A SPATIAL ECONOMETRIC INVESTIGATION OF URBAN PROXIMITY AND LABOUR MARKET BEHAVIOUR AFTER THE NEWFOUNDLAND AND LABRADOR COD MORATORIUM

JAMIE WARD

A Spatial Econometric Investigation of Urban Proximity and Labour Market Behaviour after the Newfoundland and Labrador Cod Moratorium

A thesis submitted to the

School of Graduate Studies

in partial fulfillment of the

requirements for the degree of

Master of Science

Department of Geography

Memorial University of Newfoundland

St. John's, Newfoundland and Labrador

May 2011

Abstract

Traditionally, the majority of quantitative studies that analyze labour market phenomena have utilized global data and top-down methods. Often, however, labour market dynamics observed on the national or provincial scale are the result of processes operating on a lower-level, local scale. This discrepancy between the scale at which labour market processes operate and the scale at which they are studied reduces the resolution of analysis, and may result in faulty policy development. In addition, because labour markets normally operate in discrete space, traditional local methods which operate in continuous space, such as geographically weighted regression (GWR), are not appropriate. In this thesis, a method of adapting the GWR model to discrete space is described and tested. To evaluate its effectiveness, the discrete-space GWR (DGWR) technique is applied to the aftermath of the 1992 cod moratorium in Newfoundland and Labrador. To do this, a theoretical economic model of moratorium susceptibility and impact is constructed and tested using a DGWR, ordinary least squares, and continuous GWR model. Upon conducting the analysis, it is found that the DGWR is the superior technique with respect to both model fit and the mitigation of spatial effects. In addition, the DGWR model also produces more realistic and easily applicable empirical results than the existing alternatives.

Acknowledgements

I would like to thank my supervisor, Dr. Alvin Simms, for his valuable guidance and instruction during both the analysis and writing stages of this thesis. I would also like to thank my committee members, Dr. David Freshwater and Dr. Rob Greenwood, for taking the time to provide valuable comments and insights throughout the process. In addition, I would like to thank Service Canada (Government of Canada), the Department of Human Resources, Labour and Employment (Government of Newfoundland and Labrador), and Memorial University of Newfoundland for providing financial assistance, without which none of this would have been possible.

I would like to say thank you to my fiancée, Heather, for your love, support, and commitment, and to my siblings, Jeffrey, Jennifer, and Nicholas, for your support and competitive spirit. Underlying everything is the constant support of my parents, who instilled in me the value of work and commitment from a very early age.

Table of Contents

Abstract ii

Acknowledgements iii

List of Tables vi

List of Figures vii

1.0 – Introduction 1

1.1 – Introduction 1

1.2 – Research Problem and Hypotheses 3

1.3 – Research Approach Outline 5

1.4 – Thesis Organization 5

2.0 – Literature Review 7

2.1 – Introduction 7

2.2 – Labour Market Studies 8

2.2.1 – Demand and Supply Shocks 10

2.2.2 – Mobility and Minimum Wage Effects 12

2.3 – Spatial Modelling 16

2.3.1 – Broad Methods 17

2.3.2 – Regression Analysis 19

2.4 – Newfoundland and Labrador Labour Force 22

2.4.1 – History and Context 22

2.4.2 – The Cod Moratorium in the Literature 24

2.5 – Conclusion 27

3.0 – Methodology 29

3.1 – Introduction 29

3.2 – Theoretical Adaptation 30

3.2.1 – Demand Shock 31

3.2.2 – Spatial Structure 35

3.3 – Geographically Weighted Regression Model 40

3.3.1 – Spatial Structure Adaptation 42

3.3.2 – Application of Theory 46
3.4 – Model Inputs and Outputs 49
3.4.1 – Required Data 49
3.4.2 – Model Adaptations 55
3.4.2.1 – Community List 56
3.4.2.2 – Spatial Weight Function 56
3.4.2.3 – Model Condition 58
3.4.2.4 – Variable Exclusion 59
3.4.3 – Parameters and Interpretation 60
3.5 – Conclusion 64
4.0 – Results 65
4.1 – Introduction 65
4.2 – Model Fit Results 65
4.2.1 – Global Fit Statistics and Parameter Stability 66
4.2.2 – Spatial Statistics 69
4.2.3 – Model Condition and Pre-Analysis Results 72
4.3 – Empirical Results 79
4.3.1 – Overview 79
4.3.2 – Example 97
4.4 – Conclusion 100
5.0 – Conclusion 101
5.1 – Discussion of Results 101
5.2 – Policy Implications and Conclusions 102
5.3 – Future Research 103
6.0 – Bibliography 108

List of Tables

Table 4.1 Model Fit Results 68
Table 4.2 Parameter Stationarity Results, Quality Model (p-value) 68
Table 4.3 Parameter Stationarity Results, Utilization Model (p-value) 69
Table 4.4 Moran's I Results (Index value and z-score) 70
Table 4.5 White Test Results (p-value) 72
Table 4.6 Outlier Analysis Results, Quality Model 78
Table 4.7 Outlier Analysis Results, Utilization Model 78
Table 4.8 Community Comparison Results (pseudo-t score) 98

List of Figures

Figure 2.1 Labour Market Supply and Demand 11

Figure 2.2 Anticipated Effect of a Sudden Influx of Workers to a Stable Economy 12

Figure 2.3 Demand Shock with a Minimum Wage 16

Figure 2.4 Northern Cod Adjustment and Recovery Program (NCARP) Funding per Capita, 1993/1991 24

Figure 2.5 Rural vs. Urban Migration Prevalence, 1991-1996 26

Figure 3.1 Wage Pressure to Migrate from One Labour Market to Another 40

Figure 3.2 Straight Line versus Road Network Distance Example, Bonavista to St. John's 44

Figure 3.3 Example of Community vs. CSD Location and Size 53

Figure 3.4 Newfoundland and Labrador Commuting Distance Decay 57

Figure 4.1 Local Moran's I results for Labour Market Quality Change model, 1991-1996 71

Figure 4.2 Communities Included in or Excluded from the Analysis 73

Figure 4.3 Community Sample Size 74

Figure 4.4 Community Condition Number, 1991 (Quality Model) 75

Figure 4.5 Community Condition Number, 1991-1991 (Quality Model) 75

Figure 4.6 Community Condition Number, 1996-2001 (Quality Model) 76

Figure 4.7 Community Condition Number, 1991 (Utilization Model) 76

Figure 4.8 Community Condition Number, 1991-1996 (Utilization Model) 77

Figure 4.9 Community Condition Number, 1996-2001 (Utilization Model) 77

Figure 4.10 Percent Part Year Low (Utilization) Model Results, 1991 (pseudo t-score). 81

Figure 4.11 Percent Part Year Medium (Utilization) Model Results, 1991 (pseudo t-score) 81

Figure 4.12 Percent Part Year High (Utilization) Model Results, 1991 (pseudo t-score) 82

Figure 4.13 Percent Full Year (Utilization) Model Results, 1991 (pseudo t-score) 82

Figure 4.14 Percent Part Year Low (Utilization) Model Results, 1991-1996 (pseudo t-score) 83

Figure 4.15 Percent Part Year Medium (Utilization) Model Results, 1991-1996 (pseudo t-score) 84

Figure 4.16 Percent Part Year High (Utilization) Model Results, 1991-1996 (pseudo t-score) 84
Figure 4.17 Percent Full Year (Utilization) Model Results, 1991-1996 (pseudo t-score) 85
Figure 4.18 Percent Part Year Low (Utilization) Model Results, 1996-2001 (pseudo t-score) 86
Figure 4.19 Percent Part Year Medium (Utilization) Model Results, 1996-2001 (pseudo t-score) 86
Figure 4.20 Percent Part Year High (Utilization) Model Results, 1996-2001 (pseudo t-score) 87
Figure 4.21 Percent Full Year (Utilization) Model Results, 1996-2001 (pseudo t-score) 87
Figure 4.22 Percent Middle Age (Quality) Model Results, 1991 (pseudo t-score) 89
Figure 4.23 Percent Older Age (Quality) Model Results, 1991 (pseudo t-score)..... 89
Figure 4.24 Percent Middle Age (Quality) Model Results, 1991-1996 (pseudo t-score) . 91
Figure 4.25 Percent Older Age (Quality) Model Results, 1991-1996 (pseudo t-score) ... 91
Figure 4.26 Percent Middle Age (Quality) Model Results, 1996-2001 (pseudo t-score) . 92
Figure 4.27 Percent Older Age (Quality) Model Results, 1996-2001 (pseudo t-score) ... 92
Figure 4.28 Percent Medium Skill (Quality) Model Results, 1991 (pseudo t-score) 94
Figure 4.29 Percent High Skill (Quality) Model Results, 1991 (pseudo t-score)..... 94
Figure 4.30 Percent Medium Skill (Quality) Model Results, 1991-1996 (pseudo t-score) 95
Figure 4.31 Percent High Skill (Quality) Model Results, 1991-1996 (pseudo t-score) ... 95
Figure 4.32 Percent Medium Skill (Quality) Model Results, 1996-2001 (pseudo t-score) 96
Figure 4.33 Percent High Skill (Quality) Model Results, 1996-2001 (pseudo t-score) ... 96
Figure 4.34 Example Communities 97

1.0 – Introduction

1.1 – Introduction

July 2, 1992 was a pivotal day in the history of Newfoundland and Labrador. It was on this date that the venerable cod fishery, which had supported a large portion of the province's population for hundreds of years, was first closed (Murray *et al.*, 2005). Although initially intended to be a temporary measure to allow the overfished and dwindling fish stocks time to recover, almost twenty years later the fishery has still not returned to its prior size and importance. Before the fishery closure, the labour force of rural Newfoundland and Labrador was predominantly low-skilled and highly dependent upon the fishery and other resource extraction industries (Feehan, 1991). Labour markets in urban areas, however, displayed some measure of diversification beyond the fishery, to industries such as retail and government services (Schrank, 2005). Although this spatially heterogeneous exposure to the shock of the moratorium indicates that directly affected workers would have felt different pressures and made different decisions depending on which area of the province they happened to reside in, to this point this has not been a topic of an in-depth investigation.

The study of labour market dynamics is a relatively popular area of research in both economics and economic geography (Glendon and Vigdor, 2003; Coffey and Shearmur 2002), and offers a potential pathway for the study of the effects of the cod moratorium. In particular, there has historically been a focus upon quantitative methods and sound theoretical reasoning within the economics literature (Overman, 2004). This approach prefers the search for overall consistency of behaviour in different times and areas on a fundamental level over the documentation of behaviours that are idiosyncratic

and difficult to classify. For this reason, much work in this field is based on the functioning of idealized actors producing in idealized economies, and real-world implications are often difficult to discern (Martin, 1999). On the other hand, many recent studies originating in geography tend to take an opposite approach to studying labour market dynamics. Instead of focusing on generating theories and laws of behaviour, these studies are often based on qualitative data, and are more interested in the local and unique (Overman, 2004). This focus on the local and unique makes it difficult to efficiently generate policy for individual regions, however, because simplifying assumptions that allow for comparisons to other locations cannot be made.

The field of regional science (and the sub-field of spatial econometrics) provides a middle ground between the theories of economics and the empirical observations of geography (Anselin, 1988). Generally, the aim of many regional science-based studies is to apply the same theoretical rigour and quantification found in economics to the complexity of the real world. Often, this is accomplished through the use of regression-based techniques that have been tailored to study the interactions of variables on a local or regional scale. For studies of real-world locations, this method is an improvement over approaches originating in both traditional economics and geography.

Although studies utilizing the regional science/spatial econometrics-type approach have been gaining in popularity recently (Anselin, 2007), in many cases researchers fail to understand the true nature in which space affects the system that they are studying (Huang and Leung, 2002). Compromises between empirical validity and model simplicity often result in unrealistic assumptions regarding the nature of spatial interactions of actors in a system (Fotheringham *et al.*, 2002), which reduce the

usefulness of the model. For example, while it is often assumed that workers generally interact in discrete space by commuting along roadways (Johannsson *et al.*, 2002), most available models only allow this interaction to be modelled in continuous space.

In Newfoundland and Labrador, the fishery workforce has traditionally been concentrated in small communities dispersed along the coastline, between which straight line travel is often impossible. For this reason, particular attention needs to be made to the construction of an accurate representation of the spatial system in which the cod moratorium unfolded before its effects are modelled in this thesis.

1.2 – Research Problem and Hypotheses

The primary question which this thesis will attempt to answer is, "Does proximity to an urban centre have an influence on a local labour market's ability to adapt to changes in economic structure?" Specifically, the focus will be upon the labour market effects of the cod moratorium in rural Newfoundland and Labrador. While some broad demographic and social implications have been studied in past literature (Davis, 2006; Hamilton and Butler, 2001), there is a complete lack of information available concerning the spatial distribution of these effects. Theoretically, knowledge of the spatial distribution of labour market shock sensitivity could be used to assist in the delimitation and classification of labour markets in the province, and help design policies that are sensitive to the needs of local areas.

To effectively determine the answer of the research question and address the overall research problem, a number of hypotheses need to first be defined, and then tested. These hypotheses are concerned primarily with the behavioural aspects of the model, and are essential to evaluating the results. While a more elaborate discussion of

the theoretical model and its implications will be provided in section 3.2, the hypotheses generated by the model include:

(i) Before the cod moratorium in 1992, it is hypothesized that remote rural communities would have generally had a greater proportion of their workforce employed in the cod fishery than similar urban-adjacent rural communities, because of reduced economies of scale and, by extension, fewer opportunities for economic diversification.

(ii) As a result of the different levels of effort required, it is expected that the effects of the moratorium will occur in a systematic way, with actions represented by movements along the supply curve (e.g. declines in hours worked) occurring before supply curve shifts (e.g. outmigration).

(iii) It is hypothesized that post-moratorium declines in labour market utilization (movement along the supply curve) will be more uniform across rural fishing communities than declines in labour market quality (supply curve shifts). Generally, while the universally applied moratorium on cod fishing would initially reduce the work available to fishers and fish processing workers in all regions, increased employment alternatives in urban-adjacent rural areas would allow more of them to remain in their respective local labour markets.

(iv) Concerning the nature of labour market quality declines, it is expected that age-biased outmigration will be more systematic than skill-biased outmigration. While broadly substitutable forces are expected to act on workers of different ages, workers of different skill levels would be subject to different types of forces (e.g. the cod moratorium would more directly affect lower skilled workers, while higher skilled workers would have greater opportunities elsewhere). The expected result is that while local labour markets in the province would generally become older as a result of the cod moratorium, any changes in skill levels would be less predictable.

(v) Due to the non-uniform exposure to the direct force of the cod moratorium, it is expected that the relevance of the model constructed in this thesis to communities will vary according to pre-moratorium fishery involvement. Regardless of spatial location, communities with higher dependence on the cod fishery should, *ceteris paribus*, be more likely to display the behaviour predicted in the preceding hypotheses than communities involved in other industries.

1.3 – Research Approach Outline

To answer the primary question of this thesis and in turn evaluate the different hypotheses, the following steps will be utilized:

(i) A theoretical model of the dynamics in the economy will be constructed, based on commonly accepted economic principles, and empirical observations in other locations.

(ii) The primary elements of the theoretical model will be organized into a series of independent and dependent variables, to allow empirical evaluation of the model within a modelling framework.

(iii) A realistic, locally sensitive and discrete geographically weighted regression (DGWR) model will be utilized to compare variations in variable relationships at the community level, using real-world separation distances.

(iv) The performance of the DGWR model relative to more traditional techniques will be determined using a mixture of standard fit and spatial statistics, and the significance of the empirical results will be evaluated using maps of pseudo-t scores generated by the model.

1.4 – Thesis Organization

This thesis contains six chapters. Following this introduction to the research problem and the objectives of this research, the second chapter will provide an introduction to the background of this topic and the primary research methods. Specifically, past efforts in studying labour market dynamics will be profiled, followed by an overview of research concerning some broad methods used to empirically model labour markets and labour market dynamics. The review of methods will focus upon regression methods in general and geographically weighted regression (GWR) methods in particular. The second chapter will close with a brief introduction to the Newfoundland and Labrador labour force, which forms the context of this thesis. The third chapter will detail the methodology of this thesis, including a discussion of the

theoretical model being utilized and the discrete GWR model used to empirically evaluate it. In addition, the third chapter will contain a discussion of the model's particularities, including the data requirements and statistical limitations. The fourth and fifth chapters will contain the results of the model and some discussions and conclusions of the results, respectively. The cited references will be listed in the bibliography.

2.0 – Literature Review

2.1 – Introduction

Economic health is a topic which is almost universally relevant amongst regions. The relative strength and attractiveness of a region's labour market is often intricately related to the strength and attractiveness of its overall economy (Mathur and Song, 2000). Practically, a great deal of development relies upon the ability of an economy to efficiently match qualified workers with productive jobs. An overabundance or misplacement of either one can necessitate significant shifts, with widespread negative consequences (Bound and Holzer, 2000). On a personal level, the availability of appropriate and meaningful employment is commonly cited as amongst the most important components of personal happiness and well being (Korpi, 1997). As such, researchers in a number of academic disciplines are concerned with studying different aspects of labour markets, and labour market strength. This chapter will be primarily concerned with exploring some of the past efforts in this area, with a particular focus upon methods used to theoretically model and empirically test the functionality and implications of local labour market demand and supply shocks.

The first section in this chapter will explore some of the approaches researchers have used in the past to theoretically represent and study labour markets. Specifically, the history of some of the significant schools of thought within the field will be examined, with a particular focus upon studies that acknowledge the importance of space in regional dynamics. In addition, a framework used in the literature to represent labour market dynamics will also be discussed, along with some of the previous efforts made to study sources of labour market heterogeneity.

The second section of this chapter will provide an introduction to some common methods used in modelling spatial economic phenomena. Although a broad introduction will be provided, studies that utilize regression methods capable of evaluating the relationships between different labour market variables will be specifically focused upon. In addition, the relative merits and drawbacks of the locally focused geographically weighted regression method will be explored, within the context of its applicability to a study of local labour markets.

Finally, the third section of this chapter will focus on literature relating to the 1992 cod moratorium in Newfoundland and Labrador, a labour market shock that is a potential modelling opportunity for this project. This section will provide a general introduction to the structure of the Newfoundland and Labrador labour force, and any previous efforts to systematically study and model the impacts of the moratorium.

2.2 – Labour Market Studies

Generally, the nature and health of people's work habits are regarded as closely related to the state of the economy overall (Mathur and Song, 2000). Thus, much attention has been paid to the health of labour markets, and how it relates to economic development (Jones, 2004; Mathur and Song, 2000). Studies of labour markets are often concerned with such topics as unemployment statistics, participation rates, age structure dynamics, commuting tendencies, and similar information (Jones, 2004). The study of labour markets is a largely multidisciplinary effort, with papers in the past originating from an array of disciplines, including economics (Krugman, 1991), economic geography (Martin, 2000), and regional science (Mathur and Song, 2000). The wide range of

questions posed and methodologies applied in the study of labour markets has resulted in a rich and diverse literature, with often conflicting approaches and interpretations.

The role of space in labour market dynamics has long been a subject of study in both economics and geography, stretching back to at least the late 19th century (Ravenstein, 1885). Historically, the most significant distinction between labour market studies originating within economics and geography has been the treatment of space. Although a number of geography-based studies have embraced basic structures and theories of economics (particularly since the 'quantitative revolution' of the 1960s), their primary focus has often been on investigating the role that space plays in creating and affecting labour markets (Overman, 2004). Conversely, most economics-based studies have adopted frameworks that are distinctively less empirical, and more theoretical, with a focus upon consistency with other types of markets and market behaviour. According to Martin (1999), many economists assume that space plays no role in labour market dynamics, and that actors exist in dimensionless markets, devoid of the effects of distance decay. As a result, economists generally deem the primarily empirical work of many economic geographers to be too idiosyncratic and empirical, without a sufficient theoretical base (Overman, 2004). The approach taken within most regional science/spatial econometrics papers is often located between the mainstream economics and the mainstream geography approaches, where authors attempt to reconcile the theoretical regularities of economics with the seeming idiosyncrasies of the empirical world (Anselin, 1988). These papers generally value both empirical validity and theoretical rigour, and try to reconcile the two by modifying standard economic techniques (such as linear regression) with distance effects (Brunow and Hirte, 2006).

The diverse approaches used to study the labour market phenomena in different fields often are the result of differing priorities between the fields. For example, economic geographers are often more concerned with issues of local sustainability and viability in real-world locations, where tailoring policies for specific situations is a priority (Overman, 2004). Alternatively, economists are often more focused on general goals of understanding relatively global processes and explaining more broadly relevant economic behaviour (Krugman, 1991). One advantage of a regional science approach is that it is able to engage with both local and global phenomena at the same time, by attempting to place local behaviour within a more general context (Fingleton, 2004). By testing the relationship between local behaviour and theoretical constructs, this approach is able to provide locally relevant labour market information, while also providing a mechanism for assumption testing and comparison with other locations and times (Anselin, 1988).

2.2.1 – Demand and Supply Shocks

The first step in evaluating the relative condition of economic phenomena is often to establish a benchmark to evaluate the area or event in question (Jones, 2004). When markets are concerned, this benchmark is represented with the use of supply and demand curves (Nickell and Bell, 1995). These curves allow the change in one, or both, of the supply and demand of a particular commodity to be placed in the context of its impact upon the quantity and price of that commodity. In the case of labour markets, the commodity is labour, with the supply of labour being workers, and the demand being jobs (**Figure 2.1**).

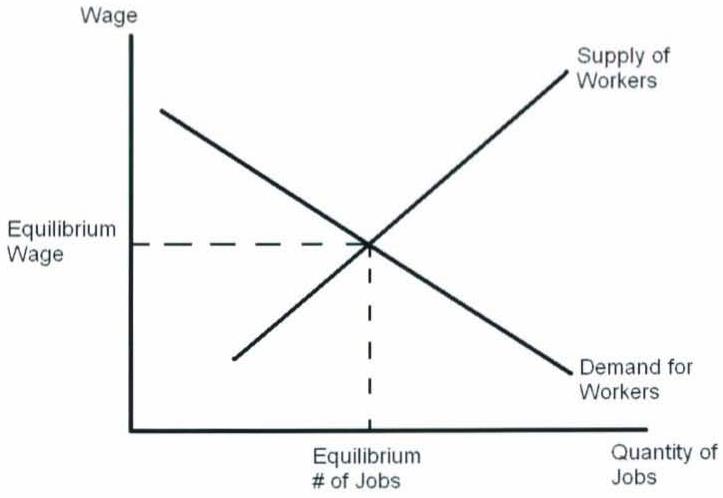

Figure 2.1 Labour Market Supply and Demand

According to this approach, a labour market will more or less function in a similar manner to any other market, with a tendency to be drawn towards a particular 'equilibrium point' where the supply curve and the demand curve meet. Any drastic shifts to either the supply of workers or jobs in a particular labour market will manifest themselves through shifts in the supply or demand curve that creates pressure on the position of the equilibrium point. For example, an unexpected increase in labour supply would theoretically trigger a downward pressure on wage that would increase the demand for labour, thus alleviating the pressure of the supply change, and shifting the system to a new equilibrium (**Figure 2.2**).

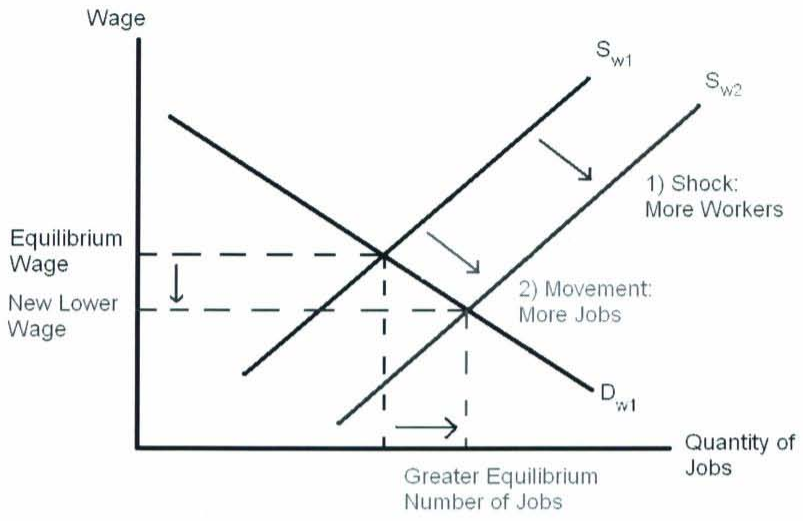

Figure 2.2 Anticipated Effect of a Sudden Influx of Workers to a Stable Economy

This framework for theoretical labour market dynamics allows for regional delineation by enabling the specification of regionally differentiated equilibrium wage/labour levels through region-specific demand and supply levels. Migration and other interaction effects between labour markets will be accounted for in this model by allowing the existence of labour markets with different equilibrium wage levels, between which some limited form of interaction (migration) is possible. This phenomenon will be explained in chapter 3.

2.2.2 – Mobility and Minimum Wage Effects

The incorporation of interaction effects in the study of labour markets allows for the modelling of a number of interesting phenomena not possible with the use of more rigid and simplified models (Krugman, 1991). For instance, it enables the incorporation of the extensive body of mobility literature available in economics and geography, as

well as the investigation of the effects of regionally-varying administrative policies (such as price floors and ceilings) on various labour market figures.

The study of how mobility changes across different demographic and labour market groups is an important activity for researchers and policy makers in many areas (Brunow and Hirte, 2009; Decressin and Fatas, 1995). An understanding of what types of people are the most likely to either enter or leave a workforce is fundamental to the implementation of any labour market policy. With this in mind, there have been efforts made in the past to evaluate differential mobility levels by studying labour markets differentiated according to such criteria as age, skill level, ethnicity, and income level (Goetz, 2001). When combined with an understanding of the different roles that these groups play in an economy, knowledge of mobility proves to be an effective tool for diagnosing labour market deficiencies and recommending the most appropriate course of action.

According to Goetz (2001), the demographic criterion which displays the most universal mobility change is age. Generally, the younger the worker, the more mobile he/she is (Bound and Holzer, 2000). Younger workers are theorized to be relatively more mobile than older workers because they are in a different stage in their lives (Goetz, 2001). It is generally theorized "that people move when they marry and when they have children" (Goetz, 2001). After people begin their careers, their propensity to migrate decreases until their children have grown, and they retire. Upon retirement, migration rates for people in northerly areas of North America rise again, because they move to warmer areas (Goetz, 2001). However, while retirement migration is often a significant

component of overall migration patterns, empirical evidence indicates that in most areas it occurs in relatively lower quantities than young worker migration.

A second demographic grouping that is often tested for differential mobility is education levels. Empirically, there have been times and places in the past where either less educated or more educated workers have been the more mobile group (Bound and Holzer, 2000). However, in the majority of cases, higher educated workers have been found to be the more mobile cohorts (Goetz, 2001). According to Beine *et al.* (2008), workers with relatively more education participate in industries that exhibit greater agglomeration economies than less educated workers. In terms of income, this means that more educated workers have relatively more to gain from migrating than do less educated workers. Of course, these studies assume *ceteris parabus* conditions between education groups, and it is not unusual for lower skilled migration to exceed higher skilled migration, if a differential shock has occurred (Bound and Holzer, 2000).

Aside from evidence that not all demographic groups are equally mobile, research has also shown that not all demographic groups are equally productive. Within the framework of regional economic development, a number of studies have focused upon the relative importance of different demographic and labour market groups to the well being of an economy (Feyrer, 2007; Rupasingha *et al.*, 2002; Bhatta and Lobo, 2000). These studies have found that, in a relatively free market economy, productivity and income levels differ from location to location due to factors such as age, education, and social capital development. The relative balance of these factors in an economy plays a significant role in determining its rate of growth and level of development. For example, a number of studies have found middle-aged workers in the 30-44 (Brunow and Hirte,

2006) or 40-49 cohort (Feyrer, 2007) to be associated with the highest productivity, so, everything else being equal, economies comprised of high percentages of middle-aged workers should be more productive than economies with lower percentages. Also, economies with high rates of well educated workers (Bhatta and Lobo, 2000) and social involvement (Rupasingha *et al.*, 2002) have been found to be relatively more productive than economies with lower educated workers and less social involvement, respectively.

Mobility effects are not the only differential forces that can affect the relationship between separate labour forces. Government-legislated price floors (minimum wages) also play a role in labour market adjustments (Mejean and Patureau, 2010). A minimum wage is the minimum allowable wage level permitted within an administrative region. What this means for labour market theory is that it is possible for markets to be held out of equilibrium for extended periods of time, resulting in market inefficiencies (Mejean and Patureau, 2010). For example, suppose a region were to suddenly experience a large lay-off of workers in an important industry, and the demand curve was to be shifted downwards. At some point, various workers unable to find work would be pressured to exit the labour force, while others would accept lower pay for the same work, and the supply curve would in turn attempt to shift backwards. However, the supply curve has to shift back past the point where the equilibrium wage level is no longer under the minimum wage level (**Figure 2.3**). Practically, this phenomenon acts to reduce flexibility within local labour markets, and to increase the attractiveness and arbitrage potential of other areas.

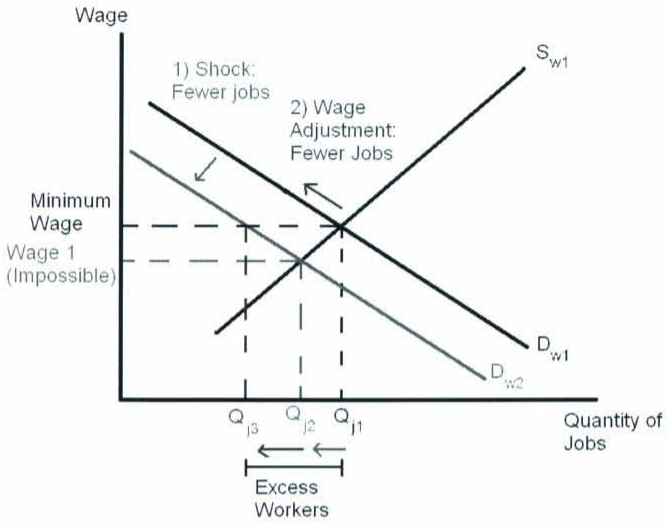

Figure 2.3 Demand Shock with a Minimum Wage

2.3 – Spatial Modelling

The creation and specification of models of economic behaviour is an activity that is at the heart of much economic policy evaluation (Romer, 1994). Without models, it would be difficult for economists to evaluate theoretical predictions concerning phenomena in one time or location, or to relate these phenomena to events occurring at different times, or similar locations. As such, there is a large and varied literature concerning empirical economic modelling available, originating largely in various fields of economics, as well as regional science and economic geography (Anselin and Griffith, 1988). Historically, the majority of studies concerned with quantitative modelling of various aspects of economic behaviour have been concerned with empirically testing theoretical frameworks originating in the economics literature. As a result, the acceptable methodological framework within which models are constructed and results are

interpreted is often derived from mainstream economics, which generally places greater importance upon theoretical consistency than empirical validity (Martin, 1999).

Despite the fact that the economics-based modelling literature appears to address many areas of concern relating to empirical economic modelling quite well, it struggles when faced with the fundamental problem of scale (Lindh and Malmberg, 1999). Although it may be theoretically justifiable and empirically more straightforward to utilize mainstream economics' theories of consumer behaviour on a global (worldwide) scale, some issues may arise in both respects if the study area of concern is somewhat smaller (Rey and Montouri, 1999). From a geographer's perspective, the ability to not only study large scale, macro effects, but to also increase resolution and explore more localized micro effects is desirable for any model. In any study which is focused on a sub-national geographical unit (such as a Canadian province) this ability becomes not only desirable, but essential to the empirical usefulness of the model. In response, researchers, particularly in the fields of economic geography and regional science, have developed a number of different methods of incorporating the effects of space into essentially non-spatial economic models (Anselin, 1988; Fotheringham *et al.*, 2002; Brunow and Hirte, 2006). The ability of any of these methods to successfully remove spatial effects from a model, however, is dependent upon how space affects the phenomena in question. As such, it is important to obtain a clear understanding of the strengths and limitations of potential methods before the model is constructed.

2.3.1 – Broad Methods

Generally, spatial non-stationarity of variable relationships in economic systems is most often approached within a regression framework (Griffith and Paelinck, 2007).

However, depending upon the type of investigation sought, and how much is known about the role of space in a particular economic system, a variety of other techniques are also available (Engelen *et al.*, 1995). Two somewhat common alternatives to regression that are used to explore the spatial structure of economic systems in the literature are hierarchical linear models (Varga, 2000) and complex systems (Engelen *et al.*, 1995).

Applications of complex system-based methods of economic modelling have grown considerably in frequency and variety in the literature since they became commonplace in the 1970's and 80's (Martin and Sunley, 2007). Common types of 'complex methods' include cellular automata and agent-based modelling, both of which have been used at some time to study economic behaviour (Tesfatsion, 2002; Engelen *et al.*, 1995). Generally, studies utilizing these methods tend to focus on simulating economies by creating virtual environments analogous to the areas being studied, and then interpreting the results (Martin and Sunley, 2007). Whereas regression analysis as a technique is generally used to test relationships between variables and make predictions based on those relationships, complex methods are more often used to test how local level decisions interact and propagate outward to form higher level societies and economies. Although complex methods offer an ability to study the processes behind economic phenomena that regression methods cannot match, they often require more data, and their results are more difficult to directly relate to real-world behaviour (Santé *et al.*, 2010).

Hierarchical linear models (HLM) are distinguished in the economic modelling literature from more mainstream regression models mainly through their treatment of the hierarchical aspects of some spatial phenomena (Varga, 2000). Generally, HLMs allow processes that operate on different spatial scales to be represented within the same model,

thus enabling a more accurate empirical representation in some situations. For example, using a HLM structure, labour market phenomena which depend upon interactions through commuting could be separated from phenomena that result from migration patterns. While a HLM approach may be more capable of accurately modelling labour phenomena in a particular location, it is still a regression approach, so all preparation necessary for an ordinary regression is still required (Varga, 2000). HLM therefore would seem to be a more viable option as an attempted improvement to a simpler regression approach, rather than as a primary model.

Although complex systems and hierarchical linear models offer intriguing possibilities for modelling the spatial aspects of economic systems, their acceptance and utilization in the literature is not as widespread as more common regression models. While complex systems and hierarchical linear models offer a degree of spatial flexibility that cannot easily be approached with mainstream regression methods, in contrast, spatial regression models are fairly straightforward and intuitive, and are more commonly utilized in existing regional science and spatial econometric literature (Anselin, 1988). Two spatial regression methods that will be explored in greater depth in the next section are spatial regression models in the regional science subfield of spatial econometrics and localized regression methods developed in geography.

2.3.2 – Regression Analysis

Perhaps the most commonly accepted method of explicitly accounting for space in the economic modelling literature is through the medium of spatial econometrics. As a sub-field of regional science, spatial econometrics seeks to, "...deal with the particularities caused by space in the statistical analysis of regional science models..."

(Anselin, 1988). These particularities are classed into one of two forms of 'spatial effect': spatial dependence or spatial heterogeneity. Spatial dependence, or spatial autocorrelation, occurs when the value of a variable at one location depends in some way on the values of that variable at other locations (Cliff and Ord, 1970). Spatial heterogeneity, on the other hand, occurs when the trend of the values of some variable varies inconsistently over space. These spatial effects can have profound impacts on the empirical veracity of any economic model, unless they are accounted for (Rey and Montouri, 1999). A number of methods have been developed that address the issue of spatial effects in regression models, dealing most often with the presence of spatial dependence (Brunow and Hirte, 2006). This spatial dependence is usually representative of the degree of interaction that takes place between agents (people, communities, labour markets, etc.) in a spatial economy (Rey and Montouri, 1999). Methods used to account for this interaction often function by adapting typical regression models to explicitly account for the spatial dependence (Anselin, 1988). A list of some of the models resulting from efforts in this area includes relatively popular spatial lag and spatial error formulations, and other, less well known types, such as spatial regressive models (Brunow and Hirte, 2006). While broadly similar, these models differ depending on how the spatial dependence is theorized to enter the system. For example, in a spatial lag model (**Equation 2.1**), the dependence enters the system via interactions and spillovers between actors in a system, and as such is represented, to some degree, by spatially lagging the dependent variable:

$$y = \rho W y + \varepsilon \qquad \textbf{(2.1)}$$

Where ρ = an autoregressive parameter, W = a spatial weights matrix, and ε = the error. In a spatial error model (**Equation 2.2**), the spatial dependence is assumed to occur via random regional shocks across the study area, and is accounted for in the error term:

$$\varepsilon = \lambda W\varepsilon + \varepsilon \qquad (\mathbf{2.2})$$

Where λ = the spatial autoregressive error parameter.

A second way in which researchers have previously dealt with space has been to utilize a technique known as geographically weighted regression (GWR) (Fotheringham *et al.*, 2002). What separates GWR from other methods used to deal with the effects of space is that its coefficients are allowed to vary locally. Effectively, this ability to vary the coefficients localizes the analysis by removing the assumption held in other techniques that global processes are at work in the system. This is a clear advantage in situations where such an assumption is difficult to make or defend, such as when modelling an economy that is composed of several different labour markets. In the literature, GWR is utilized in modelling a number of socio-economic phenomena in space, including such things as house prices in the greater London area (Fotheringham *et al.*, 2002), illness in the UK (Brunsdon *et al.*, 1998), and regional industrialisation in China (Huang and Leung, 2002). Generally, GWR differs from the spatial econometric methods of dealing with space in that GWR does not explicitly include a term to represent the global effect of space on the dependent variable (**Equation 2.3**). Instead, as seen in **Equation 2.4**, the effect of space is included implicitly in a GWR model, by

weighting the influence that specific cases of the independent variables have on the dependent variable, according to spatial proximity.

$$y_i = \beta_0(u_i, v_i) + \sum_k \beta_k(u_i, v_i)x_{ik} + \varepsilon_i \qquad (2.3)$$

Where u_i, v_i = the spatial coordinates of the i^{th} point, and x_{ik} = the value of independent variable k at that i^{th} point. These terms are calibrated via the expression:

$$\widehat{\boldsymbol{\beta}}(u_i, v_i) = (\boldsymbol{X}^T\boldsymbol{W}(u_i, v_i)\boldsymbol{X})^{-1}\boldsymbol{X}^T\boldsymbol{W}(u_i, v_i)\boldsymbol{y} \qquad (2.4)$$

Where $\boldsymbol{W}$ = the spatial weights matrix, $\boldsymbol{X}$ = the matrix of independent variables, and $\boldsymbol{y}$ = the dependent variable (bold type indicates a matrix).

The specific method used to derive the regression point weights varies depending on the nature of the study being considered, and could entail the use of either a fixed or adaptive kernel (Fotheringham *et al.*, 2002). A fixed kernel utilizes a fixed neighbourhood threshold distance, while an adaptive kernel uses a threshold distance that varies according to the number of neighbours a particular case has.

2.4 – Newfoundland and Labrador Labour Force

2.4.1 – History and Context

The cod fishery has long held a central role in the Newfoundland and Labrador economy. A resident fishery was in place since at least the late 1700's, and persisted as the economic mainstay of many rural communities for two hundred years (Perry and Ommer, 2003). With the technological advances of the mid-twentieth century, however, fishing productivity increased, and by the mid-1960's landings were reaching all time highs (Sinclair, 1996). Despite the high catches and increased income reported by offshore trawlers in the late 20th century, the inshore fishery remained a seasonal industry

and incomes remained low. To increase the attractiveness of the fishery as a source of employment for rural workers, the federal government began to supplement market income for fishery workers by reducing the employment insurance qualification requirements (Hamilton and Butler, 2001). In effect, this action reduced the efficiency of the fishery, and led to a greater capacity than likely would have existed in free market conditions.

Although the exact cause is disputed, some combination of these increased fish landings, climate change, and habitat damage (Perry and Ommer, 2003) ultimately proved unsustainable, and the fishery collapsed in the early 1990's. The Federal Government instituted a moratorium on fishing starting on July 2, 1992 (Murray *et al.*, 2005), and as many as 30,000 people were immediately thrown out of work (Davis, 2006). Many of the affected communities depended almost entirely upon the cod fishery as an economic base, with heavy involvement in both the harvesting and processing of fish often occurring in the same community. As such, they were unprepared for the surge in unemployment that accompanied the closure of the fishery (Perry and Ommer, 2003). As seen in **Figure 2.4**, the effects of the cod moratorium were not equally felt in all areas of the province. Through a combination of resource proximity, government initiatives, and lack of alternative economic opportunities, most of the impacted communities (as measured by government aid, per capita) were located on the northeast coast of Newfoundland. For many of these communities, the period following the end of government support programs in the late 1990's were characterized by large amounts of outmigration, as unemployed former fishers and fish processing workers left to find employment elsewhere (Davis, 2006).

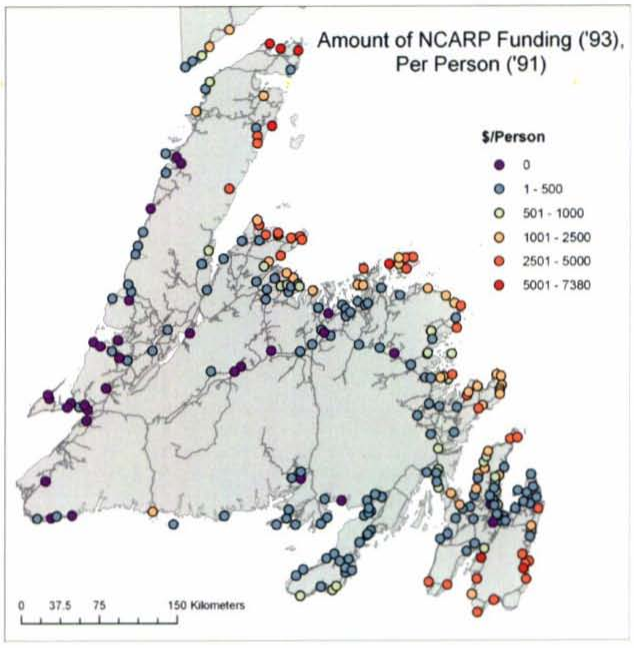

Figure 2.4 Northern Cod Adjustment and Recovery Program (NCARP) Funding per Capita, 1993/1991 (Government of Newfoundland and Labrador, 2010)

2.4.2 – The Cod Moratorium in the Literature

To understand why many rural communities were not prepared for the cod moratorium, one must first understand that the economy of Newfoundland and Labrador as a whole has long been highly dependent upon resource extraction industries (and the fishing industry in particular) as a source of output and employment (Sinclair, 1996). This situation was most profound in coastal rural areas of the province, where incomes have historically been lower and unemployment rates higher than the more diversified urban areas of St. John's and Corner Brook (Schrank, 2005). While these urban centres maintained a measure of economic diversity due to various types of manufacturing and government services, for many rural areas there was only the fishery (Schrank, 2005). In

addition, while traditionally a number of alternative species have been caught for subsistence in rural areas, the cod fishery alone represented almost the entire commercial enterprise of these communities (Perry and Ommer, 2003).

The high reliance of rural areas on the fishery exposed them to the full force of the cod moratorium of the early 1990's, and the resulting negative adjustments (Davis, 2006). Although the effects of the cod moratorium affected all areas of the province to some degree, the impact on young people in rural labour markets has been especially pronounced (**Figure 2.5**). According to Davis (2006), upwards of a third of the population of young people in rural Newfoundland and Labrador migrated in some form between 1991 and 1996. Ten years later, the net effect of years of such migration is that many rural labour markets now have an abundance of dependent, mainly retired members of society, and a deficit of working-age members of the workforce (Murray *et al.*, 2005). As discussed in section 2.2.2, this differential effect is expected, because of the simultaneous increased mobility and decreased job security of young people.

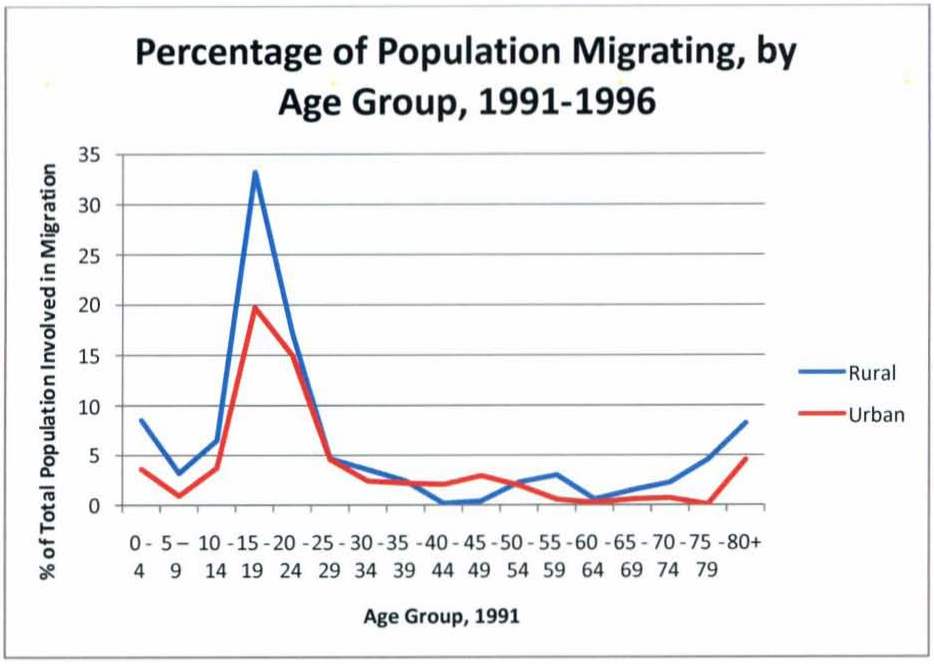

Figure 2.5 Rural vs. Urban Migration Prevalence, 1991-1996 (Government of Newfoundland and Labrador, 2009)

To reduce or delay the exposure of rural areas to the full force of the cod moratorium, the federal government instituted aid programs in the years following 1992 (Hamilton and Butler, 2001). Immediately following the announcement of the cod moratorium, the government created the Northern Cod Adjustment and Recovery Program (NCARP). NCARP was created to provide affected workers with income to substitute the loss from the fishery collapse (Murray *et al.*, 2005), which would reduce the incentive for the unemployed to migrate out of their communities in search of work. Initially, like the moratorium itself, NCARP was intended to be temporary. Once the fish stocks had regained their strength and fishing could resume, NCARP would end. However, when it became apparent that the fish stocks were not returning to their

expected levels, the government was forced to reconfigure the purpose of the aid they were providing. To this end, The Atlantic Groundfish Strategy (TAGS) was launched in 1994 to replace NCARP. The initial goal of TAGS was to reduce the capacity of the fishing industry by funding licence buyouts and continuing education for fishery workers. However, for many people, TAGS continued the tradition of income substitution (and migration deterrence) started by NCARP, and the degree to which TAGS achieved its initial goals is debatable. Regardless, with the ending of TAGS in 1998, the government's direct efforts to reduce the effects of the moratorium ended, as well. After this date, people in the fishing industry were once again required to largely rely upon the traditional avenue of employment insurance for supplemental income.

While it is not difficult to find literature relating to the descriptive strengths and weaknesses of the Newfoundland and Labrador economy, it is considerably harder to find literature based on more in-depth studies or modelling efforts. In fact, despite a search for such studies, not one could be located that empirically investigated the nature of relationships between various types of labour market reactions, in the spatial context of the Newfoundland and Labrador economy. While this is unfortunate from the literature search perspective, it does indicate a hole in the existing literature to be filled with this research.

2.5 – Conclusion

The systematic study of labour market dynamics is an activity engaged with a number of different approaches and from a variety of perspectives. Depending upon the characteristics of the study area, the information being sought, and the preferences of the researcher, the frameworks and assumptions considered reasonable can vary widely

(Overman, 2004). Despite being fragmented across a number of different fields, the literature concerning the theoretical operation of regional labour market shocks and their implications is relatively well developed. The use of supply and demand function shifts to represent labour market dynamics is fairly well established (Nickell and Bell, 1995), and evidence of heterogeneous productivity and mobility amongst different demographic groups has been found for a number of different regions (Bhatta and Lobo, 2000; Goetz, 1999). In addition, the literature concerning different types of modelling frameworks capable of studying labour market behaviour is also well developed. Over time, the regression framework has become the most popular and commonly accepted method of evaluating relationships between variables in economic research. Recently, locally-derived methods such as geographically weighted regression (GWR) have been used to study relationships and relationship stability on a smaller scale than is possible with mainstream regression methods. However, due to their basis in continuous space, GWR models have to this point had limited applicability to phenomena based on human interactions, which largely operate in discrete space. Therefore, an opportunity exists to modify the existing GWR structure to facilitate discrete space interactions, and evaluate any performance improvements in an empirical study. As a potential study area, Newfoundland and Labrador is an excellent candidate, due to the combination of the spatially distinct local labour markets that result from the relative isolation of most of the fishery settlements, and the labour demand shock it experienced in 1992 with the closure of the cod fishery.

3.0 – Methodology

3.1 – Introduction

A primary goal in the study of regional economic development is to understand the processes which underlie economic phenomena, so that appropriate policy for situations involving these phenomena can be determined. While theoretical developments consistent with idealized behaviour are important for relating events in one location or time with events in another, knowledge of the degree to which this simplified behaviour represents real-world locations is essential to developing strategies for actual locations. Most of the methods currently used for this purpose fail to adequately capture the true, discrete nature of labour market interaction behaviour. For this reason, this thesis will focus upon developing a model capable of empirically testing local labour market effects of the 1992 cod moratorium in Newfoundland and Labrador, which reflects workers' commuting behaviour. Aside from gaining an accurate account of the extent of structural changes that followed the moratorium, the model described in this chapter will allow a theoretical framework of the cod moratorium to be evaluated. The following sections will detail the rationale supporting the local model, as well as information regarding its construction, interpretation, strengths and limitations.

Section 3.2 will introduce a theoretical model that has been constructed for the behaviour of the cod moratorium in Newfoundland and Labrador. Specifically, the moratorium will be interpreted as a sharp retraction of demand for fisheries workers, and any effects resulting from it will be interpreted through shifts in the labour supply and demand functions that occur when they adjust. This structure will allow space to play a

role in the functioning of the economy because it will allow exposure to the effects of the moratorium to differ between communities.

Section 3.3 will explore the local empirical model that will be used to explore the actual effects of the cod moratorium, and the relevance of the theoretical model. To capture local effects, a modified version of the geographically weighted regression model (GWR) will be used. In most applications of the GWR model, it is assumed that any interaction effects between actors in a system occur along straight lines (Fotheringham *et al.*, 2002). However, the long distances involved and lack of public transportation infrastructure means that workers who commute between communities in Newfoundland and Labrador probably most often do so along discrete road networks. For this thesis, the standard GWR procedure will be modified to accommodate discrete distances between communities, which will allow a more realistic analysis that is better equipped to comment on the usefulness of the theoretical model constructed in the first section.

Finally, section 3.4 will introduce and explain some of the more technical aspects of this thesis. The empirical data required to accurately represent the various aspects of the theoretical model will be outlined, along with the technical steps required to modify the GWR model so that it can operate in discrete space. In addition, some statistics used in novel ways which are important for the interpretation of different aspects of the model will be discussed.

3.2 – Theoretical Adaptation

To gain a sufficient understanding of the significance of any empirical events that occur in an economy, it is often necessary to relate the events to an established theoretical framework. The fact that economies in diverse locations can share broadly similar

institutional arrangements and market structures indicates that lessons learned or theorized in one location can often be transferred and applied (at least in part) to another location. This is a common practice particularly in studies relating to western-style 'market' economies, which are regulated by a mixture of free market and government controls (Pita *et al.*, 2010; Rupasingha *et al.*, 2002). Generally, people in these types of economies have broadly similar ambitions and goals, which are often measured with quantitative measures such as income and output per worker (Romer, 1994), as well as less tangible measures such as quality of life (Hudson, 2007). The Newfoundland and Labrador economy, being one of these mixed market economies, can therefore be compared to other economies, in other locations. By extension, because the Newfoundland economy features relatively mobile factors of production (being a Canadian province, workers are free to move to and from other provinces), it is well suited for the application of various theoretical economic frameworks. In particular, this analysis will examine the applicability of a supply and demand curve framework to a demand shock in the Newfoundland and Labrador labour force, as well as the potential implications that would result from such a shock, according to this framework. It should be noted that the primary purpose of this formal modelling exercise is to provide context for the results of the empirical model presented in section 4.

3.2.1 – Demand Shock

As introduced in section 2.1, a common method of representing shocks to a labour market is through the use of supply and demand curves. The shock is represented in this framework as a force which pressures the supply or demand curve to shift, and the impacts of the shock are represented as the adjustments that occur when the system is

pressured to return to an equilibrium state (see **Figure 2.2**). These adjustments can take the form of a change in the equilibrium wage, a response shift in the curve not under direct pressure from the shock, and/or an efficiency/elasticity related response, where the pressure from the shock is at least partially resisted by the targeted curve through the elasticity of the curve itself. The equilibrium wage or number of workers would change when one of the curves is somewhat resistant to change, and movement along that curve happens instead. For example, in the event of a large local lay-off (demand shock), should some workers refuse to exit the local labour market, the supply curve would not shift fully in response. Instead of a shift, this situation would result in some workers moving down the supply curve and accepting lower wages to stay in the labour market, which would cause the equilibrium wage level to decrease. If the curve not under direct pressure from the shock shifts fully in response, then the equilibrium wage level would not decrease. In many real-world examples, however, the labour market does not behave this cleanly. For example, when under pressure to reduce the number of workers, a business or government may decide to allow workers to work fewer hours, or convert to seasonal work in order to collect employment insurance benefits (Hamilton and Butler, 2001). This would reduce the size of the demand curve shift, but also reduce the efficiency of the economy at the same time.

In Newfoundland and Labrador, a method of representing and testing the effects of a labour market shock is perhaps more important than in other Canadian provinces. The cod moratorium which began in 1992 represents a significant demand shock to the province's labour force because it removed thousands of jobs from the economy. Not all workers in the province were equally exposed to the moratorium, however, because its

direct impact was limited to the cod fishery, which was primarily rural and staffed with low skilled workers (Hamilton and Butler, 2001). The limited scope of the cod moratorium means that it is the ideal demand shock to be represented with supply and demand curve behaviour. Any resulting effects of the demand shock will be assumed to have occurred within a fairly simplified framework, because of the relatively homogeneous nature of the affected labour force. The lack of diversity in education and age requirements, in addition to the general lack of extensive organizational hierarchy means that most of the workers and jobs in the industry will be represented with a single supply and demand curve, respectively. Also, the lack of stringent entrance requirements to the cod fishery and employment diversity in the rural labour markets would have limited the options of many of the affected workers after the moratorium, thereby restricting the range of potential effects, and simplifying the model necessary to evaluate them.

Another empirical reality which restricts the potential effects of the cod moratorium is the provincial minimum wage. The minimum wage is the minimum legal hourly wage that employers can pay workers in the province, regardless of location (Baker *et al.*, 1999). Practically, the minimum wage acts as a price floor, below which labour cannot be acquired, regardless of the economic conditions. This restriction may be important for an analysis of the impacts of the cod moratorium, because any desire of firms to provide work for less than the minimum wage would not be tolerated without government subsidies and transfer payments. Therefore, if the demand shock was severe enough to cause downward pressure on the equilibrium wage level, and that wage level was at or near the minimum level initially, an additional downward shift in labour supply

would be necessitated, barring any government intervention. It will be assumed that the equilibrium wage level may have been at risk of crossing the minimum wage threshold after the moratorium, and these forces may have been active, because of the low income levels historically associated with cod fishing and the lack of substitutable work in fishing communities (Feehan, 1991).

The role of government intervention in the supply and demand curve framework is essential to the empirical specification of the model. Upon closing the fishery, the federal government implemented two distinct but related programs specifically targeting the east coast fishing industry. The 'Northern Cod Adjustment and Recovery Program' (NCARP) was in place from 1992 through 1994, and was followed by 'The Atlantic Groundfish Strategy' (TAGS), which existed from 1994 through 1998. By design, NCARP functioned as a mechanism for income substitution, to keep the affected workers in their respective labour markets until the moratorium was over, and they could return to work. Once this was ruled out as a possibility, the government implemented TAGS, which was designed as a program to assist the affected workers in getting out of the fishing industry, through retraining and licence buybacks. While not as direct as NCARP, for some people TAGS functioned as an income substitution program, as well (Sinclair, 1996). Practically, these income substitution programs enabled workers to earn less than a subsistence wage from market sources, but still receive enough income to make staying in their community a possibility. During the time these programs were in place, the labour market for these workers would have been held out of equilibrium, with a surplus of supply. In essence, these programs increased the reservation wages of the workers, as the wage required to entice them to move and acquire a job was inflated

while they were in place. Once the income substitution programs ended, the market would have exerted downward pressure on the labour supply. Therefore, one could reasonably expect any resulting negative change in the supply of labour in the affected communities to not necessarily have occurred immediately after the demand shock, but rather to have been delayed and drawn out over a period of five or ten years. Any impacts relating to the utilization or efficiency of the workforce would not be delayed, however, because such decreases were practically encouraged under NCARP and the national employment insurance program (Schrank, 2005).

3.2.2 – Spatial Structure

The reality in which labour markets operate is that geography matters. The willingness of workers to commute to a particular job or labour centre will generally decrease as the distance required to commute to it increases (Karlsson and Olsson, 2006). The distance a worker is willing to travel varies based on variables such as pay, ease of travel, the availability of substitutable intervening employment opportunities, and personal preferences (Karlsson and Olsson, 2006). Complicating this further is the fact that the types of jobs available for workers vary spatially, because of regional differences in both industrial make-up and market and supplier accessibility. These types of forces reduce to some degree the ability of industries to operate freely, and for labour markets to compete equally. The result is that sources of real-world labour supply and demand tend not to be uniformly distributed, which means that the spatial distribution of economic opportunities and risks tends to vary as well. When factored together, these forces of agglomeration and dispersion often lead to labour markets taking on hierarchical structures, with labour demand being focused primarily upon one or a couple of

communities in a labour market, and the labour supply radiating outwards to the commuting limit of most workers (Coffey and Shearmur, 2002). Neighbouring labour markets will often differ somewhat with respect to indicators such as economic structure, income, and culture because of the commuting distance limits felt by most workers, and their tendency to select the nearest available job (everything else being equal) (Karlsson and Olsson, 2006).

As a result of the heterogeneous functions provided by different services and industries (Wensley and Stabler, 1998), hierarchical rankings and structures can frequently be found on the level of labour markets themselves (and their centres). Often, high level services benefit from agglomeration effects to a greater degree than lower level services, so there is a benefit to centralizing the service (Coffey and Shearmur, 2002). For example, university research benefits to a greater degree from daily contact between faculty members than does primary education. The presence of both Marshall-type 'localization externalities', (where agglomeration benefits are generated through same-industry agglomerations) (Henderson, 1997) and Jacobs-type 'urbanization externalities' (where agglomeration benefits are generated simply through the presence of a large and diverse local market) within labour market systems often results in a given system containing one or a few primary centres, a few more secondary centres, a few more tertiary centres and so on (Mulligan, 1984). The presence of this hierarchy means that the effects of any large scale economic perturbations will often be different in one part of a region versus another. Factors used to measure shock preparedness for a local labour market, such as size, diversity, and utilization can often be traced in part to the market's

relative spatial location, and its proximity to higher order centres and their agglomeration effects.

The spatial distribution of the labour markets of Newfoundland and Labrador has always played an important role in determining local labour market structure throughout the province. Overall, the population of the province in 1991 was fairly urbanized, with 304,450 (representing 53.5% of the total population) of people living in urban areas containing more than 1000 people (Statistics Canada (n.d.a)). Of these urban areas, St. John's was by far the largest in the province, because it contained the highest population, the largest workforce, and a concentration of retail and public services unmatched by any other centre (Hamilton and Butler, 2001). On the eve of the cod moratorium in 1991, the city of St. John's, with 95,770 people (not including its metropolitan area) contained more than four times the number of people of the next largest regional centre, Corner Brook, which had 22,410 people (Statistics Canada (n.d.b)).

In addition to population, the rural areas of the province also had economic structures broadly different from those found in urban and urban-adjacent areas of the province prior to the moratorium, due to the agglomeration forces introduced above. As a result of the clustering of government and retail services in the province's large urban centres, the rural areas of the province were relatively dependent upon primary resource extraction for employment. While the labour markets of urban areas such as St. John's, Corner Brook, Gander, and Grand Falls-Windsor involved a mixture of services, manufacturing and some resource extraction activities, rural labour markets often depended to a greater degree upon primary activities for employment (Hamilton and Butler, 2001). For dozens of isolated small communities along the northeast coast of the

island, the cod fishery was the primary economic base for many years due to its labour-intensive nature, lack of significant agglomeration benefits, and dispersed resource distribution (see **Figure 2.4**). Although the dispersed nature of the small communities was ideal for harvesting cod when the cod were plentiful, it left the communities overly exposed and under prepared for the shock that occurred when the cod moratorium was imposed.

One implication of the spatial distribution of communities for the aftermath of the cod moratorium is that those communities located within commuting distance of an urban centre would have experienced less of an impact, for a shorter time period than those communities beyond commuting distance. Generally, affected fishers and fish processing workers located in urban-adjacent communities would have had greater access to substitutable employment than those located in more remote communities. It should be noted, however, that there is no reason to assume that all affected workers in urban-adjacent labour markets would have been able to find substitutable employment, but rather that it would be relatively easier for them, on an aggregate level, than for those in more remote environments.

A second way in which the spatial system of the province likely affected the impacts of the cod moratorium is a result of the distribution of the resource. Instead of being uniformly dispersed throughout Newfoundland and Labrador, the communities most heavily involved in the fishery before the moratorium were concentrated along the northeast coast of the island (see **Figure 2.4**). The primary implication of this non-uniform distribution of vulnerable fishing communities is that not all communities in the province will conform to the theoretical model to the same degree. Relevance is expected

to be highest along the northeast coast where cod fishing was the most viable, and lower in regions where viability lagged, such as the western/interior portions of the island, and in Labrador.

Beyond the direct moratorium impacts experienced by a community, spatial location also potentially influenced the distribution of some of the more structural forces underlying their economic well-being. One of the primary incentives for arbitrage between labour markets is through wage differentials (Gabriel *et al.*, 1993). To expand on **Figure 2.2**, the wage a worker can obtain by relocating to a different labour market regardless of conditions in their own labour market is ever present, and possibly different from their current wage. In general, the more positive the difference between external wage and local wage, the more the worker is pressured to migrate to the other labour market. This pressure is demonstrated in **Figure 3.1**. The pressure to migrate resulting from differential wage rates is another force expected to vary with urban proximity. The expected equilibrium wage in urban-adjacent labour markets should not have decreased to the same relative degree as the expected equilibrium wage in smaller and more remote labour markets because of the higher degree of diversity and competition found in the larger urban labour markets. As a result, the pressure to migrate out of their labour market would have been relatively higher for workers living in remote areas, and, everything else being equal, the supply shift would have been greater. This means that while initial utilization declines immediately following the cod moratorium would have been almost universally severe in both remote and urban-adjacent fishing communities, the more permanent quality declines associated with the outmigration of workers would have been more significant for remote communities.

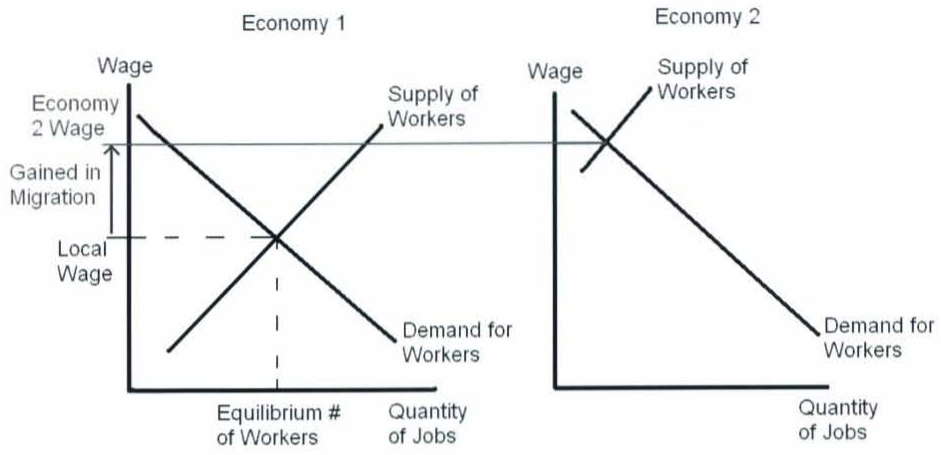

Figure 3.1 Wage Pressure to Migrate from One Labour Market to Another

3.3 – Geographically Weighted Regression Model

The analysis of different aspects of real-world labour market behaviour is often difficult to structure. Model selection is a particularly important and sensitive process because of the convergence of theoretical expectations, empirical realities, and technological limitations. The literature that explores the study of labour markets contains a variety of approaches, both quantitative and qualitative, each with a different set of advantages and disadvantages. For this reason, a necessary step prior to attempting a labour market study is to evaluate the nature of the study area and available data, as well as the demands of the research question. This way, it is assured that any desired information is obtained in the most efficient way possible, by maximizing the interpretable results for a given level of cost and effort.

The primary concern with this research project is with the determination of relative effects of the cod moratorium. From a modelling perspective, this analysis would

ideally involve testing relationships between different variables, in different locations, over time. Perhaps the most common technique used in the literature for the evaluation of relationships is regression. As introduced in section 2.3, regression analysis involves the combination of one or more 'independent' variables to create or predict a separate 'dependent' variable. To evaluate stability over time, separate regressions are completed for different time periods, creating a different set of relationship data for each period, which are then compared. Evaluating temporal change is relatively simple, due to the singular direction and constant speed in which time travels. The multiple directions and rates associated with spatial change, however, present more difficult challenges. Depending upon the nature of spatial interaction in a system, the regional science/spatial analysis literature offers a number of potential solutions that allow the incorporation of space into a regression framework, as explored in section 2.3.

In this thesis, it is theorized that differential labour market behaviour in Newfoundland and Labrador causes the nature of spatial interactions to vary from community to community. Therefore, the best method to study these interactions is one able to accommodate locally-varying relationships. Geographically weighted regression (GWR) is one such method that meets the criteria of this thesis by combining the ability to evaluate both temporal and spatial stability simultaneously at a community level. For this reason, GWR was selected as the primary method of analysis for this thesis. Modifications to the standard GWR framework required for this analysis will be explored in section 3.3.1, and an explanation of how the model will be accommodated within the existing theoretical structure, along with a discussion of specific hypothesis generation, will be provided in section 3.3.2.

3.3.1 – Spatial Structure Adaptation

The primary advantage that the geographically weighted regression (GWR) framework holds over more conventional regression frameworks is that it allows relationships between variables to vary over space. GWR is able to evaluate variable relationships through the use of a 'moving' regression kernel, which is a local regression window that evaluates variable relationships within a limited spatial neighbourhood. This local window moves throughout the entire study area, testing local sets of points until all points are accounted for, and the entire study area has been covered. Upon completion, the local parameter estimates are combined to create a continuous map of variable relationships, which indicates where a particular relationship is positive and where it is negative, as well as where the magnitudes differ. In its unmodified state, however, the ability of the GWR framework to sufficiently account for the real-world spatial structures and interactions existing in the moratorium-era Newfoundland and Labrador economy is questionable. The general structure of GWR is vague and not committed to any particular incarnation of spatial structure (Fotheringham *et al.*, 2002). Despite this, most studies situate their analysis in continuous space, where spatial interaction is allowed to occur unimpeded in all directions. Although an assumption of continuous space is only reasonable in areas of complete isotropic connectivity, it is assumed in all available previous applications of GWR, which range from investigating regional industrialization (Huang and Leung, 2002) to estimating net primary production in a forest ecosystem (Wang *et al.*, 2005). This precedent of utilizing continuous space for interaction represents a challenge for this thesis, which is concerned with labour market interactions that are constrained to a discrete road network. This is a problem

because the distance between communities in this province along the road network is often very different from the straight line distance, due to the irregular coastline and large distances involved. For example, **Figure 3.2** illustrates the large difference between the straight line and road distance between Bonavista and St. John's. At 124 km, the straight line distance in this instance is less than half the distance people actually travel between the two communities, which is 293 km. For the model, this situation has two potential outcomes if the GWR framework is utilized. The first option would be to assume the status quo, and complete the analysis using the standard continuous GWR. The second option is to modify the GWR framework so that it is able to accommodate discrete distances, which would allow the analysis to be more theoretically acceptable. To utilize either option, a trade-off would be required. Whereas the first option would undoubtedly be simpler and easier to apply, the second option would require fewer assumptions and offer easily interpreted results. It should be noted, however, that the empirical validity of the second method at this point is unproven, because the nature of spatial interaction can only be theorized prior to analysis. For this reason, both methods will used in this project. A method of incorporating discrete space into a GWR model will be developed and tested against the pre-existing continuous model for accuracy and efficiency. Whichever model performs better will be the one used for interpretation. A non-spatial ordinary least squares linear regression will also be performed and evaluated, to ensure a fixed point of reference for both models.

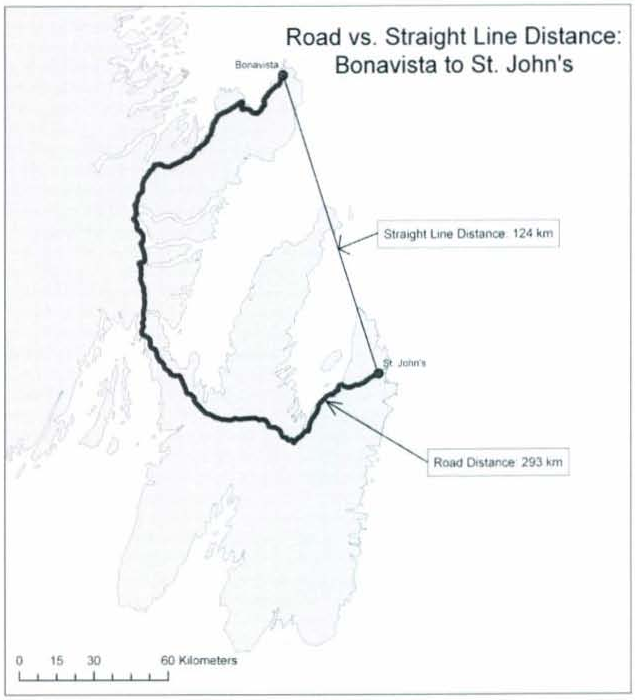

Figure 3.2 Straight Line versus Road Network Distance Example, Bonavista to St. John's

Presently, there are two widely available pieces of software that are capable of implementing GWR with a minimum of background or effort. These programs are the Environmental Systems Research Institute's (Esri's) GWR tool for its ArcMap geographic information system (GIS) software, and a standalone program known as 'GWR' (Fotheringham *et al.*, 2002). Although these programs offer modellers a variety of customization options and statistics to use when implementing a GWR, they unfortunately only allow the use of continuous spatial weight functions. The software used for this project should be able to utilize both discrete and continuous spatial weight functions, so for this reason neither of the existing options is sufficient. To perform a

satisfactory analysis, it is necessary to look beyond pre-packaged software tools to programs that offer greater flexibility and a more generalized approach.

Despite the fact that the two primary software programs used to construct a GWR are based in a GIS environment, GWR is not a strictly geographical technique. Instead, like all multiple regression techniques, GWR is primarily based upon linear algebra (Fotheringham *et al.*, 2002). The only input for space in the GWR algorithm is through the use of a spatial weights matrix, which is used to represent the degree of interaction between different locations in the model (Brunsdon *et al.*, 1998). Generally, this spatial weights matrix is an n x n matrix, where n is the number of events being studied (in this case communities), and the elements represent the degree or probability of interaction between two separate communities. As such, the spatial weights matrix (and by extension the entire GWR) will be created within the framework of standard mathematics software with linear algebra functionality.

One mathematics software package that allows a number of benefits not available with more traditional, purpose-built software is MATLAB. For this analysis, the most important benefit of using MATLAB is the ability to exert total control over the choice of spatial weight function for the model. In MATLAB, distance is treated in the same manner as any other parameter, unlike in GIS-based software, where it is often calculated internally in the program. This means that while GIS-based analyses are restricted to basing their spatial weight function on the distance model included in the software, studies based in MATLAB are able to customize their weight function to fit their needs. A second area in which MATLAB allows a degree of flexibility over pre-packaged software is calculable statistics. For example, although the program 'GWR' allows one

to test for global spatial autocorrelation within the residuals of a GWR model, it does not allow the calculation of local spatial autocorrelation statistics (e.g. local Moran's I (Anselin, 1995)), as well as any tests for spatial heteroskedasticity (e.g. the White test (White, 1980)). The open-ended framework of MATLAB, however, allows one to program tests for whichever phenomena one wishes, provided the equation can be found and programmed. A third advantage to using MATLAB for this analysis is that, as a result of MATLAB's transparent interface, one can follow the equations directly, and determine where unusual or unexpected results originate. Although it is only strictly required for the discrete GWR model, all three models evaluated in this thesis (the continuous GWR, the discrete GWR, and the OLS model) are coded in MATLAB, both for the reasons cited above and to ensure consistency in model structure and results.

3.3.2 – Application of Theory

A number of aspects need to first be clarified before the theoretical model of the Newfoundland and Labrador labour force constructed in section 3.2 can fit within a geographically weighted regression (GWR) framework. While the implications of the cod moratorium in various regions of the province may be relatively simple to anticipate within the simplified structure of a theoretical model, they are not so clear empirically. The various simplifying assumptions implicit in the theoretical construction of the model tend to remove the applicability of the model results from reality to some degree. The extent to which the results are or are not applicable to the real world depends upon how reasonable the assumptions are, and the degree to which they apply to actual behaviour. The only way to be sure whether an assumption is reasonable, and if model interpretations can be applied or not, is through the use of an empirical model and

hypothesis test of some kind. Prior to performing the empirical analysis for this thesis, the spatial relationships need to be defined, and the supply and demand curves, the shocks, and the various shifts need to be incorporated into the GWR framework. The fact that GWR is a regression model means that these parameters will be represented in the model through the interplay of a dependent variable and a number of independent variables. Events and effects can then be evaluated by testing the stability and interaction of these relationships in space. The theoretical model in this thesis will be tested by evaluating the nature of these anticipated effects in different regions of the province.

Knowledge of the practicality of any assumptions is essential to gauge the usefulness of the results in any abstract economic theory-based study. A trade-off between theoretical consistency and empirical relevance is often required to some degree, based on the goals of the particular project. For this thesis, the goal is to create a theoretical model that is consistent with real-world local labour market behaviour using idealized supply and demand curve shifts. The purpose is not to replicate reality exactly, but rather to create an avenue through which general labour market behaviour will be explored. To accomplish this, a number of generalizing assumptions are required. The first assumption is that workers with a particular skill level can fill any job requiring that or a lower skill level, so that all fishery jobs are able to be staffed by the lowest skilled workers. This assumption is reasonable due to the limited entry requirements of the fishing industry in the years before the moratorium. A second simplifying assumption is that all workers in the economy possess perfect information regarding job opportunities and potential wages in their own and other labour markets. Like the first assumption, this one is required to simplify the representation of arbitrage potential in the analysis. For

the purpose of simplifying the explanation of results from the regression model, a third assumption is that all spatial dependence between labour markets in the province results from people commuting to work. Although this is certainly a simplification, commuting is widely regarded to be the most important method by which communities interact (Karlsson and Olsson, 2006). In addition, this assumption allows the use of a single spatial weight function for the entire analysis, which is fundamental for maintaining simplicity of the model and comparing results.

To sufficiently evaluate the effects of the demand shock using the GWR framework, a number of different models will be created. These models will all share the same dependent variable, but have different independent variables. The dependent variable for all model runs will be the initial exposure to the cod moratorium, and the independent variables will each represent the different potential effects of the cod moratorium, in different time periods. A total of six different sets of independent variables will be created and tested, consisting of tests for utilization and quality declines from the years 1991 through 1996, and 1996 through 2001, as well as tests for the initial relationships in 1991. If the relationships behave the way they are anticipated to in the theoretical model, then the results of the empirical analysis will be that the theoretical model, with its assumptions, is a reasonable approximation of the Newfoundland and Labrador labour force in the 1990's. Both short and long term implications of the cod moratorium for the economic geography of the province's labour force will then be determined, and recommendations based on the model will be made.

3.4 – Model Inputs and Outputs

Like any empirical analysis, the usefulness of this thesis will not only rest upon the strength of the theoretical representation of the Newfoundland and Labrador economy, but also upon the modelling framework used to study it. While important, the theoretical base of the project ultimately relies upon the availability of a number of different data sets and statistics, and their ability to accurately represent the phenomena being studied. In this case, these 'phenomena' refer to both the theoretical concepts that are being evaluated in the model, as well as the empirical indicators of model performance that are necessary to interpret the results. For the modelling exercise to have any application to real-world policy, it needs to have tangible inputs and outputs. For instance, if the theoretical concepts discussed in section 3.2 could not be represented with actual data, it would be impossible to test the implications of the model. Similarly, if there was no standardized, replicable way to evaluate the results of the model, it would be impossible to know their significance, and to relate the study area to other locations or times that endure similar economic conditions.

An examination of the data required for this analysis, along with the appropriateness of the data that is available, will be provided in section 3.4.1, followed by a discussion of some ways in which the data, model, and theory are required to adapt to each other in section 3.4.2. This section will conclude with a summary of some of the parameters used in the model, and a discussion of their interpretation, in section 3.4.3.

3.4.1 – Required Data

Prior to being incorporated into the model, different items that constitute the framework of the economy need to be represented with specific and tangible pieces of

data. If a reasonable real-world indicator cannot be found to act as a surrogate for a particular concept, then the concept cannot be represented with the geographically weighted regression (GWR) framework introduced in section 3.3. However, data are rarely gathered to specifically represent individual theoretical concepts, so it may not be possible to match a concept exactly. In this case researchers are forced to approximate the representations of their theoretical concepts, an activity that weakens their model and narrows its applicability (e.g. using returns to schooling to represent returns to human capital (Lindh and Malmberg, 1999)). Therefore, for this thesis it is essential to carefully evaluate both the requirements of the theoretical model, and the data that are available, prior to performing any analysis or drawing any conclusions.

As previously discussed, the explicit phenomena being investigated in this model are the relationships between the cod moratorium labour demand shock in 1992, and the changes in labour market utilization and quality in different communities around Newfoundland and Labrador. These relationships will be tracked initially from the moratorium in 1992 to the end of specific government support during the period from 1996-1998, and then for a time following the end of government support, when the markets were theoretically forced to endure more severe adjustments. To gather relevant data to represent these phenomena, there are two practical options: direct fieldwork (*i.e.* primary data) or statistical data collected by others (*i.e.* secondary data). Fieldwork can immediately be eliminated as an option for this thesis, because the events being studied are in the past, and people may have moved and perceptions may have changed a great deal since they occurred. Secondary data, on the other hand, are relatively easy to access, are available concerning a wide variety of topics covering the appropriate geographical

areas and resolution, and stretch back well before the cod moratorium occurred in the early 1990's. For these reasons, data sourced from different government agencies will be utilized in this analysis.

Prior to selecting data for this analysis, it is important to determine specifically what sources may be used, what data are available, and how closely the concepts used in the theoretical model will be approximated. Perhaps the most visible source of data pertaining to the Newfoundland and Labrador labour market is Statistics Canada (www.statcan.gc.ca), which is responsible for a large amount of statistical data gathering and dissemination, including the national census and the CANSIM socioeconomic database. Although Statistics Canada offers a wide variety of employment and other labour market statistics, they are only available for specific, standardized spatial units. The smallest spatial unit available for most of Newfoundland and Labrador is the polygon-based census subdivision (CSD), which in many cases is analogous to the community level. However, for some rural or sparsely populated areas, CSD's consist of a number of communities merged together, and in some cases are hundreds of kilometres across (**Figure 3.3**). This is an issue in this study because the focus is on rural versus urban behaviour, and a consistent, community-based spatial unit is a priority. Fortunately, a community-based spatial unit does exist, in data disseminated by a number of provincial government departments. This unit, the government community list (GCL) unit, is a standard format utilized by the provincial government for community-based data (Government of Newfoundland and Labrador, 2008). There are 742 entries in the GCL, each corresponding to an individual community, and they are represented with points instead of polygons. The departments which publish labour market data using the GCL

system that is relevant to this thesis are the Department of Human Resources, Labour, and Employment (HRLE), and the Newfoundland and Labrador Statistics Agency (NLSA), located in the Department of Finance. In the context of this thesis, employment and labour market statistics are available from HRLE, whereas distance related data and spatial datasets are available from the NLSA. Generally, the labour market data published by HRLE is census-based data that has been aggregated by the provincial government into communities according to the GCL. The temporal resolution is therefore the same as the census, so it is only available for one in every five years. In this thesis, this means that data are available for the years 1991, 1996, and 2001, which can serve as approximations for the theoretical benchmarks of (i) the period immediately prior to the moratorium, (ii) the point when direct government support was ending, and (iii) an adjustment period shortly afterward.

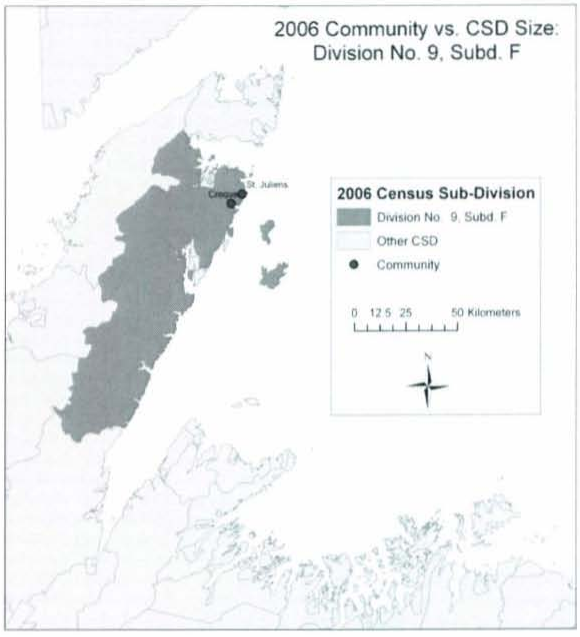

Figure 3.3 Example of Community vs. CSD Location and Size (Government of Newfoundland and Labrador, 2008; Statistics Canada (n.d.b))

The data available for this thesis is classified into two groups: employment and labour market data intended to represent the components of the relationships being modelled, and supporting data, such as distance information and spatial data. The labour market data from HRLE ranges from such topics as employment by industry, education, and income, to work activity (months worked per year), and labour force activity (participation rates, unemployment rates). The dependent and independent variables for the models will be created using this data. The distance information and spatial data, on the other hand, are required to situate the model in space. Data of this type for Newfoundland and Labrador is generally available from the NLSA, as well as federal government sources such as GEOBASE (www.geobase.ca) or GEOGRATIS (geogratis.cgdi.gc.ca).

Theoretically, the dependent variable should be the initial exposure of a community's labour market to the cod moratorium. To represent this in the model, the percentage of each community's labour force working in fishery-related occupations (fishers and fish processing workers) in 1991 will be used. It is expected that the more dependent upon the fishery a community is in 1991, the more susceptible that community will be to the effects of the cod moratorium.

The potential independent variables for the model discussed in section 3.3 consist of some anticipated effects of the moratorium on local labour markets. In particular, changes in labour market utilization and labour market quality are expected. Based on the available data, labour market utilization is represented in this model by the work activity, or the percentage of workers in the labour force who worked different numbers of weeks during the year. Due to its seasonal nature, a high degree of fishing employment in a community is expected to be represented by a relatively high percentage of the labour force working for part of the year. Likewise, high unemployment would be represented by a high percentage of workers who did not work, and non-fishing, service employment would be represented by a high percentage of full year workers. To represent these dimensions, the work activity variable is split into five different classes: did not work (0 weeks), part year low (1-14 weeks), part year medium (15-20 weeks), part year high (21-48 weeks), and full year worked (48 weeks plus).

Two different datasets need to be used to represent quality impacts, because of their classification into two separate effects. The first area in which labour market quality is anticipated to have been affected is through the age structure of the workforce. Data in which the workforce of each community is classified into different age cohorts are used

to represent the age structure of the workforce. Although the dataset acquired from HRLE consists of the entire workforce classified into four separate nine-year cohorts and one upper cohort (55+), this thesis only requires distinctions between life stages. For this reason, the cohorts will be merged to form three broad groups: young (15-34), middle aged (35-54), and older (55+). The second area in which labour market quality may have been affected is through the skill level of the workforce. To represent skill level, education data is used. The education data available from HRLE consists of the percentages of the workforce with different education levels. Again, the data are available with more divisions than are required for this thesis, so they will also be merged to form three broad classes representing: low skill (no high school diploma), medium skill (high school diploma), and high skill (post secondary education (either university or college/trade school)).

A discrete version of the GWR as described in section 3.3 requires data able to describe the discrete distances between locations. For this analysis, the locations are communities, and the distances are discrete because they are along roads. ArcMap's network analyst will be used to calculate the road distance between communities, using a NLSA GCL point file as the origin and destination communities, and the provincial road network file from GEOBASE.

3.4.2 – Model Adaptations

Beyond the formulation of theory and collection of data, a number of essential manipulations are required before the analysis for this thesis is performed. While the selection of an appropriate model and relevant data are necessary for a strong analysis, it is important to ensure that any standardization or adaptation required by either the

theoretical framework or the geographically weighted regression (GWR) model is performed. For this study, data need to be standardized according to communities included and model sample size, the created road network distances need to be converted into a spatial weights matrix, and model adaptations such as variable exclusion need to be performed. Only after the data are standardized and the GWR model properly constructed can the results of the study be evaluated and interpreted.

3.4.2.1 – Community List

The first class of adaptations required for this thesis relate to the community data. In their original state, the datasets intended to represent the dependent and independent variables in the GWR model are fairly heterogeneous with respect to the communities contained within them. Changes in community definitions and settlement patterns between 1991 and 2001 have resulted in some communities appearing in or disappearing from the datasets during the time span of the analysis. To proceed with the analysis, the list of communities contained in every dataset need to be standardized. In this project, this is accomplished by excluding any communities not listed in every dataset from the analysis. This initial standardization of the data creates a universal list of 589 communities upon which the rest of the analysis is based.

3.4.2.2 – Spatial Weight Function

The adaptation of a broadly applicable modelling framework such as GWR to an empirical study area such as Newfoundland and Labrador requires some modifications specific to the model. In particular, the discrete GWR (DGWR) model requires a spatial weight function to be defined, and the data to be specified according to this function. The purpose of the spatial weight function is to approximate the mechanism responsible for

spatial dependence in the study area (Fotheringham *et al.*, 2002). This project is concerned with labour markets and labour market behaviour, so any spatial dependence occurring in the study area will be assumed to be the result of commuting (Johansson *et al.*, 2002). Therefore, the spatial weight function is based upon commuting behaviour in the province. The first step in creating the spatial weight function is to determine the cumulative distance decay function for commuting between communities, which is calculated based on optimized (1996-2006) data acquired from the Newfoundland and Labrador Statistics Agency (Government of Newfoundland and Labrador, 2008) (**Figure 3.4**).

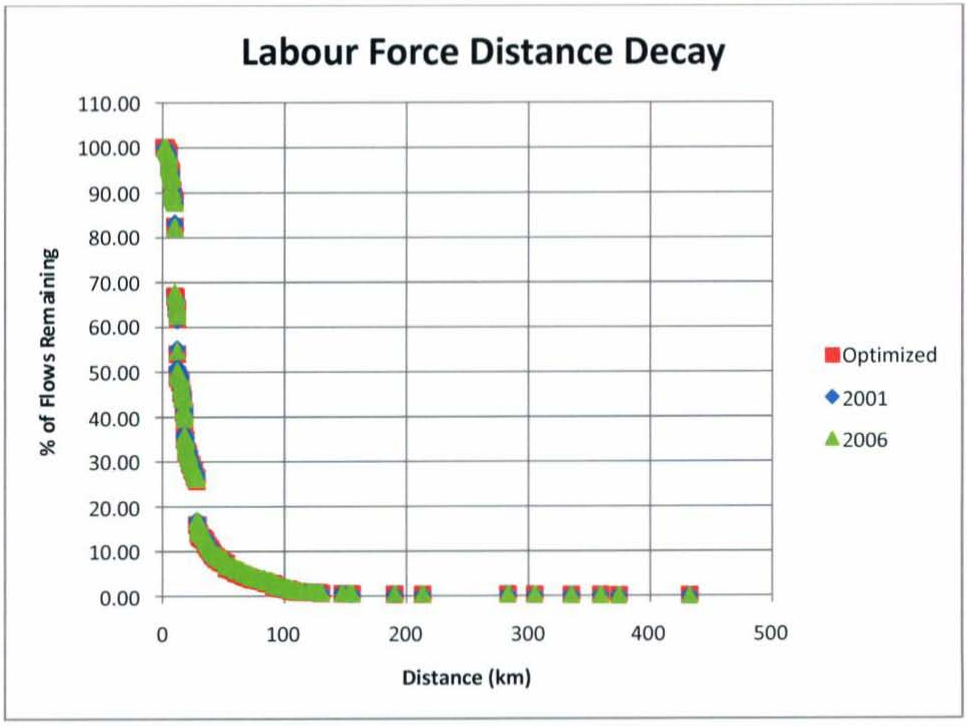

Figure 3.4 Newfoundland and Labrador Commuting Distance Decay (Government of Newfoundland and Labrador, 2008)

In the province, approximately 90% of workers commute less than 40 km every day, and 99.5% of workers commute less than 151 km. Accordingly, a fixed kernel spatial weights function is calculated by applying fuzzy set theory to benchmarks derived from the commuting decay function. Mathematically, this fuzzy function is represented in **Equation 3.1**:

$$M = \begin{cases} 1 & \text{for } d < 40 \text{ km} \\ \left(1 - \left(\frac{(d-40)}{(151-40)}\right)\right) & \text{for } 40 \leq d \leq 151 \text{ km} \\ 0 & \text{for } d > 151 \text{ km} \end{cases} \qquad \textbf{(3.1)}$$

Where M = the spatial weight value and d = the distance between two communities.

According to this equation, two communities receive a membership of 100% if they are within 40 km of each other, and 0 % if they are beyond 151 km. This fuzzy membership function may be used as the spatial weight function in a GWR model, if the percentage membership between two communities is used as the percentage of consideration of one community in the local regression of the other. For example, if two communities are separated such that there is a 50% membership probability (chance of commuting) between them (i.e. they are located 96.5 km apart), then each of those communities will be included in the other's regression 50%. In other words, they will count as half of a neighbour for each other.

3.4.2.3 – Model Condition

A properly functioning regression model requires a sample size large enough to ensure that individual cases do not have especially large influences upon the analysis.

This is an issue when attempting to conduct a local analysis in a largely rural province such as Newfoundland and Labrador, which is composed of both densely and sparsely populated regions. If a community is so remote that it has too few neighbours for the function to be properly specified, it can have an impact on the global statistics used to evaluate the overall model, in addition to the reliability of local estimates. For this analysis, the condition number (see section 3.4.3) is used to determine sample size sufficiency for the different areas of the study area. Individual communities that are found to have fewer than ten neighbours and condition numbers greater than thirty are considered to be too unstable for this analysis, and are excluded. This step removes a further eleven communities from the list of included communities, reducing the coverage of the model from 589 to 578.

3.4.2.4 – Variable Exclusion

The final adaptation required for the GWR model in this thesis concerns data exclusion. In general, when conducting a regression analysis on percentage data in which all possibilities sum to one, all possibilities do not need to be included in the analysis (Feyrer, 2007). Instead, in this case, one possibility is excluded, because it can be calculated from the others. For example, a work activity division can be excluded from the utilization change analysis, because together they sum to the entire workforce. If the 'did not work' class was excluded from the model, for instance, the coefficients of the other classes would then be calculated in relation to the percentage of workers who did not work. A coefficient of +5 for 'full year worked' would therefore mean that for every percentage increase in the dependent variable, there would be an increase in full year workers five times that of workers who did not work. The youngest (15-34) age structure

and lowest (no high school) education class are excluded from the quality model, and the 'did not work' class is excluded from the utilization model.

3.4.3 – Parameters and Interpretation

A sound understanding of the structure and significance of the outputs of any study is necessary for the results of the study to be properly utilized. This is particularly true for quantitative and statistical studies which have a diverse range of outputs used to track a number of different aspects of the model. Commonly used groups of statistics include tests for model fit or comparison and tests for model conditioning or stability, in addition to statistics used to relay the empirical results and significance of the model (Brunow and Hirte, 2006). Although this thesis mainly utilizes statistics for fit, model conditioning, and results that are standard across many different modelling exercises and regression models, it does use some modified statistics.

As a regression model, the discrete geographically weighted regression (DGWR) model produces a number of statistics that are common to more standard techniques. A list of these statistics include the coefficient of determination (r^2) and the Akaike Information Criterion (AIC), which are used to determine overall model fit, and White's heteroskedasticity test and the global Moran's I test, which are used to compare the ability of each of the three models to remove the spatial effects of spatial dependence and spatial heterogeneity respectively from the data (Anselin, 1988).

While the functioning of the r^2 statistic is fairly well understood, the AIC is somewhat obscure. Unlike the r^2 statistic, the AIC does not compare a given distribution to an ideal case (Fotheringham *et al.*, 2002). Instead, the AIC is generally calculated by measuring, for a given model, the distance between the model distribution and the

distribution of the actual phenomena being studied. This procedure is repeated for each model candidate, and the model with the lowest AIC value is generally considered to be the best.

Although the fact that GWR and OLS are both regression methods means that the model results are broadly similar, they differ in one important way, because of the local nature of GWR. In general, both types of models report results in the form of variable coefficients and standard errors. However, while GWR produces a new, separate, estimate and standard error for each community, OLS only produces one of each for the entire study area. Also, while the significance of estimates in the OLS model is determined through the use of a standard t-score, calculated by dividing the estimate by the standard error, the local nature of GWR necessitates the use of a 'pseudo'-t score, which is not as statistically sound. The reasons for this, as well as the limitations it imposes on the interpretability of results, will be further explored later in this section.

Once the optimum model type is selected and empirical results determined, it is important to conduct some tests of the model conditioning and efficiency. It is difficult to ensure that the results of a model are their most accurate if limited information is available concerning the ability of the model to satisfy its assumptions in the study area. For this reason, tests for sample size, condition number, and both univariate and multivariate outliers are used in this analysis. Local Moran's I is also used to test the ability of the model to remove systematic spatial autocorrelation. To ensure that the extra effort required to incorporate space into the GWR model is warranted, a test for parameter stationarity originating in Leung *et al.* (2000) is used.

Although the interpretation of results in this thesis relies, for the most part, on standard, commonly available statistics, the discrete nature of the GWR and the often sparsely populated study area means that some dimensions of the output need to be tailored specifically for the technique. The most important such statistic with respect to the direct interpretation of the model results is the t-type score used to determine significance of model coefficients (Wimberly *et al.*, 2008). While the significance of estimates resulting from global OLS models is readily determined and interpreted with a standard t-score, the same formula is more difficult to apply to local GWR models, because the GWR estimation procedure violates the assumption of independence required for the t-test (LeSage, 1999). Generally, to perform a t-test on two separate parameter estimates that would enable the significance of the estimates to be compared, the two estimates need to be independent of each other. Any overlap between the areas from which the estimates have been calculated would mean that the independent/dependent variable relationships being modelled would not be independent for the two regression points, and therefore any tests conducted on the point assuming independence between them would be technically invalid. This is a potential issue because if regression estimates are not independent, then the estimates themselves may become unstable if a different sampling design is used (Wheeler, 2007). The estimates would only become unstable if the data underlying them was of inconsistent quality, and the overall model was not strongly conditioned. Whereas the estimates of an OLS model are assumed to be independent (because they are only calculated once), the estimates of a GWR model cannot. For this reason, the results of t-score calculations for GWR estimates have to be labelled 'pseudo' t-scores (Wimberly *et al.*, 2008). As explained in section 2.3, GWR

parameter estimates are calculated though the use of a 'moving window' regression, where a fresh regression is performed at each regression point. For this thesis, although this means that results are now available at a community-level resolution, it also creates a resampling issue whereby the variable relationships for a particular community may be considered in the regressions of a number of other communities. If, in any single community, the quality of the model were to fall below that which is commonly acceptable in a regression model, through an issue such as low sample size, it could potentially impact the estimates for other, otherwise well specified, communities. Therefore, it is important to ensure that all areas of a GWR model are well specified prior to testing for parameter significance. While the regression points in a GWR can technically never be independent of each other, the pseudo t-score may be interpreted as practically equivalent to a t-score, if the model is stable and well conditioned (Wimberly *et al.*, 2008).

For this study, a condition number is used to quantify the stability of local regressions, as an alternative to the individual sample size. Mathematically, the condition number is calculated as the largest division of the relative error in the regression coefficient and the relative error in the model estimate (Velleman and Welsch, 1981). It measures the amount of change in either the coefficient or the model estimate that is expected to occur relative to a unit of change in the other. The literature suggests that the lower the condition number, the better specified the model, and values of less than 30 are ideal (Wheeler, 2007). To ensure that the results of an analysis are as accurate and well-specified as possible, the condition number for each location is evaluated prior to interpretation of the pseudo t-scores.

3.5 – Conclusion

In this chapter, a method of adapting a local geographically weighted regression (GWR) to study a labour market shock was described. Although most efforts in this area to date have neglected to recognize any significant role for spatial location, a discrete GWR (DGWR) model will allow a shock and its effects to be studied in a realistic manner, with empirically-derived anticipated interaction effects. Analysts utilizing the DGWR model can also draw upon the existing regression literature to evaluate its performance relative to other modelling frameworks, in addition to the applicability of alternative theoretical approaches to a defined empirical problem. In this thesis, the DGWR model is used to test for effects of the 1992 cod moratorium in a realistically-structured Newfoundland and Labrador labour market. It allows the distances between actors in the economy to be accurately simulated to a degree not possible with continuous or non-spatial models.

4.0 – Results

4.1 – Introduction

The primary focus of this study is to determine the applicability of the theoretical model described in the previous chapter to the reaction of local labour markets to the cod moratorium in Newfoundland and Labrador. The complex nature of the moratorium means that a number of indicators are required to be evaluated before the degree of this applicability is known. For this thesis, these indicators are classified into two broad groups: statistics relating to ensuring the specification and quality of the analysis, and statistics relating to the performance and interpretation of the overall model. This chapter will report on the specific results of these two groups.

The first section will contain the results of the pre-analysis stage of the analysis. In this section, the optimum method for conducting the analysis will be determined using a combination of model fit, parameter stationarity, and spatial effects tests. Following the selection of the best model, the data will be tested using condition numbers and outlier analysis and adapted to ensure that the model is well specified, and that the results are empirically sound.

The second section of this chapter will present the empirical results of the model. The success of the theoretical model in representing the empirical reality will be evaluated using inferential statistics that are applied to each model run.

4.2 – Model Fit Results

The first step in this empirical analysis is to select the most appropriate model. Based on a review of the literature and the particularities of the study area, three different modelling frameworks were considered for this thesis. These methods consist of an

ordinary least squares (OLS) model, a continuous geographically weighted regression (GWR) model, and a discrete GWR (DGWR) model. For the DGWR model, it has been theorized that labour market phenomena in Newfoundland and Labrador operate between communities through commuting behaviour along the roadways. One implication of this hypothesized interaction is that commonly used OLS and continuous GWR models will fail to accurately represent economic behaviour, and will therefore not be able to satisfactorily evaluate any moratorium impacts. However, the DGWR model is unproven, and the extent of any advantages over traditional models is unknown. For these reasons, the first step in this analysis is to compare the overall performance of the three frameworks, and quantitatively determine the degree to which the DGWR model outperforms the traditional approaches. The approaches will be assessed according to both traditional goodness of fit statistics, and the models' ability to represent the spatial economic system.

4.2.1 – Global Fit Statistics and Parameter Stability

It is a common practice to evaluate the overall performance of a regression model through the use of statistics (Fotheringham, 1998). Although a number of alternative options for gauging global fit exist, a combination of the coefficient of determination (r^2) and the Akaike Information Criterion (AIC) statistics were used in this thesis. In general, r^2 is a measure of the degree to which a model can replicate a set of empirical data; it is a measure of the 'goodness of fit' of the model (Fotheringham *et al.*, 2002). The maximum possible value of r^2 that a model can achieve is one, which would indicate a perfect fit between model and data. The AIC, on the other hand, is a relative measure of model performance. Unlike the r^2 statistic, the AIC does not have an optimum value. Instead, to

determine which of a group of models provides the best fit, an AIC is calculated for each model being tested, and the model with the lowest AIC value is interpreted as having a relatively better fit than the others.

For this study, r^2 and AIC statistics were calculated for each of the six model runs (representing quality (Q) and utilization (U) in the 1991, 1991-1996, and 1996-2001 time periods), to compare the performance of the three modelling frameworks in every relevant situation. The results are summarized in **Table 4.1**. It is immediately evident when analyzing these results that the performance of the discrete geographically weighted regression (DGWR) model far exceeds that of both the ordinary least squares (OLS) and the continuous GWR model, according to both r^2 and AIC statistics. In every case, the DGWR model obtains the highest r^2 value (e.g. $r^2 = 0.74$ for the 1991 utilization model) and the lowest AIC value (e.g. AIC = -1013.05), while the OLS model obtains the lowest r^2 value (e.g. $r^2 = 0.42$), and the highest AIC value (e.g. AIC = -723.06). The continuous GWR model is consistently ranked in the middle (e.g. $r^2 = 0.63$ and AIC = -932.51). The primary implication of this comparative analysis is that, of the three models tested, the DGWR is the most statistically appropriate modelling method for this thesis. In addition, the OLS model represents the least appropriate modelling approach. Therefore, there is evidence that discrete space represented by road network distances provides a better approximation of the functioning of Newfoundland and Labrador labour markets than existing methods.

Table 4.1 Model Fit Results

Model	Year	AIC DGWR	AIC GWR	AIC OLS	AdjRsq DGWR	AdjRsq GWR	AdjRsq OLS
Q	**91**	-837.45	-775.04	-533.93	0.64	0.52	0.19
Q	**91-96**	-762.02	-653.94	-420.34	0.60	0.41	0.01
Q	**96-01**	-947.44	-827.97	-429.91	0.69	0.56	0.03
U	**91**	-1017.56	-932.51	-723.06	0.74	0.63	0.42
U	**91-96**	-874.74	-806.67	-570.46	0.67	0.55	0.24
U	**96-01**	-1013.05	-844.21	-670.32	0.72	0.57	0.36

While a local model such as GWR may provide a better overall fit for a set of data, it also requires a great deal of additional work to construct and interpret. If some phenomenon were to be largely stationary in space, it is conceivable that any additional interpretation ability provided by a GWR would not outweigh the costs of constructing the model. For this reason, a test for parameter stationarity used by Leung *et al.* (2002) was applied to the data. In general, a significant ($p >= 0.95$) result for this test indicates that the tested parameter does vary to a significant degree over space, and that a local method would be the most appropriate to use. For this thesis, both the quality and utilization models contained a majority (n = 11/15 and 15/15, respectively) of significant spatially unstable relationships (see **Tables 4.2** and **4.3**).

Table 4.2 Parameter Stationarity Results, Quality Model (p-value)

Model	Year	Constant	Middle Aged Change	Older Change	Medium Skill (Education) Change	High Skill (Education) Change
Q	**91**	0.93	0.67	0.68	1.00	1.00
Q	**91-96**	1.00	1.00	1.00	1.00	0.91
Q	**96-01**	1.00	1.00	1.00	1.00	1.00

Table 4.3 Parameter Stationarity Results, Utilization Model (p-value)

Model	Year	Constant	Worked Part Year (Low)	Worked Part Year (Medium)	Worked Part Year (High)	Worked Full Year
U	**91**	1.00	1.00	1.00	1.00	1.00
U	**91-96**	1.00	1.00	1.00	1.00	1.00
U	**96-01**	1.00	1.00	1.00	1.00	0.95

4.2.2 – Spatial Statistics

To understand how Newfoundland and Labrador labour markets function, it is important to evaluate how different spatial incarnations perform with respect to spatial effects in the economy. Although r^2 and AIC statistics allow the selection of the best modelling framework overall, tests for residual spatial dependence and heterogeneity are useful to specify the nature and extent of any discrepancies between the framework and any spatial structure which exists in the data. In this study, statistics were used to test for heteroskedasticity, as well as spatial autocorrelation amongst the residuals of each framework. Moran's I (both local and global (Anselin, 1995)) was used to test for spatial autocorrelation, and the White test was used for heteroskedasticity.

Like the tests for model fit, both types of tests for spatial effects were applied for each of the three model frameworks, to each of the six model runs. The results of the global Moran's I tests are shown in **Table 4.4**. The discrete GWR (DGWR) model is able to remove significant ($z >= 1.96$) global spatial autocorrelation amongst its residuals in all but one model run (labour market quality change from 1991 through 1996 ($z = 2.78$)), while significant spatial autocorrelation still exists in four runs of the continuous GWR model, and all runs of the OLS model.

Table 4.4 Moran's I Results (Index value and z-score)

Model	Year	MoransI DGWR	MoransI GWR	MoransI OLS	MoransIz DGWR	MoransIz GWR	MoransIz OLS
Q	**91**	0.01	0.02	0.35	1.40	3.78	42.83
Q	**91-96**	0.02	0.02	0.35	2.78	4.48	42.72
Q	**96-01**	0.01	0.01	0.33	1.09	3.12	40.72
U	**91**	-0.01	0.00	0.21	-0.69	1.27	26.93
U	**91-96**	0.00	0.00	0.29	-0.06	1.31	36.03
U	**96-01**	0.01	0.01	0.13	1.59	2.14	16.99

As seen in **Figure 4.1**, the local Moran's I test indicates that the communities contributing the most to the significance of the global Moran's I statistic for the 1991-1996 quality change model are in urbanized areas of central Newfoundland and the northeast Avalon. This initial stability of the moratorium exposure/quality change relationship in non-fishing areas is accounted for in the theoretical framework detailed in section 3.2.2, so the significance of the global Moran's I statistic for the 1991-1996 time period is expected.

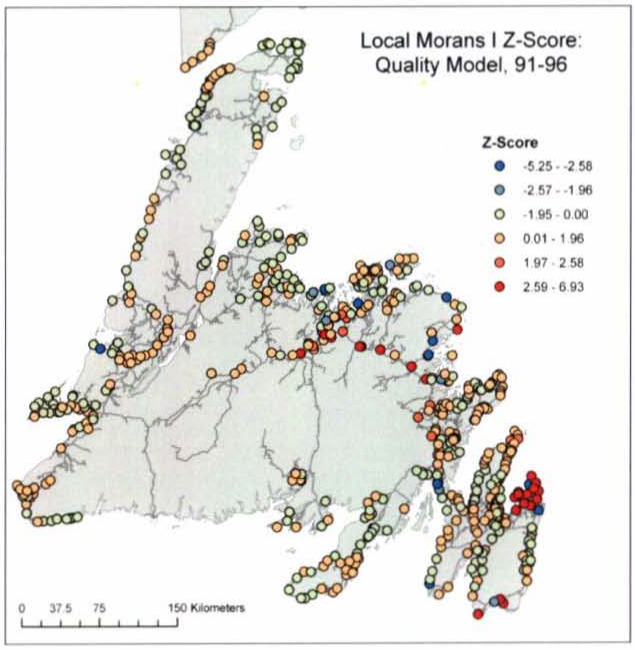

Figure 4.1 Local Moran's I results for Labour Market Quality Change model, 1991-1996

The primary implication of the spatial statistics utilized in this study is that, overall, the DGWR model is able to represent more of the interactions and spatial dependency that exists in the Newfoundland and Labrador labour force than either of the two alternative methods. This means that, in general, the combination of road network distances and commuting-derived distance decay effects is the best approximation of labour market structure tested in this thesis, and that local analysis on its own cannot capture the spatial dependence in the system.

While the DGWR is clearly the best model regarding model fit and spatial dependence, it only managed to outperform the continuous GWR with respect to heteroskedasticity by a slim margin. According to the White test (**Table 4.5**), significant heteroskedasticity ($p >= 0.95$) is present in all but two DGWR model runs, in all but one

continuous GWR run, and in every OLS run. A possible explanation for this is that the labour markets in the province exhibit a hierarchical structure, where the reach of higher order centres for some interactions extends farther than either lower order or supply centres. However, incorporating hierarchical effects into the model is beyond the scope of this thesis. The only significant implication of the White test results for this research is that the DGWR did outperform the two existing methods with respect to heteroskedasticity, and therefore better and more efficiently represents the spatial dynamics of the study area.

Table 4.5 White Test Results (p-value)

Model	Year	White GWR(dsc)	White GWR	White OLS
Q	**91**	0.67	0.74	1.00
Q	**91-96**	1.00	1.00	1.00
Q	**96-01**	0.94	0.97	1.00
U	**91**	0.97	0.99	1.00
U	**91-96**	1.00	1.00	1.00
U	**96-01**	1.00	0.97	1.00

4.2.3 – Model Condition and Pre-Analysis Results

According to the global fit, parameter stability and spatial effect statistics, the discrete GWR (DGWR) model is more appropriate for this analysis than either the OLS or continuous GWR model. However, to evaluate relationships at a local level, the data requirements for GWR are higher than for OLS. In a sparsely settled province such as Newfoundland and Labrador, it can often be difficult to obtain a sufficient sample size within the specified regression window. As such, communities with fewer than ten neighbours and which consistently produce condition numbers larger than thirty in the

models are excluded from the model. **Figure 4.2** is a map of all communities which have been included in and excluded from this analysis. The primary concern of this thesis is with evaluating regional labour market trends, so the exclusion of isolated communities is a necessary step toward conducting a focused and relevant analysis. In addition, because the absence of the effects of commuters may cause isolated communities to react to labour market shocks differently than communities located in regions (Glendon and Vigdor, 2003), effects in one are not comparable to effects in the other, and require different methods to accurately evaluate.

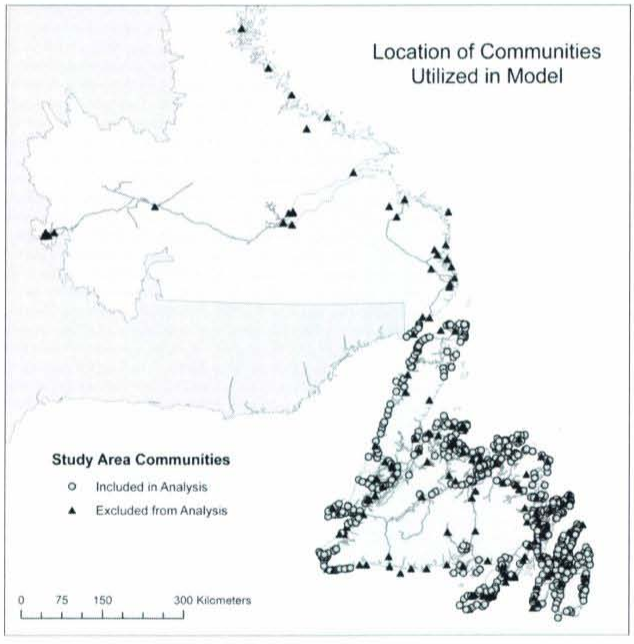

Figure 4.2 Communities Included in or Excluded from the Analysis

Sample sizes for each included community are represented in **Figure 4.3** and condition numbers for the performed models are displayed in **Figures 4.4-4.9**. All communities in Labrador (excluding the Labrador Straits area), and some communities

located along the south coast of Newfoundland are excluded from this analysis because of the degree of isolation calculated using these methods. To preserve overall sample size, however, communities possessing a sample size below ten and/or condition numbers above thirty in one or two model iterations are included in the analysis, if they remain well specified according to the other indicators.

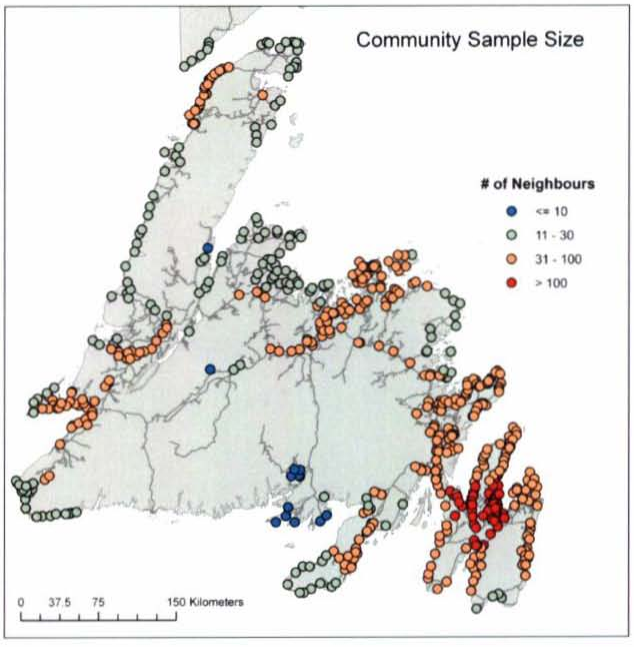

Figure 4.3 Community Sample Size

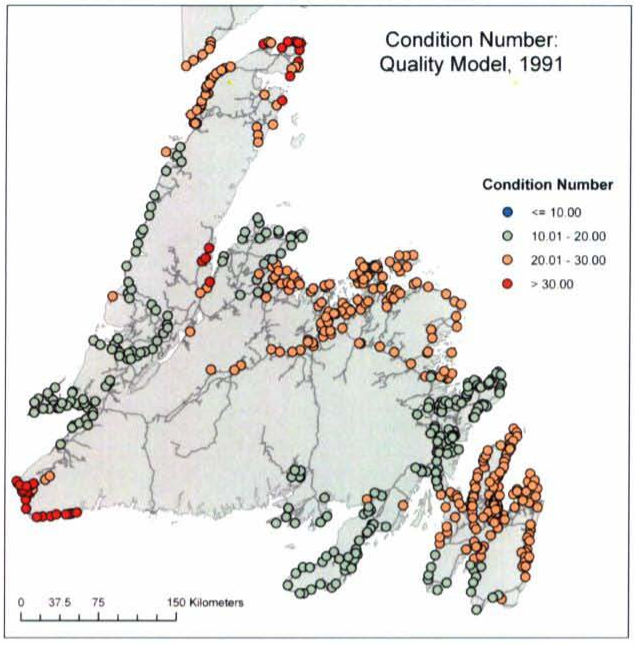

Figure 4.4 Community Condition Number, 1991 (Quality Model)

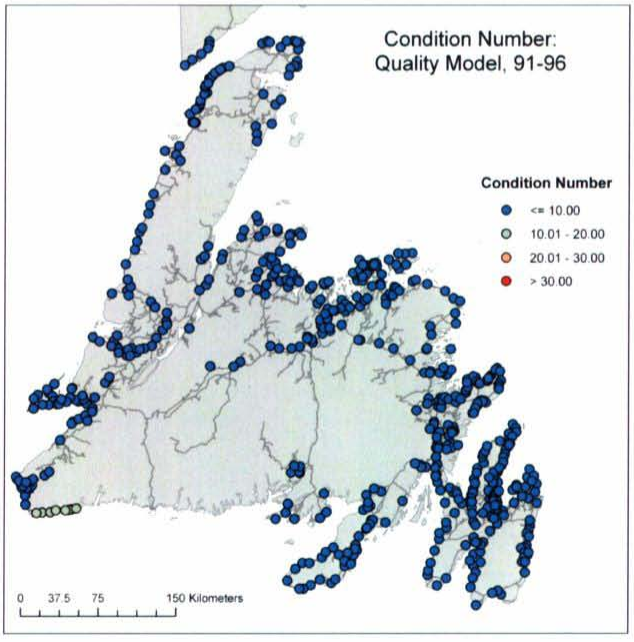

Figure 4.5 Community Condition Number, 1991-1991 (Quality Model)

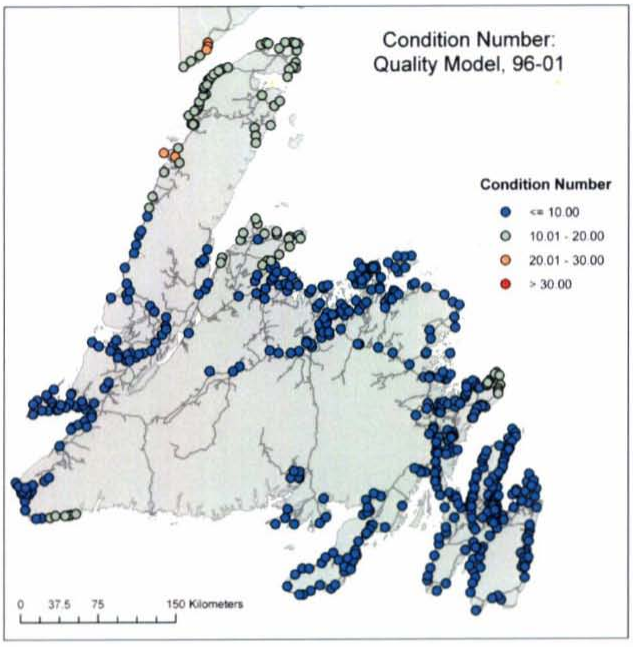

Figure 4.6 Community Condition Number, 1996-2001 (Quality Model)

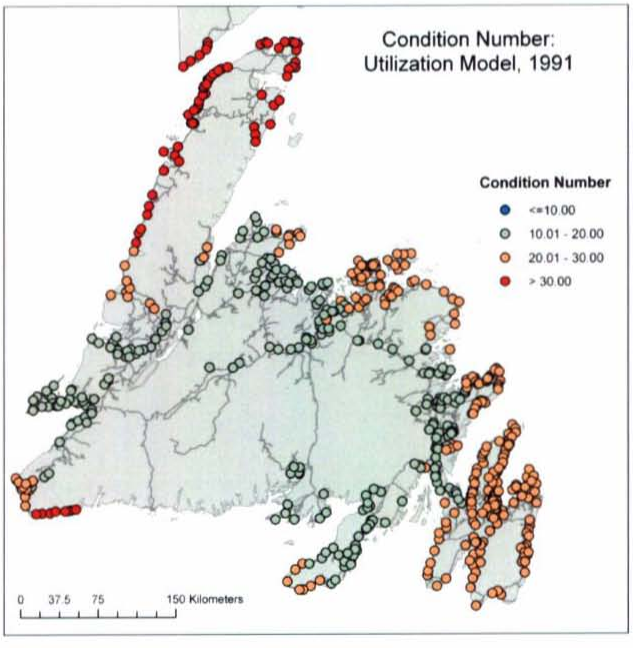

Figure 4.7 Community Condition Number, 1991 (Utilization Model)

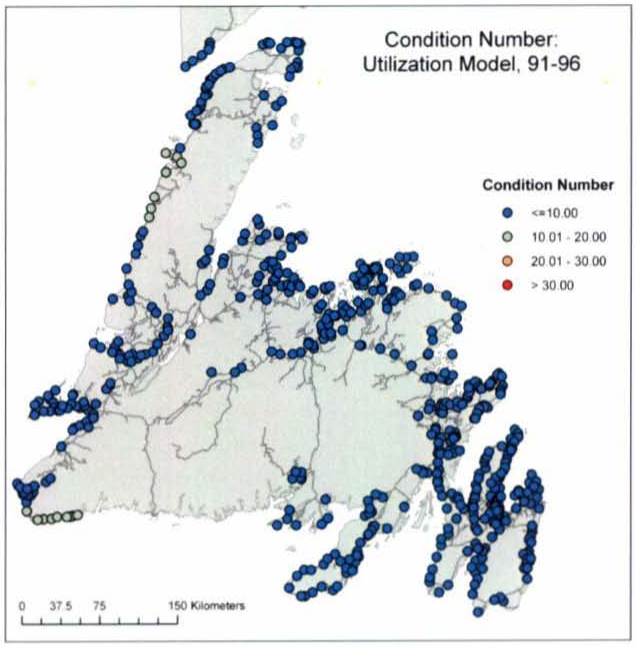

Figure 4.8 Community Condition Number, 1991-1996 (Utilization Model)

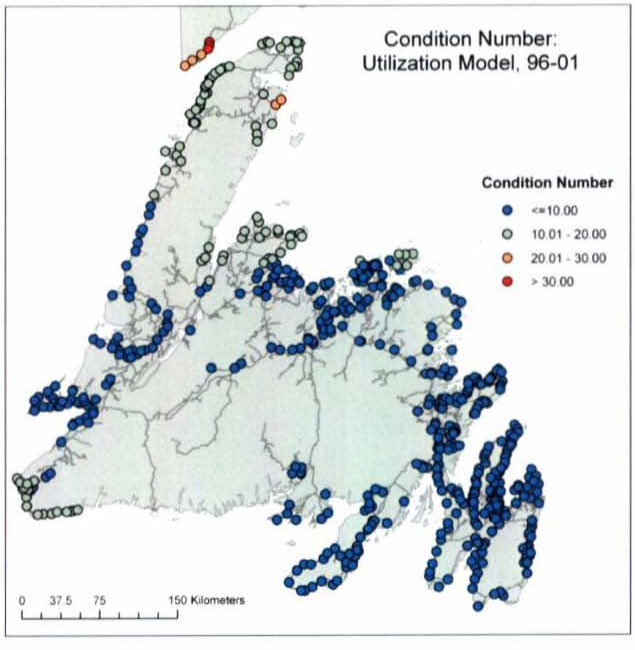

Figure 4.9 Community Condition Number, 1996-2001 (Utilization Model)

The final set of results related to the pre-analysis of this thesis is for the outlier analysis. In order to make inferences from the pseudo-t statistic, the underlying data are required to be normally distributed. To ensure this, the quantities of both univariate and multivariate outliers were determined using z-scores and Mahalanobis distance (both $p < 0.05$) respectively. With normally distributed data, the expected number of outliers for a sample size of 578 is approximately 29. While all 6 model runs contain more than 29 multivariate outliers, only 6 of the 30 variables have more than 29 univariate outliers (see **Tables 4.6** and **4.7**). The higher than expected number of multivariate outliers may be explained by the localized nature of interactions between labour market variables, in which case they should not interfere with the results of the DGWR model.

Table 4.6 Outlier Analysis Results, Quality Model

Model	Year	Fishery Dependence (91)	Middle Aged Change	Older Change	Medium Skill (Education) Change	High Skill (Education) Change	Multi-variate Outliers
Q	**91**	17	24	14	31	26	42
Q	**91-96**	17	23	29	30	20	45
Q	**96-01**	17	22	28	24	17	36

Table 4.7 Outlier Analysis Results, Utilization Model

Model	Year	Fishery Dependence (91)	Worked Part Year (Low)	Worked Part Year (Medium)	Worked Part Year (High)	Worked Full Year	Multi-variate Outliers
U	**91**	17	32	28	24	26	45
U	**91-96**	17	21	20	28	27	46
U	**96-01**	17	44	34	32	20	40

4.3 – Empirical Results

Due to its better model fit (AIC = -1017.56 through -762.02 versus -932.51 through -653.94 for the continuous geographically weighted regression (GWR) and -723.06 through -420.34 for the ordinary least squares (OLS) models), and ability to remove more spatial autocorrelation (n = 1 significant (z >= 1.96) model run versus 4 for the continuous GWR and 6 for the OLS models) and spatial heteroskedasticity (n = 4 significant (p >= 0.95) model runs versus 5 for the continuous GWR and 6 for the OLS models), the empirical testing of the theoretical model will be conducted using a discrete GWR (DGWR) framework. In addition to being a more accurate and better specified model than other techniques, the results of the DGWR model are easier to interpret because distances between communities are more realistic, and regression windows therefore contain more logical groupings of communities.

To allow overall trends and local applicability to be explored simultaneously, the empirical results of the model will be featured first as aggregated maps in section 4.3.1, and second as an illustrated example of two rural communities (one remote and one urban-adjacent) in section 4.3.2.

4.3.1 – Overview

To assess the accuracy of the constructed theoretical model, its idealized theoretical relationships will be compared to empirical relationships between cod moratorium exposure and quality- and utilization-based labour market effects for Newfoundland and Labrador communities in the years following the 1992 fishery closure. The pseudo-t statistic will be the quantity used to make this comparison. Specifically, a pseudo-t score representing the significance of each relationship will be

generated for every community and time period, and then displayed in individual maps. Once displayed, the patterns generated in the maps will be analyzed, and their resemblance to the anticipated theoretical patterns tested.

Prior to interpreting any empirical results of this study, it should be recalled that affiliation with the cod fishery, and therefore applicability to this model, is not expected to be constant throughout the study area. Instead, as stated in section 3.2.2, the predictions of the theoretical model are expected to be generally more applicable to the northeast coast of the island, so interpretation should in turn focus more on this area.

According to the theoretical model, the first effect that the cod moratorium should have on the Newfoundland and Labrador labour force is a change in utilization. Initially, because of the seasonal nature of the fishery, it is expected that coastal communities will have relatively under-utilized labour markets, with a higher percentage of their labour force working for less of the year. This occurs as expected, with noticeable coastal/inland differences in labour market utilization for percentage employed in the fishery in 1991 versus both 'part year high' and 'part year medium' (in relation to 'did not work') (see **Figures 4.10 – 4.13**). Communities with a high degree of dependence upon the fishery in coastal areas have historically higher percentages of workers working for part of the year, versus not at all.

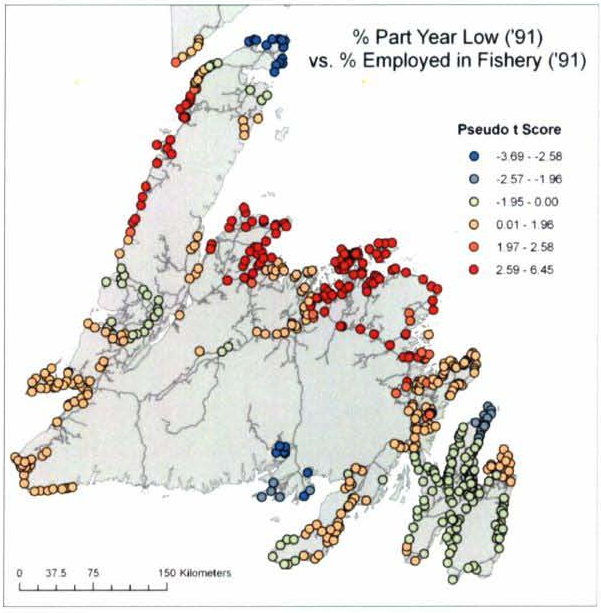

Figure 4.10 Percent Part Year Low (Utilization) Model Results, 1991 (pseudo t-score)

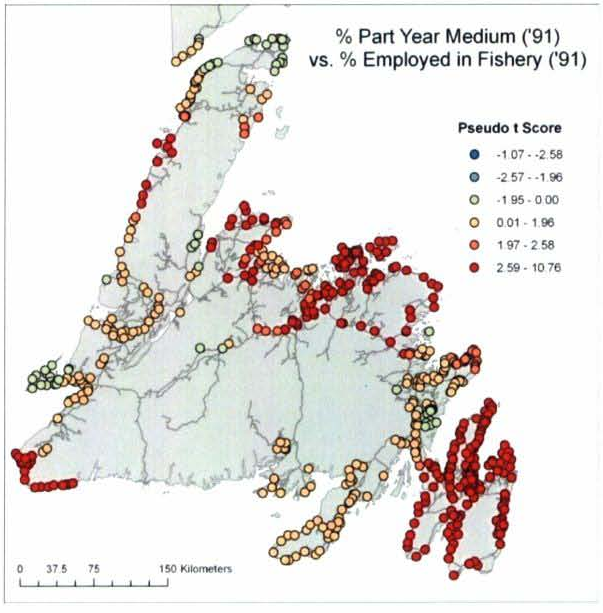

Figure 4.11 Percent Part Year Medium (Utilization) Model Results, 1991 (pseudo t-score)

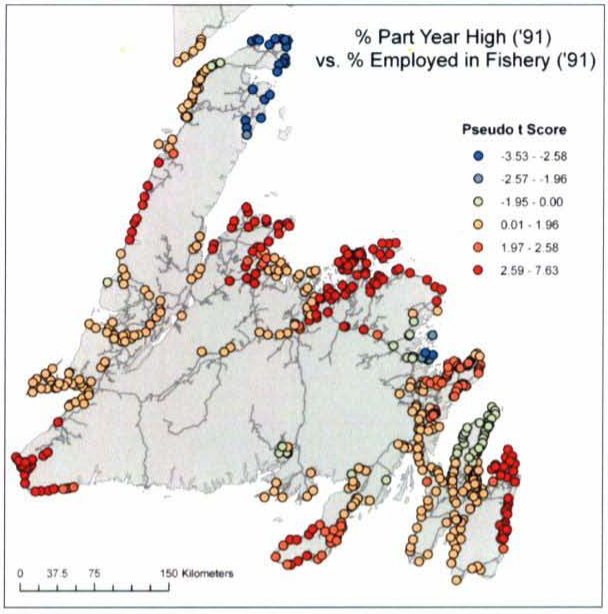

Figure 4.12 Percent Part Year High (Utilization) Model Results, 1991 (pseudo t-score)

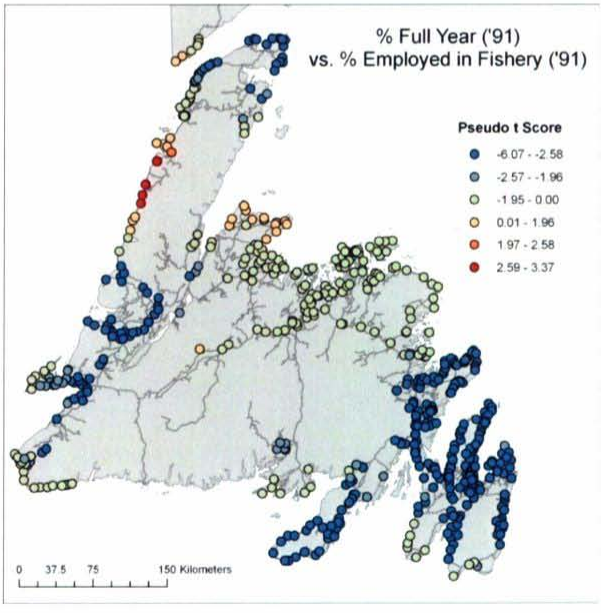

Figure 4.13 Percent Full Year (Utilization) Model Results, 1991 (pseudo t-score)

With the suspension of the cod fishery in 1992, the trend of widespread seasonal employment is expected to reverse for the first few years, when fishers and fish processing workers who formerly worked seasonally are temporarily (at first) placed on government support such as NCARP and TAGS to wait for the cod stocks to rebound. This trend occurs as expected, as shown in **Figures 4.14 – 4.17**. Again, the forces impacting labour market utilization appear to be primarily concerned with coastal workers employed for 'part year medium' and 'part year high', because these are the cohorts that experience the most widespread significant declines (pseudo t < -1.96) in relation to the 'did not work' cohort from 1992-1996.

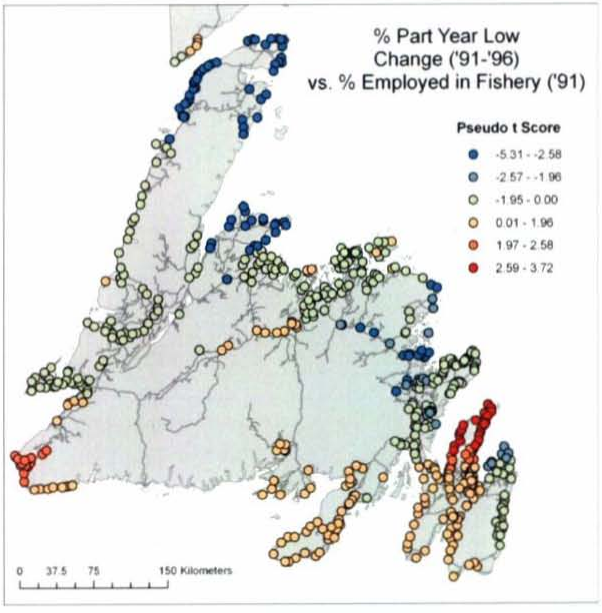

Figure 4.14 Percent Part Year Low (Utilization) Model Results, 1991-1996 (pseudo t-score)

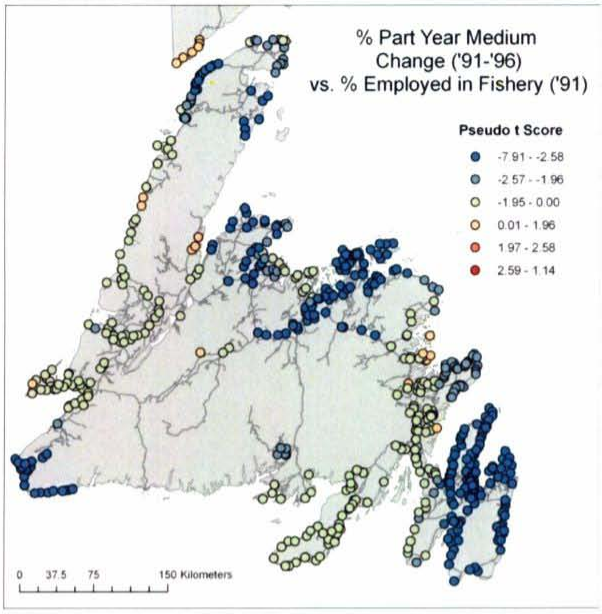

Figure 4.15 Percent Part Year Medium (Utilization) Model Results, 1991-1996 (pseudo t-score)

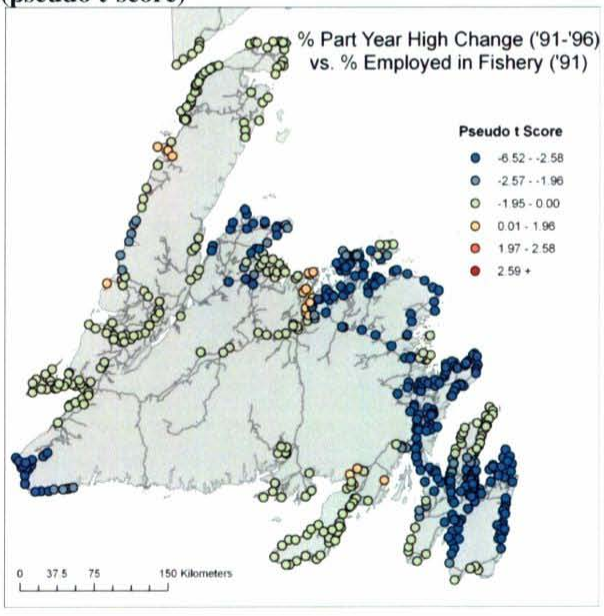

Figure 4.16 Percent Part Year High (Utilization) Model Results, 1991-1996 (pseudo t-score)

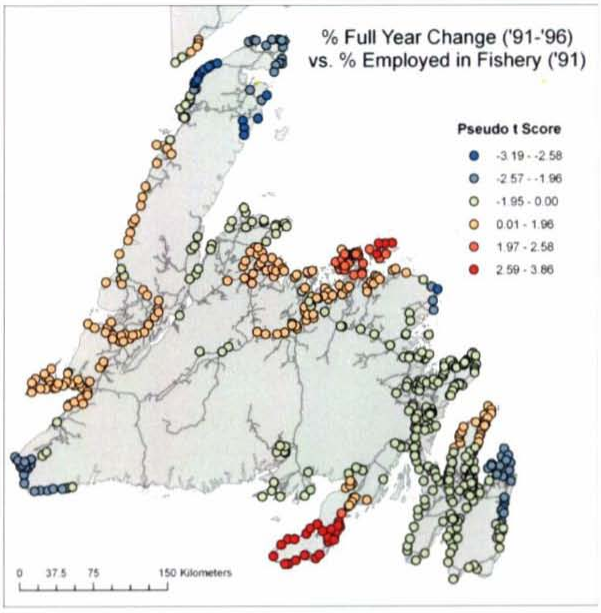

Figure 4.17 Percent Full Year (Utilization) Model Results, 1991-1996 (pseudo t-score)

Finally, as the initial support programs run out in the mid-1990's, the proportion of workers not working in coastal communities is expected to decline. Theoretically, the lack of continued government support, coupled with the rise in availability of seasonal employment in alternative (but relatively capital intensive) fisheries such as crab and shrimp (Hamilton and Butler, 2001) would cause some workers to resume their previous seasonal employment patterns, and others to outmigrate. Once again, as shown in **Figures 4.18-4.21**, this pattern emerges as expected, as there are significantly positive (pseudo $t > 1.96$) relationships between pre-moratorium fishery dependence and the change in percentage of workers working 'part year medium', 'part year high' and 'full year' in a number of (predominantly coastal) communities.

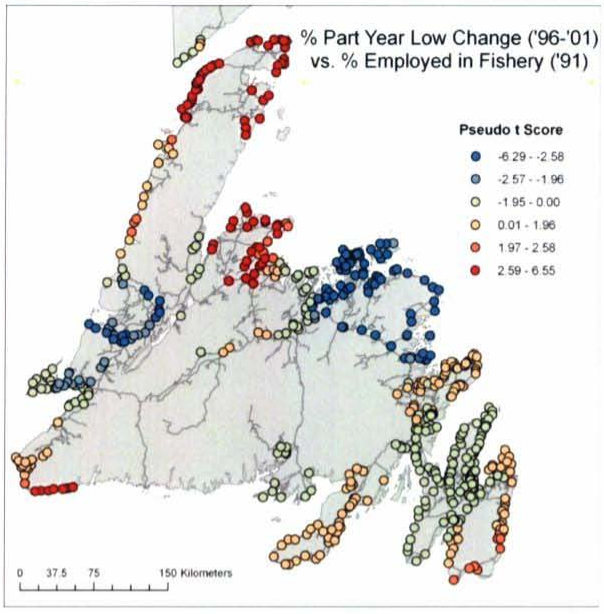

Figure 4.18 Percent Part Year Low (Utilization) Model Results, 1996-2001 (pseudo t-score)

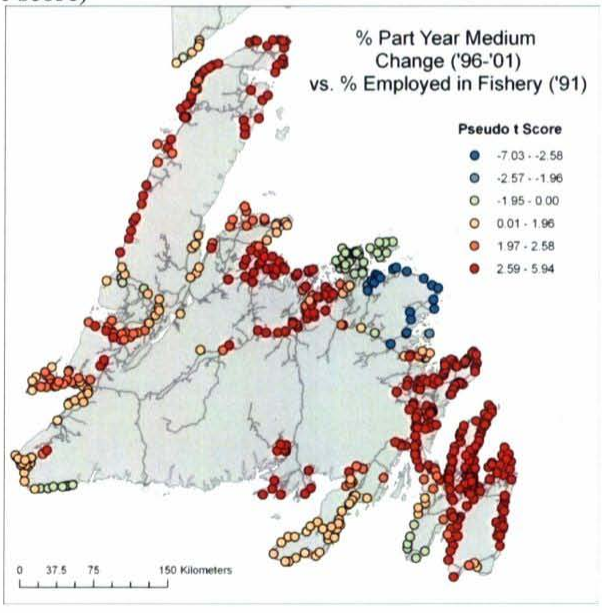

Figure 4.19 Percent Part Year Medium (Utilization) Model Results, 1996-2001 (pseudo t-score)

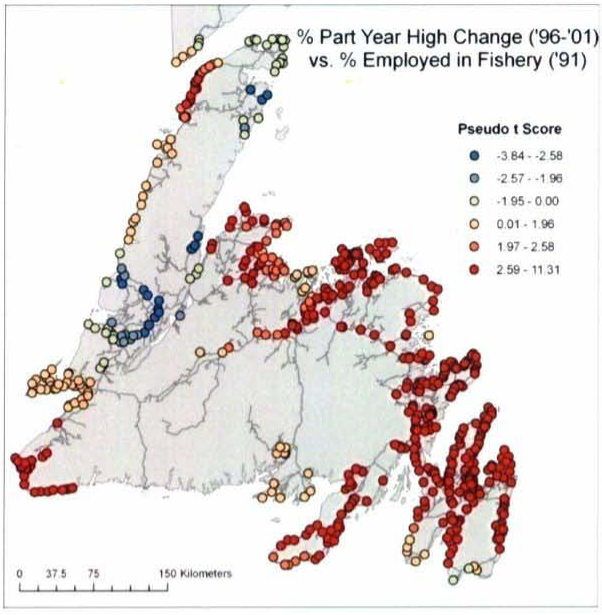

Figure 4.20 Percent Part Year High (Utilization) Model Results, 1996-2001 (pseudo t-score)

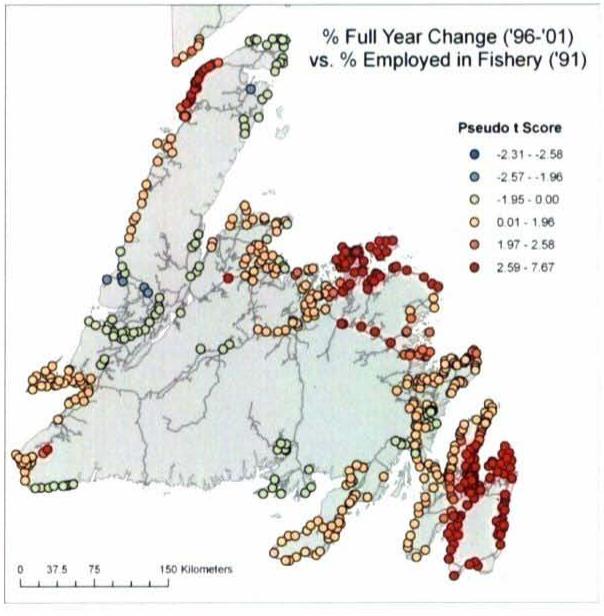

Figure 4.21 Percent Full Year (Utilization) Model Results, 1996-2001 (pseudo t-score)

The model to test for quality changes over time in relation to pre-moratorium fishery dependence is structured slightly differently than the model for utilization changes. Instead of testing for changes in a single variable, the quality model is used to test for changes in two quality indicators: age structure and education (skill) structure. Generally, the theoretical expectations for changes in age structure are fairly straight forward. The combination of the fishery's history of comparatively low wages and the fact that young people are relatively more mobile than older people means that the fishery in most areas is expected to have been ageing prior to the cod moratorium. While not as consistently significant as utilization-related relationships, many coastal areas do indeed feature positive (pseudo $t > 0$) relationships between percentage employed in the fishery in 1991 and the proportion of the workforce that is 'middle aged' and the proportion of the workforce that is 'older workers' in 1991 (versus the proportion of the workforce that is 'younger workers') (see **Figures 4.22** and **4.23**).

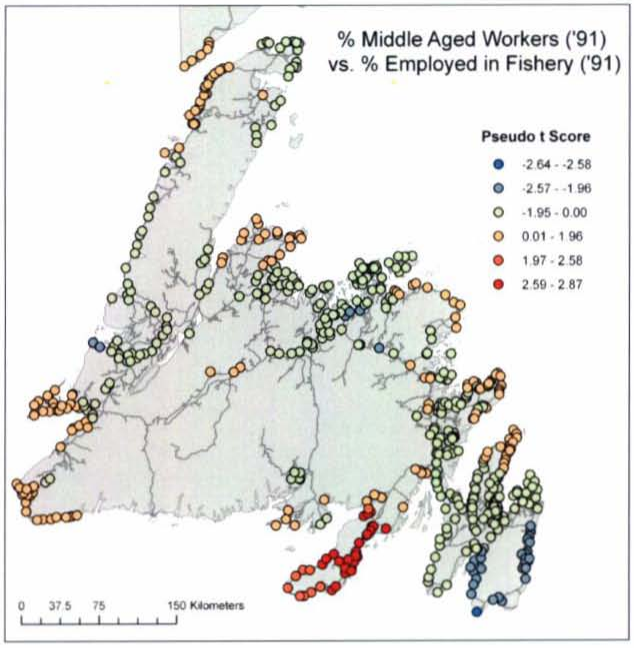

Figure 4.22 Percent Middle Age (Quality) Model Results, 1991 (pseudo t-score)

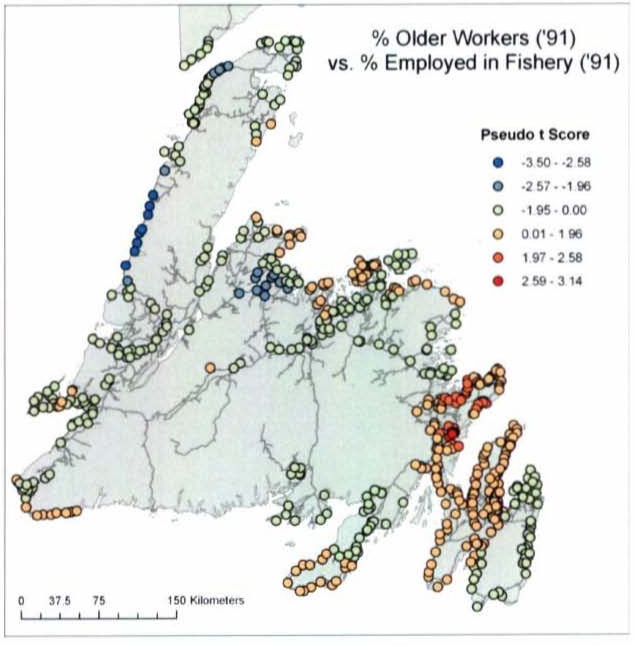

Figure 4.23 Percent Older Age (Quality) Model Results, 1991 (pseudo t-score)

After the fishery closure in 1992, the age structure is expected to change in a different manner than the utilization structure. While a downward shift in labour market utilization may be a temporary adjustment which rebounds after a short period, a change in age structure requires a more permanent decision, and a longer period of time to rebound. Practically, this means that after all shifts in labour market utilization have been accounted for, any downward shifts in labour market quality would represent the deeper and more permanent effects of the labour market shock. Previous studies have demonstrated that young people are generally more mobile than older people (Goetz, 2001), so in the event of a labour market shock, young people are expected to leave first, followed by older people. In this thesis, this means that the extent of damage suffered by a labour market will be gauged by the relative loss of middle aged and older workers in the years following the fishery closure. It is expected that, given their initial conditions, only the most rural areas will lose relatively more older workers in the initial (1991-1996) time period. However, it is also expected that the relative predominance of middle aged and older worker outmigration will extend into more urban-adjacent fishery-dependent communities and labour markets once the government support programs run out during the second (1996-2001) time period.

As anticipated, the empirical relationships between initial fishery dependence and proportion of both middle-aged and older workers in the initial 1991-1996 time period are only significantly negative (pseudo $t < -1.96$) for a few remote coastal communities (see **Figures 4.24 -4.25**). In the later 1996-2001 time period, however, the range of communities with significantly negative relationships expand to include much of central and northeastern Newfoundland (see **Figures 4.26-4.27**).

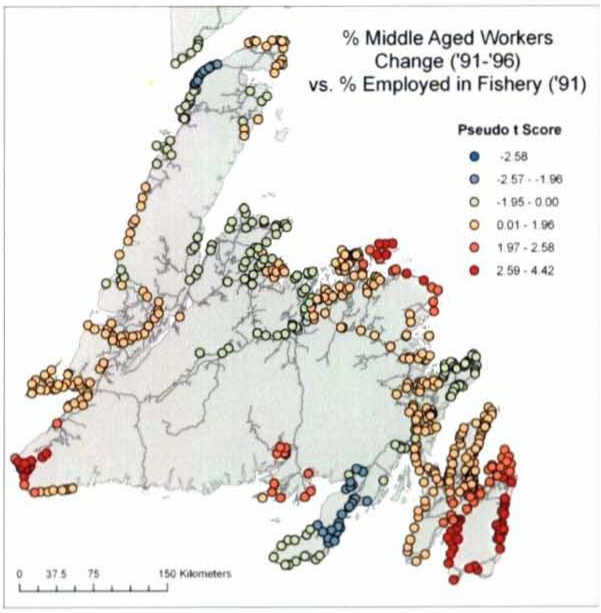

Figure 4.24 Percent Middle Age (Quality) Model Results, 1991-1996 (pseudo t-score)

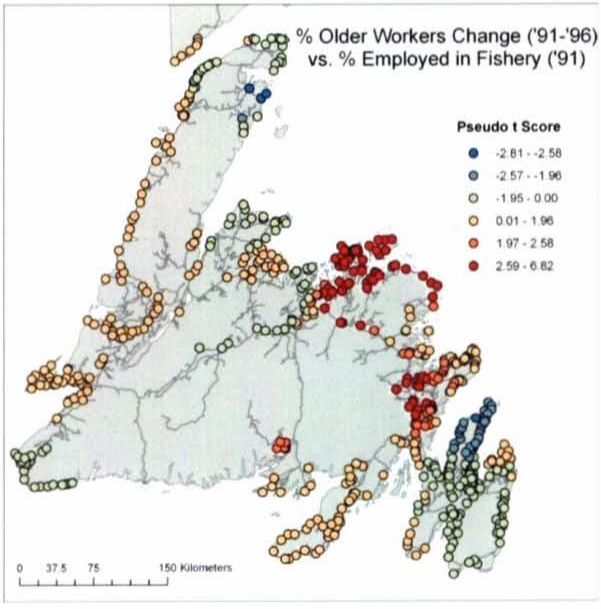

Figure 4.25 Percent Older Age (Quality) Model Results, 1991-1996 (pseudo t-score)

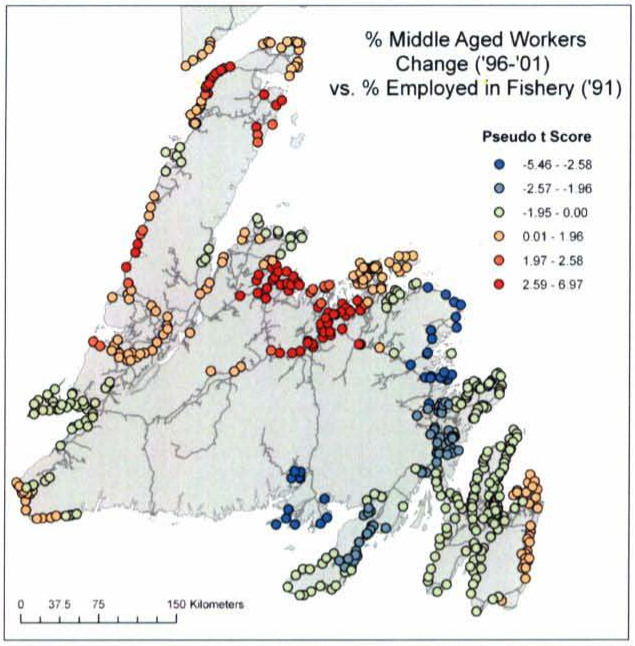

Figure 4.26 Percent Middle Age (Quality) Model Results, 1996-2001 (pseudo t-score)

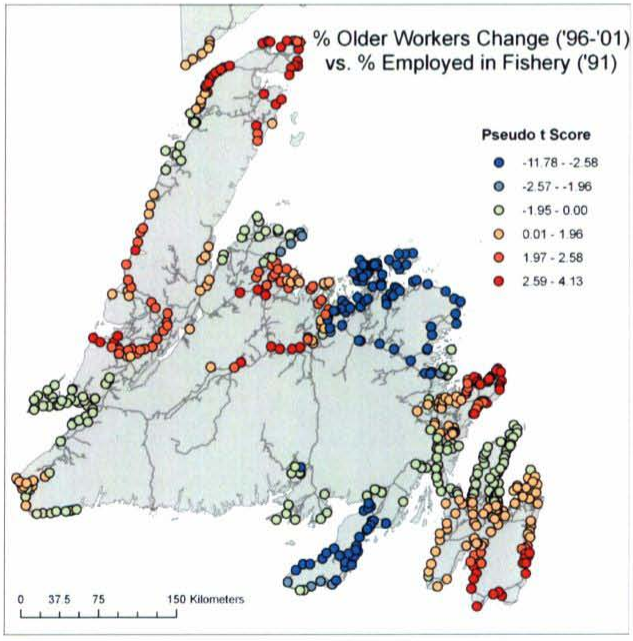

Figure 4.27 Percent Older Age (Quality) Model Results, 1996-2001 (pseudo t-score)

The results of the tests concerning the second quality indicator, education level, differ from those concerning age structure. Prior to the moratorium, it is expected that high association with the fishery would be positively correlated with the lowest education cohort, 'less than a high school diploma' because of the low entrance requirements. Therefore, the closure of the cod fishery in 1992 would be expected to impact the lowest educated members of the communities the most. According to studies conducted in other areas, however, mobility generally increases with education (Goetz, 2001). Any net change in education levels in affected communities would therefore be the result of separate push and pull factors on higher and lower educated workers. Unlike the forces acting upon different age cohorts, the forces acting upon education cohorts are non-substitutable between workers at different levels. As a result, the spatial pattern of the relationship between 1991 fishery dependence and the percentage change in different education levels is expected to be unclear and unsystematic during the periods following the fishery closure.

The empirical results confirm the predictions of the theoretical model, with many communities initially featuring significantly negative (pseudo $t < -1.96$) relationships between pre-moratorium fishery dependence and proportion of workforce with either a high school diploma (see **Figure 4.28**) or a college or university education (see **Figure 4.29**) The observed relationships between pre-moratorium fishery dependence and post-moratorium workforce education change also correspond to the theoretical predictions, as little systematic change is detected in either the 1991-1996 (see **Figures 4.30 - 4.31**) or the 1996-2001 (see **Figures 4.32 – 4.33**) periods.

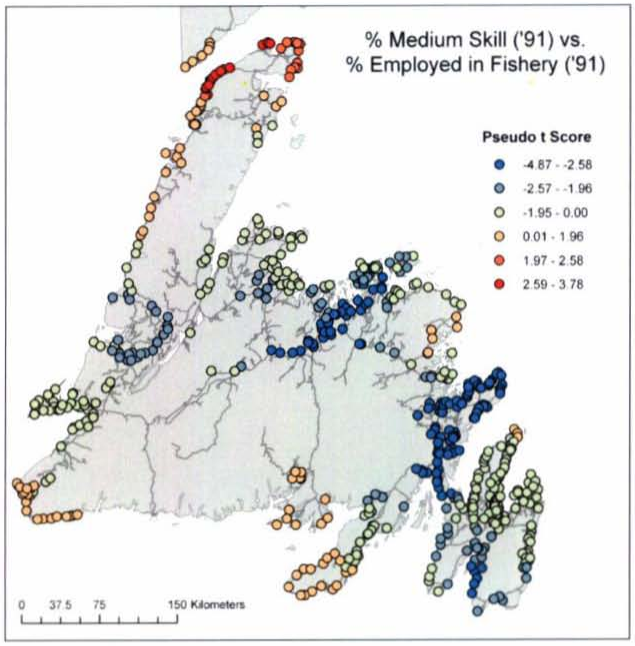

Figure 4.28 Percent Medium Skill (Quality) Model Results, 1991 (pseudo t-score)

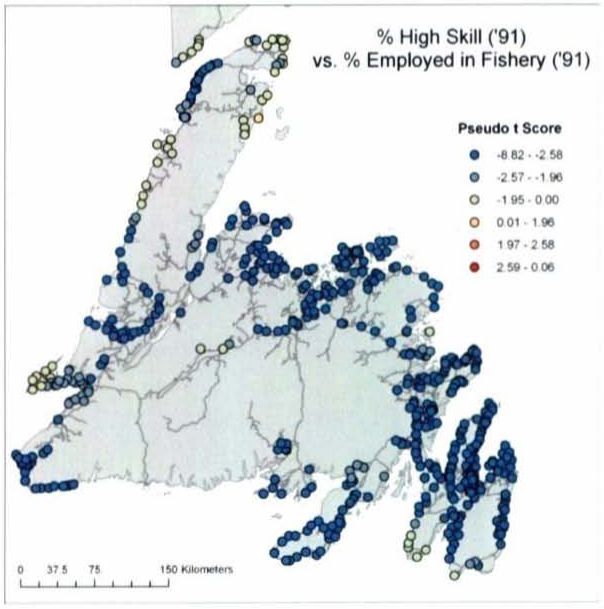

Figure 4.29 Percent High Skill (Quality) Model Results, 1991 (pseudo t-score)

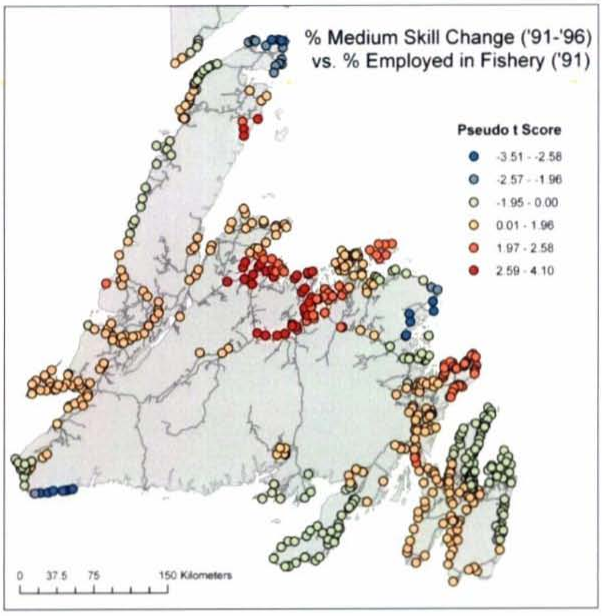

Figure 4.30 Percent Medium Skill (Quality) Model Results, 1991-1996 (pseudo t-score)

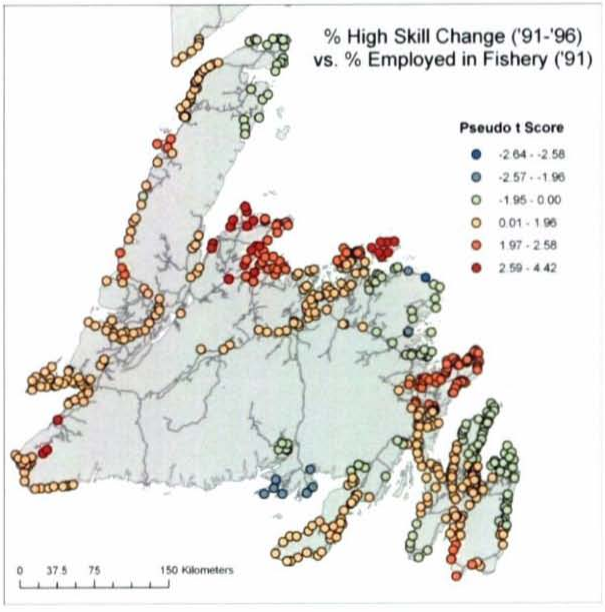

Figure 4.31 Percent High Skill (Quality) Model Results, 1991-1996 (pseudo t-score)

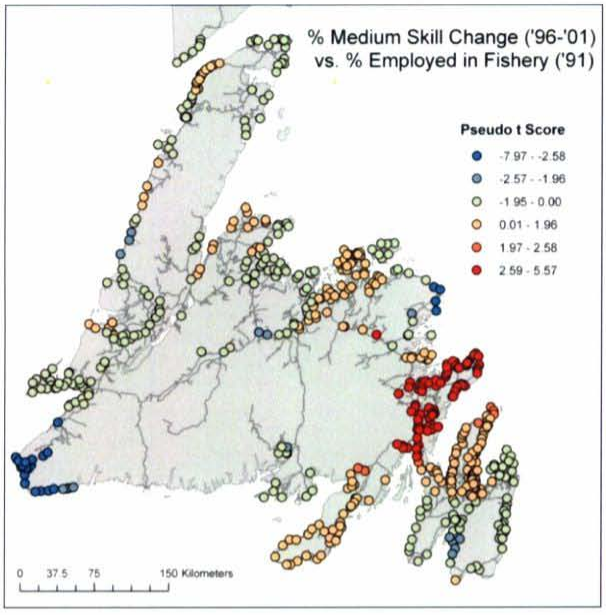

Figure 4.32 Percent Medium Skill (Quality) Model Results, 1996-2001 (pseudo t-score)

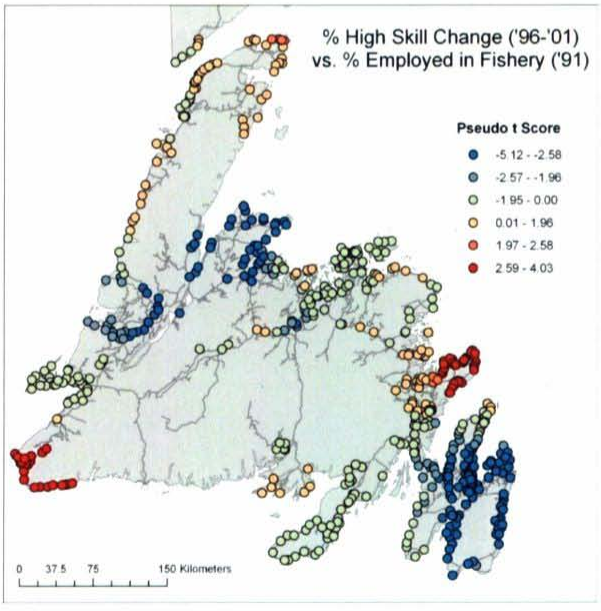

Figure 4.33 Percent High Skill (Quality) Model Results, 1996-2001 (pseudo t-score)

4.3.2 – Example

As an empirical endeavour, one of the goals of this thesis is to be relevant to the real world. While the maps in section 4.3.1 do succeed in displaying the results of the analysis in an efficient manner, it is difficult to compile them all and generate a comprehensive narrative that, overall, supports or refutes the theoretical model. To assist in this role, the results of two rural communities, one remote (St. Bride's) and one urban-adjacent (Fermeuse) (**Figure 4.34**) will now be compared and analyzed (**Table 4.8**). While they are located different distances from the nearest major urban centre (167 and 89 km from St. John's, respectively), both St. Bride's and Fermeuse are located in the same census division (Statistics Canada (n.d.b)), and have similar 1991 populations of 586 and 505 respectively (Statistics Canada (n.d.b)).

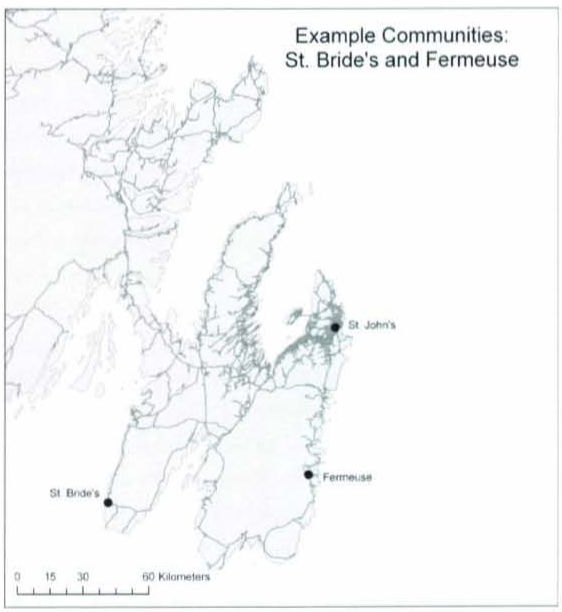

Figure 4.34 Example Communities

Table 4.8 Community Comparison Results (pseudo-t score)

	Pseudo t-Score	Fermeuse	St. Bride's	Fermeuse	St. Bride's	Fermeuse	St. Bride's
		1991		1991-1996		1996-2001	
Utilization	PartYear Low	-0.18	0.20	-0.97	0.73	2.22	-0.34
	PartYear Med	6.81	3.21	-7.82	-1.76	5.76	-0.95
	PartYear High	2.76	1.53	-3.38	-1.41	5.34	1.52
	WorkFull Year	-3.20	-0.29	-1.43	-1.16	6.54	0.04
Quality	MidAge	-2.28	-1.21	3.91	1.16	0.79	-0.44
	OldAge	-1.23	0.15	-1.30	0.35	3.35	-0.07
	MedSkill	-2.03	-1.59	-0.41	1.78	-1.61	-1.35
	HighSkill	-4.60	-1.13	0.34	-0.50	-5.08	-1.40

During the initial (1991) time period, Fermeuse obtained significant ($t > 1.96$) and relatively more positive scores than St. Bride's with regards to the 'part year medium' ($t = 6.81$ vs. $t = 3.21$) and 'part year high' utilization variables, and significant ($t < 1.96$) and relatively more negative scores with regards to the 'work full year' utilization variable, and the 'middle age', 'medium skill', and 'high skill' quality variables. This means that Fermeuse initially had a younger and lower skilled fishery than St. Bride's, with relatively more people working part of the year (15-48 weeks) than either for the full year or not at all. These results indicate that the fishery was likely more important to the St. Bride's economy, with Fermeuse probably featuring more diversified options for their middle-aged and medium-skilled workforce.

During the first time period after the moratorium (1991-1996), the declines for 'part year medium' workers are negative for both example communities, with Fermeuse

declining relatively more for that statistic (t = -7.82 vs. t = -1.76), as well as 'part year high' (t = -3.38 vs. t = -1.41). Regarding the more permanent quality declines, the relative share of middle-aged workers grows in both communities. However, while the cohort change is significantly positive (t > 1.96) in Fermeuse (t = 3.91), it is not in St. Brides (t = 1.16). These results indicate that while declines in utilization significantly affected both economies initially, outmigration in Fermeuse was centred more on the youngest workers than in St. Bride's, where the situation was more mixed.

Finally, during the 1996-2001 time period, Fermeuse begins to experience a rebound in all of its workforce utilization indicators, with universally significant (t > 1.96) and relatively more positive values than St. Bride's ((t = 2.22 vs. t = - 0.34) for 'part year low', (t = 5.76 vs. t = - 0.95) for 'part year medium', (t = 5.34 vs. t = 1.52) for 'part year high', and (t = 6.54 vs. t = 0.04) for 'worked full year', respectively). Regarding changes in workforce quality, the significant loss of young people experienced by Fermeuse from 1991-1996 does not appear to have spread to the rest of the workforce, because the change in both middle (t = 0.79) and older aged (t = 3.35) workers remains positive. While they are not significant, the pseudo-t scores for age in St. Bride's now dip into negative values (t = -1.35 and -1.40 for middle-aged and older workers, respectively), indicating that the relative share of mature workers in the town's workforce has now begun to drop.

Overall, the results of the comparative example verifies, on a local level, the hypotheses generated in the theoretical model. St. Bride's, being the more remote community, appears to have initially been more dependent upon the fishery than

Fermeuse, and suffered more extensive utilization and quality declines after the cod moratorium. As expected, both communities experienced utilization declines initially, with the less remote community (Fermeuse) rebounding more strongly and suffering less extensive age-biased outmigration in the longer term. Also as expected, quality effects related to worker skill level were mixed with respect to urban proximity.

4.4 – Conclusion

The results of this study were presented in this chapter. First and foremost, it was found that the local, discrete space approach offered by the discrete geographically weighted regression (DGWR) model is better suited for this particular study of labour market effects than either the standard ordinary least squares or continuous GWR model. Using the appropriate statistical tests, it was found that the DGWR approach offers significant improvements over the more traditional methods in both overall model fit and accounting for spatial autocorrelation. Using the DGWR model, it was found that the theoretical model constructed in chapter 3 forms a sound basis for assessing the effects of the fishery closure on Newfoundland and Labrador labour markets, as, for the most part, the dynamics of the theoretical model have been verified empirically. The primary implication of this is that the significance of the dynamics of the province's labour markets can now be assessed, with the assistance of studies in other locations that utilize similar theoretical models.

5.0 – Conclusion

5.1 – Discussion of Results

The Newfoundland and Labrador labour force has endured a great deal of upheaval and change in the last two decades. Through the closure of the once dominant cod fishery, many rural areas of the province have seen the composition of their labour markets completely restructured in less than a generation. However, despite a relatively large amount of qualitative research and anecdotal evidence concerning other aspects of the moratorium (Davis, 2006; Murray *et al.*, 2005), prior to this thesis no systematic effort had been made to quantitatively measure the impact of location on the nature of labour market effects to communities. The discrete geographically weighted regression (DGWR) method provides an ability to study these effects with a combination of local scale and realistic spatial weights that to date has not been possible.

Underlying all empirical results is the fact that the DGWR model outperforms both the ordinary least squares (OLS) and continuous GWR frameworks with the data used in this project. This observation is not surprising, given the superior ability of the DGWR model to more closely approximate the real-world commuting interactions that are fundamental to the theoretical functioning of the province's labour markets. It is evident from their failure to remove spatial dependence that both the non-spatial framework utilized by the ordinary least squares model, and the straight-line distance framework utilized by the continuous GWR model do not sufficiently represent the interactions between labour markets in Newfoundland and Labrador. Therefore, to attempt to interpret the results of either of these models in this or a similar study would be incorrect. Of the three tested models, the DGWR model removes the greatest amount of

significant spatial dependence from the study area, and therefore provides the most realistic modelling environment.

In the theoretical model developed for this thesis, it was hypothesized that the cod moratorium would constitute a demand shock to the province's labour markets, and that remote rural areas would be less sheltered from any potential effects than urban-adjacent rural areas. The order of effects was theorized to progress from short run labour utilization changes to longer run quality (labour movement) shifts as the severity of the moratorium increased. The final results of this modelling process indicate that this scenario did occur as predicted. While most areas experienced severe utilization declines relative to moratorium exposure, the more remote regions experienced deeper cuts to their age structure (i.e. relatively more middle aged workers left the workforce) than urban-adjacent regions. However, the magnitude of these cuts was, as expected, not uniform across all remote regions. The distance decay effect appears to be the strongest in the areas nearest to St. John's, and weaker in central and western Newfoundland labour markets. Whereas the concentration of the cod fishery off the northeast coast of the island likely explains most of this discrepancy (see **Figure 2.4**), literature from the European Union indicates that the degree of outmigration due to downturns in the fishing industry can also be significantly altered by differences in culture and personal preference across regions (Pita *et al.*, 2010).

5.2 – Policy Implications and Conclusions

The structure of this study generates results in two distinct areas. The first area from which conclusions are drawn is through the empirical analysis, which was structured around the hypotheses introduced in section 1.2, and creates practical

implications for policy in the study area and similar locations. Conclusions drawing from this area are as follows:

(i) In general, remote communities were more dependent upon the cod fishery during the year immediately preceding the cod moratorium than urban-adjacent communities.

(ii) Once the moratorium occurred, its effects did, as hypothesized, follow a common trend, with almost universal and immediate utilization declines occurring during the 1991-1996 time period followed by more substantial age structure declines in relatively more remote rural communities from 1996-2001.

(iii) Regarding the quality changes, age effects are more distinct than skill (education) effects in almost every region.

(iv) The observed trends stated in conclusions (i) and (ii) above are generally most applicable to communities located along the northeast coast of the island.

In addition to evaluating the stated hypotheses, conclusions are also drawn from inferences made while conducting the analysis. Some of these conclusions are:

(i) The DGWR model outperforms both an unmodified ordinary least squares model, and a continuous GWR model, with respect to both spatial autocorrelation and traditional model fit statistics.

(ii) The workforce involved in the fishery has a legacy of relative under-utilization, likely due to its history of seasonal work.

(iii) Although the quality declines of the most vulnerable resource extraction industry workforce in many rural areas are likely due to outmigration, the lack of symmetry indicates that the greatest pull locations have been in markets outside the province instead of urban areas in the province (see **Figures 4.24-4.27**).

5.3 – Future Research

The results of this thesis indicate that its primary questions have been resolved. However, throughout the course of conducting this research, a number of additional

questions suitable for future study concerning both the empirical analysis and the modelling framework have been raised. Through the exploration of these questions, the significance of the research detailed in this thesis will be determined, and a context for any supplementary analysis will be provided.

One of the conclusions of this research is that potential appears to exist for the classification of Newfoundland and Labrador labour markets based on their respective proximity to the province's major urban centres. With everything else equal, the relative degree of response in some variables to the demand shock of the cod moratorium generally decreases in severity the closer a community is to St. John's or Corner Brook. The access to larger markets and the increased economies of scale available within regular commuting distance to a larger service centre theoretically provides an advantage to urban-adjacent rural communities that more remote communities do not have access to. Based on this observation, it would seem reckless to design and implement policy that is insensitive to the real costs associated with location. According to this research, however, not all economic variables are as sensitive to location and the forces of agglomeration as others. For example, while it is clear that utilization declines affected almost every area of the province uniformly after the closure of the cod fishery, the effects on age structure were more severe and occurred earlier in the most remote areas. Additional research could investigate this further and attempt to identify the sensitivity of these processes to changes in scale and time.

A potential future area of research related to the classification of labour markets based on the spatially systematic nature of shock effects found in this project concerns the effects themselves. The priority in this thesis has been to evaluate the effect that the

cod moratorium has had on strictly economic aspects of the Newfoundland and Labrador economy in the decade following the initial shock. However, in the economic geography literature, the concepts of regional growth and well-being often have definitions that encompass more than economic or employment aspects (Hudson, 2007). Instead, these studies feature more holistic definitions of regional well-being, and often include aspects of social and cultural development that are difficult to quantify. While beyond the scope of this thesis, an exhaustive investigation into the effects of the cod moratorium on the people of Newfoundland and Labrador would provide some commentary on fishery closure impacts on the qualitative aspects of their lives, to complement the aggregate quantitative indicators. Some examples of qualitative or non-economic changes that could potentially be studied in relation to urban proximity include changes in health status (Murray *et al.*, 2005), risky behaviour (Davis, 2006), cultural identity, and changes in the "social economy" (Hudson, 2007).

A third area in which the analysis explained in this thesis could potentially be expanded is to include additional time periods. Although the model as it exists captures the initial conditions and the most direct effects of the cod moratorium, it would be interesting to extend the analysis to time periods both before the moratorium occurred, and after this analysis ends. The inclusion of time periods prior to 1991 in the analysis would help establish the context for the moratorium, by determining whether the demand shock associated with the closure of the cod fishery was a single, sudden event, or representative of the final step of a long, drawn out process of decline. On the other hand, evaluating effects after 2001 would enable the analysis to gain a more long-term

perspective with respect to factor movements and labour market adjustments, which would allow for the generation of more long-term policy implications.

A further potential avenue of research stemming from this project would be to apply the discrete GWR (DGWR) model to other situations, thereby evaluating its versatility and its stability. At this stage, it is essential to ascertain the general conditions in which a DGWR would be the best modelling option, and those in which it would not. To do this, the model should be applied to labour market behaviour in different political and cultural systems, as well as to phenomena that operate on different scales than commuting behaviour. However, the nature of relationships being modelled would need to be known beforehand, because to this point a method of optimizing bandwidth size with only a weight matrix has not been developed.

One final area of potential future research identified through the results of this thesis involves the spatial relationships being modelled. Although the DGWR model as used in this analysis is superior to existing ordinary least squares and continuous GWR modelling techniques according to the statistics used in this thesis, it still may be able to be improved. It is evident from the presence of heteroskedasticity in a number of the model runs that variance in the error terms is not necessarily constant throughout the study area. According to Anselin, (1988), this spatial heterogeneity could be caused by the presence of different spatial effects, such as "central place hierarchies, the existence of leading and lagging regions, vintage effects in urban growth, etc…", all operating in the same region. As a result, to be considered exhaustive, an analysis would be required to have any number of spatial effects, all operating at potentially different scales, incorporated within it. To accomplish this, models that allow hierarchical analysis, such

as hierarchical linear modelling, could be used. Hypothetically, the model could be constructed in the same manner as described in this thesis, but with spatial effects that operate on spatial scales different from commuting behaviour potentially included in the analysis. Any improvements over the model executed in this thesis would likely be minor in most respects, however, because of the theorized predominance of commuting in determining labour market interaction in the study area.

Although this thesis has raised a significant number of additional potential research avenues, this fact alone does not mean that the analysis contained within is lacking. The procedure developed to incorporate distance metrics alternative to continuous space represents a significant step toward the creation of models that are more meaningful and locally relevant than existing methods. In addition, the identification of spatially heterogeneous relationships between exposure to the 1992 cod moratorium in Newfoundland and Labrador and changes in regional age structure provides quantitative proof of the need for locally distinct labour market policies that recognize the importance of urban proximity.

6.0 – Bibliography

Anselin, L. (1995). Local Indicators of Spatial Association – LISA. *Geographical Analysis, 27*(2), 93-115.

Anselin, L. (1988). *Spatial Econometrics: Methods and Models*. Kluwer Academic, Boston, MA.

Anselin, L., & Griffith, D. (1988). Do spatial effects really matter in regression analysis? *Papers of the Regional Science Association,* 65, 11-34.

Baker,M., Benjamin, D., & Stanger, S. (1999). The highs and lows of the minimum wage effect: A time-series cross-section study of the Canadian law. *Journal of Labor Economics,* 17 (2), 318-348.

Beine, M., Docquier, F., & Rapoport, H. (2008). Brain drain and human capital formation in developing countries: Winners and losers. *The Economic Journal*, 118, 631-652.

Bhatta, S., & Lobo, J. (2000). Human capital and per capita product: A comparison of US states. *Papers in Regional Science*, 79(4), 393.

Bound, J., & Holzer, H. (2000). Demand shifts, population adjustments, and labor market outcomes during the 1980s. *Journal of Labor Economics*, 18(1), 20-54.

Brunow, S., & Hirte, G. (2009). The age pattern of human capital and regional productivity: A spatial economic study on German regions. *Papers in Regional Science*, 88(4), 799-823.

Brunow, S., & Hirte, G. (2006). Age structure and regional economic growth. *Jahrbuch Für Regionalwissenschaft*, 26(1), 3-23.

Brunsdon, C., Fotheringham, S., & Charlton, M. (1998). Geographically weighted regression – modelling spatial non-stationarity. *The Statistician*, 47(3), 431-443.

Cliff, A., & Ord, K. (1970). Spatial autocorrelation: A review of existing and new measures with applications. *Economic Geography*, 46, 269-292.

Coffey, W., & Shearmur, R. (2002). Agglomeration and dispersion of high-order service employment in the Montreal metropolitan region, 1981-96. *Urban Studies*, 39(3), 359-378.

Davis, R. (2006). All or nothing: Video lottery terminal gambling and economic restructuring in rural Newfoundland. *Identities: Global Studies in Culture and Power*, 13, 503-531.

Decressin, J., & Fatas, A. (1995). Regional labor market dynamics in Europe. *European Economic Review*, 39, 1627-1655.

Engelen, G., White, R., Uljee, I., & Drazan, P. (1995). Using cellular automata for integrated modelling of socio-environmental systems. *Environmental Monitoring and Assessment*, 34, 203-214.

Feehan, J. (1991). Learning from studies of Newfoundland's economic problems. *Newfoundland Studies,* 7(1), 1-22.

Feyrer, J. (2007). Demographics and productivity. *The Review of Economics and Statistics*, 89(1), 100.

Fingleton, B., (2004). Some alternative geo-economics for Europe's regions. *Journal of Economic Geography*, 4, 389-420.

Fotheringham, A. S. (1998). Trends in quantitative methods II: stressing the computational. *Progress in Human Geography*, 22(2), 283-292.

Fotheringham, A., Brunsdon, C. & Charlton, M. (2002). *Geographically weighted regression: the analysis of spatially varying relationships*. West Sussex: John Wiley & Sons, Ltd.

Gabriel, S., Shack-Marquez, J., & Wascher, W. (1993). Does migration arbitrage regional labor market differentials? *Regional Science and Urban Economics*, 23, 211-233.

Glendon, S. & Vigdor, J. (2003). Thy neighbor's jobs: geography and labor market dynamics. *Regional Science and Urban Economics,* 33, 663-693.

Goetz, S. (2001). Migration and local labor markets. In: Loveridge, S. (Ed.), *The Web Book of Regional Science.* Regional Research Institute, West Virginia University, Morgantown, WV. Retrieved April 12, 2010, from http://www.rri.wvu.edu/WebBook/Goetz/contents.htm

Government of Newfoundland and Labrador. (2010). *Government Transfers by Source.* Retrieved June 26, 2010 from: http://www.communityaccounts.ca/communityaccounts/onlinedata/display_table.asp?_=0bfAjIydpaWrnbSTh5-FvKFteKWik8e4tWSpyZp4WIs_

Government of Newfoundland and Labrador. (2009). *Census 1986-1996: Migration.* Retrieved July 24, 2009 from: http://www.communityaccounts.ca/communityaccounts/onlinedata/display_table.asp?p=%D1%B7%C0%8C%8C%9D%A5%A5%AB%9D%B4%93%87%9F%C5%BC%9Cmx%A5%A2%93%C7%B8%B5d%B6%CC%9C%BAY

Government of Newfoundland and Labrador. (2008). *Optimal to Flows 2006* [Data File].

Griffith, D., and Paelinck, J. (2007). An equation by any other name is still the same: on spatial econometrics and spatial statistics. *Annals of Regional Science*, 41, 209-227.

Hamilton, L., & Butler, M. (2001). Outport adaptations: Social indicators through Newfoundland's cod crisis. *Research in Human Ecology*, 8 (2), 1-11.

Henderson, V. (1997). Externalities and industrial development. *Journal of Urban Economics*, 42(3), 449-470.

Huang, Y., & Leung, Y., (2002). Analysing regional industrialisation in Jiangsu province using geographically weighted regression. *Journal of Geographical Systems,* 4, 233-249.

Hudson, R. (2007). Regions and regional uneven development forever? some reflective comments upon theory and practice. *Regional Studies*, 41(9), 1149-1160.

Johansson, B., Klaesson, J., & Olsson, M. (2002). Time distances and labor market integration. *Papers in Regional Science,* 81, 305-327.

Jones, M. (2004). The dynamic benchmarking of labour markets. *Regional Studies,* 38(5), 493-504.

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

Korpi, T. (1997). Is utility related to employment status? Employment, unemployment, labor market policies and subjective well-being among Swedish youth. *Labour Economics,* 4, 125-147

Krugman, P. (1991). Increasing Returns and Economic Geography. *Journal of Political Economy*, 99(3), 483-499.

LeSage, J. (1999). Spatial Econometrics. In: Loveridge, S. (Ed.), *The Web Book of Regional Science.* Regional Research Institute, West Virginia University, Morgantown, WV. Retrieved May 18, 2010, from http://www.rri.wvu.edu/WebBook/LeSage/etoolbox/index.html

Leung, Y., Mei, C., & Zhang, W. (2000). Statistical tests for spatial nonstationarity based on the geographically weighted regression model. *Environment and Planning A*, 32, 9-32.

Lindh, T., & Malmberg, B. (1999). Age structure effects and growth in the OECD, 1950-1990. *Journal of Population Economics*, 12, 431-449.

Martin, R., (2000). Local labour markets: Their nature, performance, and regulation. In: Clark, G., Feldman, M., & Gertler, M. (Eds.), *The Oxford handbook of economic geography*. Oxford University Press, Oxford, 455-476.

Martin, R., (1999). The new 'geographical turn' in economics: some critical reflections. *Cambridge Journal of Economics*, 23, 65-91.

Martin, R., & Sunley, P. (2007). Complexity thinking and evolutionary economic geography. *Journal of Economic Geography,* 7, 573-601.

Mathur, V., & Song, F. (2000). A labor market based theory of regional economic development. *The Annals of Regional Science*, 34, 131-145.

Mejean, I., & Patureau, L. (2010). Firms' location decisions and minimum wages. *Regional Science and Urban Economics,* 40, 45-59.

Mulligan, G. (1984). Agglomeration and central place theory: A review of the literature. *International Regional Science Review*, 9 (1), 1-42.

Murray, M., MacDonald, D., Simms, A., Fowler, K., Felt, L., Edwards, A., & Gates, K. (2005). *Community resilience in Newfoundland: The impact of the cod moratorium on health and social well-being*. St. John's, NL: Memorial University of Newfoundland.

Nickell, S., & Bell, B. (1995). The collapse in demand for the unskilled and unemployment across the OECD. *Oxford Review of Economic Policy*, 11(1), 40-62.

Overman, H. G. (2004). Can we learn anything from economic geography proper? *Journal of Economic Geography*, 4(5), 501-516.

Perry, R., & Ommer, R. (2003). Scale issues in marine ecosystems and human interactions. *Fisheries Oceanography*, 12(4), 513-522.

Pita, C., Dickey, H., Pierce, G., Mente, E., & Theodossiou, I. (2010). Willingness for mobility amongst European fishermen. *Journal of Rural Studies,* 26(3), 308-319.

Ravenstein, E., (1885). The Laws of Migration. *Journal of the Statistical Society of London*, 48(2), 167-235.

Rey, S., & Montouri, B. (1999). US regional income convergence: A spatial economic perspective. *Regional Studies,* 33 (2), 143-156.

Romer, P. (1994). The origins of endogenous growth. *The Journal of Economic Perspectives*, 8(1), 3-22.

Rupasingha, A., Goetz, S., & Freshwater, D. (2000). Social capital and economic growth: A county-level analysis. *Journal of Agricultural and Applied Economics*, 32(3), 565-572.

Santé, I., García, A., Miranda, D., & Crecente, R. (2010). Cellular automata models for the simulation of real-world processes: A review and analysis. *Landscape and Urban Planning*, 96, 108-122.

Schrank, W. E. (2005). The Newfoundland fishery: Ten years after the moratorium. *Marine Policy*, 29(5), 407-420.

Sinclair, P. (1996). Sustainable development in fisheries dependent regions? Reflections on Newfoundland cod fisheries. *Sociologia Ruralis*, 36(2), 224-235.

Statistics Canada. (n.d.a). *Selected population characteristics, Canada, provinces and territories, every 5 years (number unless otherwise noted).* Retrieved May 18, 2010, from the CANSIM database, through the ESTAT distributor: http://estat.statcan.gc.ca/cgi-win/cnsmcgi.pgm?regtkt=&C2Sub=&LANG=E&C2DB=EST&C2USER=&C2PASS=&C2APASS=&C2USEWRK=&SDDSLOC=//www.statcan.gc.ca/imdb-bmdi/*.htm&ROOTDIR=ESTAT/&VEC=&RESULTTEMPLATE=ESTAT/CII_ASUM&ARRAY_SUMM=1&SDDSID=&DRILLFILE=&ARRAYID=1530037

Statistics Canada. (n.d.b). *1991 (2A) basic questionnaire, Provinces to Municipalities.* Retrieved May 18, 2010, from the ESTAT distributor: http://estat.statcan.gc.ca/cgi-win/CNSMCGI.PGM?Lang=E&C91SubDir=ESTAT\&DBSelect=PSD2A

Tesfatsion, L. (2002). Agent-based computational economics: Growing economies from the bottom up. *Artificial Life*, 8, 55-82.

Varga, A. (2000). Local academic knowledge transfers and the concentration of economic activity. *Journal of Regional Science*, 40(2), 289-309.

Velleman, P., & Welsch, R. (1981). Efficient computing of regression diagnostics. *The American Statistician*, 35(4), 234-242.

Wang, Q., Ni., J., & Tenhunen, J. (2005). Application of geographically-weighted regression analysis to estimate net primary production of Chinese forest ecosystems. *Global Ecology and Biogeography*, 14, 379-393.

Wensley, M., & Stabler, J. (1998). Demand-threshold estimation for business activities in rural Saskatchewan. *Journal of Regional Science*, 38(1), 155-177.

Wheeler, D. (2007). Diagnostic tools and a remedial method for collinearity in geographically weighted regression. *Environment and Planning A,* 39, 2464-2481.

White, H. (1980). A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica*, 48(4), 817-838.

Wimberly, M., Yabsley, M., Baer, A., Dugan, V., & Davidson, W. (2008). Spatial heterogeneity of climate and land-cover constraints on distributions of tick-borne pathogens. *Global Ecology and Biogeography*, 17 (2), 189-202.