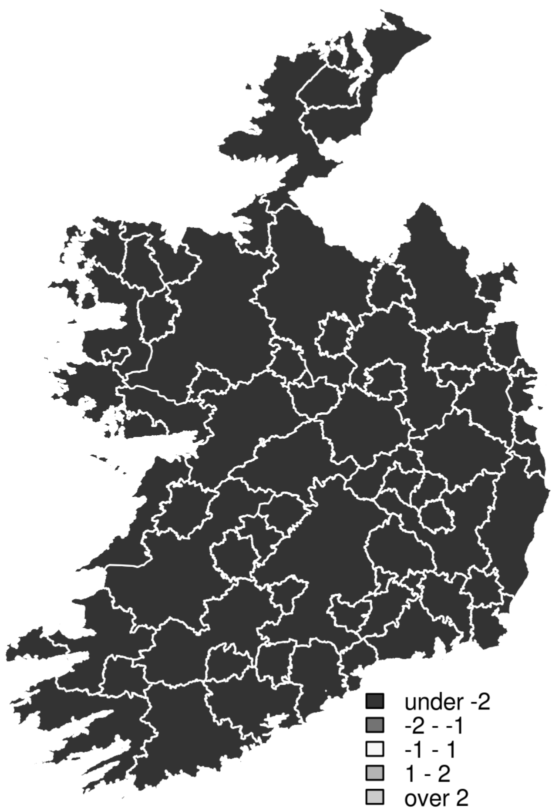

(a) distance

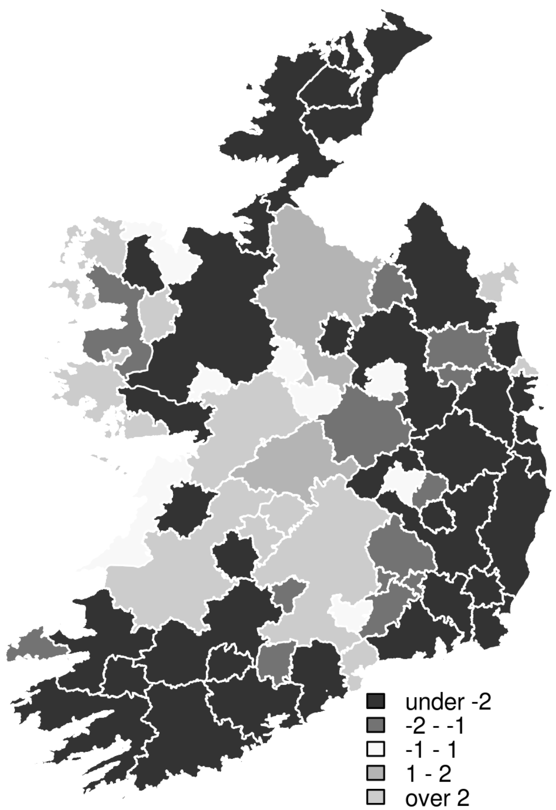

(b) accessibility

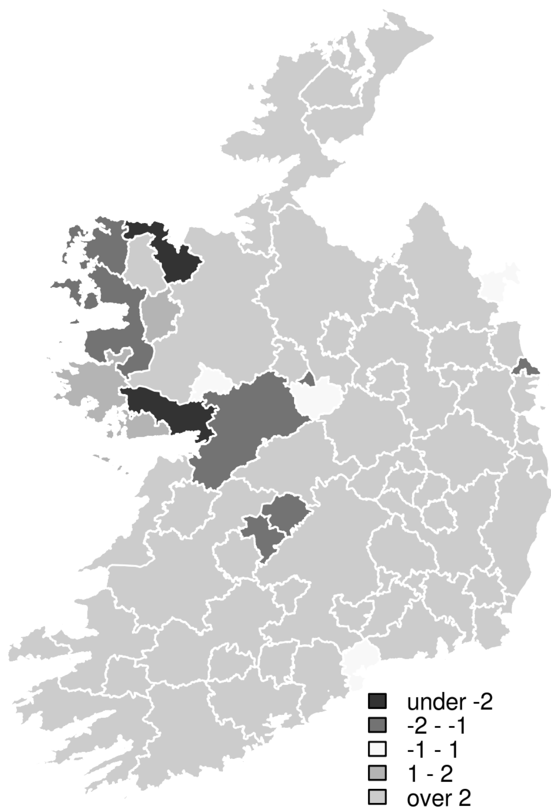

(c) workers

Figure 7.7: Local labour market-specific z-values for the 6 variables in our model(s). Non-significant values (values between -1 and +1) are less opaque than significant values [continued on following page].

Figure 7.7: [continued from previous page]

Looking at education, we see that it has the strongest impact on commuting flows in the GDA and Waterford. In terms of significance, it appears to have a measurable impact on commuting flows only in regions of relatively high or low education levels. Regions with a more general mix of education levels do not display significant education parameters. This is not entirely unexpected as seemingly random mixtures of education levels, as is the case in many of the periphery regions, provide very little in terms of explanatory power. The housing variable displays the largest range in parameter values, ranging from highly negative in and around Ballina and between Waterford and Cork, to highly positive in Skerries and Swords in the GDA. Isolated large parameters in Limerick, Cork, and north of Carlow are also visible. The spatial distribution of housing *values* is relatively uniform across the study region, which may explain the relatively uniform distribution of parameter estimates and lack of significance throughout most of the country.

In the current context (i.e., an exploration of variations in model parameters due to local labour market effects), the parameters for our attributes of space (distance and accessibility) are perhaps most relevant. Beginning with accessibility, we note that the regions with the most significant parameters are largely associated with accessibility 'extremes' (i.e., highly isolated and/or central EDs). This observation is most noticeable when we compare the maps of accessibility values from Appendix B with the corresponding significance map in Figure 7.7. What these maps indicate is that accessibility (i.e., origin-centric destination competition) is most relevant in areas where there is either very strong competition, or significantly less competition to begin with. This finding makes intuitive sense: where there is strong competition, a centrally located ED can expect far fewer commuters than a similar, but isolated ED. Conversely, where there is minimal competition, differences between isolated and central EDs will be less apparent.

Distance is the only parameter for which all regions display statistically significant parameter estimates. Estimates range from -3.16 to -1.28, and are strongly associated with the structure of the local labour markets. By this we mean that strong distance decay is associated with the periphery local labour markets as defined in Section 5.4.2, whereas weaker distance decay is associated with large hub regions, such as those surrounding Cork, Limerick, Galway, and the GDA. This is a highly intuitive result and suggests that local labour markets with a single, large central city will attract commuters from further afield than periphery local labour markets with multiple smaller cities/towns. A comparison of distance decay parameters with observed mean distances for each origin in Figure B.1 does not suggest a clear correspondence between the two maps. This indicates that the structure of the local labour markets may indeed explain a large component of distance decay, quite possibly in the form of functional distance decay as alluded to previously. Additionally, all local distance decay parameters were substantially smaller than the estimated inter-regional distance decay parameter (INTER-REGION model from Section 7.2), again validating our theory of a functional distance decay. In explanation, since within-region distance decay is less than between-region distance decay across the board, the situation appears to be quite similar to that depicted in Figure 7.3, thus lending credence to our stated assumptions. It is also interesting to note that the magnitude of distance decay for each local labour market is almost directly proportional to the order in which the local labour markets were 'found' by the regionalisation algorithm; which is in turn proportional to the strength of its internal connections (i.e., flows).

The general trends and observations touched on in the preceding paragraphs highlight the importance of considering local variations in commuting trends. Furthermore, the relatively large range of parameter estimates for each of the 6 variables considered attests to the complex nature of travel-to-work patterns/behaviour; a complexity that is difficult to disentangle at the global level. While some of the local parameter estimates were contrary to our original expectations, several of these estimates were found to be non-significant and therefore unlikely to be representative of actual local conditions. Conversely, many of the unintuitive variations in parameter values *do* appear to indicate important differences in commuting processes, perceptions, and behaviours. With this in mind, it is clear that until we consider the context within which commuting takes place - the local labour market - we cannot properly understand the extent to which local labour markets influence, and are influenced by, travel-to-work.

7.4 Moving forward

Building on the insights offered from previous chapters, in this chapter we have implemented a more explicit treatment of 'local' versus 'global' commuting patterns. This was designed to facilitate a more thorough understanding of variations in commuting flows. With this in mind, we arrived at two additional modelling viewpoints; the first being a single 'inter-regional' commuting model designed to examine commuting flows *between* local labour markets and the second, a set of 'region-specific' commuting models designed to explore commuting patterns *within* local labour markets. Both models helped us to explore the spatial relationships between travel-to-work patterns and local labour markets and yielded important insights into the effects of local labour markets on commuting.

Based on the empirical analyses from this chapter, we have found that the concept of functional distance decay provides a useful explanation for the significant changes in distance decay between global and local commuting flows. Furthermore, we have found that commuting flows *between* local labour markets may be characterised by significantly different relationships than commuting flows *within* a single local labour market. As such, it is possible that when operating at the global level, local parameters become 'averaged', leading us to under- and over-estimate their effects in certain regions. Indeed, contrary to the assumptions of a global model, the relationship between commuting flows and the various parameters in our models do not appear to be invariant across space. While there does not appear to be any sweeping spatial trends across the study region, there are clear spatial variations in parameter estimates that appear to be related to local labour markets, particularly for the distance, accessibility, and housing variables.

In parallel to these findings, we note that while at the *global* level all parameters were significant, this is not necessarily the case at the *local* level. In fact, most parameters were only significant in particular local labour markets and, in some cases, there was a clear relationship between parameter significance and the local attributes of origins. Furthermore, where parameters were significantly related to commuting flows, they did not always agree across space. For example, while accessibility may hinder commuting in some regions (e.g., around Letterkenny), it may actually facilitate commuting in others (e.g., Skerries and regions between Limerick and Kilkenny). This observation of non-stationary parameter estimates is not surprising and attests to the complex nature of the commuting processes. In particular, it advances the theoretical argument that commuting is partly contingent on the structure of local labour markets.

In the following chapter, we take a slightly different approach to examining travel-to-work patterns than previous chapters; this time focusing on the general commuting process, rather than local labour market effects specifically. The idea is to explore the concept of choice set generation and provide a theoretically and empirically valid means of integrating additional behavioural concepts into spatial interaction modelling. More specifically, we introduce the concept of excess zeros in commuting data and present a class of models designed specifically to deal with this issue of zero-augmentation. Furthermore, we provide a theoretical interpretation of excess zeros and zero-augmented processes in the context of travel-to-work and demonstrate the superiority of zero-augmented models when dealing with (a subset of) the commuting data used in this thesis.

Chapter 8

Choice set integration

8.1 Introduction

In the previous three chapters we developed, calibrated, and explored several different spatial interaction models of commuting. These models were found to fit well to the Irish commuting data presented in Chapter 4, particularly when over-dispersion in the commuting flows was explicitly considered using a negative binomial model. Furthermore, local labour market effects were explored and modelled by considering aggregate, internal, external, inter-regional, and intra-regional commuting flows separately. What we have found is that by exploring the concept of functional distance implicit in all five of these model types, we can begin to explain the extent to which local labour markets influence commuting flows and how commuting flows enforce the boundaries of local labour markets. There is however, an additional component to any spatial interaction setting - the behavioural component - or more specifically, destination choice.

In this chapter, we focus on the concept of choice set definition. While not directly contingent on previous chapters, the work presented in this chapter is intended to provide a 'way forward' in terms of spatial interaction modelling research. As such, the models presented in this chapter provide a powerful and theoretically pleasing interpretation of spatial interaction models of commuting and lend a theoretical interpretation of the large number of zeros in many travel-to-work datasets. The primary goal of this chapter is therefore to provide an intuitive and previously unavailable means of integrating choice set generation into existing spatial interaction models.

8.2 Excess zeros

In the previous two chapters, the over-dispersed nature of our commuting data strongly favoured the use of a negative binomial-based spatial interaction model of commuting. However, it is important to keep in mind that over-dispersion, while consistent with the negative binomial specification, does not necessarily mean that the negative binomial model will be entirely adequate (Cameron & Trivedi, 1998). Indeed, it has been suggested that the negative binomial model sometimes spuriously indicates over-dispersion, when in fact the underlying process may actually consist of a two-part process generating excess zeros (Shankar et al., 1997). In Chapter 6, the inability to properly predict large flows may be an indication that problems with our negative binomial commuting model still exist. The likely culprits for these potential problems are a possible misspecification of the conditional mean (i.e., commuter choice process) and unobserved heterogeneity in the flow generating process (i.e., variations in commuter choices). Both these issues may be leading to the excess zeros and over-dispersion observed in the commuting flow data. While the negative binomial model appears to be able to address this latter problem (at least at a global level), in the following section, we provide alternative modelling approaches designed to take into account the issue of misspecification of the commuter choice process.

In addition to regular over-dispersion, commuting data may be overdispersed due to excess zeros. It is important in these cases to separate excess zeros from regular over-dispersion, because these two forms of over-dispersion are likely generated from different underlying processes (Greene, 1994). As we have alluded to previously, over-dispersion due to unobserved heterogeneity in the flow generating process is common in commuting data, however, overdisperion due to excess zeros may also occur when the incidence of zero-flows is greater than that expected from the Poisson (or related) base-model. For example, in travel-to-work situations, it is possible that there will be zero commuting between two regions i and j, due to limitations imposed by a lack of transportation, infrastructure, or jobs, or simply because the attributes of the origin, destination, or their separation, are not conducive to commuting (i.e., the distance is simply too far for most commuters). In the first case, zero flows are inevitable (i.e., the probability of commuting from i to j is essentially *a priori* zero), whereas in the second situation, zero flows *may* arise, but positive flows are also possible (i.e., the probability of commuting from i to j is greater than zero). These two types of zeros may be termed *structural* and *observational* zeros respectively[1].

According to Cameron & Trivedi (1998), in the above situation of two zero-generating processes, it would be a misspecification to assume that the zeros and non-zeros arise from the same underlying process. Furthermore, to ignore zero-inflation would lead to similar problems arising from regular over-dispersion, such as biased and/or inconsistent parameter estimates, inflated standard errors, and subsequent inference problems (Miller, 2007). As such, models designed to explicitly consider excess zeros in count data have been developed; these models are termed *zero-augmented* models. The most commonly applied forms of zero-augmented models are the hurdle and zero-inflated mod-

[1]These different types of zero counts are also sometimes referred to as *structural* and *sampling* zeros.

els (Mullahy, 1986; Lambert, 1992). Hurdle models are two-component models that combine a truncated-at-zero (left-truncated) count component with a right-censored 'hurdle', or zero component. Similarly, a zero-inflated count model is a finite mixture model that combines a count component with a point mass at zero (Zeileis et al., 2008). Excess zeros are assumed to arise differently in each case and, as a result, they have different theoretical interpretations.

It is easiest to think of these two models as modelling a situation in which there are two processes controlling commuting; one that controls whether or not commuting between i and j is possible and another controlling the *positive* (hurdle), or *non-negative* (zero-inflated) number of commuters travelling from i to j in the case that commuting *is* possible. In other words, both models assume there is an additional unknown processes generating excess zeros, however, in the case of zero-inflated model, there are *two* possible ways that zero-flows may arise, whereas in the hurdle model, there is only one process generating zero-flows. In the following two sections, we introduce these alternative models in more detail and provide theoretical interpretations which are relevant to commuting behaviour.

8.2.1 Hurdle model

As alluded to above, the hurdle model treats zero and non-zero flows separately. It is essentially a conditional Poisson model[2], or finite mixture model with two components. The most common formulation of the hurdle model is one in which a binomial probability model governs the binary outcome of whether a zero or positive flow is realised (i.e., the transition stage) and a truncated-at-zero count data model governs the conditional distribution of the positives (i.e., the events stage) (Mullahy, 1986; Miller, 2007). The term 'hurdle' comes

[2]The hurdle and zero-inflated models were originally formulated using a Poisson count component, although, as we will show later on, other count models such as the negative binomial model are equally valid.

from the idea that if the realisation at the transition stage is positive, a *hurdle* is crossed and only then are positive counts possible. Normally a truncated Poisson or negative binomial model is used to address the positive flows and the log-likelihoods for the two separate models are estimated separately. This is a beneficial feature of the hurdle model, as it makes calibration much faster and more stable than alternatives. Formally, the probability that T_{ij} is equal to the observed flow t_{ij}, is given as

$$\Pr(T_{ij} = t_{ij}) = \begin{cases} f_{\mathrm{z}} & t_{ij} = 0 \\ (1 - f_{\mathrm{z}})\dfrac{f_{\mathrm{c}}(t_{ij})}{1 - f_{\mathrm{c}}(0)} & t_{ij} > 0 \end{cases} \tag{8.1}$$

where f_{c} is the count, or events stage model and f_{z} is the zero, or transition stage model. As such, the event stage model is essentially the probability for a positive realisation, multiplied by the probability for the counts (Miller, 2007). The most common/intuitive practical formulation of this model consists of a binomial distribution with logit link function for the transition stage and a truncated-at-zero Poisson distribution with log link function for the event stage:

$$\Pr(T_{ij} = t_{ij}) = \begin{cases} p & t_{ij} = 0 \\ \dfrac{(1 - p)e^{-\lambda_{ij}}\lambda_{ij}^{t_{ij}}}{(1 - e^{\lambda_{ij}})t_{ij}!} & t_{ij} > 0 \end{cases} \tag{8.2}$$

where λ_{ij} is now the truncated Poisson mean for counts greater than zero, p is the probability of a zero count (often modelled using logistic regression) and $p = 1/(1 + \exp(\mathbf{z}'_{ij}\boldsymbol{\beta}))$. The two separate models can be implemented as GLMs and, as mentioned, need not have the same independent variables. This provides a powerful means of modelling two-part commuting processes and can be extended to take into account additional over-dispersion by replacing the truncated Poisson component with a truncated negative binomial model,

$$\begin{aligned}\Pr(T_{ij} = t_{ij}, t_{ij} > 0) &= \frac{\Gamma(\theta^{-1} + t_{ij})}{\Gamma(\theta^{-1})\Gamma(t_{ij} + 1)} \left(\frac{1}{(1 + \theta\lambda_{ij})^{1/\theta} - 1}\right)^{-\theta^{-1}} \\ &\quad \times \left(\frac{\lambda_{ij}}{\lambda_{ij} + \theta^{-1}}\right)^{t_{ij}}. \end{aligned} \tag{8.3}$$

This allows for over-dispersion at the event stage, providing an intuitive interpretation of heterogeneity in the commuting flow process, which we will discuss further in Section 8.2.3.

8.2.2 Zero-inflated model

Another way to model excess zeros is given by the zero-inflated count model. In this case, the zeros are modelled separately, as well as along with positive values, providing two possible sources of zero flows. The zero-inflated model may also be considered a mixture model, with a point mass at zero mixed with a Poisson or negative binomial distribution (or other count model). As such, the probability of observing a zero in a zero-inflated model is given by the probability of observing an excess zero (i.e., p in Equation 8.3), plus the probability of observing a zero in the count model. In the case of a zero-inflated Poisson model then, we have

$$\Pr(T_{ij} = t_{ij}) = \begin{cases} p + (1-p)e^{-\lambda_{ij}} & t_{ij} = 0 \\ (1-p)\dfrac{e^{-\lambda_{ij}}\lambda_{ij}^{t_{ij}}}{t_{ij}!} & t_{ij} > 0 \end{cases}, \tag{8.4}$$

where all variables are defined as before. The event stage of the above model is quite similar to Equation 3.22 in Chapter 3 and, in fact, for $p = 0$, (8.4) reduces to the regular Poisson model. Similarly, the zero-inflated negative binomial model is given by

$$\Pr(T_{ij} = t_{ij}) = \begin{cases} p + (1-p)\left(\frac{\theta^{-1}}{\theta^{-1}+\lambda_{ij}}\right)^{\theta^{-1}} & t_{ij} = 0 \\ (1-p)\frac{\Gamma(\theta^{-1}+t_{ij})}{\Gamma(\theta^{-1})t_{ij}!}\left(\frac{\theta^{-1}}{\theta^{-1}+\lambda_{ij}}\right)^{\theta^{-1}}\left(\frac{\lambda_{ij}}{(\lambda_{ij}+\theta^{-1})}\right)^{t_{ij}} & t_{ij} > 0 \end{cases}. \tag{8.5}$$

Notice that, contrary to the hurdle model, the two components of the zero-inflated model are not functionally independent and, as such, their likelihood functions cannot be maximised separately. This being the case, the hurdle model is sometimes preferred over the zero-inflated model due to its orthogonal parametrisation, making it both simpler to fit and easier to interpret. However, from a theoretical point of view, there are distinct differences between these two models which make them more, or less appropriate for modelling commuting flows.

8.2.3 Theoretical interpretations

The aforementioned zero-augmented models, while useful in a purely mathematical sense for accounting for excess zeros and over-dispersion, also offer some potential benefits from a theoretical perspective; particularly in the context of travel-to-work behaviour/models. As mentioned previously, zero-augmented models can model situations in which there are two underlying processes controlling commuting. The first process controls whether or not commuting between i and j is *possible* and the second process controls the *volume* (i.e., number of commuters travelling from i to j) of commuting given that commuting is possible. We have previously provided empirical evidence to suggest that this type of two-stage process is indeed driving commuting in the Irish commuting dataset (see the large number of zero flows in Figure 6.1a for example). There are many possible reasons for this type of two-stage commuting process to occur, however, the most likely explanations revolve around the notion of choice set generation.

Choice set generation is most frequently discussed in the context of spatial and aspatial choice modelling and is linked to the early work of Manski (1977). Essentially, a choice set refers to the group of alternatives that are evaluated by an individual when making a spatial choice and consists of a subset of the universal choice set that describes all alternatives available to an individual (Pellegrini et al., 1997). Manski's two-stage discrete choice paradigm has remained a popular framework for empirically determining the set of available alternatives to an individual, both in a deterministic sense and via probabilistic approaches. In Section 3.4.4 we presented the competing destinations framework in the context of spatial choice, highlighting how it can be used to model hierarchical information processing. In this sense, the competing destinations framework is a probabilistic approach to choice set definition whereby the probability of a destination being in the true choice set of an individual is determined by a measure of its accessibility relative to alternative destinations (see Fotheringham (1988), Section 3.4.4 and references therein). The zero-augmented models of this section provide a means to further refine choice set definition in a spatial interaction context, combining the hierarchical information-processing assumptions of the competing destinations framework, with a two-stage choice set definition process which is able to account for additional information processing strategies which may or may not be explicitly spatial.

In a spatial interaction setting, the set of destinations available to a particular origin is generally assumed to be equal to the universal choice set. That is, all alternatives have an equal probability of being evaluated. Indeed, this is a fundamental assumption in Poisson spatial interaction models: that there is a constant probability of an individual in i commuting to j. However, in reality, this is often not the case, as there may be a myriad reasons why interaction between two locations may not be possible. For example, it may be that the transportation network between a set of cities makes it impossible for workers living at origin i to travel to destination j within a single work-day, making commuting impossible. Alternatively, there may be some unknown socio-political or religious reason why employees living in origin i would never consider/be hired to work in a particular destination j. In these cases, the zero flows between i and j could be considered structural zeros, in that they are inevitable consequences of the particular spatial interaction setting. Thus, our zero-augmented models may be interpreted as modelling a two-part decision making process, whereby the transition component of the model defines the choice set (m_i) generation stage and the event component defines the (conditional) level of commuting.

Differences in interpretation between the hurdle and zero-inflated models may be attributed to differences in assumptions regarding the explicitness of the choice set. In other words, the hurdle model assumes the choice set is explicitly defined by the transition component, such that positive flows *must* occur between i and all destinations in its choice set ($k \in m_i$). Conversely, the zero-inflated model assumes the choice set is approximately defined and allows for additional observational zeros which may be the result of unobserved limitations on commuting flows. In the following section, we empirically evaluate these zero-augmented models in the context of Irish commuting patterns and provide an interpretation of the results which may provide valuable recommendations for future spatial interaction models of commuting.

8.3 Model results

Building on the models from Section 6.4, we continue to develop our commuting spatial interaction model, this time exploring the potential improvements offered by the zero-augmented models. In all cases, we use a binomial GLM with logit link function for the transition component, with all the same independent variables as in the event components[3]. For ease of computation and interpretation, we calibrate our zero-augmented models (and recalibrate the Piosson and negative binomial models for comparison) on data for the Dublin region only, which provides a good mix of zero and positive flows, encompassing approximately 3% of the full ~ 11 million records. An initial negative binomial model fit to the Dublin data (not shown) indicated that unemployment was not significant. Dropping this variable and refitting the model provided an improvement in likelihood and AIC and BIC measures, whereas any changes in the remaining parameter estimates were negligible. As such, we move forward with this modified model.

Table 8.1 provides parameter estimates and various model diagnostics for evaluating the four zero-augmented models. Firstly, we consider two hurdle models, based on the Poisson and negative binomial models as the event components. For both models, the variables and data are the same as for the initial Poisson and negative binomial models, however, both count components are now truncated for $flows < 1$, with the transition components modelling zeros versus counts. In addition to the hurdle models, we examine two zero-inflated models, again based on the Poisson and negative binomial models previously presented. The additional probability weight for zero flows (transition component) is fitted via a binomial GLM and the event component is essentially the usual (non-truncated) Poisson or negative binomial model.

In terms of likelihoods, the negative-binomial-based models appear to provide better fits than their Poisson counterparts and the information criteria measures corroborate this statement. Additionally, the zero-inflated variants appear to fit better than their hurdle counterparts, suggesting that the zero-

[3]Note that the destination fixed effect variable is *not* included in the transition component, as it is not needed to enforce our destination constraints.

Table 8.1: Model outputs from zero-augmented models.

	ZIP	ZINB	HURDP	HURDNB
	Event component			
(Intercept)	13.7534	14.9203	14.4931	15.0982
	(0.2322)**	(0.2492)**	(0.165)**	(0.776)**
distance	-1.3559	-1.6644	-1.3147	-1.532
	(0.0015)**	(0.0052)**	(0.0016)**	(0.0069)**
accessibility	-0.6237	-0.6258	-0.594	-0.5626
	(0.0016)**	(0.0036)**	(0.0016)**	(0.0048)**
workers	0.504	0.5601	0.5069	0.5857
	(0.0017)**	(0.0048)**	(0.0017)**	(0.006)**
education	0.4559	0.4839	0.4573	0.5747
	(0.0037)**	(0.0089)**	(0.0037)**	(0.0115)**
housing	1.9821	1.6732	1.9167	2.1352
	(0.0142)**	(0.0366)**	(0.0144)**	(0.0475)**
	Transition component			
(Intercept)	-18.7455	-35.8691	16.3199	16.3199
	(0.2391)**	(0.857)**	(0.113)**	(0.113)**
distance	1.7659	3.5692	-1.7717	-1.7717
	(0.0156)**	(0.0653)**	(0.0074)**	(0.0074)**
accessibility	0.1521	-0.0897	-0.4536	-0.4536
	(0.0079)**	(0.0251)**	(0.0048)**	(0.0048)**
workers	-0.4429	-0.7826	0.5899	0.5899
	(0.0089)**	(0.024)**	(0.0061)**	(0.0061)**
education	0.2985	2.0454	0.2697	0.2697
	(0.0247)**	(0.0841)**	(0.0124)**	(0.0124)**
housing	2.799	6.13	0.7645	0.7645
	(0.1029)**	(0.2902)**	(0.0508)**	(0.0508)**
-log Likelihood	442260	270855	469442	306480
AIC	885728	542921	940093	614170
BIC	892233	549437	946598	620687
pseudo-R^2	0.4029	0.6743	0.4276	0.7673

*$p < 0.05$; **$p < 0.01$

inflated negative binomial model (ZINB) provides the best fit of the four zero-augmented models. The pseudo-R^2 measures are less clear than the other model diagnostics, indicating that the ZINB model may not be the clear choice it appears to be. The parameters (and standard errors) in the event components are similar between all four models and major differences between these four models are not entirely clear until we begin to examine their respective transition components. In particular, the zero-inflated negative binomial model appears to differ most in terms of both parameter sign and magnitude. As expected, the parameters (and standard errors) for the hurdle models are identical, due to the fact that the same logit model is used to model zero-flows in both cases.

Interpretation of the parameters reveals some interesting aspects of the models, particularly when we examine the transition components. As before, the distance and accessibility parameters (for the event components) are negative and are within the expected range. In the transition components however, the parameter magnitudes have increased and in some cases, the signs have reversed. This sign reversal is due to the fact that the transition component in the zero-inflation models describes the probability of observing a *zero* flow, whereas the transition component in the hurdle model describes the probability of observing a *positive* flow. Moving forward with the ZINB model, we note that with distance, a 1% increase leads to the probability of a zero flow increasing by 3.57%. In other words, if the distance between i and j is increased, this will cause an decrease in the probability of positive commuting flows. Similarly, if we increase the level of accessibility by 1%, this will cause the probability of a zero flow to increase by 0.15%. Thus, the addition of a zero-inflation component into the model provides a powerful means of examining the probability of commuting between any set of origins and destinations. The interpretation of the event component parameters equates to the same interpretation as in the standard Poisson and negative binomial models, whereby the parameters control the volume of commuting between origins and destinations. In all four models, all parameters are significant, with p-values < 0.01.

Now that we have fit these zero-augmented models to the commuting data, it is of interest to compare and contrast these models with the Poisson and negative binomial models (Table 8.2). The parameters (and standard errors) in the event components of the zero-augmented models are similar to those from the modified Poisson models. There are some small differences in magnitude, however, these differences are to be expected, as the parameters for the zero-augmented models must now be interpreted in the context of the zero-augmentation (Zeileis et al., 2008). Despite this, the relative similarity of the model parameters suggests that the estimated mean functions for the various models are similar. To get a better picture of the changes induced by the zero-augmented models, we can examine the likelihood of the models. In terms of likelihoods, the basic Poisson spatial interaction model is inferior to all other models examined. We are able to improve the Piosson model fit by allowing for additional zero-inflation via the hurdle and zero-inflated Poisson models, however, this improved fit is not able to compensate for the improvements induced by the negative binomial model.

Table 8.2: Poisson and negative binomial models for the Dublin subset.

	POIS	NEGBIN
(Intercept)	13.1792	15.4843
	(0.1642)**	(0.2965)**
distance	-1.5101	-1.8591
	(0.0013)**	(0.005)**
accessibility	-0.638	-0.6064
	(0.0014)**	(0.0037)**
workers	0.6046	0.7077
	(0.0016)**	(0.0044)**
education	0.4721	0.491
	(0.0036)**	(0.009)**
housing	2.1291	1.6139
	(0.0137)**	(0.0363)**
-log Likelihood	486195	274613
AIC	973585	550424
BIC	980026	556863
pseudo-R^2	0.2686	0.6836

*$p < 0.05$; **$p < 0.01$

As mention previously, the negative binomial model provides a significant improvement in fit over the Poisson model, however, this fit can be improved further by allowing for additional zero-inflation (see -log Likelihood in Table 8.1 and 8.2). This assertion can be tested more formally using a Vuong closeness test in order to examine whether our zero-augmented models are preferred over the standard Poisson and negative binomial models. A comparison of the zero-augmented models with their non-zero-inflated counterparts (e.g., zero-inflated Poisson versus Poisson, Poisson hurdle versus Poisson) indicates that in almost all cases, the zero-augmented models are preferred (p-value < 0.01). The only exception is the negative binomial hurdle versus regular negative binomial model, in which case the regular negative binomial model is preferred (p-value < 0.01). Similarly to the information criteria presented earlier, the Vuong statistic also indicates that the zero-inflated models are preferred over their hurdle counterparts and that the negative binomial-based models are preferred over their Poisson counterparts (p-value < 0.01 in all four cases). This leads us to again favour the ZINB model.

As with previous models, it is beneficial to compared the expected probabilities of the various models with the observed probabilities of the commuting flows. In particular, we want to evaluate if the additional zero-inflation of the zero-augmented models better predicts zero flows than the previous Poisson and negative binomial models. Again, we do this by comparing the observed and predicted *counts* for zero and large flows ($flows > 250$), to get an idea of the model fit(s) at both tails of the commuting flow distribution. Contrary to previous diagnostics in this chapter, this comparison generally favours the negative binomial model: 1) the ZINB model over-predicts zero flows (270420-269289=+1131), whereas 2) the ZIP model, while much closer than the original Poisson model, still under-predicts zero flows (264269-269289=-5020). Both these predicted zero counts are in contrast to the earlier negative binomial model, which only under-predicted zero-flows by 393. For larger commuting flows, 3) the ZINB model is now slightly closer to the observed number (327-125=+202; compare with a difference of 388 without the zero-inflation), whereas 4) the ZIP model is now under-predicting large flows (86-125=-39), which was not occurring with the regular Poisson model. While 5) the hurdle models both predict large (and zero[4]) flows quite well (HURDNB: 215-125=90; HURDP: 78-125=-47), small to moderate flows are not well presented. This is much better visualised by comparing the actual predicted commuting flows versus the observed commuting flows (Burger et al., 2009), as in Figure 8.1, where we plot the residuals from the various models against the observed commuting flows.

From Figure 8.1, it is clear that the three negative binomial-based models outperform the Poisson-based models in terms of out-of-sample forecast. This is particularly evident for medium to large commuting flows, where in some cases, the Poisson-based estimators produce errors an order of magnitude larger than their negative binomial counterparts. While the negative binomial-based models do not appear to under-predict the volume of commuting at any given point, there does appear to be significant over-prediction for extremely small flows. Both the zero-inflated and hurdle negative binomial models are relatively poor predictors for small flows ($0 > flows < \sim 75$) when compared with the negative binomial model, however, in absolute terms they perform quite well for the most part. As mentioned before however, both negative binomial-based zero-augmented models are better at predicting *zero* flows than the basic negative binomial model, which accounts for a large proportion of the total flows. This is evidenced by the relative values of the squared correlation coefficients[5] between observed and predicted flows for the two best

[4]Recall that by definition, the hurdle models reproduce the observed number of zeros flows exactly.

[5]The squared correlation coefficient used here can be considered another form of a pseudo-R^2 measure.

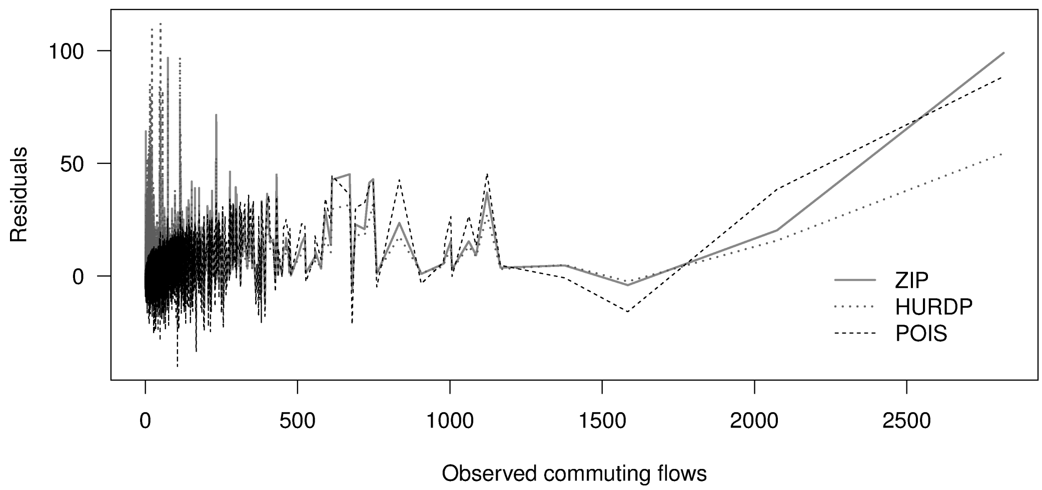

(a)

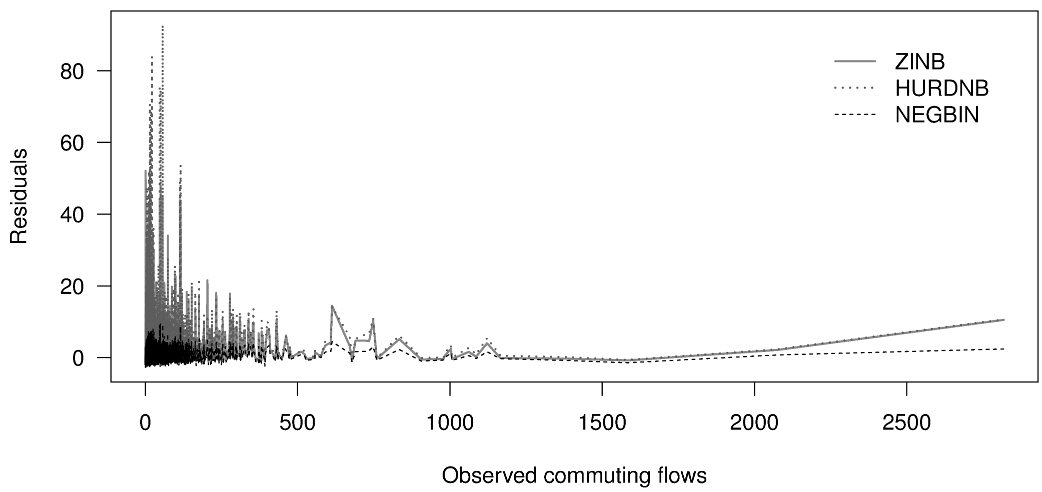

(b)

Figure 8.1: Model residuals for Poisson (a), and negative binomial (b) based models. Residuals are computed as the observed minus predicted commuting flows and are used here to examine model goodness-of-fit.

models: NEGBIN (0.1644) and ZINB (0.2737). This measure actually favours the zero-inflated model, despite the observation that, visually, the negative binomial model appears to be superior to all other models in terms of minimising residuals.

Clearly, the negative binomial and ZINB models both have merits and drawbacks; neither model is strongly preferred over the other. The negative binomial model appears to fit the observed data better, as evidenced by the smaller residuals and larger pseudo-R^2 value from Table 6.2. However, since we are using MLE, we would not expect residuals to be explicitly minimised as in OLS and, in fact, there are several problems inherent in these types of goodness-of-fit tests (see Cameron & Windmeijer (1996) and Burger et al. (2009)). Furthermore, while the negative binomial model appears to provide a better fit to the observed data, the ZINB model has been shown to have the higher likelihood. In the end however, the 'best' model will depend on the extent to which over-dispersion and excess zeros are empirically relevant (Burger et al., 2009). As such, the zero-inflated model may make more theoretical sense for our current application.

8.4 Moving forward

The zero-augmented models presented in this chapter provide a powerful and theoretically pleasing interpretation of spatial interaction models of commuting, particularly when there are a number of structural and observational zero flows. Furthermore, the interpretation of the transition component of these two-part models provides an intuitive and previously unavailable means of integrating choice set generation into existing spatial interaction models. Additionally, we have show empirically that the zero-inflated negative binomial model performs relatively well in the context of real-world commuting flows and in many respects is superior to more simplistic modified Poisson models such as the negative binomial model.

In the following chapter we present our main conclusions and recommendations derived from the culmination of our previous theoretical, methodological, and empirical findings. We summarise our main contributions, discuss the implications of our results, and detail future directions for extending this research area. In particular, we provide linkages between this and the previous seven chapters and readdress some of the key objectives and goals presented in our Introduction. Finally, specific contributions of this research to the commuting, regionalisation, and spatial interaction literature are discussed, followed by a retrospective examination of the body of work encompassed by this thesis.

Chapter 9

Conclusions

9.1 Research findings

With this thesis, we have attempted to tease out aspects of local labour market effects on commuting by asking two main questions: "to what extent does commuting enforce the spatial boundaries of local labour markets?" and "how do the boundaries of local labour markets influence commuting?". In doing so, we have addressed a series of research goals and objectives designed to guide our empirical research. These research goals were separated into three distinct but related targets: 1) define and delineate local labour markets (functional regions) for Ireland in a logical and defensible manner, 2) develop a spatial interaction model of commuting which takes into account the problems and opportunities inherent in Irish commuting patterns, and 3) incorporate goals 1 and 2 into a single modelling framework to investigate the effects of local labour markets on models of commuting flows. Coupling the theoretical and methodological concepts covered in chapters 2 and 3 with the travel-to-work data presented in Chapter 4, we have addressed these research goals via our empirical analyses in chapters 5, 6, 7, and 8.

9.1.1 Generating functional regions

We started our empirical investigation of local labour market effects on commuting patterns by defining and delineating a series of functional regions for Ireland. We did this for aggregate commuting patterns and for the specific commuting patterns of various socio-economic sub-groups, allowing us to examine how the commuting patterns of different components of the workforce contribute to the formation and maintenance of the local labour market. In doing so, we were able to develop a useful definition of a local labour market in the context of the functional regionalisation literature, having many desirable properties, including limiting the need for tuning parameters or threshold values. Furthermore, the regionalisation algorithm developed in this phase of our research (Chapter 2) has value for other types of research where interaction data are the focus, including migration, transportation, and many ecological phenomena.

Empirical evaluation of the functional regionalisation procedure and associated outputs was completed in Chapter 5. By weighting the commuting flows to take into account the geographical distances between regions, the procedure was shown to find stable, spatially-constrained local labour markets in both simulated and real-world geographical networks. In addition, the stability of the regionalisation procedure was tested using bootstrap resampling techniques originally developed by Hennig (2007) and extended to a regionalisation framework in Chapter 5. Using this technique, we were able to measure the effects of noise and random error on the algorithm's performance, ultimately leading to a more statistically viable means of evaluating algorithm performance and output suitability. Thus, from chapters 2 and 5 we were able to develop an efficient means of evaluating disaggregate functional regionalisations in order to examine the structure of the aggregate functional regions and, by association, the local labour market. With no need to specify *a priori* threshold values or self-containment criteria, we presented modularity maximisation as a useful tool for both exploring the local labour market structure and providing a point of origin for defining aggregate functional regions.

In terms of results, our regionalisation exercises revealed several important structural features. For instance, throughout the regionalisations there remained several consistent ‘periphery zones’ which tended to be much larger in size than their surrounding urban (hub) functional regions. Areas where these periphery functional regions were most common include the north- and south-west corners of central Ireland, including portions of counties Mayo, Cork, Galway, and Limerick. Additionally, the sub-group regionalisations presented in conjunction with the aggregate analysis showed clear differences in functional region characteristics between population sub-groups. While several of these differences were associated with specific spatial locations, such as periphery regions, others were less obvious at the aggregate level; highlighting the importance of considering the structure of local labour markets with respect to the intricacies of sub-group commuting behaviour. Our final conclusions from this phase of the research suggested that while no single functional regionalisation - whether it be based on aggregate or disaggregate data - can capture the true structure of complex commuting patterns, it is clear that structure does indeed exist and that this structure may be used to explain a significant amount of variation in commuting flows.

9.1.2 Modelling commuting flows

Our second primary goal was to develop spatial interaction models of commuting which are able to accommodate the unique context of travel-to-work and, in particular, travel-to-work in Ireland. We addressed this goal by developing a base-model for commuting in Ireland from modern spatial interaction principals. As part of this goal, we provided accumulated evidence for the superiority of negative binomial-based spatial interaction models over standard Poisson spatial interaction models in the context of sparse commuting flows. Empirical findings from this stage of analysis revealed that in addition to distance and working population size, the spatial structure of origins and destinations and a number of non-spatial attributes such as unemployment, housing density, and education, all significantly affected commuting flows. Additional improvements in model performance were supported by statistical results and spatial interaction theory and provided the necessary evidence to move forward with our selected empirical model, leading to more fundamental discoveries at later stages of our research.

These more fundamental discoveries were made clear in chapters 6, 7, and 8 and include the fact that local labour markets were found to play a key role in differentiating between different types of commuting patterns. This is corroborated by previous work by Thorsen & Gitlesen (1998) and Gitlesen & Thorsen (2000), who find that special care should be taken regarding the benefits of residing and working in the same local region. What our research indicates is that a failure to account for local labour market conditions may seriously hinder the applicability of models of commuting. Indeed, parameter estimates from models that do not take into account local labour market effects may misrepresent spatial effects such as distance decay because the effects of the local labour market become entangled with the attributes of origins and/or destinations. As such, models which take into account the effects of local labour markets, as well as local models at the level of the local labour market, may offer benefits over more simplistic models of commuting flows.

The combination of our first two primary goals provided strong evidence that there are important differences between commuting at the national and local levels and, in particular, that there are important local labour market effects that must be taken into account when considering travel-to-work behaviour. In general, we found strong evidence that commuting flows decrease rapidly with distance and that origins surrounded by a large number of potential destinations will tend to export fewer commuters to a given destination than their more isolated neighbours *ceteris paribus*. Furthermore, we found that when over-dispersion is accounted for via the negative binomial model, the effects of unemployment and housing density on commuting flows decreases, indicating that these variables may be associated with heterogeneity in commuting behaviour. When accounting for local labour market effects, we also found a significant positive increase in internal commuting flows and furthermore, by considering internal and external commuting separately, we were also able to significantly improve the fit of our models. From these findings, we can confidently say that, at least at the global level, local labour markets have a *significant* impact on observed and predicted commuting flows.

9.1.3 Local labour market effects

Our third and final research goal - examining local labour market effects on commuting flows and, more specifically, to implement local spatial interaction models at the level of the local labour market - was addressed in Chapter 7. In this chapter, we took our analysis of local labour markets to the next level, examining both inter- and intra-regional flows. Here, we examined inter-regional commuting patterns as a separate process and found that there are major differences in commuting behaviour/patterns when we look at commuting from the perspective of travel-to-work *between* local labour markets. In particular, we found that the concept of functional distance decay once again provided a viable explanation for the changes in distance decay between global, internal, external, and inter-regional commuting flows. Following our examination of inter-regional commuting flows, we progressed our evaluation of local labour market effects one step further by evaluating *local* spatial interaction at the level of the local labour market. From this, we explored parameter trends/differences across Ireland and exposed the extent to which local labour markets have an influence on travel-to-work.

This empirical stage allowed us to 'tie together' all previous findings in a unified manner and solidified our theories of functional distance decay. Important findings from this phase of the analysis included the fact that commuting flows *between* local labour markets were characterised by significantly different variable responses than commuting flows *within* a single local labour market. Furthermore, we found that there are clear spatial variations in parameter estimates that appear to be related to local labour markets, particularly for the distance, accessibility, and housing variables. Similarly, while most parameters were only significant in particular local labour markets, in some cases there was a clear relationship between parameter significance and the (spatially) local attributes of origins. Moreover, we noted clear heterogeneity in the parameter estimates, indicating that commuting processes vary across the study region and that the nature of travel-to-work is more complex than one might assume, especially where there is the potential for multiple interacting local labour markets.

9.1.4 Choice set integration

A further improvement to the modelling efforts in this thesis stems from the work on excess zeros and choice set generation presented in Chapter 8. In this work, we suggested that the negative binomial model sometimes spuriously indicates over-dispersion, when in fact the underlying process may actually consist of a two-part process generating excess zeros (Shankar et al., 1997). Building on this notion, we explored the concept of choice set generation as a viable theoretical interpretation of a two-part commuting process. More specifically, we modelled travel-to-work as a two-part process, whereby the first process controls whether or not commuting between i and j is *possible*, and the second process controls the *volume* (i.e., number of commuters travelling from i to j) of commuting given that commuting is possible. Both theoretical and empirical results provided strong evidence for the superiority of a zero-augmented approach to spatial interaction models of commuting. Additionally, we found that the concept of zero-inflation, whereby zero flows may arise from two sources[1], provided a superior theoretical interpretation *and* empirical fit to our Dublin commuting data.

While not directly extending the work of the previous three empirical chapters, the work presented in Chapter 8 provides significant improvements in model performance and interpretation over many of our previous models. Furthermore, we found that while a more accessible origin (i) will have a slightly higher *probability* of sending commuters to a given destination j, the actual *volume* of commuters it sends will be less than its more isolated neighbours *ceteris paribus*. In addition, results from Chapter 8 indicate that distance, education, and housing all have a strong impact on the probability of commuting between any origin and destination pair. Most importantly however, our use of zero-augmented models in this chapter provides an intuitive link to the concept of choice set generation in commuting models, and may provide an important linkage between models of spatial choice and spatial interaction in general.

9.2 Future directions

The research presented in this thesis provides many interesting and important avenues for future research. In general, the methods and models applied here could be extended to other regions and datasets and the findings presented in this conclusion could be put to use in future analyses of inter-/intra-regional

[1]As opposed to the concept of a hurdle, where there is only one source of zeros.

commuting patterns. Additionally, the general work-flow adhered to throughout this thesis (i.e., from regionalisation to modelling at global and local levels) could be implemented in extremely different settings with the goal of answering fundamentally different questions. In particular, we envisage this type of research being applied to such broadly different areas as animal home-range research and retail and housing market analyses. The specific individual components of this thesis also provide additional avenues for developments, whether they be focused on improving the current implementation or being applied in entirely new directions.

9.2.1 Local labour markets

The modularity maximisation procedure presented in chapters 2 and 5 is effective in delineating functional regions, although some limitations do arise in practice. According to Fortunato & Barthélemy (2007), the original modularity formulation suffers from a resolution limit, such that functional regions which are smaller than some threshold value r, may be combined to form larger functional regions, thereby missing out on potentially important substructures. This value of r will depend on both the size of the overall network, as well as the interconnectedness of the functional regions within the network (Porter et al., 2009). There are several potential solutions to this problem, some of which require a resolution parameter to be specified (e.g., Arenas et al., 2008; Blondel et al., 2008). However, the use of additional parameters is not ideal in the context of functional regionalisation, as it is difficult to know beforehand which value of r to use. Future work will explore related methods (using alternate quality functions for example), which could be employed in a similar context to the work presented here to limit the need for resolution parameters (see Porter et al. (2009) for some potetially viable solutions/examples).

Furthermore, because our regionalisation procedure employs a divisive, hierarchical heuristic to optimise modularity, it is possible that the final regionalisation may be sub-optimal in terms of the objective function. This is a common problem with divisive hierarchical methods, which is one of the reasons the current procedure employs an intermediate optimisation step in the regionalisation algorithm. The procedure outlined in Section 5.2.2 is designed to test for the occurrence of these types of errors and, in the current case, these problems were shown to be minimal. The hierarchical nature of the proposed procedure also provides a convenient means of evaluating functional regions at multiple scales, leading to potentially useful insights into the multi-scale nature of the labour market, which is currently another open research topic.

A third avenue worth exploring in terms of regionalisation research involves the 'scale' of analysis. In our current implementation, the final regionalisation is ultimately linked to the choice of bandwidth h. Indeed, preliminary tests indicate that by altering the bandwidth and keeping all else fixed, the number of functional regions generated as h is increased (from 1 kilometre to 100 kilometres), approaches that of the non-weighted regionalisation after a bandwidth of approximately 20 kilometres. Therefore, in order to produce functional regions which are operationally valid (i.e., relatively small and coherent), a value of h between the 2nd and 3rd quartiles should be used. This being the case, further research into the implications of differing values of h, and whether this geographical weighting can be altered to avoid possible complications due to bandwidth size, will be explored.

It is also important to note that the resampling strategy presented in Section 5.2.2 is not limited to any particular functional regionalisation procedure and can therefore be used to assess the stability of any regionalisation procedure. It is this wide-ranging applicability that makes it a particularly useful method for comparing the stability of different functional regionalisation procedures, which is itself is a topic for further study.

9.2.2 Spatial interaction

The modelling exercises performed in this thesis also provide some interesting avenues for further research, including as a means to improve our current model performance and as a means to further tie together the substantive topics covered in this thesis. In terms of improving model performance, for example, the negative binomial model used in this thesis is not necessarily a 'true' fixed effects model and, as such, the constraints imposed on the negative binomial model are not as strict as they are for the Poisson model. Hausman et al. (1984) have presented a conditional negative binomial model which estimates parameters via conditional maximum likelihood methods and can be relatively easily calibrated using standard numerical maximisation routines. This model should provide the means to easily estimate destination constraints, allowing us to take advantage of the power of the negative binomial model while enforcing stricter constraints than those currently implemented in this thesis. However, recent evidence suggests that the Hausman model may only apply in relatively restricted circumstances, namely when the constraints are directly related (in a very specific way) to the parameter of over-dispersion (Allison & Waterman, 2002; Guimarães, 2008). Incidentally, there is evidence to suggest that applying an unconditional negative binomial regression estimator with dummy variables to represent the fixed effects (as done in this thesis) may in fact yield results as viable as those based on the conditional model (Allison & Waterman, 2002).

In travel-to-work situations, it is possible that there will be no commuters travelling between two regions i and j for two completely separate reasons: 1) due to limitations imposed by a lack of transportation infrastructure or jobs, and/or 2) simply because the attributes of the origin, destination, or their separation are not conducive to commuting. There are many possible reasons for this type of two-stage commuting process/behaviour to occur; however, the most likely explanations revolve around the notion of choice set generation, as presented in Chapter 8. While the work presented in this chapter provided strong evidence in favour of a two-part commuting process, further work is required to conclusively present zero-augmented models as a viable alternative to more traditional spatial interaction models. As such, future research will further explore the potential applicability of zero-augmented models and provide a more solid theoretical foundation upon which to build this potentially powerful spatial interaction framework.

When working with highly sparse spatial interaction datasets, as is the case for large-area commuting dataset like the one used in this thesis, local models can become unstable; particularly when origin/destination constraints are involved. As we observed in Section 7.3, this can lead to areas where parameters are unreliable (i.e., due to a lack of degrees of freedom) or simply unavailable (i.e., due to rank deficient matrices). In these situations, it may be useful to explore techniques designed to "borrow strength" from surrounding regions to account for small samples . This is common in small-area population forecasting and disease mapping applications. Techniques such as conditional autoregressive models, multilevel models, and/or random effects models should be explored in this context. Additionally, this type of problem would also likely benefit from the work on geographically weighted spatial interaction models currently being undertaken here at the National Centre for Geocomputation.

Finally, whether explicitly or implicitly, the concept of 'functional distance' has played a key role in the various methods and results produced throughout the successive phases of this research. This is in fact one of the key findings of this thesis, although the *extent* to which crossing the boundaries of local labour markets decreases commuting flows has not been fully solved. As such, another important avenue of future research based on the modelling exercises initiated in this thesis is to try to extricate the actual effect of crossing into another local labour market from the overall distance decay. In other words, we aim to tease out the difference between *functional* and *actual* distance decay, and potentially utilise this information to refine the boundaries of local labour markets. This final step would provide another means of integrating the two primary streams of research presented in this thesis.

9.3 General conclusions

The four preceding empirical chapters, coupled with the theoretical, conceptual, and methodological issues discussed throughout this thesis, have ultimately lead to a better understanding of the importance of the local labour market as a spatial entity unto itself. Clearly the spatial structure of local labour markets has a strong connection to observed and predicted commuting patterns and this connection is undoubtedly one that is reciprocal in nature: while local labour markets influence the patterns of travel-to-work, so too does travel-to-work influence the distribution of local labour markets. The interplay between these two complementary views on local labour market effects is difficult to disentangle from what is observed 'on the ground'; however the evidence presented in this thesis provides a step towards this goal.

We have shown in this thesis that network-based partitioning methods lend themselves well to the delineation and analysis of functional regions and in many respects may offer benefits over legacy regionalisation methods from the geographical literature. Furthermore, the theory underlying travel-to-work areas (TTWAs) and local labour markets provides an intuitive link to the concept of modularity and, as a result, methods which explicitly consider the modular nature of travel-to-work data should be further explored. Attention should also be paid to examining the ‘stability’ of these and previous regionalisation methods. How stable are these functional regions and how much confidence can one have in their delineation? Would the same or similar functional regions be observed under different amounts of error or noise? These are important questions to consider when choosing a functional regionalisation procedure and statistically exploring these questions will ultimately provide the justification necessary for moving forward with a particular representation of local labour markets.

The boundaries of local labour markets are an integral component in the characterisation of inter- and intra-regional commuting patterns. It is not until we begin to separate commuting based on the milieu in which it takes place - local labour markets - that we begin to see the extent to which local labour markets influence commuting patterns. Furthermore, it is not until we examine commuting at the level of the local labour market, be it between or within them, that we realise the extent to which commuting influences the shape and configuration of local labour markets. As such, studies which explicitly consider the local labour market as an important spatial entity are required if we are to properly understand the behavioural, economic, and social factors at play in travel-to-work systems. As has been noted by spatial interaction researchers in the past, distance decay is inextricably linked to the spatial structure of origins and destinations. However, a key finding of this thesis is that distance decay is not only dependent on the configuration of origins and destinations, but also on the spatial structure of local labour markets, or more generally, the totality of surrounding conditions in which spatial interaction takes place.

9.4 Moving forward

There has been a surprising dearth of work on spatial interaction in geography over the last few decades, particularly in the context of urban spatial processes. This is bewildering, given the wide ranging applicability of spatial interaction modelling. As Olsson (1970) and Fotheringham & O'Kelly (1989, p. 233) note:

> The concept of spatial interaction is central for everyone concerned with theoretical geography and regional science ... Under the umbrella of spatial interaction and distance decay, it has been possible to accommodate most model work in transportation, migration, commuting, and diffusion, as well as significant aspects of location theory.

More recently however, the need to explain and understand complex spatial interactions in both human and physical environments has reinvigorated interest in spatial interaction modelling. This is likely due, in part, to the renewed interest in network analysis and related techniques in physics and sociology (see for instance the many recent publications in Science, Social Networks, Proceedings of the National Academy of Sciences (PNAS), and Physical Review (E, X, Letters), as well as several recent papers by Mark Newman, and Albert-László Barabási), as well as the continued interest in understanding the patterns of emergence, complexity, and interactions of modern urban systems. Furthermore, the now ubiquitous availability of new forms of spatial interaction data, such as mobile phone communication, user generated content and data (UGC), and social network services such as Facebook, Twitter, Academia.edu, *ceteris paribus*, means that new questions can be asked and answered using spatial interaction modelling techniques and theories.

9.4.1 Ultima ratio

In closing, spatial analysis and spatial interaction are relevant to many areas of geographical inquiry. The power and possibilities when spatially-enabled data are applied to real-world problems is potentially infinite. The role of geographers is to work to develop new approaches, foster expanded use of spatial analysis in new fields, and educate others in the use of these methods. Due to rapidly increasing spatial data availability and the multidisciplinary nature of many of today's problems, there is a growing niche for spatial research that is able to integrate data and theories across systems. Our hope is that the research presented in this thesis provides a step towards filling this niche. As we move forward, a deeper understanding of the complex nature of urban systems, including travel-to-work, will undoubtedly become central in characterising urban processes and, this understanding will ultimately come at the intersection of regional science, complexity, and spatial interaction.

Bibliography

Agresti, A., 1990. *Categorical data analysis*. New York, NY: Wiley.

Allaway, A. W., Berkowitz, D., & D'Souza, G., 2003. Spatial diffusion of a new loyalty program through a retail market. *Journal of Retailing*, *79*(3), 137–151.

Allen, W. B., 1972. An economic derivation of the "Gravity law" of spatial interaction: A comment on the reply. *Journal of Regional Science*, *12*(1), 119–126.

Allison, P. D., & Waterman, R. P., 2002. Fixed–Effects negative binomial regression models. *Sociological Methodology*, *32*(1), 247–265.

Alonso, W., 1978. A theory of movement. In N. M. Hansen (Ed.) *Human Settlement Systems*, (pp. 197–211). Cambridge: Ballinger.

Anas, A., & Chu, C., 1984. Discrete choice models and the housing price and travel to work elasticities of location demand. *Journal of Urban Economics*, *15*(1), 107–123.

Anselin, L., 1995. Local indicators of spatial association-LISA. *Geographical Analysis*, *27*(2), 93–115.

Anselin, L., & Getis, A., 1992. Spatial statistical analysis and geographic information systems. *Annals of Regional Science*, *26*(1), 19.

Arenas, A., Fernández, A., & Gómez, S., 2008. Analysis of the structure of complex networks at different resolution levels. *New Journal of Physics*, *10*(5), 053039.

Ball, R. M., 1980. The use and definition of travel-to-work areas in Great Britain: Some problems. *Regional Studies: The Journal of the Regional Studies Association*, *14*(2), 125–139.

Bartley, B., & Kitchin, R. (Eds.) , 2007. *Understanding Contemporary Ireland*. Pluto Press.

Batten, D. F., & Boyce, D. E., 1986. Spatial interaction, transportation and interregional commodity flow models. In P. Nijkamp (Ed.) *Handbook of Regional and Urban Economics*, Regional Economics, (pp. 357–406). Amsterdam: North-Holland Publishing, 1 ed.

Baxter, M., 1982. Similarities in methods of estimating spatial interaction models. *Geographical Analysis*, *14*(3), 267–272.

Baxter, M., 1984. A note on the estimation of a nonlinear migration model using GLIM. *Geographical Analysis*, *16*(3), 282–286.

Baxter, M. J., 1985. Quasi-likelihood estimation and diagnostic statistics for spatial interaction models. *Environment and Planning A*, *17*(12), 1627–1635.

Ben-Akiva, M., & Lerman, S. R., 1985. *Discrete choice analysis: Theory and application to travel demand*. MIT Press Cambridge, Mass.

Benjamini, Y., & Yekutieli, D., 2001. The control of the false discovery rate in multiple testing under dependency. *The Annals of Statistics*, *29*(4), 1165–1188. ArticleType: research-article / Full publication date: Aug., 2001 / Copyright © 2001 Institute of Mathematical Statistics.

Bhat, C. R., & Guo, J., 2004. A mixed spatially correlated logit model: Formulation and application to residential choice modeling. *Transportation Research Part B: Methodological*, *38*(2), 147–168.

Birkin, M., Clarke, G., & Clarke, M., 2010. Refining and operationalizing entropy-maximizing models for business applications. *Geographical Analysis*, *42*(4), 422–445.

Blondel, V. D., Guillaume, J., Lambiotte, R., & Lefebvre, E., 2008. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, *2008*(10), P10008.

Borgers, A., & Timmermans, H., 1986. A model of pedestrian route choice and demand for retail facilities within inner-city shopping areas. *Geographical Analysis*, *18*(2).

Breathnach, P., 1998. Exploring the 'Celtic tiger' phenomenon: Causes and consequences of ireland's economic miracle. *European Urban and Regional Studies*, *5*(4), 305–316.

Breslow, N., 1990. Tests of hypotheses in overdispersed poisson regression and other quasi-likelihood models. *Journal of the American Statistical Association*, *85*(410), 565–571.

Brown, L. A., & Holmes, J., 1971. The delimitation of functional regions, nodal regions, and hierarchies by functional distance approaches. *Journal of Regional Science*, *11*(1), 57–72.

Brown, L. A., & Horton, F. E., 1970. Functional distance: An operational approach. *Geographical Analysis*, *2*(1), 76–83.

Brunsdon, C., Fotheringham, A., & Charlton, M., 1996. Geographically weighted regression: A method for exploring spatial nonstationarity. *Geographical Analysis*, *28*(4), 281–298.

Brunsdon, C., Fotheringham, A. S., & Charlton, M., 1998. Geographically weighted regression - modelling spatial non-stationarity. *Journal of the Royal Statistical Society (Series D): The Statistician*, *47*(3), 431–443.

Burger, M., Oort, F. V., & Linders, G., 2009. On the specification of the gravity model of trade: Zeros, excess zeros and zero-inflated estimation. *Spatial Economic Analysis*, *4*(2), 167–187.

Cameron, A. C., & Trivedi, P. K., 1990. Regression-based tests for overdispersion in the poisson model. *Journal of Econometrics*, *46*(3), 347–364.

Cameron, A. C., & Trivedi, P. K., 1998. *Regression analysis of count data*. Cambridge University Press.

Cameron, A. C., & Windmeijer, F. A., 1996. R-squared measures for count data regression models with applications to health-care utilization. *Journal of Business & Economic Statistics*, (p. 209–220).

Carr, K., 2008. Qualitative research to assess interest in public transportation for work commute. *Journal of Public Transportation*, *11*(1), 1–16.

Casado-Díaz, J. M., 2000. Local labour market areas in Spain: A case study. *Regional Studies*, *34*(9), 843–856.

Clark, W. A. V., & Onaka, J. L., 1985. An empirical test of a joint model of residential mobility and housing choice. *Environment and Planning A*, *17*(7), 915–930.

Clarke, G., & Clarke, M., 2001. Applied spatial interaction modelling. In G. Clarke, & M. Madden (Eds.) *Regional science in business*, Advances in spatial science, (pp. 137–157). Germany: Springer-Verlag.

Clarke, G., Langley, R., & Cardwell, W., 1998. Empirical applications of dynamic spatial interaction models. *Computers, Environment and Urban Systems*, *22*, 157–184.

Clarke, G., & Madden, M., 2001. *Regional science in business*. Advances in spatial science. Germany: Springer-Verlag.

Clauset, A., Newman, M. E. J., & Moore, C., 2004. Finding community structure in very large networks. *Physical Review E*, *70*(6), 066111–1–066111–6.

Clifton, K. J., & Handy, S. L., 2003. Qualitative methods in travel behaviour research. In P. R. Stopher, & P. M. Jones (Eds.) *Transport Survey Quality and Innovation*, (pp. 283–302). Oxford: Elsevier Science Ltd.

Commins, N., & Nolan, A., 2010a. Car ownership and mode of transport to work in Ireland. *The Economic and Social Review*, *41*(1), 43–75.

Commins, N., & Nolan, A., 2010b. The determinants of mode of transport to work in the Greater Dublin Area. *Transport Policy*.

Congdon, P., 1993. Approaches to modelling overdispersion in the analysis of migration. *Environment and Planning A*, *25*(10), 1481 – 1510.

Convery, F. J., Mcinerney, D., Sokol, M., & Stafford, P., 2006. Organizing space in a dynamic economy: Insights for policy from the Irish experience. *Built Environment*, *32*(2), 172–183.

Coombes, M., & Bond, S., 2007. Travel-to-Work Areas: The 2007 review. Final project report, Office for National Statistics, UK.

Coombes, M., & Casado-Díaz, J. M., 2005. The evolution of local labour market areas in contrasting regions. *ERSA conference papers, European Regional Science Association*.

Coombes, M. G., Dixon, J. S., Goddard, J. B., Openshaw, S., & Taylor, P. J., 1982. Functional regions for the population census of Great Britain. *Geography and the Urban Environment: Progress in Research and Applications*, *5*, 63–112.

Coombes, M. G., Green, A. E., & Openshaw, S., 1986. An efficient algorithm to generate official statistical reporting areas: The case of the 1984 Travel-to-Work Areas revision in Britain. *The Journal of the Operational Research Society*, *37*(10), 943–953.

Coombes, M. G., Green, A. E., & Owen, D. W., 1988. Substantive issues in the definition of "localities": Evidence from sub–group local labour market areas in the West Midlands. *Regional Studies*, *22*(4), 303.

Coombes, M. G., & Openshaw, S., 1982. The use and definition of travel-to-work areas in Great Britain: Some comments. *Regional Studies: The Journal of the Regional Studies Association*, *16*(2), 141–149.

Cordey-Hayes, M., & Wilson, A. G., 1971. Spatial interaction. *Socio-Economic Planning Sciences*, *5*, 73–95.

Cörvers, F., Hensen, M., & Bongaerts, D., 2009. Delimitation and coherence of functional and administrative regions. *Regional Studies*, *43*(1), 19–31.

CSO, 2006a. Census 2006. Place of work census of anonymised records (POWCAR), Central Statistics Office Ireland, Ireland.

CSO, 2006b. Census 2006. Small area population statistics (SAPS), Central Statistics Office Ireland, Ireland.

CSO, 2007. Census 2006. Volume 12 - Travel to Work, School and College, Central Statistics Office Ireland, Ireland.

CSO, 2011. Live register. April 2011, Central Statistics Office Ireland, Ireland.

Cubukgil, A., & Miller, E., 1982. Occupational status and the journey-to-work. *Transportation*, *11*(3).

Curry, L., 1972. A spatial analysis of gravity flows. *Regional Studies*, *6*(2), 131–147.

Daganzo, C. F., 1979. *Multinomial probit: The theory and it's application to demand forecasting*. New York, NY: Academic Press.

Danon, L., Díaz-Guilera, A., Duch, J., & Arenas, A., 2005. Comparing community structure identification. *Journal of Statistical Mechanics: Theory and Experiment*, (09), P09008.

Davies, R. B., & Guy, C. M., 1987. The statistical modeling of flow data when the poisson assumption is violated. *Geographical Analysis*, *19*(4), 300–314.

Dean, C., & Lawless, J. F., 1989. Tests for detecting overdispersion in poisson regression models. *Journal of the American Statistical Association*, *84*(406), 467–472.

Dean, C. B., 1992. Testing for overdispersion in poisson and binomial regression models. *Journal of the American Statistical Association*, *87*(418), 451–457.

Dodd, S. C., 1950. The interactance hypothesis: A gravity model fitting physical masses and human groups. *American Sociological Review*, *15*(2), 245–256. ArticleType: primary_article / Full publication date: Apr., 1950 / Copyright © 1950 American Sociological Association.

Drobne, S., Konjar, M., Lisec, A., Milanovi, N. P., & Lamovšek, A. Z., 2010. Functional regions defined by urban centres of (inter) national importance: The case of Slovenia. In *Cities For Everyone: Liveable, Healthy, Prosperous*. Reed Messe Wien (Austria/Vienna).

Duch, J., & Arenas, A., 2005. Community detection in complex networks using extremal optimization. *Physical Review E*, *72*(2), 027104.

Efron, B., & Tibshirani, R. J., 1993. *An introduction to the bootstrap*. Chapman & Hall.

Elhorst, J., & Oosterhaven, J., 2006. Forecasting the impact of transport improvements on commuting and residential choice. *Journal of Geographical Systems*, *8*(1), 39–59.

Eurostat, 1992. Study on employment zones. Eurostat (E/LOC/20), Luxembourg.

Farmer, C. J. Q., Nelson, T. A., Wulder, M. A., & Derksen, C., 2010. Identification of snow cover regimes through spatial and temporal clustering of satellite microwave brightness temperatures. *Remote Sensing of Environment*, *114*(1), 199–210.

Felsenstein, J., 1985. Confidence limits on phylogenies: An approach using the bootstrap. *Evolution*, *39*(4), 783–791.

Findlay, A. M., Stockdale, A., Findlay, A., & Short, D., 2001. Mobility as a driver of change in rural Britain: An analysis of the links between migration, commuting and travel to shop patterns. *International Journal of Population Geography*, *7*(1), 1–15.

Fischer, M. M., 2001. Neural spatial interaction models. In M. M. Fischer, & Y. Leung (Eds.) *Geocomputational modelling: Thechniques and appliactions*, Advances in Spatial Science, (pp. 195–219). Berlin: Springer.

Flórez-Revuelta, F., Casado-Díaz, J., & Martínez-Bernabeu, L., 2008. An evolutionary approach to the delineation of functional areas based on travel-to-work flows. *International Journal of Automation and Computing*, *5*(1), 10–21.

Flowerdew, R., & Aitkin, M., 1982. A method of fitting the gravity model based on the Poisson distribution. *Journal of Regional Science*, *22*(2), 191–202.

Flowerdew, R., & Lovett, A., 1988. Fitting constrained poisson regression models to interurban migration flows. *Geographical Analysis*, *20*(4), 297–307.

Fortunato, S., & Barthélemy, M., 2007. Resolution limit in community detection. *Proceedings of the National Academy of Sciences*, *104*(1), 36–41.

Fotheringham, A., Brunsdon, C., & Charlton, M., 2002. *Geographically Weighted Regression: The Analysis of Spatially Varying Relationships*. London: John Wiley & Sons.

Fotheringham, A. S., 1980. *Spatial structure, spatial interaction, and distance-decay parameters*. Open access dissertations and theses, McMaster University, McMaster, Ontatrio, Canada.

Fotheringham, A. S., 1981. Spatial structure and distance-decay parameters. *Annals of the Association of American Geographers*, *71*(3), 425–436.

Fotheringham, A. S., 1983a. A new set of spatial-interaction models: The theory of competing destinations. *Environment and Planning A*, *15*(1), 15–36.

Fotheringham, A. S., 1983b. Some theoretical aspects of destination choice and their relevance to production-constrained gravity models. *Environment and Planning A*, *15*(8), 1121–1132.

Fotheringham, A. S., 1984. Spatial flows and spatial patterns. *Environment and Planning A*, *16*(4), 529–543.

Fotheringham, A. S., 1986. Modelling hierarchical destination choice. *Environment and Planning A*, *18*(3), 401–418.

Fotheringham, A. S., 1987. Hierarchical destination choice: Discussion with evidence from migration in The Netherlands. Working Paper no. 69, Netherlands Interuniversity Demographic Institute, The Hague, Netherlands.

Fotheringham, A. S., 1988. Consumer store choice and choice set definition. *Marketing Science*, *7*(3), 299–310.

Fotheringham, A. S., 1997. Trends in quantitative methods I: Stressing the local. *Progress in Human Geography*, *21*(1), 88–96.

Fotheringham, A. S., 2001. Spatial interaction models. In N. J. Smelser, & P. B. Baltes (Eds.) *International Encyclopedia of the Social and Behavioral Sciences*, (pp. 14,794–14,800). Oxford: Elsevier Science.

Fotheringham, A. S., 2006. Quantification, evidence and positivism. In S. Aitken, & G. Valentine (Eds.) *Approaches to Human Geography*, (pp. 180–220). London: Sage Publications.

Fotheringham, A. S., Brunsdon, C., & Charlton, M., 1996. The geography of parameter space: An investigation of spatial non-stationarity. *International Journal of Geographical Information Systems*, *10*, 605–627.

Fotheringham, A. S., Brunsdon, C., & Charlton, M., 2000a. *Quantitative geography: Perspectives on spatial data analysis*. London: Sage Publications.

Fotheringham, A. S., Champion, T., Wymer, C., & Coombes, M., 2000b. Measuring destination attractivity: A migration example. *International Journal of Population Geography*, *6*(6), 391–421.

Fotheringham, A. S., & Dignan, T., 1984. Further contributions to a general theory of movement. *Annals of the Association of American Geographers*, *74*(4), 620–633.

Fotheringham, A. S., & O'Kelly, M. E., 1989. *Spatial interaction models: Formulations and applications*. Dordrecht: Kluwer Academic Publishers.

Fotheringham, A. S., & Rogerson, P. A., 1993. GIS and spatial analytical problems. *International Journal of Geographical Information Science*, *7*(1), 3–19.

Fotheringham, A. S., & Trew, R., 1993. Chain image and store-choice modeling: the effects of income and race. *Environment and Planning A*, *25*(2), 179–196.

Fotheringham, A. S., & Webber, M. J., 1980. Spatial structure and the parameters of spatial interaction models. *Geographical Analysis*, *12*(1), 33–46.

Fotheringham, A. S., & Williams, P. A., 1983. Further discussion on the poisson interaction model. *Geographical Analysis*, *15*(4), 343–347.

Fotheringham, S., Nakaya, T., Yano, K., Openshaw, S., & Ishikawa, Y., 2001. Hierarchical destination choice and spatial interaction modelling: A simulation experiment. *Environmenta and Planning A*, *33*, 901–920.

Frost, M., Linneker, B., & Spence, N., 1998. Excess or wasteful commuting in a selection of British cities. *Transportation Research Part A: Policy and Practice*, *32*(7), 529–538.

Gabriel, S. A., & Rosenthal, S. S., 1996. Commutes, neighborhood effects, and earnings: An analysis of racial discrimination and compensating differentials. *Journal of Urban Economics*, *40*(1), 61–83.

Gardner, W., Mulvey, E. P., & Shaw, E. C., 1995. Regression analyses of counts and rates: Poisson, overdispersed poisson, and negative binomial models. *Psychological Bulletin*, *118*(3), 392.

Gehlke, C., & Biehl, K., 1934. Certain effects of grouping upon the size of the correlation coefficient in census tract material. *Journal of the American Statistical Association*, *29*(185), 169–170.

GeoDirectory, 2011. GeoDirectory. http://www.geodirectory.ie/.

Geographical Analysis, 2010. Special issue: A 40th anniversary celebration of alan wilson, 1970, entropy in urban and regional modelling. *Geographical Analysis*, *42*(4), 361–488.

Gerard, R., 1958. Commuting and the labor market area. *Journal of Regional Science*, *1*(1), 124–130.

Getis, A., & Ord, K., 1992. The analysis of spatial association by use of distance statistics. *Geographical Analysis*, *24*(3), 189–206.

Girvan, M., & Newman, M. E. J., 2002. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences of the United States of America*, *99*(12), 7821–7826.

Gitlesen, J. P., & Thorsen, I., 2000. A competing destinations approach to modeling commuting flows: a theoretical interpretation and an empirical application of the model. *Envionrment and Planning A*, *32*(11), 2057–2074.

Giuliano, G., & Narayan, D., 2003. Another look at travel patterns and urban form: The US and Great Britain. *Urban Studies*, *40*(11), 2295–2312.

Giuliano, G., & Small, K. A., 1993. Is the journey to work explained by urban structure? *Urban Studies*, *30*(9), 1485–1500.

Gobillon, L., & Selod, H., 2007. The effect of segregation and spatial mismatch on unemployment: Evidence from france. CEPR Discussion Paper DP6198, Centre for Economic Policy Research, London.

Goldner, W., 1955. Spatial and locational aspects of metropolitan labor markets. *The American Economic Review*, *45*(1), 113–128.

Goodman, J. F. B., 1970. The definition and analysis of local labour markets: Some empirical problems. *British Journal of Industrial Relations*, *8*, 179–196.

Gordon, P., Richardson, H. W., & Jun, M. J., 1991. The commuting paradox: Evidence from the top twenty. *Journal of the American Planning Association*, *57*(4), 416–420.

Green, A., Coombes, M., & Owen, D., 1986. Gender-specific local labour market areas in England and Wales. *Geoforum*, *17*(3-4), 339–351.

Green, A. E., Hogarth, T., & Shackleton, R. E., 1999. Longer distance commuting as a substitute for migration in britain: a review of trends, issues and implications. *International Journal of Population Geography*, *5*(1), 49–67.

Green, A. E., & Owen, D. W., 1990. The development of a classification of travel-to-work areas. *Progress in Planning*, *34*(1), 1–92.

Greene, W. H., 1994. Accounting for excess zeros and sample selection in poisson and negative binomial regression models. Working Paper 94-10, Leonard N. Stern School of Business, New York University.

Griffith, D. A., 2010. Celebrating 40 years of scientific impacts by alan wilson. *Geographical Analysis*, *42*(4), 361–363.

Griffith, D. A., & Jones, K. G., 1980. Explorations into the relationship between spatial structure and spatial interaction. *Environment and Planning A*, *12*(2), 187–201.

Grosvenor, T., 2000. Qualitative research in the transportation sector. In *Proceedings of an International Conference on Transport Survey Quality and Innovation May 24-30, 1997*. Grainau, Germany.

Guimarães, P., 2008. The fixed effects negative binomial model revisited. *Economics Letters*, *99*(1), 63–66.

Guimerà, R., & Amaral, L. A. N., 2005. Functional cartography of complex metabolic networks. *Nature*, *433*(7028), 895–900.

Guimerà, R., Mossa, S., Turtschi, A., & Amaral, L. A. N., 2005. The worldwide air transportation network: Anomalous centrality, community structure, and cities' global roles. *Proceedings of the National Academy of Sciences of the United States of America*, *102*(22), 7794–7799.

Guldmann, J. M., 1999. Competing destinations and intervening opportunities interaction models of inter-city telecommunication flows. *Papers in Regional Science*, *78*(2), 179–194.

Guy, C. M., 1987. Recent advances in spatial interaction modelling: An application to the forecasting of shopping travel. *Environment and Planning A*, *19*(2), 173–186.

Hansen, W. G., 1959. How accessibility shapes land use. *Journal of the American Planning Association*, *25*(2), 73–76.

Hanson, S., & Pratt, G., 1988. Reconceptualizing the links between home and work in urban geography. *Economic Geography*, *64*(4), 299–321.

Hanushek, E. A., 1981. Alternative models of earnings determination and labor market structures. *The Journal of Human Resources*, *16*(2), 238–259.

Hausman, J. A., Hall, B. H., & Griliches, Z., 1984. Econometric models for count data with an application to the Patents-R&D relationship. *National Bureau of Economic Research Technical Working Paper Series*, *No. 17*.

Haynes, K. E., & Fotheringham, A. S., 1984. *Gravity and spatial interaction models*. California: Sage Publications.

Hennig, C., 2007. Cluster-wise assessment of cluster stability. *Computational Statistics & Data Analysis*, *52*(1), 258–271.

Hensen, M., & Cörvers, F., 2003. The regionalization of labour markets by modelling commuting behaviour. In *43rd Congress of the European Regional Science Association*. Jyväskylä, Finland.

Hollingsworth, T. H., 1971. Gross migration flows as a basis for regional definition: An experiment with Scottish data. In *Proceedings, IUSSP Conference, London, 1969*, vol. 4, (pp. 2755–2765).

Horner, A., 1999. The tiger stirring: Aspects of commuting in the Republic of Ireland 1981–1996. *Irish Geography*, *32*(2), 99–111.

Horner, M. W., 2004. Spatial dimensions of urban commuting: A review of major issues and their implications for future geographic research. *The Professional Geographer*, *56*(2), 160–173.

Horowitz, J., 1980. A utility maximizing model of the demand for multi-destination non-work travel. *Transportation Research Part B: Methodological*, *14*(4), 369–386.

Huber, P. J., 1967. The behavior of maximum likelihood estimates under nonstandard conditions. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, vol. 1, (p. 221–33).

Huff, D. L., 1959. Geographical aspects of consumer behavior. *University of Washington Business Review*, *18*, 27–37.

Huff, D. L., 1963. A probabilistic analysis of shopping center trade areas. *Land Economics*, *39*(1), 81–90.

Huff, D. L., 1964. Defining and estimating a trading area. *The Journal of Marketing*, *28*(3), 34–38.

Hunter, L. C., 1969. Planning and the labour market. In S. C. Orr, & J. Cullingworth (Eds.) *Regional and Urban Studies*. George Allen & Unwin, Glasgow.

Jaccard, P., 1901. Distribution de la flore alpine dans le bassin des dranses et dans quelques régions voisines. *Bulletin de la société vaudoise des sciences naturelles*, *37*, 241–272.

Johnston, R. J., 1973. On frictions of distance and regression coefficients. *Area*, *5*(3), 187–191.

Jones, C., Marron, J. S., & Sheather, S. J., 1996. Progress in data-based bandwidth selection for kernel density estimation. *Computational Statistics*, (11), 337–381.

Jones, P. M., Dix, M. C., Clarke, M. I., & Heggie, I. G., 1985. *Understanding Travel Behaviour*. Aldershot, England: Gower.

Kain, J. F., 1968. Housing segregation, negro employment, and metropolitan decentralization. *The Quarterly Journal of Economics*, *82*(2), 175.

Kain, J. F., 1992. The spatial mismatch hypothesis: Three decades later. *Housing policy debate*, *3*(2), 371–460.

Karlsson, C., & Olsson, M., 2006. The identification of functional regions: Theory, methods, and applications. *The Annals of Regional Science*, *40*(1), 1–18.

Kauermann, G., & Carroll, R. J., 2001. A note on the efficiency of sandwich covariance matrix estimation. *Journal of the American Statistical Association*, *96*(456), 1387–1396.

Kernighan, B. W., & Lin, S., 1970. An efficient heuristic procedure for partitioning graphs. *Bell System Technical Journal*, *49*(2), 291–307.

Kerr, M. K., & Churchill, G. A., 2001. Bootstrapping cluster analysis: Assessing the reliability of conclusions from microarray experiments. *Proceedings of the National Academy of Sciences of the United States of America*, *98*(16), 8961–8965.

Kitchin, R., Gleeson, J., Keaveney, K., & O'Callaghan, C., 2010. A haunted landscape: Housing and ghost estates in Post-Celtic tiger ireland. Working Paper Series No. 59, NIRSA.

Konjar, M., Lisec, A., & Drobne, S., 2010. Methods for delineation of functional regions using data on commuters. In *Geospatial Thinking, Proceedings of AGILE*, (pp. 1–11). Guimaraes, Portugal.

Kuppam, A. R., Pendyala, R. M., & Rahman, S., 1999. Analysis of the role of traveler attitudes and perceptions in explaining mode-choice behavior. *Transportation Research Record*, *1676*, 68–76.

Lambert, D., 1992. Zero-inflated poisson regression, with an application to defects in manufacturing. *Technometrics*, *34*(1), 1–14.

Lawless, J. F., 1987. Negative binomial and mixed poisson regression. *Canadian Journal of Statistics*, *15*(3), 209–225.

Leicht, E. A., & Newman, M. E. J., 2008. Community structure in directed networks. *Physical Review Letters*, *100*(11), 118703–1–118703–4.

Leinbach, T. R., 1973. Distance, information flows, and modernization: Some observations from West Malaysia. *The Professional Geographer*, *25*(1), 7–11.

Lerman, S. R., & Liu, T. K., 1984. Microlevel econometric analysis of retail closure. In D. Pitfield (Ed.) *Discrete Choice Models in Regional Science*, (p. 181–201). London, UK: Pion Ltd.

Lloyd, C. D., 2011. *Local Models for Spatial Analysis*. CRC Press.

Longley, P. A., 1984. Comparing discrete choice models: Some housing market examples. In D. Pitfield (Ed.) *Discrete Choice Models in Regional Science*, (p. 163–180). London, UK: Pion Ltd.

Louviere, J. J., & Timmermans, H., 1990. Stated preference and choice models applied to recreation research: a review. *Leisure Sciences*, *12*(1), 9–32.

Lovett, A., & Flowerdew, R., 1989. Analysis of count data using poisson regression. *The Professional Geographer*, *41*(2), 190–198.

Manski, C. F., 1977. The structure of random utility models. *Theory and Decision*, *8*(3), 229–254.

Masser, I., & Brown, P. J. B., 1975. Hierarchical aggregation procedures for interaction data. *Environment and Planning A*, *7*(5), 509–523.

Masser, I., & Scheurwater, J., 1980. Functional regionalisation of spatial interaction data: An evalution of some suggested strategies. *Environment and Planning A*, *12*(12), 1357–1382.

Mathur, V. K., 1970. An economic derivation of the "Gravity law" of spatial interaction: A comment. *Journal of Regional Science*, *10*(3), 403–405.

McArthur, D. P., Kleppe, G., Thorsen, I., & Ubře, J., 2010. The spatial transferability of parameters in a gravity model of commuting flows. *Journal of Transport Geography*.

McCullagh, P., & Nelder, J. A., 1989. *Generalized linear models*. London: Chapman Hall.

McFadden, D., 1980. Econometric models for probabilistic choice among products. *The Journal of Business*, *53*(3), 513–529.

McKenna, J. E., 2003. An enhanced cluster analysis program with bootstrap significance testing for ecological community analysis. *Environmental Modelling and Software*, *18*(3), 205–220.

McLafferty, S., 1997. Gender, race, and the determinants of commuting: New York in 1990. *Urban Geography*, *18*(3), 192–212.

McLafferty, S., & Preston, V., 1991. Gender, race, and commuting among service sector workers. *The Professional Geographer*, *43*(1), 1–15.

MECLG, 2002. National spatial strategy for Ireland 2002-2020. People, places and potential, Ministry for the Environment, Community and Local Government, Ireland.

Miller, J. M., 2007. *Comparing Poisson, Hurdle, and ZIP model fit under varying degrees of skew and zero-inflation*. Dissertation, University of Florida, Florida, U.S.A.

Milligan, G., & Cooper, M., 1985. An examination of procedures for determining the number of clusters in a data set. *Psychometrika*, *50*(2), 159–179.

Morgenroth, E., 2002. Commuting in Ireland: An analysis of inter-county commuting flows. Papers WP144, Economic and Social Research Institute (ESRI).

Moss, J. E., Jack, C. G., & Wallace, M. T., 2004. Employment location and associated commuting patterns for individuals in disadvantaged rural areas in Northern Ireland. *Regional Studies*, *38*(2), 121–136.

Mullahy, J., 1986. Specification and testing of some modified count data models. *Journal of econometrics*, *33*(3), 341–365.

Nakaya, T., Fotheringham, A. S., Hanaoka, K., Clarke, G., Ballas, D., & Yano, K., 2007. Combining microsimulation and spatial interaction models for retail location analysis. *Journal of Geographical Systems*, *9*, 345–369.

Nelder, J. A., & Wedderburn, R. W., 1972. Generalized linear models. *Journal of the Royal Statistical Society. Series A (General)*, *135*(3), 370–384.

Nemec, A. F. L., & Brinkhurst, R. O., 1988. Using the bootstrap to assess statistical significance in the cluster analysis of species abundance data. *Journal Canadien des Sciences Halieutiques et Aquatiques*, *45*(6), 965–970.

Newman, M. E. J., 2003. The structure and function of complex networks. *SIAM Review*, *45*(2), 167–256.

Newman, M. E. J., 2004a. Analysis of weighted networks. *Physical Review E*, *70*(5), 056131.

Newman, M. E. J., 2004b. Fast algorithm for detecting community structure in networks. *Physical Review E*, *69*, 066133.

Newman, M. E. J., 2006a. Finding community structure in networks using the eigenvectors of matrices. *Physical Review E*, *74*(3), 036104–19.

Newman, M. E. J., 2006b. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, *103*(23), 8577–8582.

Newman, M. E. J., & Girvan, M., 2004. Finding and evaluating community structure in networks. *Physical Review E*, *69*(2), 026113–1–026113–15.

Ng, A. Y., Jordan, M. I., & Weiss, Y., 2002. On spectral clustering: Analysis and an algorithm. In T. Dietterich, S. Becker, & Z. Ghahramani (Eds.) *Advances in neural information processing systems*, vol. 14, (pp. 849–856). Cambridge: MIT Press.

Niedercorn, J. H., & Bechdolt Jr, B. V., 1969. An economic derivation of the "Gravity law" of spatial interaction. *Journal of Regional Science*, *9*(2), 273–282.

Niedercorn, J. H., & Bechdolt Jr, B. V., 1970. An economic derivation of the "Gravity law" of spatial interaction: Reply. *Journal of Regional Science*, *10*(3), 407–410.

Niedercorn, J. H., & Bechdolt Jr, B. V., 1972. An economic derivation of the "Gravity law" of spatial interaction: A further reply and a reformulation. *Journal of Regional Science*, *12*(1), 127–136.

Nielsen, T. A. S., & Hovgesen, H. H., 2008. Exploratory mapping of commuter flows in england and wales. *Journal of Transport Geography*, *16*(2), 90–99.

Noronha, V. T., & Goodchild, M. F., 1992. Modeling interregional interaction: Implications for defining functional regions. *Annals of the Association of American Geographers*, *82*(1), 86–102.

O'Kelly, M., 2009. Applied retail location models using spatial interaction tools. In A. S. Fotheringham, & P. Rogerson (Eds.) *The SAGE handbook of spatial analysis*. London, UK: Sage Publications.

O'Kelly, M. E., 1981. A model of the demand for retail facilities, incorporating multistop, multipurpose trips. *Geographical Analysis*, *13*(2), 134–148.

O'Kelly, M. E., 2004. Isard's contributions to spatial interaction modeling. *Journal of Geographical Systems*, *6*(1), 43–54.

O'Kelly, M. E., & Lee, W., 2005. Disaggregate journey-to-work data: Implications for excess commuting and jobs-housing balance. *Environment and Planning A*, *37*(12), 2233–2252.

O'Kelly, M. E., & Niedzielski, M. A., 2008. Efficient spatial interaction: Attainable reductions in metropolitan average trip length. *Journal of Transport Geography*, *16*(5), 313–323.

O'Kelly, M. E., & Niedzielski, M. A., 2009. Are long commute distances inefficient and disorderly? *Environment and Planning A*, *41*(11), 2741–2759.

Olsson, G., 1970. Explanation, prediction, and meaning variance: An assessment of distance interaction models. *Economic Geography*, *46*, 223–233.

OMB, 2000. Standards for defining metropolitan and micropolitan statistical areas. Federal register, Office of Management and Budget (OMB), Office of Information and Regulatory Affairs, USA.

Onaka, J., & Clark, W. A. V., 1983. A disaggregate model of residential mobility and housing choice. *Geographical Analysis*, *15*(4), 287–304.

Openshaw, S., 1977. Optimal zoning systems for spatial interaction models. *Environment and Planning A*, *9*(2), 169–184.

Openshaw, S., Charlton, M., Wymer, C., & Craft, A., 1987. A mark 1 geographical analysis machine for the automated analysis of point data sets. *International Journal of Geographical Information Systems*, *1*(4), 335–358.

Openshaw, S., & Taylor, P. J., 1979. A million or so correlation coefficients: Three experiments on the modifiable areal unit problem. *Statistical Applications in the Spatial Sciences*, *127*, 144.

Ord, J., & Getis, A., 1995. Local spatial autocorrelation statistics: Distributional issues and an application. *Geographical Analysis*, *27*(4), 286–306.

Owen, D., & Green, A. E., 2000. Estimating commuting flows for minority ethnic groups in England and Wales. *Journal of Ethnic and Migration Studies*, *26*(4), 581–608.

Palla, G., Derényi, I., Farkas, I., & Vicsek, T., 2005. Uncovering the overlapping community structure of complex networks in nature and society. *Nature*, *435*(7043), 814–818.

Papps, K. L., & Newell, J. O., 2002. Identifying functional labour market areas in New Zealand: A reconnaissance study using travel-to-work data. Discussion Paper 443, Institute for the Study of Labor (IZA), Bonn.

Pellegrini, P. A., & Fotheringham, A. S., 1999. Intermetropolitan migration and hierarchical destination choice: A disaggregate analysis from the US Public Use Microdata Samples. *Environment and Planning A*, *31*, 1093–1118.

Pellegrini, P. A., & Fotheringham, A. S., 2002. Modelling spatial choice: A review and synthesis in a migration context. *Progress in Human Geography*, *26*(4), 487–510.

Pellegrini, P. A., Fotheringham, A. S., & Lin, G., 1997. An empirical evaluation of parameter sensitivity to choice set definition in shopping destination choice models. *Papers in Regional Science*, *76*(2), 257–284.

Pendyala, R. M., Goulias, K. G., & Kitamura, R., 1991. Impact of telecommuting on spatial and temporal patterns of household travel. *Transportation*, *18*(4), 383–409.

Peterson, G. L., Dwyer, J. F., & Darragh, A. J., 1983. A behavioral urban recreation site choice model. *Leisure Sciences*, *6*(1), 61–81.

Pillar, V. D., 1999. How sharp are classifications? *Ecology*, *80*(8), 2508–2516.

Pisarski, A., 2002. Transportation trends and smart growth. In *Presentation to Transportation Research Board Workshop, Providing a Transportation System to Support Smart Growth: Issues, Practice and Implementation, September*, (pp. 8–11).

Polena, C., & Glazer, L. J., 1991. Examination of 11 guaranteed ride home programs nationwide. *Transportation Research Record*, *1321*, 57–65.

Porter, M. A., Onnela, J. P., & Mucha, P. J., 2009. Communities in networks. *Notices of the AMS*, *56*(9), 1082–1166.

Prashker, J., Shiftan, Y., & Hershkovitch-Sarusi, P., 2008. Residential choice location, gender and the commute trip to work in Tel Aviv. *Journal of Transport Geography*, *16*(5), 332–341.

Ravenstein, E. G., 1885. The laws of migration. *Journal of the Royal Statistical Soceity*, *48*, 167–235.

Redmond, L. S., & Mokhtarian, P. L., 2001. The positive utility of the commute: modeling ideal commute time and relative desired commute amount. *Transportation*, *28*(2), 179–205.

Reichardt, J., & Bornholdt, S., 2006. Statistical mechanics of community detection. *Physical Review E*, *74*(1), 016110.

Reilly, W. J., 1929. Methods for the study of retail relationships. Bulletin 2944, University of Texas, Austin.

Renkow, M., & Hoover, D., 2000. Commuting, migration, and Rural-Urban population dynamics. *Journal of Regional Science*, *40*(2), 261–287.

Robinson, D., 1968. Wage drift, fringe benefits and manpower distribution: a study of employer practices in a full employment labour market. Tech. rep., Organisation for Economic Co-operation and Development.

Rodriguez, D. A., 2004. Spatial choices and excess commuting: A case study of bank tellers in Bogota, Colombia. *Journal of Transport Geography*, *12*(1), 49–61.

Røe, P. G., 2000. Qualitative research on intra-urban travel: An alternative approach. *Journal of Transport Geography*, *8*(2), 99–106.

Rogers, E. M., 1993. The diffusion of innovations model. In I. Masser, & H. J. Onsrud (Eds.) *Diffusion and Use of Geographic Information Technologies*, (pp. 9–24). London.

Rosenbloom, S., & Burns, E., 1993. Gender differences in commuter travel in Tucson: Implications for travel demand management programs. *Transportation Research Record*, *1404*, 82–90.

Rouwendal, J., & Rietveld, P., 1994. Changes in commuting distances of Dutch households. *Urban Studies*, *31*(9), 1545 –1557.

Roy, J. R., 2004. *Spatial interaction modelling: A regional science context*. Advances in spatial science. Berlin Heidelberg New York: Springer.

Roy, J. R., & Thill, J., 2003. Spatial interaction modelling. *Papers in Regional Science*, *83*(1), 339–361.

Sayer, R. A., 1992. *Method in social science: A realist approach*. Routledge.

Shankar, V., Milton, J., & Mannering, F., 1997. Modeling accident frequencies as zero-altered probability processes: an empirical inquiry. *Accident, Analysis and Prevention*, *29*(6), 829–837.

Shannon, C. E., 1948. A mathematial theory of communication. *Bell System Technical Journal*, *27*, 379–423 and 623–656.

Sheather, S. J., & Jones, M. C., 1991. A reliable data-based bandwidth selection method for kernel density estimation. *Journal of the Royal Statistical Society. Series B (Methodological)*, *53*(3), 683–690.

Shen, Q., 2000. Spatial and social dimensions of commuting. *Journal of the American Planning Association*, *66*(1), 68–82.

Singell, L. D., & Lillydahl, J. H., 1986. An empirical analysis of the commute to work patterns of males and females in two-earner households. *Urban Stud*, *23*(2), 119–129.

Singleton, A. D., Wilson, A. G., & O'Brien, O., 2010. Geodemographics and spatial interaction: an integrated model for higher education. *Journal of Geographical Systems*.

Slater, P. B., 1975. A hierarchical regionalization of Russian administrative units using 1966-69 migration data. *Soviet Geography*, *16*, 453–465.

Slater, P. B., 1976. A hierarchical regionalization of Japanese prefectures using 1972 interprefectural migration flows. *Regional Studies*, *10*(1), 123–132.

Slater, P. B., 1981. Comparisons of aggregation procedures for interaction data: An illustration using a college student international flow table. *Socio-Economic Planning Sciences*, *15*(1), 1–8.

Smart, M. W., 1974. Labour market areas: Uses and definition. *Progress in Planning*, *2*, 239–353.

Smith, S. P., & Dubes, R., 1980. Stability of a hierarchical clustering. *Pattern Recognition*, *12*(3), 177–187.

Sobel, K. L., 1980. Travel demand forecasting by using the nested multinominal logit model. *Transportation Research Record*, *775*, 48–55.

Spielman, D. A., & Teng, S., 2007. Spectral partitioning works: Planar graphs and finite element meshes. *Linear Algebra and its Applications*, *421*(2-3), 284–305.

Stevens, S. S., 1957. On the psychophysical law. *Psychological Review*, *64*(3), 153–181.

Stewart, J. Q., 1941. An inverse distance variation for certain social distances. *Science*, *93*, 89–90.

Stewart, J. Q., 1942. A measure of the influence of a population at a distance. *Sociometry*, *5*(1), 63–71.

Stillwell, J. C. H., 1978. Interzonal migration: Some historical tests of spatial-interaction models. *Environment and Planning A*, *10*(10), 1187–1200.

Stopher, P. R., 1992. Use of an activity-based diary to collect household travel data. *Transportation*, *19*(2), 159–176.

Stouffer, S. A., 1940. Intervening opportunities: A theory relating mobility and distance. *American sociological review*, *5*(6), 845–867.

Stouffer, S. A., 1960. Intervening opportunities and competing migrants. *Journal of Regional Science*, *2*(1), 1–26.

Sultana, S., & Weber, J., 2007. Journey-to-Work patterns in the age of sprawl: Evidence from two midsize southern metropolitan areas. *The Professional Geographer*, *59*(2), 193–208.

Taaffe, E. J., Gauthier, H. L., & O'Kelly, M. E., 1996. *Geography of transportation*. USA: Prentice-Hall, Inc.

Taylor, B. D., & Ong, P. M., 1995. Spatial mismatch or automobile mismatch? An examination of race, residence and commuting in US metropolitan areas. *Urban Stud*, *32*(9), 1453–1473.

Thorsen, I., & Gitlesen, J. P., 1998. Empirical evaluation of alternative model specifications to predict commuting flows. *Journal of Regional Science*, *38*(2), 273–292.

Thorsen, I., & Gitlesen, J. P., 2002. A simulation approach to studying the sensitivity of commuting-flow predictions with respect to specific changes in spatial structure. *Environment and Planning A*, *34*(2), 271–288.

Tiefelsdorf, M., & Boots, B., 1995. The specification of constrained interaction models using the SPSS loglinear procedure. *Geographical Systems*, *2*, 21–38.

Timmermans, H., 1984. Discrete choice models versus decompositional multi-attribute preference models: a comparative analysis of model performance in the context of spatial shopping-behaviour. In D. Pitfield (Ed.) *Discrete Choice Models in Regional Science*, (p. 88–102). London, UK: Pion Ltd.

Timmermans, H., & Golledge, R. G., 1990. Applications of behavioural research on spatial problems II: preference and choice. *Progress in Human Geography*, *14*(3), 311–354.

Train, K., 1986. *Qualitative choice analysis: Theory, econometrics, and an application to automobile demand*. Cambridge, MA: MIT Press.

Turner, T., & Niemeier, D., 1997. Travel to work and household responsibility: new evidence. *Transportation*, *24*(4), 397–419.

Ubøe, J., 2004. Aggregation of gravity models for journeys to work. *Environment and Planning A*, *36*(4), 715–729.

Ubøe, J., Petter Gitlesen, J., & Thorsen, I., 2008. Laboratory testing of spurious spatial structure in trip distribution models. *Spatial Economic Analysis*, *3*, 361–372.

Valente, T. W., 1996. Social network thresholds in the diffusion of innovations. *Social Networks*, *18*(1), 69–89.

van der Laan, L., 1991. *Spatial labour markets in the Netherlands*. Eburon.

van der Laan, L., & Schalke, R., 2001. Reality versus policy: The delineation and testing of local labour market and spatial policy areas. *European Planning Studies*, *9*(2), 201–221.

Vance, J. E., 1960. Labour shed, employment field and dynamic analysis in urban geography. *Economic Geography*, *36*(3), 189–220.

Vega, A., & Reynolds-Feighan, A., 2008. Employment sub-centres and travel-to-work mode choice in the Dublin region. *Urban Studies*, (pp. 1747 –1768).

Venables, W., & Ripley, B., 2002. *Modern Applied Statistics with S*. New York: Springer-Verlag, fourth ed.

Vermeulen, W., 2003. A model for Dutch commuting. Tech. Rep. 1, CPB Netherlands Bureau for Economic Policy Analysis, The Nederlands.

Vuong, Q. H., 1989. Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica: Journal of the Econometric Society*, (p. 307–333).

Wachs, M., Taylor, B. D., Levine, N., & Ong, P., 1993. The changing commute: A case-study of the jobs-housing relationship over time. *Urban Stud*, *30*(10), 1711–1729.

Walsh, J. A. (Ed.) , 2007. *People and Place - A census atlas of the Republic of Ireland*. Ireland: National Institute for Regional and Spatial Analysis.

Ward, J. H., 1963. Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, *58*(301), 236–244.

Warnes, A., 1972. Estimates of journey-to-work distances from census statistics. *Regional Studies: The Journal of the Regional Studies Association*, *6*(3), 315–326.

Watts, M., 2004. Local labour markets in New South Wales: Fact or fiction? Working Paper No. 04-12, Centre of Full Employment and Equity, The University of Newcastle, Callaghan NSW 2308, Australia.

Wedderburn, R. W., 1974. Quasi-likelihood functions, generalized linear models, and the Gauss—Newton method. *Biometrika*, *61*(3), 439.

Wegener, M., 2000. Spatial models and GIS. In A. S. Fotheringham, & M. Wegener (Eds.) *Spatial Models and GIS: New Potential and New Models*, no. 7 in GISDATA, (p. 3–20). Philadelphia: Taylor & Francis.

Werlen, B., 1993. *Society, action and space: An alternative human geography*. London: Routledge.

Westin, K., & Sandow, E., 2010. People's preferences for commuting in sparsely populated areas: The case of sweden. *Journal of Transport and Land Use*, *2*(3).

Wheeler, J. O., 1967. Occupational status and work-trips: A minimum distance approach. *Social Forces*, *45*(4), 508–515.

Wheeler, J. O., 1969. Some effects of occupational status on work trips. *Journal of Regional Science*, *9*(1), 69.

White, H., 1982. Maximum likelihood estimation of misspecified models. *Econometrica: Journal of the Econometric Society*, (p. 1–25).

White, M. J., 1977. A model of residential location choice and commuting by men and women workers. *Journal of Regional Science*, *17*(1), 41–52.

White, S., & Smyth, P., 2005. A spectral clustering approach to finding communities in graph. In *Proceedings of the 2005 SIAM International Conference on Data Mining*, (pp. 274–286). Newport Beach, CA, USA.

Williams, P. A., & Fotheringham, A. S., 1984. The calibration of spatial interaction models by maximum likelihood estimation with program SIMODEL. Geographic Monograph Series Volume 7, Department of Geography, Indiana University.

Wilson, A. G., 1967. Statistical theory of spatial trip distribution models. *Transportation Research*, *1*, 253–269.

Wilson, A. G., 1970. *Entropy in urban and regional modelling*. London, UK: Pion Ltd.

Wilson, A. G., 1971. A family of spatial interaction models, and associated developments. *Environment and Planning A*, *3*(1), 1–32.

Wilson, A. G., 1975. Some new forms of spatial interaction model: A review. *Transportation Research*, *9*(2-3), 167–179.

Wilson, A. G., 2000. The widening access debate: student flows to universities and associated performance indicators. *Environment and Planning A*, *32*(11), 2019–2031.

Wrigley, N., 1985. *Categorical Data Analysis for Geographers and Environmental Scientists*. London, UK: Longman Group Ltd.

Zeileis, A., Kleiber, C., & Jackman, S., 2008. Regression models for count data in R. *Journal of Statistical Software*, *27*(8), 1–25.

Zenou, Y., 2000. Urban unemployment, agglomeration and transportation policies. *Journal of Public Economics*, *77*(1), 97–133.

Zipf, G. K., 1949. *Human behavior and the principle of least effort*. Oxford, England: Addison-Wesley Press.

Appendix A

Technical Notes

Open-source software

All methods and analyses presented in this thesis can be/were performed using freely available open source software, including Quantum GIS (http://www.qgis.org/), the Python scripting language (http://www.python.org/), and the R statistical programming language (http://www.r-project.org/). The thesis itself was written using TeXstudio (http://texstudio.sourceforge.net/) and typeset using the LaTeX 2ε document preparation system (http://www.latex-project.org/). Additionally, the modelling and data exploration techniques performed in this thesis are currently being developed into a unified package for both the R statistical programming language and the Python programming language and will be made available to the wider public in the near future.

Software details

All statistical results in this thesis were obtained using R version 2.13.0 (2011-04-13), with packages rgeos_0.1-4, rgdal_0.6-33, sp_0.9-80, pscl_1.03.10, MatrixModels_0.2-1, lmtest_0.9-27, sandwich_2.2-6, MASS_7.3-13, as well as several secondary packages and custom R code (which is available upon request from the author).

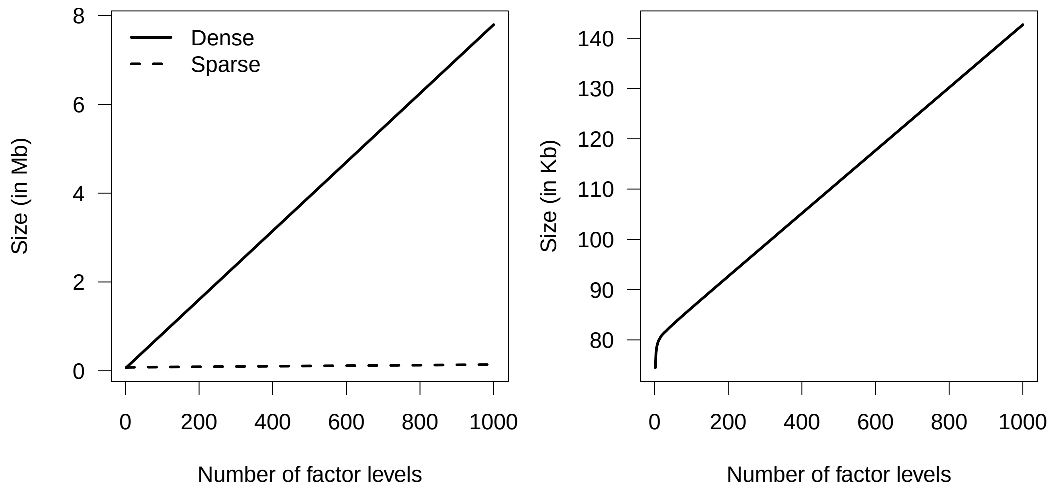

(a) (b)

Figure A.1: Memory usage of dense (a) versus sparse (b) model matrices. Note the the y-axis scales between the two figures differ (megabytes versus kilobytes), and that the sparse model trend is provided in (a) for reference. Factor levels correspond to the value of m (or n) in the context of constrained spatial interaction models.

The calibration of all statistical models in this thesis was performed using sparse matrix methods from the `MatrixModels` R package. In particular, we took advantage of the `glm4` routine, which uses sparse Cholsky decomposition in the fitting of model parameters via iterative re-weighted least squares (IRLS). This is because using regular dense matrix representations creates huge matrices that are too large to work with using conventional methods/tools (see Section 3.5.1). In terms of memory usage for example, the dense model matrices increase linearly as in Figure A.1a, whereas the sparse model matrices increase at a much slower rate (0.0625 Kbs/+level versus 7.9258 Kbs/+level), particularly after approximately 20 factor levels (i.e., m (or n) = 20) (Figure A.1b). The use of sparse model fitting routines was therefore required when working with the large commuting dataset presented in this thesis. Indeed, the model matrix used in the majority of the models explored in this thesis was 11,621,281 x 3,415 in size, which, when employing sparse matrix representations equated to approximately 1.7 Gb of memory! For further information on the `MatrixModels` package and its associated methods, see `http://cran.r-project.org/web/packages/MatrixModels/index.html`.

Grey-scale visualisations

All visualisations and figures in this thesis have been purposefully created and rendered in grey-scale. This was done for several reasons, including as an aide to printing, to promote a focus on the content of the figures/diagrams rather than their colours, to provide a unifying colour theme to the thesis as a whole, and finally, to prove that colours are often unnecessary distractions from effective visualisations and, for the most part, are not needed!

Appendix B

Supporting materials

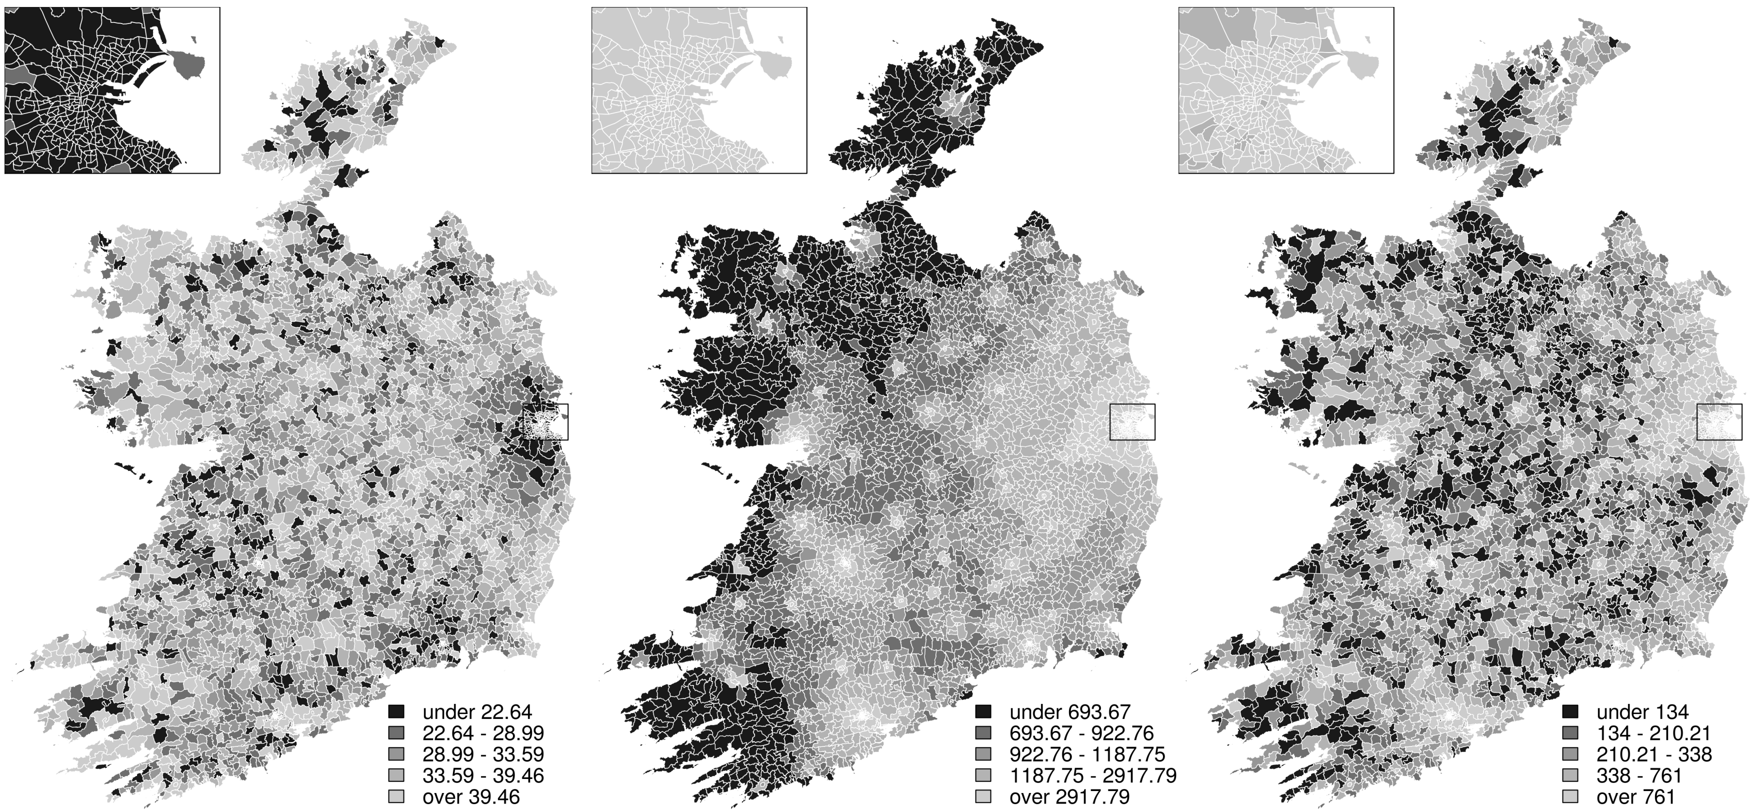

(a) distance (km) (b) accessibility (c) workers

Figure B.1: Spatial distribution of values for distance (mean distance travelled from each origin), accessibility, and the attributes of origins [continued on following page]. Note that class breaks are based on quantiles.

(d) unemployment

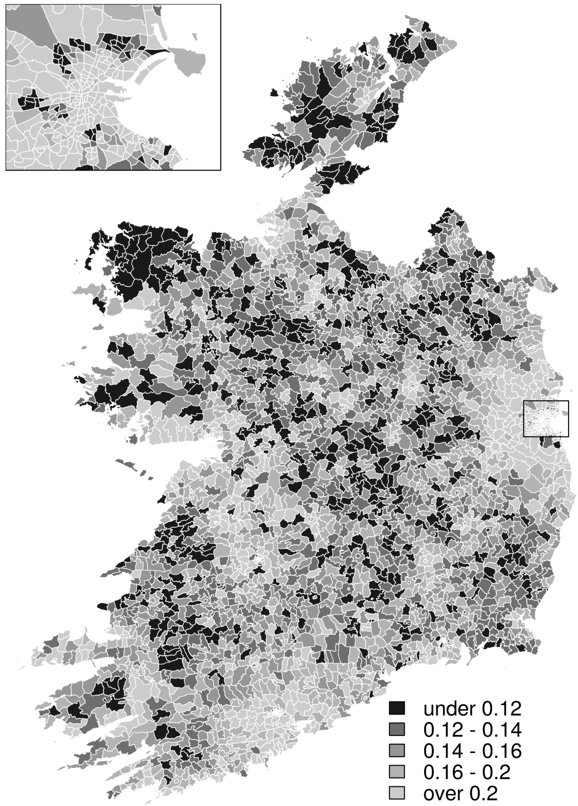

(e) education

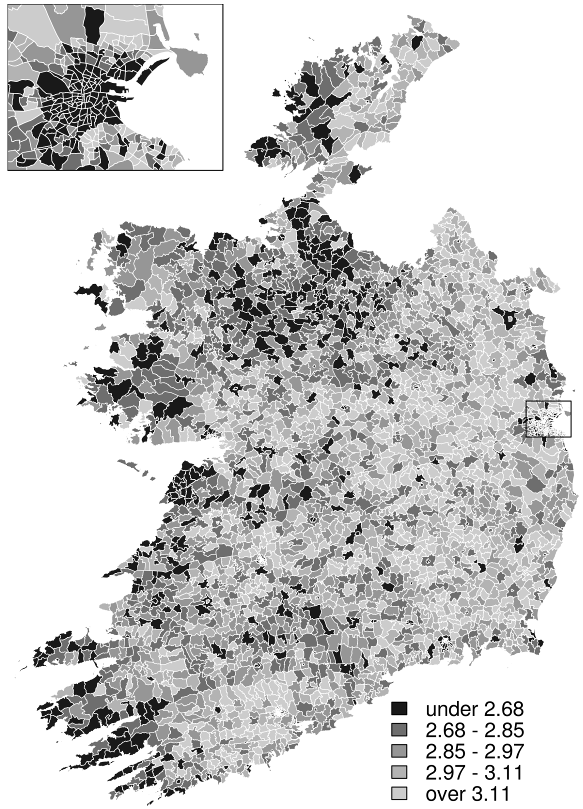

(f) housing

Figure B.1: [continued from previous page]