UNIVERSITY OF SOUTHAMPTON

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS

School of Civil Engineering and the Environment

Forecasting the Use of New Local Railway Stations and Services Using GIS

by

Simon Philip Blainey

Thesis for the degree of Doctor of Philosophy

December 2009

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS

SCHOOL OF CIVIL ENGINEERING AND THE ENVIRONMENT

Doctor of Philosophy

FORECASTING THE USE OF NEW LOCAL RAILWAY STATIONS AND SERVICES USING GIS

by Simon Philip Blainey

The aim of this thesis is to develop an integrated methodology for investigating the potential for new local railway stations within a given area, with particular emphasis on the use of Geographical Information Systems (GIS). Existing methods for assessing the case for constructing new local railway stations have often been found wanting, with the forecasts produced proving to be inaccurate.

A review of previous work in this field has been undertaken and methodologies with the potential to enhance local rail demand models have been identified. Trip rate and trip end models have been developed which are capable of forecasting usage at new station sites anywhere in England and Wales. Geographically Weighted Regression (GWR) has been used to enhance the performance of these models and to account for local variations in the effects of explanatory variables on rail demand. Flow level models have been produced for stations in South-East Wales, with a range of model formulations tested. A survey of ultimate passenger trip origins and destinations was carried out in the same area, enabling the accuracy of theoretical station catchment definition methods to be tested.

A GIS-based procedure for locating potential sites for new railway stations within a given area has been developed. This was combined with the results from the demand models and estimates of associated costs and benefits to give a synthesised appraisal procedure capable of assessing the case for constructing particular stations. This procedure was applied to 14 sites in South-East Wales and, along with trip end forecasts for 421 sites across the country, this indicated that there is almost certainly a positive case for constructing a significant number of new railway stations in the UK.

Table of Contents

List of Tables

List of Figures

DECLARATION OF AUTHORSHIP

I, Simon Philip Blainey, declare that the thesis entitled 'Forecasting the Use of New Local Railway Stations and Services Using GIS' and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly while in candidature for a research degree at this University;
- no part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published as:

    - Blainey S P & Preston J M (2010) Modelling Local Rail Demand in South Wales, *Journal of Transportation Planning and Technology*, 33(1):55-73.
    - Blainey S P (2010) Trip End Models of Local Rail Demand in England and Wales, *Journal of Transport Geography*, 18(1):153-165.
    - Blainey S P & Preston J M (2009) *Assessing the Potential Performance of New Local Railway Stations*, European Transport Conference, Leiden, 5-7 October 2009
    - Blainey S P (2009) *Modelling Local Rail Demand in South Wales*, 41st University Transport Studies Group Conference, London, 5-7 January 2009.
    - Blainey S P (2008) *Some Issues in Modelling New Local Rail Services*, 40th University Transport Studies Group Conference, Portsmouth, 3-5 January 2008.

Signed: ……………………………………………………………………………

Date:……………………………………………………………………………

Acknowledgements

This project was financed by an Engineering and Physical Sciences Research Council DTA (Doctoral Training Account) Studentship as part of Rail Research UK.

Thanks are due to a number of people without whom this thesis could not have been produced. Firstly I must thank my supervisor, Prof. John Preston, for his continued help, advice and support, and for coming up with the original idea for the project. Prof. Corinne Mulley introduced me to the wonders of GWR and explained some of the more esoteric aspects of the GWR software. Dr John Armstrong very kindly produced custom-designed Perl scripts to extract frequency and journey time data from CIF files and advised me on their use. Doug Robinson helped out with IT issues on a number of occasions. Several staff members at ITS Leeds, particularly Dr William Lythgoe, provided advice on various aspects of rail demand modelling. Derek Thompson and Danny Gilbert at Arriva Trains Wales supplied LENNON data for the project along with permission and free travel for the Rhymney Line travel survey. Ian Palmer at Greater Manchester PTE provided GMATS survey data and gave advice on its use. David Ford kindly allowed me access to facilities for backing up my files, providing a safety-net in case of computing disaster.

I am extremely grateful to the wide range of groups who have given me the opportunity to present and discuss my work, including (in no particular order) the University Transport Studies Group, the ATOC Passenger Demand Forecasting Council, Rail Research UK, the British Science Association, Science Oxford, Oxford University Transport Studies Unit, the Monday Evening Group, and my own research group, TRG. Thanks are due to everyone at TRG for their friendship and support over the last three years, and for providing a pleasant working environment when I have been in Southampton. Of course, having spent much of the last three years working from home I must also thank my housemates over that time, particularly Chris and Katherine, for providing just the right amount of distractions to allow this thesis to be written with the minimum of stress. All errors which remain in it are entirely my own responsibility.

Finally, and most importantly, a huge thankyou is due to my parents, to Anna and to Rachel not only for taking an interest in rail demand modelling, but also for reminding me in the best possible way that there are more important things in life.

Chapter One: Introduction

1.1 Forecasting Local Rail Demand: An Overview

While the rate at which new local railway stations have opened in the UK has slowed since rail privatisation in 1994, concerns over oil depletion and the effects of global warming, coupled with the fuel efficiency and versatility of rail as a transport mode, have led to a recent resurgence of interest in such projects. A detailed review of the rail demand modelling literature is given by Chapter 2, but a brief summary is provided here to explain the rationale for this project. The main source of information on rail demand modelling in the UK is the Passenger Demand Forecasting Handbook (PDFH), produced by the Association of Train Operating Companies (ATOC) (ATOC, 2002). However, the recommendations for modelling local rail demand in the PDFH are largely based on work undertaken in the 1980s by Preston (1987, 1991a, 1991b). While these models gave reasonably good results, limits on data availability and computing power meant that they were inevitably somewhat restricted and the passing of time means that their results can no longer be relied upon. Despite this these models are still widely-used, and they played a significant role in the production of the high profile ATOC report on the potential for new rail services, 'Connecting Communities' (ATOC, 2009).

Extensive data on rail usage and demographics are now available along with the computing capacity and techniques to enable it to be processed quickly. However, recent rail demand modelling research in the UK has mainly focused on inter-urban services, and while substantial progress has been made in incorporating the newly-available datasets in these models (Lythgoe, 2004; Wardman et al., 2007) and in investigating the potential of Geographical Information Systems (GIS) in this area (Whelan & Wardman, 1999b), little attention has been paid to local stations. The distinctive nature of local and suburban stations and services means that models developed for inter-urban and intercity services are unlikely to be suitable for use without major modification. There is therefore a need for a procedure which can evaluate the demand impacts of potential alterations to local rail networks, making use of the extensive datasets and computing capability which are now available, and which presents the results of this evaluation in a way which allows them to be easily communicated to stakeholders. This procedure needs to be straightforward and quick to implement so that the case for new stations can be established as cost-effectively as is possible, given the restricted funding available for new infrastructure. The research described in this thesis aims to develop such a procedure.

1.2 Rail Demand Modelling: Art or Science?

There is perhaps a tendency in any field of science which aims to provide a modelled representation of some aspect of reality for the models involved to become progressively more complicated in an attempt to improve the accuracy of their predictions. The modelling of passenger rail trips is no exception to this, and recent work has sought to consider more and more complicated representations of the explanatory variables involved. There is though sometimes a tendency to forget that there will always be an element of 'chaos' in determining the precise number and in particular the pattern of trips made by rail. A good illustration of this is given by the rail trips made by the author over a one month period, which are outlined in Table 1.1 as they would be recorded by the ticketing system.

Table 1.1: Rail trips made by author in October 2006

Date	From	To
02/10/2006	Oxford	Swaythling
02/10/2006	Swaythling	Oxford
05/10/2006	Oxford	Southampton Airport Parkway
05/10/2006	Southampton Airport Parkway	Oxford
07/10/2006	Oxford	Eltham
07/10/2006	Eltham	Oxford
07/10/2006	Sundridge Park	Lewisham
09/10/2006	Oxford	Southampton Airport Parkway
09/10/2006	Southampton Airport Parkway	Oxford
10/10/2006	Oxford	Southampton Airport Parkway
10/10/2006	Southampton Airport Parkway	Oxford
11/10/2006	Oxford	Southampton Airport Parkway
11/10/2006	Southampton Airport Parkway	Oxford
13/10/2006	Oxford	Darlington
15/10/2006	Darlington	Oxford
16/10/2006	Oxford	U1 London
17/10/2006	U1 London	Oxford
18/10/2006	Oxford	Swaythling
18/10/2006	Swaythling	Oxford
23/10/2006	Oxford	Swaythling
23/10/2006	Swaythling	Oxford
31/10/2006	Oxford	Southampton Central
31/10/2006	Southampton Central	Oxford

While the author had perfectly logical reasons for making all the trips represented by this data, a model would be unlikely to predict the seemingly random choice of Southampton Airport or Swaythling as a destination for the (slightly) more predictable trips to the Southampton area. When errors by ticketing staff (the ticket to Southampton Central) and trip chains including travel by other modes (hence the purchase of the single ticket from Sundridge Park to Lewisham) are taken into account the complications involved in attempting to forecast travel become very apparent. This means that however complex demand models become, and however many variables they incorporate, they will never be able to give an entirely accurate prediction of actual travel patterns. It is therefore impossible to escape the conclusion that there will always remain an element of 'art' in the science of rail demand forecasting. The majority of trips are though still likely to follow outwardly logical patterns, and this means that relatively simple models should be able to predict rail travel to a reasonable degree of accuracy.

1.3 What is a 'Local Station'?

The central concern of this thesis are 'local railway stations', and it is possible to outline some characteristics which will be shared by the majority of such stations. Local stations will normally only be served by stopping services to a limited range of destinations, with intercity trains not calling. Most will only have one or two platforms, although junction stations and those on multiple track main lines may have more. If station staff are provided these will tend to be present only during busier periods, and facilities available to passengers will be limited. While it may not always be the major access mode, a significant proportion of passengers will walk to and from the station and thus the station catchment will be fairly localised.

However, while it is possible to identify such characteristics, and while many sources refer to local stations as an implicitly distinct group, the author has not been able to find any comprehensive definitions of what distinguishes a 'local' station from any other station. Preston (1987) refers to 'stations serving local transport needs' which is reasonable as far as it goes, but the existence of through ticketing means that almost no stations on the National Rail network will be exclusively dedicated to local needs. One option is to treat 'local' as being synonymous with 'suburban', but this is no more helpful as a definition and is in any case inaccurate. It is not even possible to 'reverse engineer' a definition by eliminating inter-urban travel from the dataset being considered, as inter-urban trips are usually defined as those over 40 km in length (Lythgoe, 2004; Wardman et al., 2007). While it is probably safe to assume that the vast majority of inter-urban trips are indeed over 40 km in length, not all trips over 40 km in length will be interurban, as a significant number of 'local' trips will be longer than this. For example, the Heart of Wales line is almost always regarded as being local despite being 195 km long (Sphaera Interactive Media, 2006). Another possible definition of 'local' is a station served predominantly by regional (or ex-Network SouthEast) train operating companies (TOCs). However, there is a growing tendency to combine local and intercity services in a single franchise, with for example the Cornish branch lines now run by the intercity operator First Great Western. This means that the primary TOC can no longer be used as a reliable indicator of whether a station is 'local'. A final definition is provided by Network Rail, who divide the stations they manage into six categories, with Categories E (small staffed) and F (unstaffed) likely to contain the majority of local stations. This is the most useful definition for demand modelling purposes and will be used in this thesis, but the categorisation is not faultless as Categories E and F include some stations which would not normally be considered as local (such as Llandudno Junction and Newark Northgate) while omitting some which would (such as Erdington and Purley Oaks).

1.4 New Local Railway Stations in the UK

There was a boom in local railway station construction from the late 1970s until the early 1990s, following a long period when the number of stations had been steadily declining. The end of this boom has been linked to the privatisation of the British railway system from 1994 onwards, which has led to expected rates of return on investments becoming much higher and to confusion over who should pay for the construction of new stations (Preston, 2001). There has also been a massive increase in the costs of building new stations, partly because of the much more stringent legislation in place today regarding accessibility and health and safety, such as the Department for Transport (DfT) code of practice on accessible station design (DfT, 2008). The question of whether or not such legislation is to the 'greater good' is outside the scope of this thesis; suffice to say that stations will now have to demonstrate much greater revenue-earning potential to justify construction than would have been necessary 20 years ago. For example, Mitcham Eastfields (opened May 2008) cost £6 million, whereas the similarly-sized station at Filton Abbey Wood (opened March 1996) cost £1.25 million in 2008 prices. This cost inflation gives improved modelling of the demand for local railway stations increased importance, as it should be able to increase the economic viability of constructing such stations by reducing uncertainty and thus the optimism bias which the Treasury impose on transport projects. Given that the government wishes to increase passenger rail use this would seem to be a particularly desirable development. Local authorities have been instructed to promote development patterns which reduce dependency on road transport (ODPM, 2004), and new local railway stations can potentially play an important part in fulfilling this role. Preston (1987) divided new railway stations into the following categories, illustrated here with recent examples:

- Stations related to new transport systems – e.g. Croydon Tramlink
- Stations related to new services – e.g. Vale of Glamorgan Line
- Intercity parkway stations – e.g. Warwick Parkway
- Stations related to new town development – no recent examples, although for example the opening of Glasshoughton was related to the construction of new housing and leisure facilities
- Stations related to improved central area rail links – e.g. the Crossrail scheme
- Stations on existing services serving local transport needs – three categories:
    - Manned – e.g. Chandlers Ford
    - Unmanned related to major employment centres – no recent examples, but several in 1980s e.g. BSC Redcar
    - Unmanned serving residential areas – e.g. Llanharan

The work reported here will concentrate on new stations on existing services serving existing settlements or employers, but it should be possible to use the same models to predict demand from stations related to new services, and from stations serving new developments. The accuracy of forecasts in the latter category will though always depend upon the accuracy of predictions of the effects of such development.

1.5 Research objectives

The overall aim of this thesis is to develop an integrated methodology for investigating the potential for new local railway stations within a given area. This aim will be achieved by meeting the following objectives:

1) Recalibrate existing aggregate trip rate, trip end, and direct demand models of local passenger rail use using up-to-date data on levels of rail usage, rail services and competition from other modes. This is covered in Chapters 4 to 5.

2) Incorporate further explanatory variables in local rail demand models, based on readily

available data such as census Special Workplace Statistics. These should increase the models' explanatory power without making them significantly more complicated or time consuming to use. This is covered in Chapters 4 to 5.

3) Explore the potential of proprietary GIS to simplify the incorporation of geodemographic data in local rail demand models, and to enhance the visual presentation of model results. This is covered in Chapters 3.7 and 4 to 7.

4) Investigate the use of methodologies not previously applied to rail demand modelling, particularly local analysis techniques such as Geographically Weighted Regression, to enhance rail demand models. This is covered in Chapters 4.3 and 5.4.

5) Produce an automated search procedure using GIS which is capable of identifying and ranking potential new local station sites. This is covered in Chapter 6.

6) Develop an easy to use procedure for examining the demand impact of opening new local stations and services. This is covered in Chapter 7.

7) Apply the search and impact assessment procedures to a small number of British case studies. This is covered in Chapters 6 and 7.

1.6 Thesis Structure

This thesis describes how the objectives outlined in Section 1.5 were fulfilled during a PhD project which ran from October 2006 to September 2009. Following the introduction contained in this chapter, existing rail demand modelling literature is reviewed in Chapter 2. Chapter 3 contains a review of alternative methodologies, before giving details of the methodology chosen for this study and describing the various sources of data used. Chapter 4 describes the development of a wide range of trip end demand models and the use of Geographically-Weighted Regression to enhance these models. Chapter 5 is concerned with two different forms of flow-level demand model and also outlines the conduct of, and results from, a survey of station access and egress designed to assess the accuracy of theoretical station catchment definition methods. Chapter 6 describes the development of a GIS-based search procedure for new station sites. Chapter 7 brings together the trip end and flow level demand models and the site search procedure to produce demand forecasts for a large number of station sites in England and Wales. A subset of these sites is investigated in more detail with an appraisal procedure used to assess the case for station construction. Finally, the findings of the project are summarised in Chapter 8 and some possible areas for future work are outlined.

Chapter Two: Literature Review

2.1 Introduction

This chapter contains a review of previous work carried out in the field of local rail demand modelling. Section 2.2 examines procedures for determining the location of new local railway stations. Section 2.3 summarises current procedures for local rail demand modelling in the UK, and the chapter then goes on to describe the two main categories of rail demand models, with disaggregate models discussed in Section 2.4 and aggregate models in Section 2.5. The latter section outlines the differences between the three main types of aggregate models: trip rate models, trip end models and direct demand models. It also looks in detail at a number of issues connected with aggregate modelling including station catchment definition, the treatment of destinations, station access, intermodal competition and station choice. Section 2.6 compares the performance of different model types, while Section 2.7 summarises the chapter and gives some conclusions.

2.2 Identifying Potential Sites

The first stage in predicting the demand for new rail stations must be the identification of possible station sites. However, this process has been largely neglected in discussions of demand modelling, with most work assuming that sites have already been identified before the modellers arrive on the scene. The ATOC 'Connecting Communities' report identified 75 settlements which could potentially be served by new stations based on population size (ATOC, 2009), but this procedure did not identify actual sites for these stations. As far as the author is aware the only explicit procedure for identifying possible locations for new stations which takes account of local conditions was outlined by Preston (1987; 2001), and this involves the following stages, which could be automated using a GIS:

1. Define search area
2. Exclude track within set distance (e.g. 1 mile) of existing stations
3. Exclude track where engineering constraints would make station construction costs prohibitive (e.g. viaducts, tunnels)
4. Exclude sites where population density is below a set value defined as being the marginal density at which a station would be viable

Horner & Grubesic (2001) outlined a GIS based methodology for evaluating potential locations for park and ride stations in Columbus, Ohio, where there were no pre-existing local rail services. This only considered access by car, and site identification was somewhat arbitrary, but still illustrates the potential for GIS to assist in identifying locations for new stations. While there are other mathematical methods centred on 'optimum station spacings', based on the work of Vuchic (1966), these make a number of idealised assumptions and are of little use in real life situations.

Ideally the planning of new station sites should be integrated with wider land-use planning to maximise potential usage, by locating new traffic generators/attractors near to the stations. However, the short term of most passenger franchises means that train operating companies may not be resourced to interface with local authority land use and planning regimes (Heywood, 2007). It should though still be possible to facilitate some level of interaction between transport and land use planning, given the level of central control which exists via the DfT and Department for Communities and Local Government, and improved demand modelling can contribute to this by increasing understanding of the factors that affect station usage.

2.3 Rail Demand Modelling: The Status Quo

The relationship between rail demand modelling and the 'classic' four stage model of travel demand (Ortuzar & Willumsen, 2001) is a rather complex one. This 'classic' model is more a framework than a detailed model, and divides demand modelling into trip generation, trip distribution, mode choice and route choice stages. In contrast, rail demand models tend to combine the first and third stages of this process, assuming that mode choice is one of the determinants of the trip generating potential of a railway station, although mode choice is often not explicitly considered. The main exception to this are logit mode choice models (see Section 2.4), which are sometimes used to forecast changes in rail's mode share following changes to rail services. Models which predict destination station choice (stage two of the classic model) either follow on from the estimation of trip generation or use observed data on trip generation. Route choice is seldom considered in rail demand modelling (particularly for new stations) except in certain special cases, as such choice is usually extremely limited. One of the main reasons for rail demand modelling's departure from the classic framework is that in most cases rail will be a minority mode both for trips made by individual travellers and for trips made on particular flows. This means that when appraising proposed changes to rail services it is not sensible to expend effort in using general transport models to forecast large quantities of trips undertaken by modes other than rail.

The rail industry's position is that ideally a demand model would be available that embodies the choice of station and access mode alongside the decision of whether or not to make a particular journey by rail (ATOC, 2002). However, until recently most models have either concentrated on demand as a function of the journey from station to station, or alternatively have forecast station and access mode choice without estimating how changes in accessibility would affect the demand for rail travel (ATOC, 2002). Recently attempts have been made to incorporate a representation of station and mode choice in rail demand models (Lythgoe, 2004), although this has inevitably increased model complexity. While the resulting forecasts can be more accurate, this increase in accuracy is not so large that simpler models can automatically be discarded, with the fit of Lythgoe's (2004) best model ($R_{adj}^2 = 0.590$) only slightly superior to that of Preston's much simpler model ($R^2 = 0.539$). If what is required is a tool to provide an overview of the potential for new stations, simpler models are potentially the most cost-effective solution as long as forecasting accuracy is not unduly compromised.

The main source for information on rail demand modelling in the UK is the Passenger Demand Forecasting Handbook (PDFH) (ATOC, 2002). This is predominantly concerned with providing a framework for modelling incremental changes in demand as a result of changes in rail fares, timetables and service quality or in external factors, and absolute demand modelling has until recently been comparatively neglected. While the PDFH does contain two chapters on forecasting absolute demand levels, because of their complexity and lack of universal applicability the handbook recommends that expert advice should be sought before using the models described there. The models for local stations are also based on out-of-date data and therefore require updating. It is undoubtedly less straightforward to forecast absolute demand levels than to predict changes in demand for existing services. In the latter case it is usually possible to adjust existing usage figures using elasticities specified in the PDFH to reflect changes in journey characteristics. This is obviously not possible where predictions are required for proposed new services or for wholesale changes to services as no relevant current rail usage figures will exist.

Most models used to predict incremental change are calibrated on time series data, because such change occurs over time. However, when studying the demand for new stations or services it may be more sensible to use models based on cross-sectional data, although it is crucial that models for new stations can predict demand growth over time, as patronage will build up gradually after a station opens (Preston, 1987). Cross-sectional models usually either predict the number of trips made over a one year period or the number of trips on a typical weekday. The former option is preferable, as modelling the number of trips per day is complicated by seasonal variations in demand and by differences in weekday and weekend travel patterns.

When predicting absolute demand it is necessary to model the generating potential of origins and the attracting potential of destinations (virtually all sites will fall into both categories, although for new stations one or other function will usually predominate). This in turn makes it necessary to address the complex issue of station catchment areas, which are a function of rail service quality, competition from other stations and station accessibility (ATOC, 2002), and are considered in detail in Sections 5.2.4 and 5.3. It is often necessary in models to quantify the relative attractiveness of different options to travellers and this is almost always expressed in terms of 'utility' or 'generalised cost'. 'Utility' can be defined as the level of benefit (or disbenefit) derived by the traveller from making a particular journey, while 'generalised cost' can be defined as the traveller's out of pocket costs plus the travel time multiplied by the traveller's value of time.

Recent demand modelling work (e.g. Lythgoe, 2004; Whelan & Wardman, 1999b) has mostly concentrated on inter-urban journeys over 40 km in length. There have also been studies of demand at relatively uncommon station types, such as parkway stations (Lythgoe & Wardman, 2004) and airport stations (Lythgoe & Wardman, 2002). In contrast, demand modelling for suburban and local services has been rather neglected, even though Category E and F stations account for 15% of passenger trip origins on the British National Rail network (Shah, 2007). The distinctive characteristics of local rail demand mean that inter-urban models are likely to require adjustment before being used to forecast local trips.

A recent and prominent study has investigated the potential for new railway lines (in all cases essentially local) to serve communities in England with populations of 15,000 or higher which do not currently have railway stations (ATOC, 2009). However, little information was provided as to the demand modelling methodology, and as far as it was possible to establish the study did not make any innovations in this area. It made use of the rather outdated modelling methods which form the PDFH standard guidance (see above and Section 2.5), and therefore the appearance of this report does not therefore greatly diminish the need for an improved demand modelling procedure for local rail services. Furthermore, by focusing on relatively large communities (with populations greater than 15,000) the report will have omitted a large number of potential sites for local stations serving smaller communities, and there is therefore still a need for a wide-ranging assessment of the potential for new stations, particularly on existing lines where project costs will be lower.

Rail demand models can be divided into two categories, aggregate models and disaggregate models, and a summary of the main types of model in each category is given here.

2.4 Disaggregate Models

2.4.1 Disaggregate Logit Models

Disaggregate models focus on the choices made by individual travellers rather than on cross-sectional variations in the total number of rail trips, and are thus able to consider the characteristics of trips in great detail, for example by differentiating between those made alone and as part of a group (Wardman et al., 1997). They forecast the choices that would be made following the introduction of a new station or service based on an understanding of how individuals make decisions in similar situations to the proposed new station or service (ATOC, 2002).

Disaggregate modelling is usually based on some form of logit model, with parameters to the independent variables estimated by maximising the likelihood that observations of actual travel behaviour will occur, given assumptions regarding the probability of individuals making particular travel choices in response to these independent variables taking certain values (Lythgoe, 2004). The most common type of logit model is the multinomial logit (MNL) model, which takes the following general form (Lythgoe, 2004):

$$P_{qi} = \frac{e^{U_i}}{\sum_k e^{U_k}} \qquad (2.1)$$

$U_k = f(\beta, E_k)$

Where:

P_{qi} is the probability of individual q making travel choice i

U_i is the utility of travel choice i

U_k is the utility of competing choice k

β is the vector of model parameters

E_k is the vector of independent variables for competing choice k

Logit models have a clear foundation in demand theory, being based on the relationship between the utility functions for travel by different modes or routes. Linear utility functions are most common, although Mandel et al. (1997) found that using nonlinear utility functions gave more realistic results. Wardman et al. (1997) also suggested that the use of such functions increased model accuracy by allowing elasticity variations to be dampened. This extra realism does though come at the cost of more complicated model calibration. A significant problem with MNL models is that they have the property of independence of irrelevant alternatives (IIA) which states that the relative probability of choosing any pair of alternatives is independent of the presence or attributes of any other alternatives (Koppelman, 2007). This property is likely to be violated in many transport applications, when for example the probability of choosing rail travel over bus will be affected by the availability of a car.

The most common alternative to MNL models is the nested logit (NL, also known as hierarchical logit) model, which is more complex but overcomes the IIA problem. For example, Ortuzar (1983) used a NL model to predict mode choice in the Garforth corridor of West Yorkshire, and found it produced more credible predictions than a MNL model. The structural difference between MNL and NL models is illustrated by Figure 2.1.

Figure 2.1: MNL and NL models

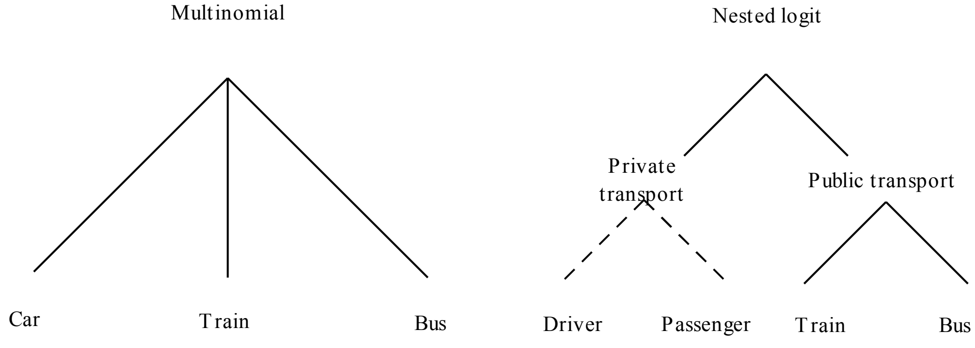

In recent years several other more complex types of logit model have been developed. Many of these are generalised extreme value (GEV) models, a family of models made up of all closed form utility maximisation formulations based on the extreme value error distribution with equal variance across alternatives (McFadden, 1978), and which includes the MNL and NL models. Examples include the paired combinatorial logit (PCL) model, the cross-nested logit (CNL) model, and the generalised nested logit (GNL) model.

The PCL model assigns equal portions of each alternative to one nest with each other alternative, with the total probability of choosing an alternative given by the sum over pairs of alternatives of the unobserved probability of the pair multiplied by the probability of the alternative given choice of that pair (Koppelman, 2007). However, the equal allocation of each alternative to a nest with each other alternative can be a limiting factor (Koppelman & Wen, 2000). The CNL model allows different proportions of each alternative to be assigned to nests defined and selected by the modeller, with each nest having the same structural parameter (Voshva, 1997). The implied correlation and substitution between alternatives is determined by the fractions of each alternative included in one or more common nests (Koppelman, 2007). The GNL model is similar to the CNL model but allows nests to have different structural parameters (Wen & Koppelman, 2001). These three models, together with the NL model, are two-level GEV models, but multi-level GEV models have also been developed allowing progressively more complex nesting structures to be represented.

GEV models as a class still have limitations, in that they are consistent with utility maximisation only under strict, and often empirically violated, restrictions to the dissimilarity and allocation parameters, and they do not allow for sensitivity or taste variations to attributes such as travel time due to unobserved individual characteristics (Bhat et al., 2007). Further (even more complex) models have been developed to overcome these problems, such as the heteroscedastic extreme value (HEV) model which avoids the IIA property of MNL models by allowing different scale factors across alternatives (Bhat, 1995). Another alternative is the mixed multinomial logit model (MMNL) which relaxes the assumption of response homogeneity (allowing for individual variations in taste) and is produced by integrating the MNL formula over the distribution of unobserved random parameters (Bhat et al., 2007). While these complex disaggregate models may give more realistic results than the more traditional MNL and NL forms, calibration of such models is far from straightforward, requiring specialised computer software and a high level of mathematical knowledge, and they do not therefore fulfil the criteria for this study of producing easy to use techniques for rail demand modelling.

A much greater quantity of data is required for disaggregate models than for aggregate models as they work at an individual level and need information on the number of trips by all modes rather than just by rail. Disaggregate logit models can be estimated against either revealed preference or stated preference data. Both have their advantages and disadvantages (see Ortuzar & Willumsen (2001) for a discussion of their relative merits), but obviously only stated preference data can be directly used to assess new stations (Lythgoe, 2004). Disaggregate models calibrated elsewhere could potentially be used to assess the demand for a proposed new station, although such models are notable for their lack of transferability (Preston, 1991a; 2001). The mode-specific constants and scaling factors will be determined by the particular population, and will thus need to be recalibrated for new locations (ATOC, 2002).

Aggregating data from disaggregate models can be problematic when data for a complete set of individuals is not available, because in many cases the function of the average of the variables is not the same as the average of the functions of the variables (Westin, 1974). One way to partially overcome such problems is to use incremental logit models, or extended incremental logit (EIL) models if a hierarchical form is required. Such models can also be used with aggregate data (aggregate logit models), and an example of an aggregate hierarchical logit model predicting the proportion of travellers using rail between a new station and a particular destination is given below. This assumes that an initial decision to use public transport over private transport is followed by a decision on which public transport mode to use.

Initial public/private transport split:

$$P'_{PT} = \frac{P_{PT}\left[\exp\left(S'_T - S_{XT}\right) + \exp\left(S'_{XT} - S_{XT}\right)\right]^{\upsilon}}{P_{PT}\left[\exp\left(S'_T - S_{XT}\right) + \exp\left(S'_{XT} - S_{XT}\right)\right]^{\upsilon} + \left[1 - P_{PT}\right]} \tag{2.2}$$

Lower split shares:

$$P'_T = \left[\frac{\exp\left(S'_T - S_{XT}\right)}{\exp\left(S'_T - S_{XT}\right) + \exp\left(S'_{XT} - S_{XT}\right)}\right] P'_{PT} \tag{2.3}$$

$$P'_{XT} = \left[\frac{\exp\left(S'_{XT} - S_{XT}\right)}{\exp\left(S'_T - S_{XT}\right) + \exp\left(S'_{XT} - S_{XT}\right)}\right] P'_{PT} \tag{2.4}$$

Where:

$P^{|}_{PT}$ (P_{PT}) is the proportion of travellers choosing public transport in the after (before) situation

$P^{|}_{T}$ is the proportion of travellers choosing rail in the after situation

$P^{|}_{XT}$ is the proportion of travellers choosing bus in the after situation

$S^{|}$ (S) is a measure of utility in the after (before) situation

XT is the old public transport mode

T is the new public transport mode (i.e. rail)

υ is the EMU parameter

Source: Preston (1987)

As well as resolving the aggregation problem such models reduce the data requirements of disaggregate approaches, but an additional model will still be required to forecast traffic generated by the new station as the EIL model only distributes journeys abstracted from other modes (Preston, 2001).

A common application of disaggregate models has been to model modal choice and station choice, where the dependent variable is discrete and linear regression can not be used if data is aggregated. For example, Whelan & Wardman (1999a) used a combination of RP and SP data to investigate station and access mode choice, using simultaneous estimation procedures to estimate a multinomial logit model. They suggest such models could be used in specifying station catchment areas by allocating individuals in a zone to competing stations given the quality and cost of rail services at each station and the availability, cost and quality of modes to access each station. However, such methods are yet to be applied to local rail services, and there may be particular difficulties involved in doing so as rail tends to be the minor mode for local journeys and therefore large samples are required if RP mode choice models are to be developed .

Logit models have been used extensively overseas for modelling rail demand, albeit mainly for interurban and intercity travel. The following examples indicate the wide variety of possible applications for such models in this field.

- Tsamboulas et al. (1992) developed disaggregate MNL models for access mode choice to metro stations in the Greater Athens area of Greece. The models were based on RP interview data, and it was found that access mode choice was significantly influenced by trip purpose.

- Mandel et al. (1997) used disaggregate mode choice models with nonlinear utility functions to model large-scale intercity passenger travel in Germany.
- Maruyama et al. (2001) combined a logit mode choice model with a network assignment model to produce an optimisation problem, which was then solved to investigate possible responses to the introduction of congestion charging.
- Marshall & Grady (2006) developed a series of binary logit models to predict inter-zonal mode choice in Washington DC, USA, based on travel to work data from the 2000 US Census. These models included a variety of land use and journey related variables, and had a very good fit with the data, with R^2 values greater than 0.9.
- Kim et al. (2007) produced a MNL model of mode choice based on survey data from St Louis, USA, and found some counterintuitive results, for example that higher income riders and those travelling in the evening were more likely to walk to the station than other travellers.
- Van Vuuren & Rietveld (2002) used a disaggregate structural demand model to estimate the elasticity of demand for train kilometres based on income and fares.

Disaggregate logit models are also used widely in more general transport models, being particularly suited to modelling mode choice as part of stage three of the four stage transport demand model (Ortuzar & Willumsen, 2001).

2.4.2 Network allocation models

Boile et al. (1995) used a novel method to model mode choice in an area of Northern New Jersey, by representing the road and rail network as a single system, and assigning travellers to routes by optimising traffic flows and travel costs over this whole system. Route choice was based on minimising traveller cost, and the framework was used to evaluate various policies aimed to induce a commuter shift to public transport. While realistic results were obtained, such modelling requires detailed and comprehensive data on precise traveller origins and destinations, on travel costs by different routes, and on background traffic flows originating outside the study area. They are therefore seldom seen as being suitable for rail demand modelling, although they are widely used in more general transport models to predict route choice for road traffic (the final stage of the four stage model) (see, for example, Yang & Bell, 1998).

2.4.3 Disaggregate regression models

Rickard (1988) used Poisson regression to model trip rates for individuals based on their personal characteristics and geographical location. The models were calibrated on data from the Long Distance Travel Survey, and as expected it was found that where a main line rail service was available within the respondent's local area they would be likely to travel by rail, but otherwise they would travel by road.

2.5 Aggregate Models

2.5.1 Overview

Aggregate models are based on grouped data as opposed to disaggregate models which are based on data on individuals. Because they are calibrated using data on actual behaviour the model parameters subsume the effects of individuals' differing perceptions of the explanatory variables (Lythgoe, 2004). They are concerned more with the end result of individuals' travel choices than with how these choices are made. There are three main types of aggregate rail demand model, all of which are detailed in the PDFH (ATOC, 2002) although the first two are comparatively lacking in sophistication.

2.5.2 Trip rate models:

Trip rate models are a very simple and pragmatic means of forecasting rail demand, which have no foundation in demand theory. They involve the observation of the number of rail trips per head of population across different stations, with the number of trips assumed to be a function of the catchment area population (Preston, 1991a). These observations are then used to forecast the number of trips generated or attracted by individual stations. Because they take no account of service levels at the new station or of the attractiveness of the destinations served it is important that the stations used to estimate trip rates are as similar as possible to the new station being modelled (ATOC, 2002). Preston (1987) suggested that while trip rate models lack a theoretical basis, are not transferable, and do not rank stations by patronage correctly, they may still provide the easiest and cheapest way of determining possible usage for one-off low cost new stations. Typical values for trip rates are given in the PDFH (ATOC, 2002), although these are calibrated on stations opened in the 1980s and are thus in need of updating. This type of model appears to have been used to analyse potential demand at new stations in the ATOC Creating Communities Report (ATOC, 2009), although the description of the methodology used in this study is unclear and it is possible that train frequencies were also considered. Trip rate (and trip end) models have also been used more generally in transport modelling to predict the total number of trips generated by all modes from particular zones (see for example Wootton & Pick, 1967).

2.5.3 Trip end models:

Trip rate models can be enhanced by including additional explanatory variables, so that the number of trips made is a function of the socio-economic composition of the population, the rail service level and the availability of competing modes as well as of the population size (Preston, 2001). Such models, calibrated on data from existing stations, are known as 'trip end models', and the PDFH suggests that they can be particularly useful when flows from the station(s) under consideration will be dominated by a single destination (ATOC, 2002). Like trip rate models, trip end models are a pragmatic approach to demand modelling with no basis in demand theory, but despite this have the potential to forecast rail use to a high level of accuracy. An example of a trip end model is given below (Preston, 1987):

$$D = \alpha + \beta_1 H_1 + \beta_2 H_2 + \gamma S + \delta J \tag{2.5}$$

Where:

D is daily usage at station i

H_1 is the number of households within 400 m of station i

H_2 is the number of households within 800 m of station i

S is the service frequency at station i

J is the journey time from station i to the nearest major employment centre

α, β, γ, and δ are constants determined by calibration

Lane et al. (2006) tested 163 explanatory variables describing the demographic and transportation attributes of station areas and of the urban area in general, and the characteristics of the rail service, for their trip end model of US local rail use. Only those variables with the most explanatory power were included in the final models, with their light rail model containing 9 variables and their commuter rail model 11 variables. Lane et al.'s work is probably the most extensive application of trip end models so far, being calibrated on a dataset of 868 commuter rail and 348 light rail stations. Like trip rate models these models tend to be context specific (ATOC, 2002) and take no account of trip destinations, and again the PDFH only gives a rather outdated example.

2.5.4 Direct demand models:

Also known as aggregate simultaneous models, these are econometric models which forecast the number of trips between specific station pairs as a function of a vector of explanatory variables. They effectively combine the first three stages of the four stage demand model (Preston, 1987), calculating trip generation, trip distribution and modal choice (although only giving forecasts for those passengers who choose rail). Such models do have some foundation in demand theory, as their general form can be derived by using the utility maximisation principle as demonstrated by Kanafani (1983) and summarised here. An individual originating from location i has a utility function (2.6) associated with the satisfaction of a particular trip purpose, and this function depends on the trips taken by the individual to the various destinations which can satisfy this purpose. These trips have associated costs, and if it is assumed that trips to different destinations are made independently and for each trip the traveller incurs a cost, then the total cost can be calculated using (2.7).

$$U = U(X_{i1}, X_{i2}, \ldots, X_{in}) \tag{2.6}$$

$$C = \sum_{j} c_{ij} X_{ij} \tag{2.7}$$

Where:

X_{ij} is the number of trips made by the individual from origin *i* to destination *j*, where there are *in* mutually exclusive destinations (*j*=1, *n*)

C is the total travel cost

c_{ij} is the cost of travelling from origin *i* to destination *j*

The utility maximisation principle states that the individual will select the values of *Xi* in order to maximise U subject to a constraint on total travel costs, which Kanafani (1983) shows leads to the general result (2.8), meaning that the marginal utility and marginal cost of additional trips to a destination should be equal. (2.8) can be seen to represent the most basic demand model of trip distribution.

$$\frac{\partial U}{\partial X_{ij}} = \frac{\partial C}{\partial X_{ij}} \quad \text{for all } j \tag{2.8}$$

To use this model it is necessary to make assumptions about the form of the utility function. Beckmann & Golob (1970) suggest that U should be assumed to be the constant elasticity function (2.9), which can be substituted into (2.8) to give the result (2.10) for X_j.

$$U = \sum_j \alpha_j X_j^{\rho} \qquad 0<\rho\leq 1 \tag{2.9}$$

$$X_j = \left(\frac{\alpha_j \rho}{c_{ij}} \right)^{1/1-\rho} \tag{2.10}$$

Where:
α and ρ are parameters

Letting the attractiveness of destination j (B_j) equal $(\alpha_j\rho)^{1/1-\rho}$ and γ equal $(1/1-\rho)$ gives the demand function (2.11) which shows that the number of trips made to destination j decreases with the cost of travel to j and increases with B_j (Kanafani, 1983).

$$X_j = B_j c_{ij}^{-\gamma} \tag{2.11}$$

If the total number of individuals at location i with similar utility functions to (2.9) is denoted by A_i, then the aggregate demand function (2.12) is obtained by adding up all the individual functions. This function (2.12) is the general form of the direct demand model.

$$T_{ij} = A_i B_j c_{ij}^{-\gamma} \tag{2.12}$$

Where:
T_{ij} is the number of trips from location i to location j

For rail demand analysis it is usually assumed that i and j are stations, with the terms within the function expanded to include the various components of the generating potential of station I and its catchment, of the attractiveness of station j and its surroundings, and of the generalised cost of the rail journey between the two stations and of other travel modes available over the same route.

Such direct demand models can be calibrated as time series models, with demand regressed against time series data for a single origin-destination pair, but to predict the demand for new stations it is necessary to use cross-sectional models with demand regressed against data for a single period across a matrix of O-D pairs (Lythgoe, 2004), or alternatively pooled time series and cross-sectional data. In recent years these models have become the preferred model form in the UK, with for example Lythgoe (2004) suggesting a number of possible refinements. The PDFH contains a range of direct demand models for interurban and parkway stations, but only gives a single model, based on data from 1981/2, for local stations (ATOC, 2002).

A number of possible functional forms can be used, with the final form determined during calibration depending on the level of data available. Direct demand models are usually calibrated using multiple regression to explain existing demand levels between a matrix of stations. In the past there have been severe problems obtaining suitable calibration data, with for example Preston (1987) having to combine information from a number of sources to produce a usable dataset. However, the availability of computerised ticket sales data through the CAPRI system and its more sophisticated successor, LENNON, mean that obtaining data on rail use is now less of a problem (Preston 2001).

Direct demand models have been widely used to predict rail demand overseas, but these applications tend to be for inter-urban rather than local travel and use time-series rather than cross-sectional or pooled data, as illustrated by the following examples.

- McDonough (1973) used a regression model to predict the short-run demand for commuter rail services in Boston, USA. No serious attempt was made to specify catchments for the stations involved.
- Talvitie (1973) used constrained linear and logarithmic cross-sectional regression models to predict the number of work trips by public transport between zones in the Chicago area. This was based on a small sample of data from a disaggregate travel survey.
- McGeehan (1984) used a simple aggregate model to forecast short-run demand changes on Coras Iomparr Eireann (CIE) interurban services in the Republic of Ireland.
- Doi & Allen (1986) used linear and logarithmic time series regression models to estimate monthly ridership levels on a single urban rapid transit line in

Philadelphia, USA, and found significant seasonal variations in demand.

- Hadj-Chikh & Thompson (1998) estimated ridership levels between commuter rail stations in South Florida using a direct demand model based on catchment populations and the distance between stations.
- Walters & Cervero (2003) developed a logarithmic regression model to forecast ridership on potential extensions to the BART network in San Francisco, USA. This was used together with a conventional 4-stage travel model to account for macro-level travel patterns and micro-area sensitivities resulting from alternative project options.
- Kuby et al. (2004) produced a multiple regression model of light rail boardings calibrated on data from 268 stations in nine US cities, using a raster-based GIS algorithm to define catchments (see Upchurch et al., 2004).

2.5.5 Reference Class Forecasting

This method of predicting demand is more of an appraisal tool then a demand forecasting procedure. It was suggested by Flyvbjerg et al. (2005) as a way of dealing with the consistently inflated forecasts produced for major rail projects in a number of countries. Reference class forecasting is based on comparisons with usage levels at completed projects, and involves the following three steps:

1. Identify relevant reference class of past projects
2. Establish probability distribution for selected reference class, which requires access to credible empirical data (Marshall & Grady, 2006)
3. Compare specific project with reference class distribution to establish most likely outcome.

At the individual station level it may prove difficult to find a completed project which is similar enough to the proposed new station to give reliable forecasts. Therefore reference class forecasting may be better used to validate the forecasts from more conventional rail demand models than as a prediction method in its own right.

2.5.6 Issues with Aggregate Modelling

2.5.6.1 Catchment Area Specification and Population

As mentioned previously, catchment area specification is a major issue for most aggregate models. The simplest models do not even consider catchment areas, instead using dummy variables to represent origin and destination populations, enabling fare and service elasticities to be estimated (Whelan & Wardman, 1999b). However, such methods are only suitable for time-series models as flows can only be forecast for stations included in the set of calibration flows. To make models transferable it is necessary to consider the origin and/or destination characteristics during calibration; in other words to use a gravity model. It is assumed when aggregating the populations that the number of decisions made per year is constant over all individuals within a zone and that the journey times to the station and thus the utility to travel are similarly constant across the zone. The most straightforward way of incorporating population is to replace the dummy variables with the population within a certain straight line distance of the station and such models can still be calibrated using linear regression. Defining catchment areas in this way is inevitably arbitrary as there is no clear boundary in reality, and is also unrealistic as in reality catchments are defined by access and egress times rather than by distance (Krygsman et al., 2004). Using two or more weighted distance bands (Preston, 1987) with differing elasticities does though allow for a limited distance decay effect.

The use of non-linear regression (summation) models allows more complex and realistic model forms to be adopted, by allocating the population around the station to a number of non-overlapping zones, and gives the following generic model form:

$$V_{ij} = \mu \left(\sum_{\alpha} P_{ai}^{\alpha} A_{ai}^{\delta} \sum_{b} P_{bj}^{\beta} E_{bj}^{\lambda} \right) GC_{ij}^{\gamma} \tag{2.13}$$

Where:

P_{ai} is the usually resident population in zone a (related to station i)

P_{bj} is the usually resident population in zone b (related to station j)

A_{ai} is the drive time from zone a to the origin station i

E_{bj} is the drive time from zone b to the destination station j

GC_{ij} is the generalised cost of rail travel from station i to station j

δ is the access elasticity

λ is the egress elasticity

α, β, γ, and μ are constants determined by calibration

Whelan & Wardman (1999b) used a GIS to allocate populations to a set of 'doughnut' zones separated by lines of equal travel time (see Figure 2.2). However, this type of zoning system implies that catchments are identical for all journeys starting or finishing at a station, when in fact they will be affected by the distance and direction of the destination (Lythgoe, 2004). If the majority of journeys from a station are to a particular destination or in a particular direction, then it may be possible to use a parabolic catchment area boundary as suggested by Farhan & Murray (2005) for park and ride services, although this still does not consider the distance to the destination. Such catchments, which may be particularly suitable for urban local stations, have been termed 'commuter-sheds' (Dickins, 1991), as illustrated by Figure 2.3.

Figure 2.2: Doughnut zones

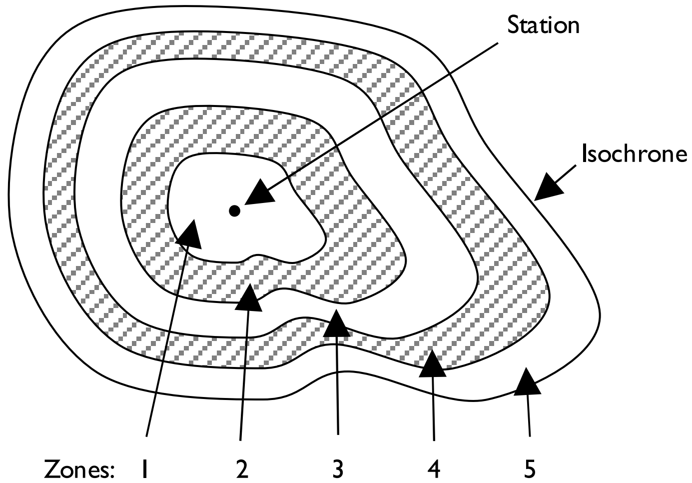

Figure 2.3: Commuter-sheds

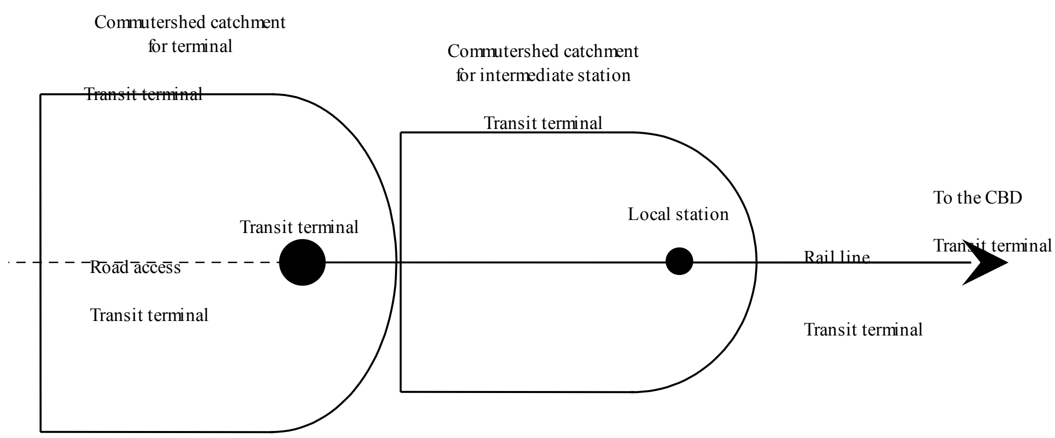

Source: Dickins (1991)

A further refinement would be to define catchments in terms of the generalised cost of access to and from the station (Sargious & Janarthanan, 1983), but while attractive in terms of realism this would complicate model calibration and would require a large amount of additional data, particularly if multimodal access was considered. Similarly the disaggregate approach developed by Farhan & Murray (2005) which considers travel times and distances for individual users requires more information than is generally available.

Lythgoe (2004) developed a new zoning system to allow models to deal more effectively with differing destination distances and directions. A grid of zonal 'seed points' is fixed in location relative to the station, arranged in a series of squares increasing in size, and zones are then defined by allocating units of population to the nearest seed point. Access times and distances from the zonal centres of population should thus give a reasonable representation of these variables for individual residents of zones to any destination, not just to a particular station as with doughnut zones (Wardman et al., 2007). The system allows larger populations with larger variations in the utility of travel in zones further away from the station, although given that the model is looking at competition between stations it is questionable how realistic this is. While this system is an improvement on previous models, there are problems with applying it to local rail journeys. Lythgoe (2004) acknowledges that the probability of individual (or population unit) *a* travelling to location *b* via stations *i* and *j* would not be constant for all individuals residing in a zone if demographic and socio-economic variables were included in the model. The inclusion of such variables would seem to be particularly desirable for local stations as they are likely to possess relatively more explanatory power at this smaller scale. This is because smaller zonal populations mean that there will be less intra-zonal variation (and correspondingly more inter-zonal variation) making it more straightforward to estimate significant parameters for such variables in demand equations. Another problem is the size of the zones used, as those defined by Lythgoe (2004) tended to be at least 6 km across, which is much too large when modelling short-distance local flows. While the zone sizes could be reduced, the aggregation of population units into zones would then become problematic as the areas covered by the population units would not be significantly smaller than the zones, with many units likely to overlap the notional zonal boundaries.

One way around this problem would be to use the population units (for example census output areas or wards) as zones in the modelling process. Despite their irregular size and shape, they might give better results than artificially created zones which do not accurately reflect the actual distribution of the population. An alternative solution would be to spatially disaggregate the population data into small regular-shaped zones, using a combination of GIS and microsimulation models. Spiekermann & Wegener (2000) produced a raster-based zonal dataset using a spatial interpolation method to create probabilistic population allocations. Another method was suggested by Zhao et al. (2003), who used extremely detailed GIS data on the spatial distribution of households to define catchments using a distance decay function. However, all these methods would be extremely complicated to apply, and the resulting improvements in accuracy might not justify the extra effort involved.

It has been argued that inappropriate catchment specification will lead to inflated rail demand elasticities, because stations with good services will have relatively large catchment areas and failing to capture this size effect will lead to the extra demand being attributed to service quality (Whelan & Wardman, 1999b). If the increased catchment size is ultimately a result of service quality this should not be a problem, but catchment size may also be affected by the quality of access and egress modes.

2.5.6.2 Destination Stations

Another shortcoming of previous modelling is that their treatment of destinations is usually far less comprehensive than that of origins. Trip end models rarely include attraction factors, and direct demand models tend to use either aggregated populations or dummy variables as a proxy for all such factors. For example, Lythgoe's (2004) origin station choice model focuses solely on the effects of population distribution to explain the relative attractiveness of different destinations. In reality it is likely that the level of economic activity or the number of jobs located in the zones around a destination station are at least as important as attraction factors (Whelan & Wardman, 1999b), particularly for suburban services which carry a large number of commuters. Therefore for demand models to be realistic they should include some measure of economic activity around destination stations, perhaps through the use of 'floorspace' variables.

2.5.6.3 Access to Stations

Most previous aggregate models have, where they consider access at all, used car drive times to and from stations as a proxy for all access modes, but for local stations walking is likely to be the dominant access mode. Ideally any model of demand for local stations should consider access mode choice as this will affect catchment specification and station choice. Including such choices can though be problematic as it is potentially very time consuming to set up access mode networks from each zone to competing stations (Whelan & Wardman, 1999b). Disaggregate logit models have been developed to model access mode choice (Tsamboulas et al., 1992), but obtaining sufficient data to calibrate such models over a large area may be problematic. While the PDFH includes some guidance on modelling the effects of different access modes , this only relates to inter-urban travel and to modelling incremental changes in demand rather than absolute demand. The same applies to modelling the effects of car parking availability and integration with other modes, which according to the PDFH can not be directly forecast (ATOC, 2002).

Modelling the effects of access on demand is further complicated because demand is unlikely to vary constantly with respect to access time. Wardman et al. (2007) attempted to overcome this by using distance decay functions to model the effects of access time within a summation model. They found that a logit decay function was most appropriate, but this function implied that willingness to use a station was effectively zero for access times in excess of six minutes, which seems extremely unrealistic. Wardman et al. (2007) suggested that this was partly due to the effect of competing stations in a dense network (unlikely for the inter-urban journeys they were considering) and that as the six minutes refers to the uncongested drive time this might represent a significant distance. However, even at a speed of 100 km/h this would give a maximum access distance of 10 km, half the radius of the catchment area specified by Lythgoe (2004) when outlining his grid-based zoning system. Even the direct demand model used by Wardman et al. (2007) to model inter-station competition allows trips to be attracted from a much wider area than this decay function. Krygsman et al. (2004) carried out a detailed investigation of access and egress times and found that while the interconnectivity ratio (the proportion of total trip time spent on access and egress) decreases with total trip distance, absolute access and egress times and consequently catchment size increase with total trip distance. Therefore while it seems likely that superior models will be obtained by specifying a flexible decay function rather than a constant elasticity, work is still needed to establish the optimal form of the function. Ironically, the function specified by Wardman et al. (2007) may be more appropriate for local journeys than for inter-urban travel, as Preston (1991a) found that trip rates for local travel fell off rapidly with distance from the station and that the dominant access mode was walking. Therefore the access distance implied by a six minute drive time seems intuitively to be a reasonable proxy for station catchment boundaries for local trips. It should though be noted that Wardman & Tyler (2000) found that in general the potential for increasing rail travel by improving accessibility to stations was limited, suggesting that the effects of accessibility on demand may have been overstated.

Some of the literature from America suggests that catchments for local rail stations there are significantly smaller than those in the UK. Zhao et al. (2003) claimed that it is commonly accepted that people are unlikely to be willing to walk more than 0.25 miles to use transit. However, Kim et al.'s (2007) study of mode choice in St Louis found that in fact the average walk distance to light rail stations was 0.47 miles (0.76 km), which is close to figures suggested for the UK by Wardman & Tyler (2000) and by the PDFH (ATOC, 2002).

2.5.6.4 Intermodal Competition

Competition with other modes is often (although not always) considered in aggregate demand models, and Wardman (1997) found that this makes an appreciable difference to the resulting demand forecasts. It may not be sufficient to simply compare the journey times and costs of the competing modes as people may have an intrinsic preference for one mode over another. For example it is commonly assumed that travellers would rather use the train than the bus, although Ben-Akiva & Morikawa (2002) found no evidence of such a preference. However, they concluded that bias can arise when rail offers a higher quality service, and it may often be the case that rail is perceived as offering such a service even if this is not actually the case. Prior commitment to specific modes also affects modal choice, as if a traveller has made an investment in using a particular mode by for example purchasing a car, or a rail season ticket, they are unlikely to immediately change their travel mode even if the relative generalised costs of alternative modes alter significantly (Simma & Axhausen, 2001). Beimborn et al. (2003) suggested that separate modal parameters should be estimated for those who are 'captive' to private or public transport and for those who make a genuine choice.

2.5.6.5 Station Choice

Travellers will often consider using more than one station as a possible origin or destination for their rail journey, but this choice is not often explicitly included in aggregate demand models. The use of stations other than the one closest to a traveller's actual trip end is known as 'railheading', and the PDFH suggests that around 20% of local passengers in the south east do this, rising to around 50% for InterCity passengers (ATOC, 2002). This indicates that the accuracy of forecasts could be significantly improved by including a station choice element in models.

Shilton (1982) identified two problems linked to station choice which affect the allocation of populations to catchment areas, and termed them the 'Bolton' and 'Brighton' effects. The Bolton effect describes travellers allocated by the zoning system to a particular station who in reality travel to a more distant station with a superior service. This is likely to be less significant for local journeys, although this will depend on station spacings and may still lead to reduced passenger numbers at local stations which are very close to large stations with a better service quality. The Brighton effect concerns travellers who are able to start their rail journey at one of a number of stations with a similar service located within the same catchment area, and tends to occur in areas served by through trains making frequent stops or in areas served by local trains connecting with a fast service (less relevant for local journeys) . This effect is though a function of poor catchment specification or of using incomplete datasets, as if all stations relevant to the study are included in the model then it is unlikely to occur.

Lythgoe (2004) developed a multinomial logit form for including a station choice decision in cross-sectional direct demand models, although this depends on a grid-based population zoning system being used, as for example with a 'doughnut zone' system such as that used by Whelan & Wardman (1999b) there will be significant intra-zone variation in journey times to competing stations. It also assumes that potential travellers choose between a number of competing stations and access them by road. This assumption is questionable even for inter-urban trips, as intuitively it seems likely that most travellers would only discern two possible origin stations: their local station, and a more distant 'hub' station enjoying a higher level of service which could justify the extra access time involved. This is supported by the findings of Whelan & Wardman (1999b), who attempted to specify catchments using a summation model which allocated five stations to each population unit. The model estimation procedure failed to converge, partly because too many minor stations were specified as competing. Therefore it seems better not to consider an arbitrary number of competing stations in demand models. This is not to say that competition between origin stations is unimportant, but only that in most cases the actual choice set will be limited.

There are certain circumstances where travellers may appear to be considering a larger choice set than they actually are, in particular where a rail trip is one element of a multi-modal trip chain involving several intermediate destinations. For example, a parent might drive their child to school some distance from their home, before leaving their car at a station close to the school to continue their journey to work by rail. Depending on the recording method used it might appear in the modelling dataset that the parent had chosen a station distant from their home because they perceived some aspect of the rail service offered there to be superior to that from their local station, even if this was not the actual reason for the station choice made. Such trip chains may also affect modal choice (Bhat & Sardesai, 2006), as for example people may drive to work instead of using the train because they 'have' to take their children to school by car (O'Fallon et al., 2004).

A key limitation of using a multinomial logit form for station choice is that it implies that the proportion new trips form of the total number of trips from a station is approximately constant, regardless of how far the competing stations are from the origin in question (Wardman et al., 2007). This appears unrealistic, as intuitively it would seem likely that demand at stations closer to the origin would be affected more by a change in generalised cost for journeys from there than demand at more distant stations. Using a cross-nested logit model might overcome this problem (Wardman et al., 2007) by relating the allocation of competing stations to nests to their distance from the new origin station.

2.5.7 Generalised cost of rail

The generalised cost of rail travel can be broken up into a number of components, as shown in equation 2.14:

$$GC_r = F + C_A + U_T + U_A + U_W \qquad (2.14)$$

Where:

F is the rail fare

C_A is the monetary cost of access to origin station and from destination station

U_T is the utility of time spent on the train

U_A is the utility of time spent travelling to and from stations

U_W is the utility of time spent waiting for the train

Establishing the value of F is reasonably simple, but the perception of this value may change if pre-paid ticketing systems are used, although the effects of such systems are not clear (Paulley et al., 2006). Calculating the value of C_A will only be straightforward if data on entire trips rather than on just the rail component of trips is available. U_T is easy to establish, and U_W tends to be included as service frequency. Evans (1969) suggested that U_T and U_W could be combined using a single measure, which he termed the mean travel time. U_A is usually calculated on the assumption that all access is by a single mode, either walk (Preston, 1987) or car (Lythgoe, 2004).

Wardman (1994) found that the demand response to changes in journey time and interchange was not constant, and that using an exponential model of generalised cost gave more realistic results, although this makes calibration much more complex. Establishing appropriate values of time can be difficult, although the PDFH gives a recommended formula for calculating the monetary cost of time spent travelling and a similar formula for wait time (ATOC, 2002), with the utility of time spent travelling assumed to be negative. Passengers may sometimes attach a positive utility to time spent travelling by train, if for example it offers an opportunity to work uninterrupted or for relaxation (Lyons et al., 2007), although this depends on the quality of the journey experience, particularly the level of crowding. Overcrowding of trains (and sometimes of stations) can be a major deterrent to train travel, particularly for business travellers who need to work during the journey (Whelan & Johnson, 2004), and this may mean new stations are less successful than expected if the services calling at them are already close to capacity. Crowding may also have effects on the wider economy, as it can reduce the productivity and efficiency of affected commuters, and can potentially have a detrimental effect on the health of passengers (Cox et al., 2006). Travellers may assign a greater negative utility to 'unscheduled' journey time resulting from unreliability and service delays. Bates et al. (2001) found that punctuality is highly valued by travellers, and it may thus have an impact on rail demand and in particular on modal choice, but including a measure of punctuality in demand models is problematic.

2.5.8 Other determinants of demand

Several other factors may potentially affect rail demand, but are seldom considered in demand modelling, and these are briefly detailed here.

2.5.8.1 Socio-economic and demographic effects

Little is known about the effects of land use patterns and the socio-economic and demographic composition of the catchment population on rail demand (Wardman et al., 2007). Most previous models have not included demographic or socio-economic variables, partly because of a lack of suitable data and partly to limit model complexity. Paulley et al. (2006) point out that care should be taken when including income and car ownership in models, as if only one is included then its effect on rail use will be ‘damped’ by the opposing effect of the missing factor, but if both are included then the collinearity between them will cause problems. Whelan & Wardman (1999a) developed a model including a number of socio-economic and demographic variables, which estimated significant elasticities for income, car ownership, and employment type, but found no significant effects for other socio-economic or demographic factors. However, it would still seem desirable to at least attempt to include such variables in any future models, particularly given the wide range of data available from the 2001 Census.

2.5.8.2 Timetable regularity

Wardman et al. (2004) investigated the impact of a regular timetable or ‘Taktfahrplan’ on demand, and found that while in general this was relatively small it could be significant where fares and journey times were low. Similar conclusions were reached by Johnson et al. (2006), who found the benefits were particularly noticeable for non-London based flows. The most significant effects were exhibited when ‘clockfaced’ timetables were introduced, and Wardman et al. (2004) suggest that current procedures underestimate these benefits.

2.5.8.3 Rail Service Quality

The general ambience of the rail service provided, or the extent to which a ‘wow’ factor is present, can affect rail demand. The best-known example is the ‘sparks effect’ causing patronage growth after electrification, but it is extremely difficult to isolate the effect of such factors. For example, Owen & Phillips (1987) suggested that the introduction of the HST had led to ‘substantial improvements in patronage’, but the model they used did not consider rail journey times, where improvements brought by the new trains may have been more responsible than their higher quality for the patronage increase. Wardman & Whelan

(2001) found that, compared to other factors, rolling stock quality had relatively little impact on demand, but that continued improvements may be necessary to maintain market share as standards of motor vehicles rise.

2.5.9 Demand build-up over time

While not a determinant of demand in itself, the lag time between the opening of a new station and demand reaching an equilibrium level in comparison with pre-existing stations may affect the perceived accuracy of model forecasts. Preston (1987) suggested that such lag periods may have a duration of several years, and this means that if model forecasts are checked too soon after stations open a false impression of the model's (in)accuracy may be obtained. Rose's (1986) study of the Chicago Transit Authority rail system found that lag times may also exist for service changes, and therefore they may affect relative elasticity models as well as absolute demand models.

2.6 Relative Performance

Preston (1987, 1991a) found that the disaggregate approach he used was highly accurate at what it did, but that what it did was very limited (predicting the number of work trips originating within a pre-defined area). It also involved approximately three times the research effort of less accurate, but more comprehensive aggregate approaches. Later reports suggest that for local stations direct demand models are in fact most accurate, followed by disaggregate models, with trip rate and trip end models least accurate (Preston, 2001). Direct demand models have certainly improved in recent years, although there tends to be a trade-off between accuracy and cost. As models get more complex they include more explanatory variables and therefore specification error is reduced, but they also require more data which is unobtainable from documentary sources, leading to increased measurement error (Preston, 2001). Much of the error in disaggregate model forecasts can be attributed to their inability to model generated travel, and to the lack of suitable calibration data for modelling non-work trips, but this should not prevent them having a role to play in decision-making if allowances are made for these limitations. Whichever type of model is used, it is important that absolute demand models are validated by forecasting demand at existing stations and comparing these forecasts with actual usage figures (ATOC, 2002).

2.7 Conclusions

This chapter has outlined current best practice in local rail demand modelling, summarised previous developments in the field and identified several areas where improvements could be made. No automated procedure for identifying new stations sites is known to exist despite the potential of GIS in this area. The absolute demand models in the PDFH are not universally applicable and those for local stations are based on outdated data, and work is therefore required to update them and to increase their transferability. Recent work has concentrated on journeys over 40 km in length, meaning that local services have been comparatively neglected. There is therefore significant potential for new data sources and methods to be applied to demand modelling for local stations and services.

The review of existing demand models showed that they can be divided into disaggregate models and aggregate models. However it should be emphasised that there is not a dichotomy between the two groups, and instead they should be viewed as forming part of a spectrum of modelling approaches. Disaggregate approaches normally use logit models to predict changes in modal choice. While they can be accurate, they tend to be limited in terms of transferability and to be computationally intensive. There are three main types of aggregate model, and these usually forecast the absolute number of trips by rail over a set period. The first type, trip rate models, forecast total demand at a station based on its catchment population. These can be extended by adding additional independent variables to give trip end models, which are calibrated using regression methods. The third type of model, direct demand models, forecast the number of trips between station pairs based on the characteristics of the stations and the journey between them, and these are also calibrated using regression analysis.

As part of this review a number of issues associated particularly with aggregate modelling were considered including station catchment definition, the treatment of destinations, station access, intermodal competition and station choice. There is potential for improvement in the way demand models deal with these issues, particularly with regard to local stations and services.

In Chapter Three some alternative methodologies not previously applied in the field of rail demand modelling will be considered and the choice of methodology for this study will be explained.

Chapter Three: Methodology and Data Sources

3.1 Introduction

This chapter starts by considering two groups of methodologies which have not previously been applied to rail demand modelling but which have the potential to enhance local rail demand models. Section 3.2 deals with the set of statistical methods known as local analysis techniques, which attempt to account for spatial variability in datasets, and Section 3.3 considers the techniques collectively known as cluster analysis, which can be used to examine and partition large datasets. Section 3.4 then reviews the use of appraisal methods in the transport sector, with a focus on cost-benefit analysis. While not a demand modelling tool in themselves, such methods play a crucial role in determining whether or not new railway stations are constructed. Methods for estimating the various costs and benefits associated with new railway stations are discussed, and the choice of methodology for this study is then outlined and explained in Section 3.5.

The chapter then moves on to describe the range of data sources considered for use in this study, in Section 3.6, and to assess their inherent advantages or disadvantages. A brief introduction to GIS is then given in Section 3.7, followed by a description of their use in this study to manage and process data and present results. The advantages and disadvantages of the GIS packages used in the study are briefly discussed. Section 3.8 explains the reasons behind the choice of case studies for the study, before Section 3.9 summarises the information presented in this chapter.

3.2 Alternative Methodologies: Local Analysis Techniques

3.2.1 Introduction to Local Analysis Techniques

While Chapter 2 described a wide range of methodologies which have been used in the past for rail demand modelling, alternative methodologies exist which could enhance rail demand models but have not so far been applied in this field. A set of methodologies which have particular promise can be collectively described as local analysis techniques.

While there is often unequivocal evidence that the individual outcomes of social processes are different in different places, the source of this variation is frequently far from clear

(Duncan & Jones, 2000). The demand for local rail services is no exception, and often appears to exhibit spatial non-stationarity. This suggests that the processes determining rail demand do not operate constantly over space, and thus that measurements of any causal relationship will depend in part on where they are made (Fotheringham et al., 2002).

Virtually all aggregate rail demand models (e.g. Preston (1987), ATOC (2002), Wardman et al. (2007)) use some form of regression for parameter estimation, but such models fail to take into account the possibility that parameters may not be constant across different points in space (Eldridge & Jones, 1991). Standard regression models where the parameter estimates are global statistics may therefore give an inadequate representation of local conditions (Fotheringham et al., 2002). Simpson's Paradox may be in operation when such models are used, in that the results obtained when data is aggregated (over space in this case) are the opposite of those that would be obtained if the data was analysed locally (Simpson, 1951). Robinson (1950) suggested that correlations computed from grouped (aggregate) data will almost certainly be significantly different from those estimated from individual level data. This is illustrated by Figure 3.1, where the trip rates at 9 imaginary local stations in three different areas are plotted against their catchment populations. When trend lines are fitted separately for each area a positive correlation between population and trip rates is obtained for each subset. However, when the data are aggregated and a trend line fitted, this gives a negative correlation between population and trip rates. While this correlation may result from the influence of other factors which mean that individuals in Area 1 are more likely to make rail trips than in Area 3, if the aggregate trend was considered on its own this possibility would be unlikely to be considered.

Figure 3.1: The individualistic fallacy

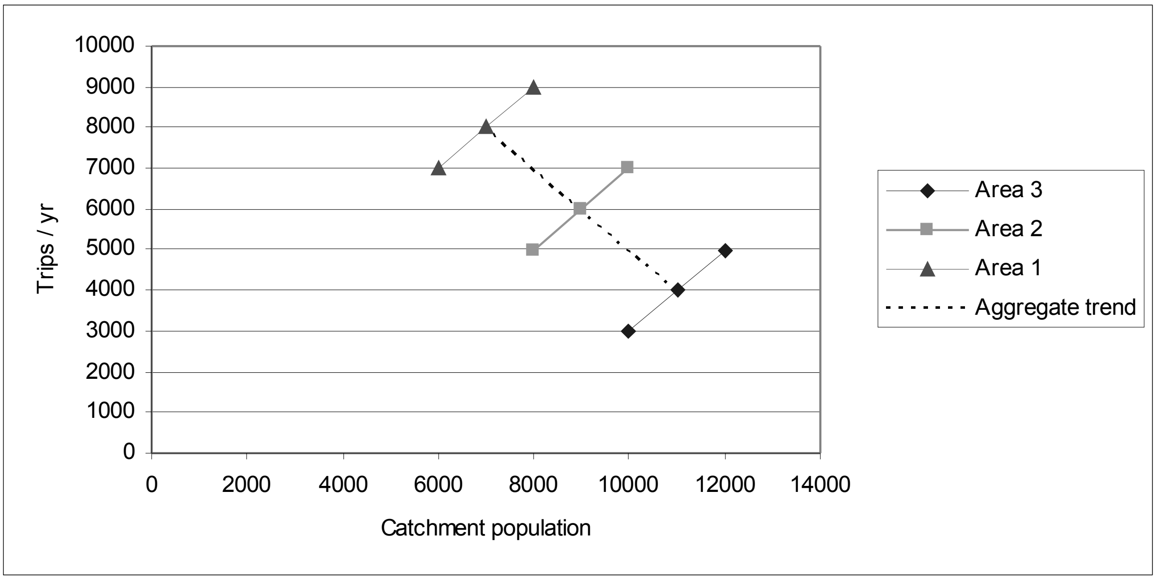

In addition to this individualistic fallacy, Alker (1969) identified several types of inaccurate generalisation which may occur when spatial variations in data are not considered, and these are illustrated schematically in Figure 3.2:

- Cross-level fallacies: Where individual relationships within a single region are extrapolated to a universe of interconnectivity relationships.
- Universal fallacies: Problems resulting from generalising relationships from a subsample.
- Selective fallacies: Problems caused by attempting to represent a universally true relationship in a subsample.
- Contextual fallacies: Where context or social structure alters the strength or form of a causal or statistical relationship.

Figure 3.2: A Typology of Ecological Fallacies

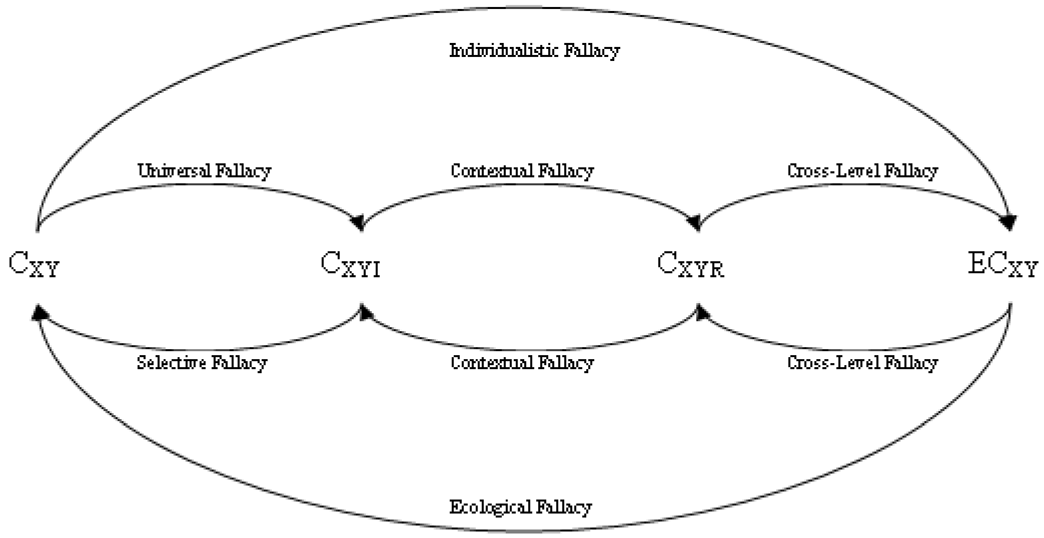

C_{XY} = universal covariance of X and Y
C_{XYI} = covariance of X and Y in subsample I
C_{XYR} = covariance of X and Y in region R
EC_{XY} = between region covariance of X and Y

Source: Alker (1969)

While global regression models have deficiencies this should not lead to them being 'demonised', and arguments in favour of local regression are sometimes rather flawed. For example, Fotheringham (1999) claims that because we would not expect an average US temperature figure to be of use in providing information on the weather in different parts of the country, we should therefore not expect global models to be useful in providing information on local spatial processes. This though overlooks the fact that while the global temperature figure is a single constant, regression models use global parameters combined with locally observed and variable data to estimate the outcomes of processes in particular places, and thus it is possible (and indeed expected) that their results will vary over space.

It is also not universally accepted that low-level spatial context affects travel and activity patterns. Timmermans et al.'s (2003) study of several cities found that variations in relative location and the characteristics of the transport system did not seem to be strongly related to characteristics of activity and travel patterns, and were far less important than differences between the cities. This seems counterintuitive, but if they are correct then the possibility of producing an accurate demand forecasting model for particular areas is increased.

Despite these caveats it is still likely that parameter non-stationarity is a major contributory factor to the imperfect fit of existing rail demand models, and it thus merits further consideration.

3.2.2 Causes of Spatial Non-Stationarity

Fotheringham et al. (2002) identified three main reasons for spatial variations in model parameters, which are sampling variation, intrinsic differences across space in relationships between variables, and poor specification of model form. The latter two illustrate a fundamental schism in social scientific theory, which merits brief consideration here as it can be argued that the reasons are in fact related.

3.2.2.1 Sampling variation

The parameters vary depending on the sample of data used. For example, a rail demand model calibrated using LENNON data might have different parameter values from one calibrated using travel survey data. Such variation is likely to result either from deficiencies in data collection or from inconsistent model implementation, and while not related to any underlying spatial processes can still complicate the identification of these processes (Fotheringham, 1999).

3.2.2.2 Relationships between variables intrinsically different across space

Variations in parameter estimates across space may result from variations in people's attitudes or preferences across space. For example, the provision of wireless internet facilities on local trains might increase patronage more in one area than in another because travellers in the first area valued this provision more highly than those in the second area, perhaps because of a higher propensity to work whilst in transit. While place is an important factor in understanding behaviour, all this factor provides on its own is a description of the spatial variation that exists. It does not give any explanation of why this variation exists, and it has been argued "that the magnitude of contextual variations will be inversely proportional to the adequacy and completeness of the underlying model of relationships between individual attributes" (Hauser, 1970). There are obviously spatial variations in the relationships between the variables affecting rail demand, but 'space' is not a causal factor in itself, merely a proxy for societal factors which are not captured by the model. To state that spatial variations in parameters occur because they are affected by indefinable differences in 'space' is to take a backward step from global models which hide these variations, as they at least attempt to explain relationships between variables at some level. This means that it is important that local analysis methods are used as a tool to identify further model variables, to explain spatial variations in parameters and to enhance model performance rather than just as a way of identifying variability. While it is impossible to identify and quantify every causal factor behind variations in rail demand, virtually all this variation does still result from a definable cause, however micro-scale this cause may be. Fotheringham et al. (2002) suggest that, within a postmodern framework, the identification of local variations in relationships would be a useful precursor to more intensive studies that highlight why such differences occur, but this makes it difficult to see how this cause of variation differs from the third reason (poor model specification).

3.2.2.3 Poor specification of model form

Variations in parameter relationships may occur because the model from which these relationships are estimated is a misspecification of reality with one or more of the independent variables affecting rail demand being omitted from the model or represented by an inappropriate functional form. This is the most likely cause of parameter variation in local rail demand models as several important causal factors are difficult to model effectively. For example, bus competition may mean that models over-predict rail demand from inner city areas where bus use is high, and under-predict demand from outer suburban areas where bus services are less attractive, but modelling this is not straightforward. Another example is micro-level access to stations, where a station might have a large catchment population and a good level of service but because the only means of access was along a poorly lit footpath or through a run-down industrial estate passenger usage would be much lower than might be expected. The fear of crime can be a significant factor in dissuading potential rail users (Cozens et al., 2003), and while such micro-level variation is almost impossible to quantify it has the potential to distort global model calibration if it is not taken into account.

While this is a positivist point of view, it is actually only an extension of the approach outlined in Section 3.2.2.2, which aims to include the factors responsible for the spatial non-stationarity of the previously considered variables in the model. If local analysis shows that there is after all some intrinsic variation in levels of rail demand over space which cannot be explained in terms of any other variable then this analysis will still increase our understanding of its precise nature. Local modelling can be seen as the statistical equivalent of a microscope, which reveals much additional detail (Fotheringham et al., 2002).

It is indeed questionable whether any rail forecasting model is truly global, as no models treat the area served by the railway network as an unvarying plane where a railway station provided with certain facilities and service frequencies would generate the same number of trips wherever it was located. Even the most basic models include the size of the population inhabiting an area, which is just as much a local feature as the area's socio-economic or cultural composition. Therefore all 'local' models actually do is include more detailed local variation than is considered by so called 'global' models.

It is important to note that the 'local modelling' discussed here refers to local adaptations of aggregate models, as opposed to disaggregate approaches which it could be argued are even more 'local'. Fairly extensive disaggregate analysis of rail demand has been undertaken (see Section 2.4), and in some cases this has included spatial stratification, with for example Sheskin & Stopher (1988) examining the difference in attitudes to transit services between rural and urban areas.

3.2.3 Types of local multivariate analysis

3.2.3.1 Geographically Weighted Regression

The traditional method for investigating errors in a spatial model is to map the residuals, but to identify which element of the model is causing the errors it is necessary to map local elements of individual parameters to establish which (if any) of them exhibit spatial non-stationarity (Fotheringham et al., 2002). While global models can be calibrated separately for different subsets of the study area, this is an unsatisfactory solution as the subset divisions would be artificial and would still conceal variation within the subsets. A better solution is to use geographically weighted regression (GWR), where each data point is weighted by its distance from the regression point by fitting a spatial 'kernel' to the data as illustrated by Figure 3.3.

Figure 3.3: Regression point kernel for GWR

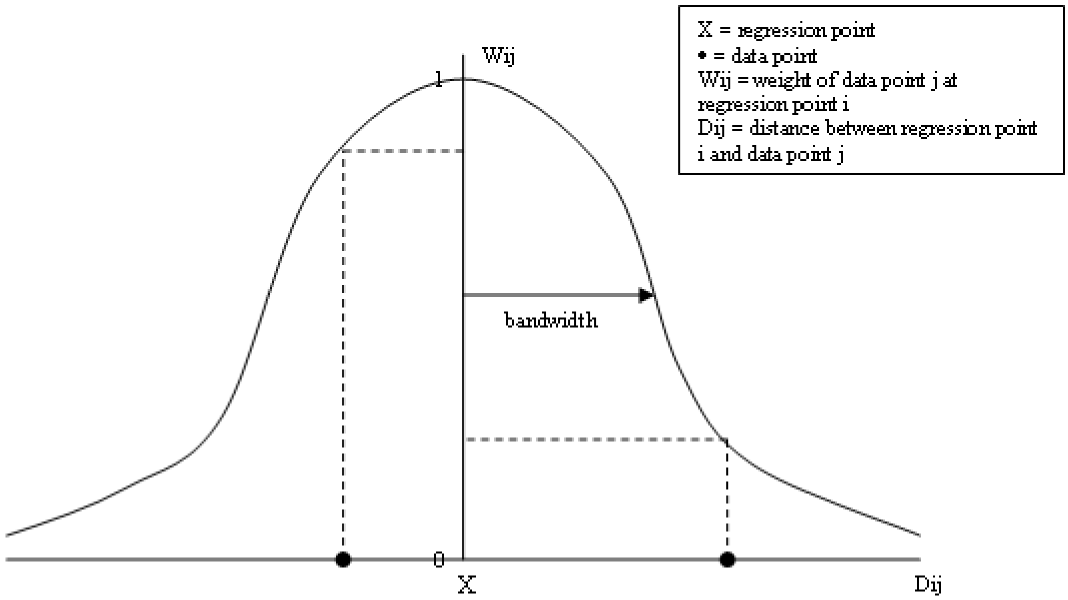

Source: Fotheringham et al. (2002)

For a 'traditional' multivariate global regression model like (3.1), the corresponding GWR model is given by (3.2) (Fotheringham et al. 2002).

$$y_i = \alpha + \sum_k \beta_k x_{ik} + \varepsilon_i \tag{3.1}$$

$$y_i = \alpha(u_i, v_i) + \sum_k \beta_k(u_i, v_i) x_{ik} + \varepsilon_i \tag{3.2}$$

Where:

(u_i, v_i) denotes the coordinates of the ith point in space

$\beta_k(u_i, v_i)$ is a realisation of the continuous function $\beta_k(u, v)$ at point i

α and β are parameters determined by calibration

The model can be calibrated locally by moving the regression point across the region, with unique results for each location due to the varying data point weightings. It is possible to make local parameter estimates at any point in space regardless of whether data had been observed there, which means that if used for modelling rail demand it should theoretically be straightforward to obtain unique parameter values for a new station in advance of its construction.

When estimating a parameter at a given location *i*, it is possible to approximate equation 3.2 in the region of *i* as equation 3.1, and perform a regression using a subset of the points in the data set which are close to *i*. Obviously more influence should be attributable to points which are closer to *i*, and thus a weighted calibration procedure is used, as shown by equation 3.3 (Fotheringham et al., 2002).

$$\hat{\beta}(u_i, v_i) = \left[\mathbf{X}^T \mathbf{W}(u_i, v_i)\mathbf{X}\right]^{-1} \mathbf{X}^T \mathbf{W}(u_i, v_i) y \tag{3.3}$$

Where:

Bold type denotes a matrix

$\hat{\beta}$ represents an estimate of β

$\mathbf{W}(u_i, v_i)$ is the matrix of weights at all data points

The model results are sensitive to the bandwidth of the weighting function, and therefore determining the optimal value for this must form part of any GWR procedure (Fotheringham et al., 2002). This calibration can be undertaken either by minimising the Akaike Information Criterion (AIC) (Hurvich et al., 1998), or by using generalised cross-validation criteria, and plotting these against the bandwidth to establish its minimum value (Fotheringham et al., 2002). Model results are relatively insensitive to the choice of weighting function, with Fotheringham et al. (2002) suggesting either an exponential function (3.4) or a bi-square function (3.5) if finite kernels are desirable.

$$w_{ij} = \exp\left[-\frac{1}{2}\left(\frac{d_{ij}}{b}\right)^2\right] \tag{3.4}$$

$w_{ij} = \left[1 - \left(\frac{d_{ij}}{b}\right)^2\right]^2$ if *j* is one of *n*th nearest neighbours of *i* and *b* is distance to the *n*th nearest neighbour, or = 0 otherwise (3.5)

Where:

w_{ij} is the weight given to the data point *j* when calibrating the parameter at point *i*

d_{ij} is the distance between *j* and *i*

b is the distance to the *n*th nearest neighbour of *i*

n is estimated during the calibration process

If a constant bandwidth was used then for some regression points where data is sparse the model might be calibrated on a small number of data points meaning that parameter estimates had large standard errors, so to overcome this issue the spatial kernels can be allowed to vary in size with the density of data points, making bandwidths larger where data is sparse and vice versa (Fotheringham et al., 2002).

While the calibration of GWR models is obviously far from straightforward, specially designed computer software (GWR 3.0) is available which makes their use relatively simple. This also allows calibration results be plotted on a map to generate surfaces of parameter estimates, illustrating where the parameters have a particularly strong effect on the dependent variable. The parameters are deterministic functions of the spatial location of the regression point, and compared to global models the residuals from GWR estimation should exhibit very little spatial autocorrelation and tend to be less extreme (Fotheringham et al., 2002).

Testing the significance of GWR models is computationally intensive, with Fotheringham et al. (2002) recommending the use of a Monte Carlo type approach, with a null hypothesis equivalent to a global model stating that any permutation of coordinates among the data points is equally likely. To test this it is first necessary to calculate the standard deviation of the GWR estimated parameters for each regression point, to give an estimate of the variability in the parameters. A number of random permutations of data points are then selected, with the GWR procedure repeated for these and their variability estimated. The set of standard deviation values is then ranked lowest to highest, with the proportion of values exceeding that from the observed data giving a measure of the probability of observing such a variation in local parameter estimates from a stationary process (Fotheringham et al., 2002).

A minor drawback of GWR models is that outliers can be harder to detect than in global regression models, because an observation only has to be unusual in its local area to have a distorting effect on the parameters. The GWR software does though include a variety of diagnostic tools which allow such observations to be identified. If some parameters considered by the model are stationary and some are non-stationary, a mixed-GWR model can be used (Fotheringham et al., 2002), but the only drawback of using a normal GWR model in this situation would be the extra calculation time involved, as if parameters are stationary the model should just produce the same estimates for each location.

While GWR was initially developed for simple linear regression models, it has been extended for use with other types of regression models. Atkinson et al. (2003) combined GWR with logistic regression to predict spatial variation in a binary response variable (the presence or absence of river bank erosion). GWR-type techniques can also be used with some non-regression models, with Paez (2006) developing a geographically weighted binary probit model to examine land use change as the result of transit developments. Similar locally varying logit models could potentially be used to enhance the disaggregate modelling of rail demand.

3.2.3.2 Moving Window Regression

Moving window regression is a less sophisticated predecessor of geographically weighted regression, which involves constructing a grid of regression points over the study area and then defining a set of regularly-shaped regions around each point (Fotheringham et al., 2002). Hagerstrand (1965) used a similar approach to model the diffusion of technological innovation over space, with Monte Carlo simulation used to model the probability in each region of adopting new technologies. Local parameter estimates can be obtained by calibrating the regression model on all data within the region around each point, and these can then be mapped to examine non-stationarity. Moving window regression is though still a discontinuous technique, with the results being affected by the size and boundaries of the regions chosen, and edge effects are likely to cause problems (Fotheringham et al., 2002). While moving window regression has not been applied to rail demand modelling, it has been used in some transport applications, for example in modelling commuting distances in Northern Ireland (Lloyd & Shuttleworth, 2005). Moving window regression should not be confused with ‘windowed models’ which extract one or more small geographic areas from the dataset, and create an additional model with added detail for those areas (e.g. Levinson & Huang, 1997).

3.2.3.3 Spatial Expansion Method

Advocates of the expansion method claim that it is both a technique for creating or modifying mathematical models, and a research paradigm (Jones & Casetti, 1992), developed in line with realist ideas. As a paradigm it suggests that important relationships based in theory should be regarded as the building blocks of more complex theoretical relationships which take account of both the basic relationships, their contexts, and the theory about the relationship between the two (Jones & Casetti, 1992).

As a technique the expansion method involves redefining the parameters of an initial model in terms of hypothesised expansion equations, and then reconstructing this model through substitution to give a terminal model capable of capturing the hypothesised drift in the original functional relationship (Jones & Hanham, 1995). Spatial expansion methods allow parameter values to vary locally by making them functions of their geographical location, which permits trends in parameter estimates over space to be measured. Such methods have been most widely applied in modelling migration patterns (e.g. Brown & Jones (1985), Brown & Goetz (1987)), but there should be no fundamental difference between this and the modelling of travel flows. Implementation of the spatial expansion method involves the following steps, based on those outlined by Jones & Casetti (1992) and Fotheringham et al. (2002):

1) Specify initial global regression model:

$$y_i = \alpha + \beta x_i + \varepsilon_i \tag{3.6}$$

2) Redefine at least one of the parameters from the initial model using an 'expansion equation', which incorporates the spatial coordinates of location *i* in the model, giving the following expanded parameters in their simplest form:

$$a_i = \alpha_0 + \alpha_1 u_i + \alpha_2 v_i \tag{3.7}$$

$$\beta_i = \beta_0 + \beta_1 u_i + \beta_2 v_i \tag{3.8}$$

Where u_i and v_i represent the spatial coordinates of location *i*

Brown & Jones (1985) suggest the following expanded form which is more mathematically correct given that the parameter is effectively a polynomial function of the (u,v) coordinates:

$$\beta_i = \beta_0 + \beta_1 u_i + \beta_2 v_i + \beta_3 u_i v_i + \beta_4 u_i^2 + \beta_5 v_i^2 \tag{3.9}$$

3) Substitute expanded parameters into initial global model to create 'terminal model':

$$y_i = \alpha_0 + \alpha_1 u_i + \alpha_2 v_i + \beta_0 x_i + \beta_1 u_i x_i + \beta_2 v_i x_i + \varepsilon_i \tag{3.10}$$

4) Calibrate terminal model using OLS to produce parameter estimates, which are then substituted back into (3.7) and (3.8) to give spatially varying parameter estimates specific to location *i*.

5) Map location-specific parameter estimates to display spatial variations in the relationships represented by the parameters.

6) If necessary iterate expansions, with terminal model produced by one expansion becoming the initial model of a subsequent expansion.

This is similar to the process followed by Lythgoe (2004) in refining his model of interurban rail demand, although the expansion equations he used were not explicitly spatial. The choice of coordinate system used in the expansion should not affect the conclusions drawn from the modelling (Fotheringham & Pitts, 1995). More complex, non-linear expansions of the global parameters can easily be accommodated (Fotheringham et al., 2002), with for example Fotheringham & Pitts (1995) using the following direction-based expansion:

$$\beta = \beta_0 + \beta_1 \cos\theta + \beta_2 \sin\theta \tag{3.11}$$

Where:

β is the parameter to be expanded

θ is the angle of destination j in a coordinate system centred on origin i with due north equal to 0 degrees.

While it is straightforward to implement, the spatial expansion method can only display relationships over space rather than actually explain them, with the complexity of the measured trends dependent upon the complexity of the expansion equations, the form of which must be assumed in advance (Fotheringham et al., 2002). This latter limitation seems likely to restrict the usefulness of the method for rail demand modelling, as the expansions in equations 3.10 and 3.11 do not allow for complex fluctuations in demand over space. The spatial expansion method can be combined with other local analysis methods such as moving window regression (McMillen, 1996), but this inevitably leads to increased model complexity and data requirements. There may in any case be interpretability problems with higher order expansions as the interactions become increasingly complex (Paez, 2006).

3.2.3.4 Spatially Adaptive Filtering

Adaptive filtering works on a 'predictor-corrector' basis, and was originally developed to compensate for the drift of regression parameters over time, with a new observation causing existing coefficients to be updated in terms of their nearest temporal neighbour (Fotheringham et al., 2002). Such methods can be used to investigate the drift of parameters over space, although as data zones tend not to have unique neighbours the coefficient estimates must be updated iteratively until convergence between neighbouring zones is achieved (Gorr & Olligschlaeger, 1994). Local and regional effects can be estimated by mapping the coefficient estimates (Fotheringham et al., 2002), and this method can therefore enhance the detection and estimation of parameters with discontinuous or sharp gradient changes over space (Gorr & Olligschlaeger, 1994). However, the results from such modelling are highly dependent on the specification and scale of the zoning system (Fotheringham et al., 2002).

3.2.3.5 Multilevel Modelling

This method models both average parameter values and the variation around these average values, in an attempt to separate the effects of personal and place characteristics on behaviour (Duncan & Jones, 2000). An individual level model representing disaggregate behaviour is combined with a macro-level model representing contextual variations in behaviour (Fotheringham et al., 2002). The modelling framework used is highly complex, and requires specialised software as ordinary least squares regression cannot be used for calibration (Fotheringham et al., 2002). Significantly different interpretations of the data may occur depending on which issues are considered in the model (Duncan & Jones, 2000), and multilevel models impose a hierarchical structure which may not exist in the process being studied (Paez, 2006). They also rely on the precise definition of a discrete set of spatial units at each level of the hierarchy, which Fotheringham et al. (2002) suggest is unrealistic because most spatial processes are continuous. However, the limitations of available data may mean that discrete boundaries have already been imposed before modelling commences, and thus will affect other forms of modelling to the same extent. Multilevel modelling has not been widely applied in transport, but has seen extensive use in educational research (for example Goldstein, 1987).

3.2.3.6 Random Coefficient Models

In random coefficient models regression parameters are assumed to vary from case to case and to be drawn from a random distribution which is either pre-specified or estimated from the data (Fotheringham et al., 2002). Such models are closely related to random intercept models where only the intercept parameter is assumed to vary (Brunsdon et al., 1999). Model calibration involves estimating the parameters of the distributions from which the casewise parameters are drawn, and then using Bayes' theorem to estimate the value of the regression coefficient actually drawn for each case (Fotheringham et al., 2002). The results from random coefficient models usually exhibit more 'noise' than those from geographically weighted regression, although if true parameter values are random this may mean the results are more accurate (Brunsdon et al., 1999). Random coefficient models are not intrinsically spatial, and because they pay no attention to the location to which parameters refer are not very suitable for modelling local rail demand. They have though seen some use in other transport applications, with for example Nielsen (2003) using them to estimate utility functions in route choice models.

3.2.3.7 Spatial Regression Models

These models assume that the error terms for observations in close spatial proximity to one another are correlated (Fotheringham et al., 2002). A set of weights are defined to represent the degree of interaction between locations, and maximum likelihood methods are then used to estimate model parameters (Ord, 1975). While the model output is still a set of global parameter estimates, local relationships are incorporated into the modelling framework through the covariance structure of the error terms (Fotheringham et al., 2002). While this may enable the accuracy of model forecasts to be improved, the diagnostic potential of such models is limited as they do not make the identification of the factors causing spatial variation any easier.

3.2.3.8 Local Spatial Interaction Models

These recognise that the global calibration of spatial interaction models hides large amounts of spatial information on interaction behaviour (Fotheringham, 1981). The severity of the misspecification bias in such parameter estimates can be shown to vary in a predictable manner with variations in spatial structure, as it depends on the pattern of accessibility existing within the spatial system (Fotheringham, 1984). This means that using localised parameters in spatial interaction models can yield much more information (Fotheringham et al., 2002). However, the disaggregation involved in such modelling is based on discrete points rather than on continuous space, and it may therefore not be applicable to modelling the demand for new stations.

3.2.4 Local Analysis and Rail Demand Modelling

It seems obvious that some form of local analysis should be incorporated in rail demand models, and Table 3.1 outlines the key strengths and weaknesses of the various techniques when applied to modelling rail demand. This comparison indicates that the two methods with most potential for enhancing rail demand modelling are geographically weighted regression and the spatial expansion method. GWR is undoubtedly a more powerful method, because it can account for much more complex spatial variation than the expansion method, and if it can be implemented successfully should give a more reliable indication of the local impact of different factors on rail demand. It is though significantly more complicated than the simple form of the expansion method as outlined in equations 3.7-3.10, and as it requires much more processing the use of both methods will be investigated to establish whether the results from GWR are sufficiently superior to justify the extra time required.

Table 3.1: Strengths and weaknesses of local analysis techniques

Technique	Strengths	Weaknesses
Geographically weighted regression	Results not affected by zoning system or edge effects Local parameter estimates can be mapped Software inexpensive and easy to use Regression points need not coincide with data points – important for modelling new stations	Outliers harder to detect than in global regression models
Moving window regression	Local parameter estimates can be mapped Relatively simple to use	Results dependent on zoning system Edge effects likely to occur
Spatial expansion method	Similar principles to some existing rail demand models Recognises explicitly that parameters in regression models can be function of context Can be calibrated using SPSS	Form of expansion equations must be determined in advance Complexity of trends identified depends on complexity of expansion equations, so local variation may be obscured
Spatially adaptive filtering	Local parameter estimates can be mapped	Results highly dependent on specification of zoning system Zones have multiple neighbours and thus processing times may be lengthy Model iterations may not always converge Spatial smoothing of parameter estimates unrealistic in some cases
Multilevel modelling	Explicitly distinguishes between personal and place characteristics	Highly complex modelling framework Unrealistic reliance on precisely defined set of spatial units Difficult to obtain suitable data for rail demand modelling at individual level
Random coefficient models	Allow high levels of variation Local parameter estimates can be mapped	Not intrinsically spatial Random coefficients pay no attention to parameter location
Spatial regression models	Local relationships incorporated into modelling framework	Limited diagnostic potential – difficult to identify factors causing variation Does not produce local parameter estimates No universally accepted method for estimating weights Calibration far from straightforward
Local spatial interaction models	Yield more information than global spatial interaction models	Disaggregation for discrete points rather than continuous space Not appropriate for trip end models

GWR has been applied extensively to explore a range of spatial phenomena, but while its use has been investigated for some transport-related applications (Du & Mulley (2006), Clark (2007)), as far as the author is aware it has not previously been applied to rail demand modelling. There are some specific problems involved in using GWR to forecast local rail demand. Data on the independent variables for such modelling (the factors determining demand) is not usually collected at the same points as the dependent variable (the number of trips from a particular station). If detailed information has been collected on the actual start and end points of trips then population units such as census output areas can be used as the data points for GWR, but such data is far from universally available.

Otherwise it will be necessary to aggregate data on independent variables for entire station catchments and use the stations themselves as the data points for regression. If this method is used it is desirable to use a large dataset so that patterns in variation between stations in different areas can be adequately illustrated, as otherwise there will be insufficient data points to give reliable results. This is not an ideal solution as it requires that some prior assumptions about the nature of station catchment areas must be made, and these assumptions will affect the parameter estimates. An alternative way around this problem would be to concentrate on investigating variations in rail trips to work, as country-wide data is available on the number of rail trips to work from individual census output areas, which are the smallest spatial units at which data on variables such as population is available, although in practice there are problems with this dataset (see Section 3.6.2.4).

Whichever solution is ultimately adopted, both the spatial expansion method and GWR should provide additional information to that provided by simply mapping the residuals from a global regression model as they make it possible to establish which parameters are responsible for variation in the size of the overall residuals. Local analysis can be seen as a model-building procedure (Fotheringham et al., 2002) in which the ultimate goal is to produce a global model of rail demand that exhibits no significant non-stationarity. By identifying variations in parameters it should be possible to hypothesise on the causes of these variations, and thus develop extra parameters to improve the fit of global demand models.

3.3 Alternative methodologies: Cluster Analysis

Whatever definition of 'local stations' is adopted, the stations described by this definition will inevitably be a diverse group. While some of this variation can be described by model parameters, partitioning the group into smaller groups of stations with similar characteristics could potentially improve model performance. One way to do this is through the use of cluster analysis, a technique for examining patterns within datasets. While it has not previously been explicitly applied in rail demand modelling, cluster analysis has been used to enhance understanding in an extremely wide range of applications. Examples range from distinguishing between different types of tissue in medical imaging (Lasch et al, 2004) to determining target groups for market research (Harrigan, 1985) and finding structural similarities between chemical compounds (Harrison, 1968).

There are three theoretical approaches to cluster analysis. The first is clustering by division, which involves starting with one cluster containing all objects, working out the best way to divide it in two, and repeating this to achieve larger numbers of clusters. This method is impractical for anything other than a tiny dataset as for a dataset of size *n* there are 2^{n-1} possible first divisions which all have to be tested to find the best division (Waterson, 2009).

The second approach, known as partition clustering, uses an iterative nearest-neighbour approach. The number of clusters is chosen in advance, and random points are then chosen to represent cluster centroids. Each data point is assigned to its nearest centroid, and the resulting clusters are then used to calculate new centroids. This process is repeated until the clusters do not change from one iteration to the next (Hawkins et al., 1982). While this method can give good results, Waterson (2009) identified several issues affecting its use. Firstly, the choice of initial centroids may affect results, as the procedure may only identify local optima rather than minimising global error. There may also be problems representing multi-dimensional observations as a single point. Finally, specifying the number of clusters in advance may result in an artificial structure being imposed on the dataset.

The third approach, hierarchical clustering, involves the creation of a series of partitions in the dataset (Everitt et al., 2001) running from a single cluster containing all individuals to n clusters each containing a single individual. This gives a graph called a dendrogram where the vertical scale represents the dissimilarity of the clusters being combined (Waterson, 2009). This method can be used to discover structure in data that is not readily apparent by visual inspection (Aldenderfer & Blashfield, 1984), an important consideration for data on phenomena such as rail demand which are determined by a large number of independent but related variables. While hierarchical clustering provides a convenient way of partitioning datasets, Hawkins et al. (1982) warn that care should be taken when applying such methods to datasets where there is not necessarily an underlying hierarchical structure. This is not to say that cluster analysis is useless in such cases, merely that the cluster solutions produced should not be reified (Aldenderfer & Blashfield, 1984).

A number of agglomerative clustering methods are available for use with hierarchical cluster analysis and Everitt et al. (2001) summarised these as shown in Table 3.2. They concluded that no one hierarchical clustering method can be recommended above others, but that different methods may give very different results on the same data. It is therefore sensible to test multiple methods to identify which gives optimal results for a particular dataset.

Table 3.2: Agglomerative clustering methods for hierarchical cluster analysis

Name	Distance between clusters defined as:	Comments
Nearest neighbour	Minimum distance between pair of objects, one in one cluster, one in the other	Tends to produce unbalanced and straggly clusters, particularly in large data sets Does not take account of cluster structure Sensitive to observational error
Furthest neighbour	Maximum distance between pair of objects, one in one cluster, one in the other	Tends to find compact clusters with equal maximum distance between objects Does not take account of cluster structure
Between-groups linkage	Average distance between pair of objects, one in one cluster, one in the other	Tends to join clusters with small variances Intermediate between nearest and furthest neighbour Takes account of cluster structure Relatively robust
Within-groups linkage	Weighted average distance between pair of objects, one in one cluster, one in the other, according to inverse of number of objects in each class	Intermediate between nearest and furthest neighbour Takes account of cluster structure Relatively robust
Centroid clustering	Squared Euclidean distance between mean vectors (centroids)	Assumes points can be represented in Euclidean space for geometrical interpretation More numerous group dominates merged cluster, subject to reversals
Median clustering	Squared Euclidean distance between weighted centroids	Assumes points can be represented in Euclidean space for geometrical interpretation. New group intermediate in position between merged groups, subject to reversals
Ward's method	Increase in sum of squares within clusters, after fusion, summed over all variables	Assumes points can be represented in Euclidean space for geometrical interpretation. Tends to find same size, spherical clusters Sensitive to outliers

A number of measures of similarity between observations are also available, and their advantages and disadvantages are summarised in Table 3.3, based on a discussion by Aldenderfer & Blashfield (1984). Once clusters have been produced, they can be examined to establish what it is about the observations within each cluster that makes them similar. This allows a decision to be made on whether or not the dataset should be partitioned for modelling purposes.

Table 3.3: Comparison of measures of similarity between observations

Method	Advantages	Disadvantages
Pearson correlation	Not affected by dispersion and size differences between variables – can be advantage or disadvantage	Two profiles may have correlation of +1.0 but not be identical Use to calculate correlation of cases does not make statistical sense
Euclidean distance	Simple to calculate Has intuitive appeal	Involves use of square root Estimation of similarity between cases strongly affected by size of variables – can be overcome by standardisation of variables
Squared Euclidean distance	Simple to calculate Has intuitive appeal Avoids use of square root	Estimation of similarity between cases strongly affected by size of variables – can be overcome by standardisation of variables
Minkowski metric distance function		Estimation of similarity between cases strongly affected by size of variables – can be overcome by standardisation of variables
Manhattan distance		May impose non-existent structure on relationships between variables
Chebychev distance	Emphasises extreme values if these are expected to be important in clustering	If extreme values unimportant will give them undue weight

3.4 Appraisal and Cost Benefit Analysis

3.4.1 Background to Appraisal and Cost Benefit Analysis

Any investment in rail infrastructure, such as a new local station, involves the commitment of resources in the present to obtain extra resources in the future. It is therefore important that investors (whether in the public or private sector) can be confident that the resources generated in the future will justify the expenditure required in the present. To establish this, appraisal of the project is carried out before construction begins, involving analysis of the resources used and generated by the project over time. The aim of appraisal is to provide prescriptive information on which course of action should be chosen from the range available. It is important that the objectives of the schemes are clear at the outset, so that appraisal criteria can follow directly on from them (Mackie & Preston, 1998).

The most commonly used appraisal procedure is Cost Benefit Analysis (CBA), which compares the costs and benefits of a project and recommends that a project should be undertaken if the total economic benefit exceeds the total cost. Extremely extensive literature exists on CBA and appraisal more generally, and this section does not set out to provide a full review, instead giving a brief summary of the application of CBA to transport, in particular in the rail sector. CBA first came to prominence with the US New Deal in the 1930s, with the transport sector one of the first to apply it in the UK, on projects such as the M1 motorway in the late 1950s and the Victoria Line in the early

1960s (Beesley & Foster, 1963). The subsequent accumulation of similar analyses led to the production of the COBA guidelines for new road projects (Nakamura, 2000). The most authoritative outline of the use (actual and potential) of CBA in the rail sector was provided by Nash & Preston (1991), although privatisation in 1994 has seen major changes. In the days of Railtrack the government was, theoretically, less directly involved in investment, with infrastructure schemes assessed on the impact they would have on Railtrack's profitability. This meant that necessary investment was not always undertaken, despite the development by the Office of Passenger Rail Franchising of a CBA framework to assess the case for subsidising investment (Vickerman, 2000). Since the collapse of Railtrack and the formation of Network Rail, control over investment has largely returned to the DfT, and rail projects therefore tend to be subject to the same criteria as other transport schemes. Current guidance is based on the appraisal criteria developed by the Strategic Rail Authority (SRA, 2003) and updated by the DfT to bring them in line with procedures for other modes (DfT, 2007c). This guidance is provided as part of the DfT's Transport Analysis Guidance website (www.webtag.org.uk), which gives information on CBA and related techniques along with guidance on estimating cost and benefit levels for projects to feed into appraisal procedures.

CBA has been commonly used to assess transport projects in many other countries, for example Germany (Rothengatter, 2000), Japan (Morisugi, 2000), and France, although in the latter country policy has oscillated between the use of CBA and of Multicriteria Analysis (MCA) (Quinet, 2000).

CBA can be seen as a pragmatic approach to a problem, as it does not attempt to optimise a situation. Instead it compares a 'do nothing' or 'do minimum' approach with one or more 'do something' scenarios. There is no need to assess anything that is in all responses, and therefore the question of what is contained in the 'do nothing' approach is very important (Mackie & Preston, 1998). For a scheme to be a success, it must pass the Kaldor-Hicks efficiency test, which means that the gainers from the scheme could hypothetically compensate the losers (Hicks, 1941; Kaldor, 1955). The main alternative to the Kaldor-Hicks test is to use the Pareto criterion for a welfare improvement, which says that a project should go ahead only if there are some gainers and no losers, with compensation paid if necessary to achieve this. In practice though there are virtually no schemes where it is possible to financially compensate all the losers (Layard & Glaister, 1994), meaning that the Pareto criterion is of limited use in appraisal.

All resources involved in the project being assessed are measured by marginal cost/benefit based on willingness to pay valuations. In common with many other projects, transport schemes provide a stream of costs and benefits over a long period of time (Nash, 1997). It is assumed that benefits or costs incurred today are worth more than the same benefits or costs incurred in the future, not because of inflation but to represent the real economic phenomenon that resources today tend to be preferred to resources in the future (McCarthy, 2001). This is accounted for in CBA by 'discounting', which adjusts future costs and benefits to give their value at the start of the project. The choice of discount rate is therefore very important in determining the financial success of a scheme, with the Treasury currently recommending a 3.5% rate for the first 30 years of a project, falling to 3% after this and to 2.5% after 75 years (HM Treasury, 2003). The discounted value of a stream of benefits or costs over time in year 0 of the project (the base year) is described respectively as the 'present value of benefits' (PVB) or 'present value of costs' (PVC). Subtracting the PVC from the PVB gives the 'net present value' (NPV) of the scheme, and if this is positive then the scheme passes the Kaldor-Hicks test. The ratio of the PVB to the PVC is the 'benefit-cost ratio' (BCR) of the scheme. The average BCR of implemented schemes is one of the four indicators against which progress in the transport sector will be measured in the UK Government's Comprehensive Spending Review 07 Period, along with journey times on main roads into urban areas, journey time reliability on strategic roads, and rail capacity and crowding levels (Dodgson, 2009). The BCR is therefore a logical choice of measure for comparing the merits of different local rail schemes.

While BCR is often used as a measure for assessing the desirability of scheme construction, a major disadvantage is that it gives no indication of the level of investment required. In the real world the expenditure budget will be constrained, and therefore if multiple schemes are proposed then a pragmatic solution is to progressively adopt the schemes with the highest BCR until either this budget is exhausted or all schemes with a positive BCR have been constructed (unlikely). In reality it may also be the case that the opportunity cost of £1 of public funds is greater than £1, in which case the minimum acceptable BCR should be greater than 1 (Layard & Glaister, 1994). A further problem is that BCRs can vary over the time the project is planned, and post-project BCRs can differ significantly from pre-project estimates (Dodgson, 2009). However, this is not unique to BCRs, being a problem common to all appraisal methodologies, and is usually a consequence of inaccuracies in the procedures used to produce the cost and benefit estimates which feed into the BCR.

An alternative to using BCRs to determine which projects should go ahead is to adopt the 'internal rate of return' (IRR) approach. The rate of return is defined as being that rate which sets the project NPV at zero, and is obtained by solving equation (3.12) (Layard & Glaister, 1994).

$$0 = B_0 - C_0 + \frac{B_1 - C_1}{1+\rho} + \ldots + \frac{B_n - C_n}{(1+\rho)^n} \qquad (3.12)$$

Where:
B_n is the project benefits in year n
C_n is the project costs in year n
ρ is the project rate of return

This approach states that a project should be undertaken if ρ is greater than the discount rate. However, there are several problems with using this approach. It is not an intrinsically correct rule for decision-making like an NPV-based approach, merely a procedure which often gives the same answer (Hirshleifer, 1958). It is also likely to provide the wrong ranking (in terms of NPV) when comparing projects of different size or length, and if there are sign changes in the stream of net returns then the rate of return calculations may give multiple solutions (Layard & Glaister, 1994).

3.4.2 Forms of appraisal

There are several forms of CBA and related appraisal techniques, and the features of these are briefly outlined here.

3.4.2.1 Full CBA

Full CBA is a theoretical ideal methodology, where the monetary values of all significant positive and negative effects, whoever they accrue to and as valued by those affected, are compared for each proposal. This is clearly impossible to implement in practice, as monetary values cannot be obtained for all effects, but this was still what early British attempts to implement CBA aspired to. These culminated in the Roskill Commission's investigations into the siting of the Third London Airport in the early 1970s, when the valuation of the Norman church at Cublington by its fire insurance value led to doubt being cast on the validity of CBA (Self, 1970). Partly as a result of such dubious valuations, it is now generally recognised that full CBA can not be implemented.

3.4.2.2 Financial appraisal

This is also known as cost revenue analysis, and strictly speaking is not a form of CBA, as it considers only financial effects and only those which affect one party (usually the operator) in isolation (Cole, 1998). It establishes the monetary effects of alternative options by using a formula similar to (3.13) to calculate their NPV to this party.

$$NPV_f = \sum_{i=0}^{N} \frac{R_i - OC_i - K_i}{(1+r)^i} \qquad (3.13)$$

Where:

NPV_f is the financial net present value of the scheme

R_i is the revenue in year i

OC_i is the operating cost in year i

K_i is the capital cost in year i

r is the interest rate

N is the project life

(Source: Holvad (2004))

Financial appraisal has been widely used in the commercial transport sector, and under British Rail was used to assess the vast majority of rail services. However, in sectors requiring subsidy (Regional and NSE), this was subject to the constraint of maintaining a ‘broadly comparable’ level of service, meaning that the appraisal aimed to find the most cost-effective way of delivering this service level (Nash & Preston, 1991).

3.4.2.3 Partial CBA

Partial CBA recognises the limitations of full CBA, and acknowledges that only some effects can be quantified and valued in monetary terms. Other effects are therefore listed alongside the results of the CBA, with the performance of each option assessed on a nominal or ordinal scale. The first major application of partial CBA in the UK transport sector was in the Leitch Committee’s 1977 recommendations on the future of trunk roads. This remains the preferred appraisal methodology in the UK, and a number of subtypes of partial CBA have evolved over the past few decades.

3.4.2.4 Social CBA

Social CBA is effectively a form of partial CBA, which aims to include all costs and benefits which can be monetarised in the appraisal, while not considering other factors. In addition to the factors included in financial appraisal, it accounts for benefits accruing to both users and non-users as a result of time savings, and also for the effects of externalities where values can be assigned. The NPV of a scheme is calculated using a formula similar to (3.14). Social CBA has seen extensive use in the road sector in the UK, and is also beginning to be applied to rail projects.

$$NPV_s = \sum_{i=0}^{N} \frac{R_i + UB_i + NUB_i + E_i - OC_i - K_i}{(1+r)^i} \tag{3.14}$$

Where:

NPV_s is the social net present value of the scheme

UB_{ia} is the user transport benefits in year i

NUB_{ia} is the non-user transport benefits in year i

E_i is the benefits external to the transport sector in year i (this might be negative e.g. environmental costs of increased car use)

3.4.2.5 Restricted CBA

Restricted CBA is in many ways similar to social CBA but differs in one important respect, in that user benefits are not included. It is assumed that these will be recovered by raising fares, and will therefore be accounted for in the revenue section of the analysis. The NPV is therefore calculated using (3.15).

$$NPV_s = \sum_{i=0}^{N} \frac{R_i + NUB_i + E_i - OC_i - K_i}{(1+r)^i} \tag{3.15}$$

Until recently this was the procedure used for local passenger rail schemes in the UK, having its origins in application for grants under Section 56 of the 1968 Transport Act (Bristow et al., 1998). It was assumed that railways were able to capture user benefits through fares and that either a financial appraisal or a restricted CBA would therefore give a fair picture of investment schemes (Vickerman, 2000). However, it is now generally acknowledged that it is invalid to disregard user benefits or to attempt to recoup them through higher fares, as raising fares to do this will reduce the social value of a scheme. Perfect price discrimination would be required to fully capture user benefit through fare revenue, and this is impossible to implement in practice. For example, it would require fares to be differentiated by station access distance, as those living closest to the station would receive the greatest benefit since they would not incur large access costs (Nash & Preston, 1991). The use of restricted CBA may have prevented some worthwhile schemes going ahead, such as the Ivanhoe line which had a positive social NPV but a negative restricted NPV (Nash & Preston, 1991).

3.4.2.6 Cost-effectiveness analysis

This is a form of partial CBA used in cases where no meaningful financial valuation of the benefits or costs can be made, and is particularly relevant for pure public goods where it is difficult to exclude people from the benefits (Layard & Glaister, 1994). The benefits of each option are held constant or expressed in non-financial units (such as the number of fatalities avoided) and analysis focuses on the relative costs of the different options in this respect, either in total or per unit. While useful in schemes involving this form of data, cost-effectiveness analysis is unlikely to be used for appraising local rail projects.

3.4.2.7 Goals-achievement matrix:

A goals-achievement matrix may be used alongside a partial CBA to account for non-quantifiable variables, but can also be used as a stand-alone method (Hill, 1968). It makes no attempt to value the effects of proposals in monetary terms, and instead weights or ranks the performance of the proposals in each of the relevant areas. The performance of proposals against each objective is then combined to give an overall performance score.

3.4.3 Application of CBA:

There are a number of stages involved in carrying out a CBA, which can be summarised as follows:

- Define appraisal case (do something), relevant alternatives (do something else) and the base case (do nothing / do minimum)
- Determine the project life
- Determine the key impacts of project

- Quantify the key impacts over time
- Monetise the key impacts over time
- Determine an appropriate interest rate
- Compare discounted costs and benefits using NPV and BCR of different options

Some items included in the CBA may be a cost to one party but a benefit to another, and therefore cancel out. These are called 'transfers' and it is necessary to remain alert to their existence when calculating the NPV (Mackie & Preston, 1998). It may be necessary to use 'shadow prices' for some resources, such as environmental externalities, if market prices do not exist or do not reflect the opportunity cost of using the resource. Finally, government guidelines suggest that an 'optimism bias correction' should be applied to capital costs, increasing them by 66% (DfT, 2004). The justification for this is that previous projects have experienced major cost overruns, but as this is effectively compensating for poor forecasting it could be argued that this cost inflation is both unnecessary and likely to mean that schemes with an otherwise positive NPV are rejected.

The range of costs and benefits included in even a partial CBA is likely to be extensive, and for a local rail scheme might account for the majority of the following factors:

- Construction costs
- Operating costs
- Revenue from new users
- User benefit (consumer surplus) for new users
- Reduced travel time and costs for new rail users
- Increased travel time for existing rail users
- Reduced congestion (and therefore travel time) for remaining road users
- Reduced accident costs (road)
- Disbenefit to remaining bus users
- Loss of revenue to bus operators
- Reduced operating costs for bus operators
- Changes to property prices and land values
- Increased employment through agglomeration benefits
- Environmental benefits and costs
- Income redistribution
- Improved mobility and social inclusion

3.4.4 Problems with CBA:

CBA has been criticised for encouraging mechanistic appraisal and for using unnecessarily econocratic language (Evans, 2004), although this is probably less the case now than it was in the past. It does still emphasise efficiency rather than equity (Layard & Glaister, 1994), perhaps inevitably given that it is an essentially capitalist tool. The development of distributional weights or the inclusion of environmental and social effects in the analysis can help to overcome this criticism, although the discounting procedures used may still fail to take into account the full impact of a project on future generations. There is still (understandably) a tendency to place a greater emphasis on things that can be measured than those which can't, and non-monetised impacts will obviously not be included in the NPV or BCR estimates, even though these may be crucial in deciding whether or not a scheme should proceed (Dodgson, 2009). CBA can be vulnerable to problems of appraisal optimism, because the calculations of the NPV and BCR may not be very transparent, but such optimism is a problem for all appraisal techniques and is more likely to result from faults in the data used to carry out the appraisal than from errors in the procedure itself.

In the UK there has been a failure to apply CBA at a strategic level, and until relatively recently the government has been wary of its use for public transport schemes. Road schemes have often been treated differently to public transport in appraisal, even though second best theory illustrates that comparability between related products in the same sector is an important issue for resource allocation. Rail expenditure in the present day in the UK tends to involve subsidising private companies, and the complex relationships between the various parties involved can make it difficult to define the baseline operation for the do-nothing option. However, this problem is again common to all appraisal techniques, and does not reflect any fundamental shortcomings of CBA.

3.4.5 Multicriteria Analysis and NATA

The problems associated with, and weaknesses of, CBA have led to a move in appraisal towards MCA. This uses a methodology similar to that used to assess non-quantifiable factors in partial CBA to analyse scheme performance. A group of impacts is defined which between them capture the performance of the different options being appraised, and the extent to which the options meet project objectives is assessed relative to these impacts, either quantitatively or qualitatively. These assessments are then transformed onto a scale giving a score for each project relative to each impact, and overall project performance is assessed by weighting these scores and combining them to give a total score (Grant-Muller et al., 2001). MCA has been widely used in Continental Europe (Quinet, 2000; Sayers et al., 2003), but has only been adopted by UK policy-makers relatively recently.

The version of MCA adopted in the UK is called the New Approach To Appraisal (NATA), and there is arguably no fundamental difference between this form of MCA and partial CBA. The preexisting UK CBA methodology (COBA) is retained as part of NATA, alongside the evaluation of elements which had previously been excluded from the analysis (Vickerman, 2000). The performance of project options is assessed based on five major criteria, which are split into 15 sub-criteria, as shown in Table 3.4.

Table 3.4: Criteria used in NATA

Criteria	Sub-criteria
Environmental impact*	Noise
	Local air quality
	Landscape
	Biodiversity
	Heritage
	Water
Safety	n/a
Economy	Journey times and vehicle operating costs
	Journey time reliability
	Scheme costs
	Regeneration
Accessibility	Access to public transport
	Community severance
	Pedestrians and others
Integration	n/a

*Environmental impact also includes data on changes in CO_2 emissions

Source: Price (1999)

The findings for each criteria are brought together in the Appraisal Summary Table (AST) as either a figure or a verbal rating, allowing them to be considered simultaneously when making a decision on the project. This AST is similar to the Benefit Incidence Table (BIT) used in appraisal in Japan (Morisugi, 2000). NATA has been applied to strategic roads reviews and multimodal studies, but is still not an ideal procedure. Nellthorp & Mackie (2000) found that decision makers were failing to consider all the sub-criteria, with their model suggesting that only eight were statistically significant in decision making. They suggested using monetary weights to represent the non-monetary impacts, which would effectively mean a return to CBA. There is also an absence of a clear procedure for ranking the criteria using the AST, which can lead to a lack of accountability in the decision-making process (Sayers et al., 2003). In common with most MCA procedures the choice and use of weights for criteria can seem arbitrary and tends to lead to ambiguity in the decision-making process (Grant-Muller et al., 2001). It is therefore unsurprising that there have been problems in putting NATA into practice, and as with any such combination of qualitative and quantitative analysis there is the danger that the results will be seen as unsatisfactory by both economists and non-economists. There are though ways around these problems, with for example Sayers et al. (2003) suggesting the use of an optimisation model to weight the different criteria, and it seems likely that NATA will remain the preferred methodology for the appraisal of large scale projects at least. However, when assessing the case for relatively small projects, like new local railway stations, Social CBA appears to provide a more pragmatic solution for initial assessment of the business case. Financial appraisal may also prove useful in some cases to establish whether or not a train operator or infrastructure provider could provide a new station without recourse to government subsidy.

3.4.6 Estimation of benefits for new station appraisal

3.4.6.1 User benefits

New rail users attracted to travel from a new station will obviously derive some level of benefit from this travel as otherwise they would not choose to travel by rail. This benefit may be particularly significant for former bus users as the journey time savings for rail compared to bus will often be large. These benefits can theoretically be quantified as the reduction in generalised costs brought about by the new services, with the 'rule of a half' applied to generated traffic (Nash & Preston, 1991). In practice two approaches to calculating user benefits are available. The first is a disaggregate approach, which involves calculating the difference between the generalised costs of rail and other modes for each flow, and multiplying this difference by the number of users abstracted to give the user benefits. However, such an approach would be extremely time-consuming for anything other than a very small dataset, and would require detailed data on factors such as car parking which was unavailable for the South-East Wales case study area. The alternative is to use an aggregate approach based on the functional relationship between rail demand and the generalised cost of rail (expressed in terms of the fare charged). This relationship is shown graphically in Figure 3.4.

Figure 3.4: Theoretical relationship between rail demand and rail fare

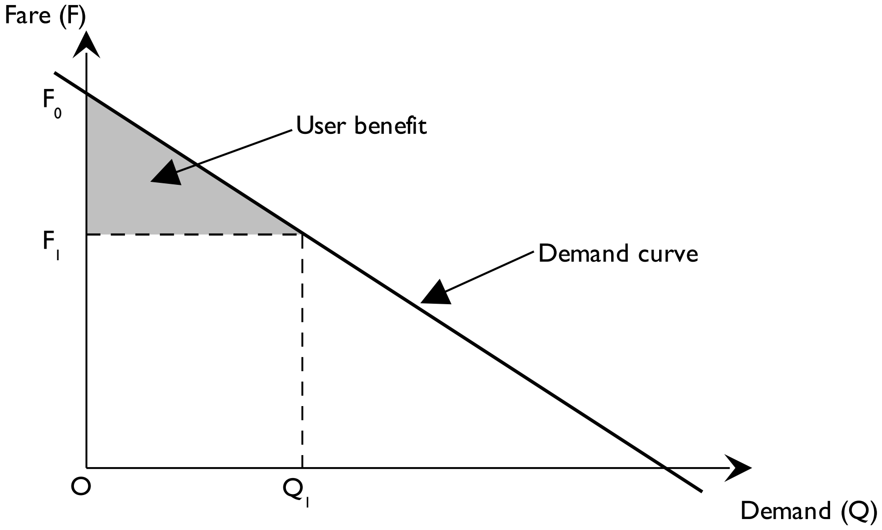

The user benefit is given by the shaded area, and can therefore be calculated by integrating the demand curve with respect to the fare between the limits F_0 and F_1.

3.4.6.2 Abstraction from other modes and non-user benefits

Many trips made from new stations will be abstracted from other modes, particularly car. This abstraction will lead to non-user benefits as a result of reduced congestion, noise, and environmental pollution, but calculating the value of these benefits is not straightforward.

It is first necessary to establish the proportion of trips from new stations which are abstracted from specific modes. Average diversion rates for growth in rail demand are given by Balcombe (2003), and are summarised in Table 3.5. While these diversion rates are based on limited evidence, there are no better data sources available. It should be noted that the urban percentages do not total to 100, presumably due to rounding during the calculation of these figures, and the additional 1% would therefore need to be distributed equally between the modes before non-user benefits were calculated. While not all local stations are urban, the urban set of diversion rates appears most realistic for local stations, as very few trips from new local stations are likely to be abstracted from air travel. However, the diversion rate from bus seems a little high for stations outside major cities.

Table 3.5: Diversion rates (%) for additional rail demand

Trip type	Bus	Car	Air	Cycle/Walk	Generated
Urban	41	33	n/a	1	24
Interurban	20	60	6	n/a	14

Source: Balcombe (2003) p105

The DfT's Transport Analysis Guidance website, WebTAG (www.dft.gov.uk/webtag), includes methodologies for calculating the monetary benefits of reductions in congestion, noise and air quality, but these are extremely difficult to implement. For congestion reductions it is necessary to have data on current traffic levels, as if traffic is being abstracted from uncongested roads then there will be no non-user benefit, but such data were not available for this study. Similarly, for noise reductions the methodology requires data on the decibel change resulting from the project and the number of households affected. While it might be possible (if time-consuming) to establish the latter figure using GIS data, no information on observed noise levels was available for this work. The local air quality methodology also requires unavailable data on pollutant concentrations, and therefore none of these WebTAG suggested procedures could be implemented.

Because rail travel is more fuel efficient than road travel, in addition to these 'local' non-user benefits the construction of new stations should lead to reductions in greenhouse gas emissions and therefore to wider non-user benefits. Guidance for estimating these reductions is also provided in WebTAG, which suggests that the first stage in the process should be to estimate fuel consumption in the before and after scenarios. However, when appraising new railway stations it is not necessary to estimate total fuel consumption, merely the change in fuel consumption, by calculating road distances for the proportion of trips from the new station abstracted from road, and using the fuel consumption formula from WebTAG to estimate the fuel saved. The reduction in carbon emissions can then be calculated using WebTAG estimates and traded off against the additional emissions generated by rail travel from the new station. Finally, the difference between these emissions figures should be monetarised using DfT estimates of the shadow price of carbon. While this methodology could have been implemented for this study, because of its complexity an alternative methodology was investigated based on a report by Sansom et al. (2001) for the DETR (a predecessor of the DfT). This expresses road sector costs and revenues in pence per vehicle km, which allows easy calculation of the benefit of a new railway station by multiplying these values by the number of vehicle kilometres removed from the highway network. The road sector marginal costs and revenues calculated by Sansom et al. (2001) are reproduced in Table 3.6.

Table 3.6: 1998 Road Sector Marginal Costs and Revenues in pence per vehicle km

Category	All vehicles		Car		PSV	
	Low	High	Low	High	Low	High
Infrastructure operating costs and depreciation	0.42	0.54	0.05	0.07	5.23	6.8
Vehicle operating costs (PSV)	0.87	0.87	n/a	n/a	79.61	79.61
Congestion	9.71	11.16	8.98	10.44	15.22	18.19
Mohring effect (PSV)	-0.16	-0.16	n/a	n/a	-14.70	-14.70
External accident costs	0.82	1.40	0.79	1.38	3.74	6.58
Air pollution	0.34	1.70	0.18	0.88	3.16	15.35
Noise	0.02	0.78	0.01	0.52	0.09	4.11
Climate change	0.15	0.62	0.12	0.47	0.56	2.24
VAT not paid (PSV)	0.15	0.15	n/a	n/a	13.44	13.44
Total of costs	12.32	17.05	10.1	13.8	106.3	131.6
Fares (PSV)	0.84	0.84	n/a	n/a	76.77	76.77
Vehicle excise duty (HGV and PSV)	0.14	0.14	n/a	n/a	0.61	0.61
Fuel duty	4.42	4.42	3.86	3.86	5.26	5.26
VAT on fuel duty	0.77	0.77	0.68	0.68	0.92	0.92
Total of revenues	6.17	6.17	4.50	4.50	83.60	83.60

Based on Sansom et al. (2001) p45 and p49
Costs exclude those attributable to pedestrians, bicycles and motorcycles
Accident costs are reported net of insurance payments

All the values in Table 3.6 required conversion into 2008 prices before they could be used in appraisal for this study, and some required further modification. It was assumed that the disbenefit of infrastructure operating costs and depreciation, vehicle operating costs, congestion, the Mohring effect, external accident costs, air pollution, noise, fares and vehicle excise duty had not changed relative to the retail price index between 1998 and 2008. In reality this may not be the case, as for example congestion is likely to have become more severe in some areas in the intervening period, but suitable data was not available to allow the estimated costs to be updated. The costs associated with climate change were increased to account for the increase in the shadow price of carbon (obtained from DEFRA (2007)). It should be noted that a second much higher set of shadow prices is also given in WebTAG guidance, but the justification for and source of these prices are not made clear. The values for VAT not paid from the PSV industry and for VAT on fuel duty were adjusted to account for the (possibly temporary) reduction in the VAT rate from 17.5% to 15%, and fuel duty values were adjust to reflect changes in tax rates. The revised figures from Table 3.6 with all values in 2008 quarter 3 prices are shown in Table 3.7. The net cost per kilometre figures from Table 3.7 can be multiplied by the vehicle kilometres removed from the highway network to give the non-user benefit of the construction of a new local railway station.

Table 3.7: 2008 Road Sector Marginal Costs and Revenues in pence per vehicle km

Category	All vehicles		Car		PSV	
	Low	High	Low	High	Low	High
Infrastructure operating costs and depreciation	0.56	0.72	0.07	0.09	6.99	9.09
Vehicle operating costs (PSV)	1.16	1.16	n/a	n/a	106.39	106.39
Congestion	12.98	14.91	12.00	13.95	20.34	24.31
Mohring effect (PSV)	-0.21	-0.21	n/a	n/a	-19.65	-19.65
External accident costs	1.10	1.87	1.06	1.84	5.00	8.79
Air pollution	0.45	2.27	0.24	1.18	4.22	20.51
Noise	0.03	1.04	0.01	0.69	0.12	5.49
Climate change	0.51	0.70	0.41	0.53	1.89	2.54
VAT not paid (PSV)	0.17	0.17	n/a	n/a	15.40	15.40
Total of costs	16.74	22.64	13.78	18.29	140.70	172.88
Fares (PSV)	1.12	1.12	n/a	n/a	102.60	102.60
Vehicle excise duty (HGV and PSV)	0.19	0.19	n/a	n/a	0.82	0.82
Fuel duty	3.53	3.53	3.08	3.08	4.20	4.20
VAT on fuel duty	0.88	0.88	0.78	0.78	1.05	1.05
Total of revenues	5.72	5.72	3.86	3.86	108.67	108.67
Net cost per km (total costs - total revenue)	11.02	16.92	9.92	14.43	32.04	64.21

3.4.7 Estimation of costs for new station appraisal

3.4.7.1 Construction costs

Data on new station construction costs was collected for the period since 1986, chosen as a cut-off date because it was 20 years before this project commenced. This data was obtained from a number of sources, including the A-Z of Rail Reopenings (Bevan, 1998) and various issues of Modern Railways. Despite an extensive search of the internet, construction costs for a number of stations opened during this period could not be established. Nonetheless, a dataset of 121 stations complete with costs was assembled for analysis. The average cost was obviously skewed by a few extremely expensive stations, specifically four large intercity parkway stations and one station constructed in a deep-level tunnel (Conway Park), which were removed from the dataset as being unrepresentative. Furthermore, the construction costs as collected are not directly comparable, as prices will have changed drastically over the study period. Costs were therefore adjusted to 2008 quarter 3 prices using the retail price index to make stations constructed at different times more comparable, and these adjusted costs were plotted against time, as shown in Figure 3.5. It seemed likely that the introduction of the DfT code of practice on access to stations published in 2002 to satisfy the requirements of the Railways Act 1993 (as amended) and Disability Discrimination Act 1995 (DDA) would have had an impact on station construction costs, as would the DDA accessibility requirements relating to station construction which became effective on 1 October 2004. These events are therefore also shown in Figure 3.5.

Figure 3.5: Construction costs in 2008 prices for stations opened 1986-2008

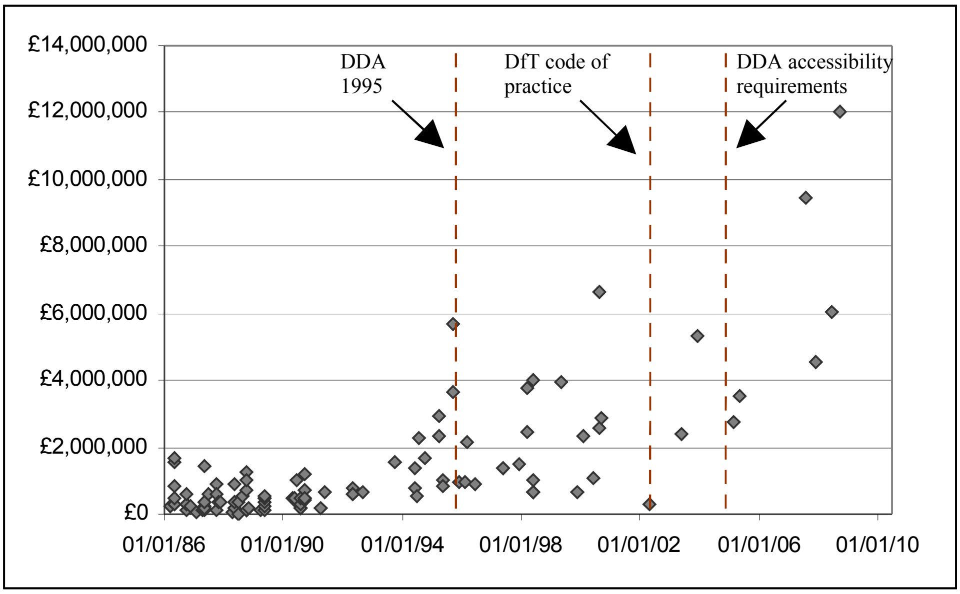

Figure 3.5 shows that even when the costs were converted into 2008 prices, the cost of constructing new stations appeared to have increased massively between 1986 and 2008. The effects of the accessibility legislation and guidelines on station construction costs are somewhat unclear, as there was no obvious step-change in costs after their introduction. This was probably because the railway industry was aware of the legislation's impending introduction well before it came into force, meaning that standards (and therefore costs) at new stations were increased in anticipation of this event.

It was possible (although unlikely) that some of the apparent increase in costs could occur because the stations constructed more recently were generally larger than those constructed in the mid-1980s. The graph was therefore redrawn with stations disaggregated by number of platforms, which was taken to be a reasonable proxy for station size, giving Figure 3.6. It should be noted that the number of platforms for earlier stations may in some cases be an overestimate, as for many of them only current data on size was available and some of these stations may have been enlarged in the intervening period. This disaggregation did not give a clear distinction between the costs for one and two platform stations. While single platform stations were in general cheaper, there are still a number of two platform stations which cost less to construct than single platform stations opened at the same time.

Figure 3.6: Construction costs in 2008 prices for stations opened 1986-2008 disaggregated by number of platforms

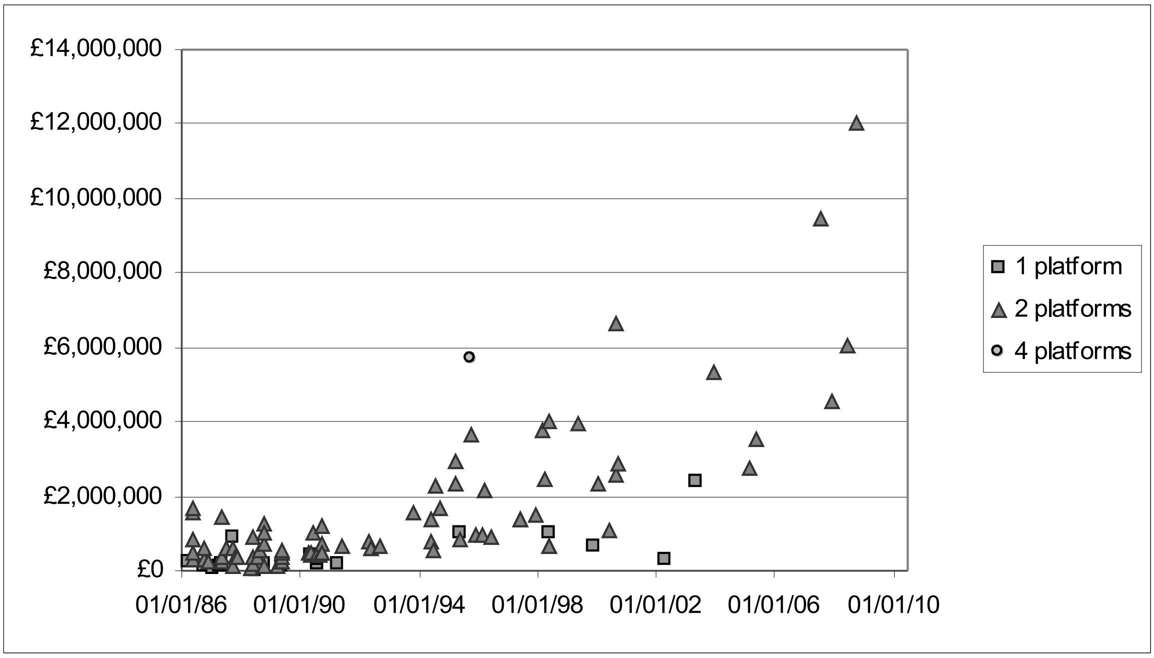

The massive increase in construction costs over the 20 year period shown in the graphs is worthy of further investigation, as it suggests that there should be major savings to be made when building new stations today, but such an investigation falls outside the scope of this project. For the purpose of appraising potential new stations it seemed sensible to concentrate on stations opened since 2000, by which time much of the cost inflation appeared to have taken place. Over the period 2000-2008 the average construction cost for two platform stations (where costs are available) was £4,926,657 in 2008 prices, whereas the average cost for single platform stations was £1,353,520. The average cost for 2 platform stations seemed likely to be representative of such stations in general, particularly given that the first example of Network Rail's (supposedly cost-saving) modular station at Mitcham Eastfields cost £6,058,523 in 2008 prices. This average cost can therefore be used as a general estimate of the cost for new 2 platform stations in appraisal. The figure for single platform stations was rather suspect, however, as one of the two stations concerned was Beauly, which has an extremely short platform that can only accommodate a single coach, and is therefore perhaps not typical of new stations that would be proposed. However, as the other station (Chandlers Ford) was perhaps built to a higher specification than would be normal for a single platform station (complete with booking office and medium-sized car park), the average cost may overall be reasonable. The wide range of variation in construction costs for stations with the same number of platforms may partly result from variations in platform length between stations. Figure 3.6 was therefore redrawn for the period 2000-2008 with the construction cost per coach length of platform

(23m) as the dependent variable, giving Figure 3.7.

Figure 3.7: Construction costs per platform unit in 2008 prices for stations opened 2000-2008

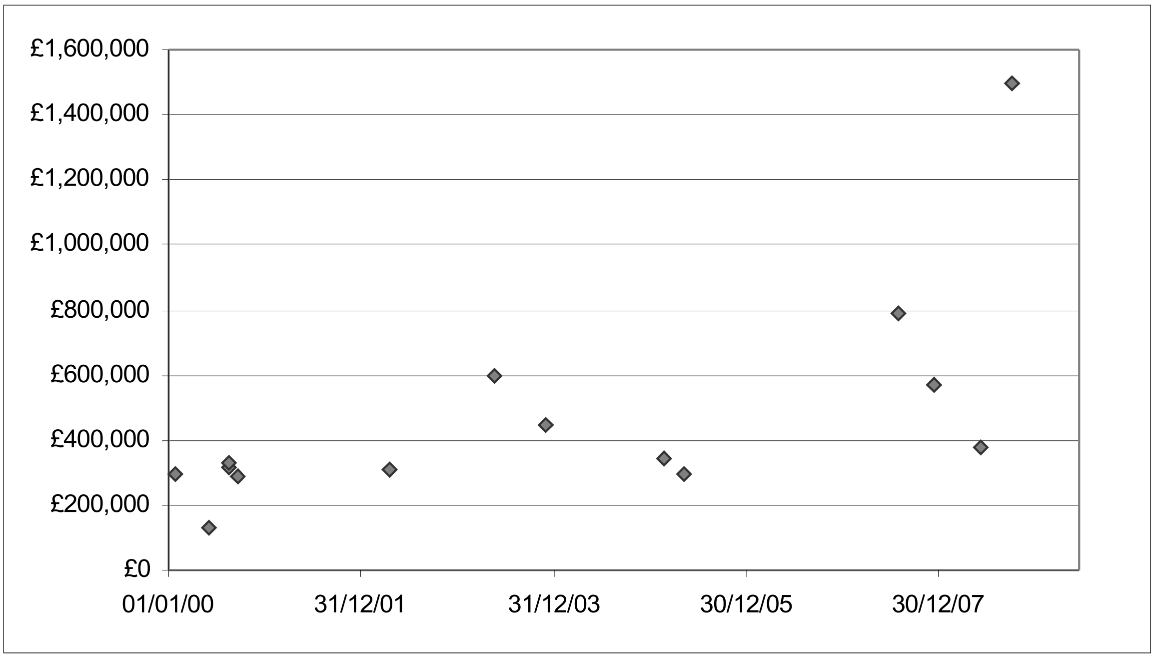

Expressing the construction cost as the cost per platform unit reduced the variation in the data quite significantly, although there still appeared to have been an increase in costs over the period covered by the graph. A regression model (3.16) was tested which aimed to explain the variation in construction costs as a function of time, the number of platform units and the size of the car park. This was calibrated on the data for the stations opened since January 2000 along with a similar model (3.17) for stations opened since the DDA became law in 1995, giving the results summarised in Table 3.8.

$$C_S = \alpha + \beta D_{00} + \gamma U + \delta P \quad (3.16)$$

$$C_S = \alpha + \beta D_{95} + \gamma U + \delta P \quad (3.17)$$

Where:

C_S is the predicted construction cost for station *S* in 2008 prices

D_{00} is the number of days between 01/01/00 and the date when station *S* opened

D_{95} is the number of days between 01/01/95 and the date when station *S* opened

U is the number of platform units at station *S*

P is the number of parking spaces provided at station *S*

Table 3.8: Summarised results from calibration of Models 3.16-3.17

	Model 3.16		Model 3.17	
	Value	t stat	Value	t stat
Intercept	-390433.526	-0.250	-1895936.241	-2.251
β parameter	1940.633	3.261	1195.028	6.284
γ parameter	145113.088	0.772	239960.807	3.294
δ parameter	5782.768	0.993	2476.860	0.886
R_{adj}^2	0.518		0.623	

Model 3.16 did not provide a good explanation for the variation in station construction costs, as it only explained just over half of the variation in the data, and only the time parameter was significant. Model 3.17 gave an improvement over Model 3.16, with all parameters now significant, but the model still only explains under two thirds of the variation in the data, and it can not therefore be relied upon to provide accurate cost estimates. Both logarithmic and translog cost models were also tested, but these gave less accurate results than the linear model and had the same problems with insignificant parameters. However, the average costs detailed above for one and two platform stations could still be used to give a rough estimate of construction costs for new stations. Alternatively, the average cost per platform unit of £470,609.97 could be used to give a more precise estimate of cost. An illustration of the cost inflation that has taken place in recent years is provided by Box's (1992) estimated average cost per platform unit of £203,405.69 (in 2008 prices) for a staffed station with brick-built shelters. However, Box also estimated a cost per platform unit of £406,811.38 for an 'intercity station without facilities', so it could be that today's local stations are being built to the intercity standards of the 1980s . Another alternative means of estimating construction costs would be to use a comparator station chosen on a case by case basis for each proposed station to be appraised. It is however unlikely that suitable comparators could be found for all proposed local stations, and construction costs for apparently similar stations can vary significantly.

3.4.7.2 Station maintenance and operating costs

Data on station maintenance and operating costs are hard to find, and only three estimates were available for this work. The first came from the Scott Wilson feasibility study for a new station at Beeston Castle in Cheshire (Scott Wilson, 2006). Operating costs were represented by Network Rail station access costs, which are meant to reflect the impact the station will have upon the network infrastructure. These will include charges for lifts and other equipment (should this be present) and also long term charges for items such as CCTV renewal (Faber Maunsell, 2007). Scott Wilson estimated that the total maintenance and operating costs (for a 2 platform station without lifts) would be £37,364 per annum (in 2008 prices), breaking down to give maintenance costs of £27,474 per annum and operating costs of £9,890 per annum. The second source of data on these costs was the Hampshire County Council bid for Rail Passenger Partnership funding for the South Hampshire Crossrail project (quoted in Halcrow Group, 2006). This estimated station access costs as being £37,937 per annum for Chandlers Ford, a single platform station with a station building. Finally, station access costs for the proposed station at Imperial Wharf in West London were estimated as £39,685 per annum for a two platform station with lifts (Faber Maunsell, 2007). Logically, the amortised construction costs should be included in the Network Rail access costs, and this introduces a risk of double counting. However, as far as it was possible to establish from the available information (Office of the Rail Regulator, 2006) this was not the case for these estimates.

3.4.7.3 Staffing costs

The question of whether or not new local stations should be staffed is a difficult one, as staff provision will obviously increase costs both through wages and through the construction of facilities such as booking offices. However, evidence suggests that the presence of staff at stations can lead to an increase in rail demand (Preston et al., 2008), and they can also improve revenue capture by reducing ticketless travel. If staff are to be provided at new stations then the cost of their employment should be included in the appraisal procedure, but only two estimates of staff costs were available for this study. The first came from the South Hampshire Crossrail study (Halcrow Group, 2006), which estimated that providing booking office staff for Chandlers Ford station would cost £17,084 per annum. Faber Maunsell (2007) gave a figure of £115,653 per annum for ‘operating costs’ at Imperial Wharf which was separate from the station access costs and therefore seems likely to be largely made up of staffing costs. The discrepancy between the two estimates is probably due to the larger staff presence planned for Imperial Wharf.

3.4.7.4 Cost of increased journey times

The construction of new railway stations will result in a small disbenefit for existing passengers on the trains which serve the new station, as the provision of an additional stop will lengthen their journey times. It is therefore necessary to calculate the monetary value of this disbenefit when appraising new station schemes. The first stage in this calculation is to establish the time taken for trains to make additional stops. This will vary from station to station as it will be affected by line speeds and gradients around the new stations, as well as by the acceleration and braking qualities of the trains concerned. However, an average figure was calculated for the purposes of this study by collecting data on journey times between stations either side of new stations before and after they opened. This data was collected for fourteen stations from various editions of the National Rail Timetable (Network Rail, 2008a; Railtrack, 1996, 2001, 2002), and is summarised in Table 3.9.

Table 3.9: Change in journey times with new station opening

New station	Date opened	Adjacent 1	Adjacent 2	Distance (m)	Time before	Time after	Difference
Horwich Parkway	01/05/1999	Blackrod	Lostock	3.5	6	7	1
Braintree Freeport	10/11/1999	Braintree	Cressing	2	4	5	1
Dunfermline Queen Margaret	25/01/2000	Cowdenbeath	Dunfermline Town	5.25	7	9	2
Wavertree Technology Park	13/08/2000	Edge Hill	Broad Green	1.75	6	6	0
Warwick Parkway	13/08/2000	Dorridge	Warwick	10.75	12	14	2
Lea Green	17/09/2000	Rainhill	St Helens Junction	3	5	6	1
Beauly	15/04/2002	Muir of Ord	Inverness	13	18	20	2
Edinburgh Park	Dec 2003	Haymarket	Linlithgow	16.25	13	16	3
Glasshoughton	21/02/2005	Castleford	Pontefract Monkhill	3.25	7	9	2
Gartcosh	09/05/2005	Greenfaulds	Stepps	8	9	11	2
Coleshill Parkway	Aug 2007	Water Orton	Nuneaton	12.25	18	20	2
Llanharan	10/12/2007	Pencoed	Pontyclun	5.25	7	9	2
Mitcham Eastfields	02/06/2008	Streatham	Mitcham Junction	2	7	9	2
Shepherd's Bush	28/09/2008	Willesden Junction	Kensington Olympia	3.25	8	9	1

This shows that, on average, 1.64 minutes was added to the journey time between the two existing stations after a new station had opened. In general this increase in journey time seems to increase with distance between the existing stations, although there are some exceptions to this. The monetary cost of the additional journey time can be calculated by multiplying it by the value of time of the existing passengers and by the number of passengers affected.

3.4.7.5 Vehicle costs

The costs relating to rolling stock which will be incurred with the opening of a new station will largely depend on whether or not the station can be served by existing services. If new services are required then the cost of acquiring and maintaining suitable trains and of providing staff to operate them will form a significant component of the total project costs (see Section 3.4.7.6). However, nearly all the stations being appraised here could be served by existing services, and therefore the only costs incurred will be those which result from these services making an additional stop. It has been suggested that such additional stops will not add any significant cost to train operators (Scott Wilson, 2006) but while it is safe to assume that these extra stops are unlikely to significantly increase maintenance costs, they will inevitably lead to increased fuel consumption because of the extra power required to accelerate away from the new station.

If the assumption is made that fuel consumption (and therefore variable rolling stock costs) is perfectly variable with respect to journey time then it is possible to estimate the extra fuel used in making an additional stop, and therefore the cost of this stop. Fuel cost estimates were given in the feasibility study for the South Hampshire Crossrail service, with expenditure on fuel forecast to be £207,100 per year (in 2001 prices) which equates to £3982.69 per week. Given that 102 services operate per week, this gives a figure of £39.04 per service. The end to end journey time was 43 minutes, but 5 minutes were spent stationary at Chandlers Ford, 1 minute at Eastleigh and 2 minutes at Southampton Central, giving an 'in motion' journey time of 35 minutes, and a fuel cost per minute of £1.12. Table 3.7 shows that on average a new station will increase journey times by 1.64 minutes, giving a fuel cost of £1.83 per additional halt per service.

This figure is in 2001 prices, and adjusting this for inflation based on the retail price index gives a cost of £2.29 per additional halt in 2008 quarter 3 prices. However, fuel prices have risen over the same period at a higher rate than the rise in the retail price index, and taking this rise in price into account gives a cost of £2.44 per additional halt. This at first sight seems to confirm the assumption of Scott Wilson (2006) that additional stops do not add significantly to train operating costs. However, while the cost per service may be small, this can still scale up to a relatively large cost per year and was therefore incorporated in the appraisal procedure.

3.4.7.6 New Service Costs

If new or extended services are to be provided to serve new stations then this will incur additional costs. The report by Sansom et al. (2001) used to calculate non-user benefits also contains estimates of marginal costs and revenues for the rail sector, and Table 3.10 shows these estimates for regional passenger services, updated to 2008 prices in the same way as the road estimates in Table 3.7.

Table 3.10: 2008 Rail Sector Marginal Costs and Revenues in £ per train km

Category	Low	High
Infrastructure	0.20	0.20
Vehicle operating	6.74	6.74
Congestion	0.12	0.12
Mohring effect	-0.90	-0.90
Air pollution	0.05	0.48
Noise	0.06	0.19
Climate change	0.10	0.14
VAT not paid	0.62	0.62
Total costs	6.99	7.58

The marginal infrastructure costs in Table 3.10 may be an underestimate, as when Wheat & Smith (2008) carried out detailed investigations of marginal rail infrastructure costs their preferred model gave an estimated cost of £1.013 per train km. However, this figure includes both passenger and freight traffic, and the figures from models which distinguish between traffic types gave estimated marginal costs of £0.187 and £0.149 for passenger traffic (Wheat & Smith, 2008), which are similar to the figure given in Table 3.10. All figures in this table are expressed per train kilometre, and therefore should not be considered if stations on existing lines are being examined, with the methodology in Section 3.4.7.5 used to estimate vehicle costs in such cases. However, if new services are being considered then these figures should be used to estimate their costs.

3.5 Choice of Methodology

3.5.1 Approaches adopted

From the review of previous work in the field of local rail demand modelling (see Chapter 2) it was obvious that so far no single methodology had been developed which could fulfil all the objectives identified for this study. A mixed-method approach was therefore adopted, with a range of modelling methods tested to identify the optimal means of assessing demand at new stations. These included both those identified as contributing to previous best practice and several techniques which had not previously been applied to rail demand modelling. While Chapter 2 showed that both aggregate and disaggregate approaches have been used for rail demand modelling, this thesis will concentrate only on aggregate approaches. This is chiefly because extensive data is available for aggregate modelling (as detailed below in Section 3.6) whereas in contrast there is little up to date disaggregate data readily available for use, meaning that a sizeable data collection exercise would be necessary before disaggregate modelling could begin. The use of disaggregate data is also complicated by concerns over confidentiality and data protection, which are less of an issue for aggregate data. Furthermore, no clear evidence was found to suggest that disaggregate approaches have any accuracy advantages over aggregate models.

The simplest and most pragmatic aggregate models are trip rate and trip end models. Current models of this type, as detailed in the PDFH (ATOC, 2002), are highly location-specific, and this study will therefore focus on generalising these models so that they can be easily applied to forecast demand at any proposed new station regardless of location. Direct demand models are theoretically capable of giving a more detailed picture of expected usage at a new station, and the general form of such models can be derived from basic demand theory, as shown in Section 2.5.4. Work in this study will focus on refining and expanding this general form to produce cross-sectional direct demand models which are capable of forecasting rail demand at the flow level in a particular area of the UK.

The study methodology can therefore be summarised as follows. Trip rate, trip end and direct demand models will be recalibrated and extended (see Sections 4.2.1, 4.2.2, and 5.2 respectively) using the extensive electronic datasets which are now available (see Section 3.6). The two most promising local analysis techniques identified in Section 3.2, GWR and the spatial expansion method, will be applied to trip end models to assess their impact on model performance (see Section 4.3.2). Cluster analysis will be used to partition the calibration dataset for trip end models to establish whether distinct subcategories of local station could be statistically identified (see Section 4.3.3). The forecasting performance of direct demand models will be compared with the combined performance of trip end models and an alternative type of flow level model, the intervening opportunity trip distribution model (see Section 5.4). An automated site search procedure using GIS will be developed to allow potential station sites to be easily identified (see Chapter 6). The various modelling methods tested will be compared, with the best models combined with the site search procedure and an appraisal procedure to give an integrated methodology for assessing the potential for new stations within a given area, forecasting demand at these stations and assessing which sites have the best case for construction using both financial appraisal and social CBA (see Chapter 7).

3.5.2 Aggregate logit models

The development of aggregate logit models of modal split was also considered, and the calibration of such models was tested for travel to work flows in the Southampton area based on 2001 census data. However, problems were encountered with the reliability of these data (see Section 3.6.2.4) and the required data on bus and car travel were extremely time-consuming to collate and of questionable accuracy. Both nested logit (NL) and multinomial logit (MNL) models were calibrated, using several different formulations of the utility functions. The NL model given by (3.18) (upper level) and (3.19) (lower level) and the MNL model given by (3.20) were the only forms that gave usable results.

$$P_{ijt|PT} = \frac{\exp \beta\, GC_{ijt}}{\exp \beta \left(GC_{ijt} + GC_{ijb}\right)} \tag{3.18}$$

$$P_{ijPT} = \frac{\exp\left(\phi\, GC_{i*}\right)}{\exp\left(\phi\, GC_{i*}\right) + \exp\left(\tau\, GC_{ijPR}\right)} \tag{3.19}$$

$$P_{ijb} = \frac{\exp U_b}{\exp U_t + \exp U_b + \exp U_c} \tag{3.20}$$

Where:

P_{ijPT} is the probability of travelling from zone i to zone j by public transport

$P_{ijt|PT}$ is the probability of travelling from zone i to zone j by rail given that the journey is made by public transport

GC_{ijt} is the generalised cost of travelling from zone i to zone j by rail

GC_{ijb} is the generalised cost of travelling from zone i to zone j by bus

GC_{ijc} is the generalised cost of travelling from zone i to zone j by car

$GC_{i*} = \ln\left(\exp\left(\beta\, GC_{ijt}\right) + \exp\left(\beta\, GC_{ijb}\right)\right)$

U_t is the utility of travelling by train

U_b is the utility of travelling by bus

U_c is the utility of travelling by car

$U_Z = \alpha_Z + \beta\, GC_Z$

U_Z is the utility of travelling by mode Z

GC_Z is the generalised cost of travel by mode Z

α, β, φ, and τ are parameters determined during calibration

The fit of both models was extremely poor, with R^2 values of 0.234 for the NL model and

0.282 for the MNL model. Because of this poor model fit aggregate logit modal split models were not investigated further. The collection of generalised cost data for bus travel in particular was extremely time-consuming, and the accuracy of the resulting datasets was questionable. The observed data on modal split from the 2001 census is now several years out of date, and more recent data would be necessary to give trustworthy predictions of rail travel from a new station. While these data problems may be the primary reason for the poor model fit, no more trustworthy datasets were available. Nonetheless, if better data did become available for both the dependent and independent variables it is possible that this model form could assist with forecasting the impact of new railway stations or services on modal split. Full details of the aggregate logit modal split models investigated for this study are given in Blainey (2009b) (see Appendix 1).

3.6 Data Sources

3.6.1 Overview

The type and quantity of data which is available obviously has a major impact on the type of model which can be used to predict demand. For example, the development of disaggregate models is still constrained by a lack of data, as it is effectively impossible to obtain accurate data on the individual behaviour of all members of a population. While data from smart cards may in theory change this situation, as they become more widely used, by providing detailed disaggregate data on trip patterns, in practice limitations of the data collected (Bagchi & White, 2005) and issues of privacy and confidentiality mean that comprehensive datasets are unlikely to appear. Even if such data did become available the sheer volume of information involved would make model calibration difficult and computationally intensive.

A wide variety of data sources were used or investigated during the development of the models described here, and these are detailed below.

3.6.2 Travel behaviour data

3.6.2.1 ORR station usage data

Before the advent of computerised ticketing there was no reliable or comprehensive source of information on ticket sales (Preston, 1987) and therefore on the number of rail journeys made and their origins and destinations. However, this situation has now changed, and extensive data on rail usage is now automatically collected and can be used as the basis for model calibration.

The Office of Rail Regulation (ORR) supplies, via their website, estimates of the total numbers of people entering, exiting and interchanging at all National Rail stations over one year periods. This is based on LENNON ticket sales data, but a significant account of data processing is undertaken before the final figures are released (see AEA Technology Rail (2006), Shah (2007), and Georgiou (2008)). This data was used as the basis for calibrating the trip end models developed during this project, where no information on trip destinations was required.

3.6.2.2 LENNON data

For direct demand models it is necessary to have information on destination stations as well as origins, and this was provided by LENNON data supplied by Arriva Trains Wales in its raw state, giving the total trips made in 2006-7 for 41,089 flows to and from local stations in South Wales. This data has shortcomings with regard to station accessibility as it does not give information on ultimate trip origins and destinations, and it is not possible to segment journeys by purpose. It can be difficult to establish which end of a flow is the trip origin (or generator) and which is the destination (or attractor). There are also problems regarding the allocation of trips between stations in geographical ticketing groups (eg tickets sold to 'London Stations' or 'Manchester Stations'), and with the representation of trips made on London and PTE travelcards. However, there are no aggregate data sources available which can overcome this problem.

3.6.2.3 Station access data

The Greater Manchester Travel Survey (GMATS) was carried out in 2002-3, and included the collection of data on rail trips to and from stations in Greater Manchester, incorporating information on access to and from the stations and ultimate trip origins and destinations. Data on 7,945 rail trips was provided by Greater Manchester PTE for this study, and was intended to be used to investigate the accuracy of station catchment definition methods. However, imperfections in the survey questionnaire meant that the accuracy of this data could not be relied upon.

It was hoped that data on ultimate trip origins and destinations could be obtained from the National Rail Travel Survey (NRTS) via the DfT (DfT, 2007a), but concerns over data confidentiality and delays in the data becoming available meant that it was not possible to gain access to the relevant information. While several formal data requests were made, unofficial advice from the Department suggested that these were not likely to meet with positive results, and this data source was not pursued further.

Because it was not possible to obtain data on ultimate trip origins and destinations from GMATS or NRTS, a survey of access and egress to stations was undertaken on the Rhymney line as part of this study. Full details of this survey are given in Section 5.4.

3.6.2.4 Census journey to work data

Journey to work data from the 2001 Census is available online via the Centre for Interaction Data Estimation and Research via the Web Interface to Census Interaction Data (WICID), Until recently this was extremely awkward to use, but it has now been partly redesigned making it somewhat easier to extract the required information. While this data should in theory provide detailed information on the size of commuting flows at inter-output area level, closer examination of the data revealed a number of discrepancies. These apparent errors are the results of modifications made as part of a disclosure control process known as ‘Small Cell Adjustment Methodology’ (SCAM). To protect individual confidentiality, table cells with a small value (0-3) were altered as shown in Table 3.11 (CIDS, 2007):

Table 3.11: SCAM adjustments to Journey To Work data

Initial Value	Adjustments
0	Initial value retained
1	Rounded to either 0 or 3, with 0 the more likely result
2	Rounded to either 0 or 3, with 3 the more likely result
3	Initial value retained

SCAM was applied independently to each output table and at each spatial level, which means that figures and totals are not consistent between tables. It also means that the numbers of people recorded as travelling to work by rail can not entirely be trusted, particularly for small flows. There is no obvious way to correct these discrepancies, as there is no way of knowing which set of figures is most accurate, although if all flows with less than 4 trips are excluded then this should assure the accuracy of the remaining dataset. There are also a few obvious errors from the SCAM procedure which can be corrected, the main one being trips to work apparently undertaken by underground in areas where no underground rail system exists.

3.6.3 Transport provision data

The most obvious source of rail and bus timetable data are the various journey planning websites which are now available. Several sites supply train times, with the Deutsche Bahn website (bahn.hafas.de/bin/query.exe/e) found to be the quickest (Deutsche Bahn AG, 2009), and this was used to obtain data on average travel times and train frequencies at individual stations and between station pairs. The Transport Direct website (www.transportdirect.info) was used to obtain similar data for bus services (Atos Origin, 2009), although it did sometimes return rather strange results. The best source of data on rail fares was found to be the National Express East Coast ticket sales website (www.nationalexpresseastcoast.com). While the data from such websites should be reasonably accurate, using them to obtain data on large numbers of stations or flows is extremely time-consuming.

Rail timetable data was also obtained as a Common Interface Format (CIF) file, which contained details of the complete National Rail timetable for 2007. This file is made up of almost 1.5 million lines of data, which means that extracting data on train frequencies and journey time is not straightforward (Armstrong et al., 2007). However, Perl scripts were developed by Dr John Armstrong which automatically interrogated the CIF data to extract the required data in a form allowing easy incorporation in the demand models. These could only extract data on flows with a direct rail service, so journey times and frequencies for flows requiring interchange still had to be calculated manually.

Data on generalised journey times (GJT) can be obtained from the MOIRA system alongside LENNON data. However, such GJT data was not included with the LENNON data supplied by Arriva Trains Wales, and therefore was not used in the calibration of the models developed here.

A final source of transport provision data is the Station Facilities Database which is freely accessible via the National Rail website (www.nationalrail.co.uk), and gives detailed information on the facilities available at all National Rail stations, including data on car park capacity.

3.6.4 Spatial data

3.6.4.1 OS Strategi and Meridian data

Spatial data on the road and rail network was obtained from the Ordnance Survey (OS) via EDINA Digimap. The use of both the Strategi and Meridian data formats was investigated. Both datasets are in vector format, as it is not possible to adequately represent networks in raster format (see Section 3.7.2). The Meridian data was found to be preferable for the applications required here, despite its larger file sizes, because it gives comprehensive coverage of the rail network and in theory also of the road network. This data is supplied in the form of tiles, and therefore the tiles for the whole area of England and Wales were progressively downloaded. In fact the coverage of the road network was found to be imperfect, with a number of road links omitted or truncated seemingly at random. While it was possible to correct some of the most obvious errors, it was not feasible to cross-check the entire Meridian road network against other more complete maps. More detailed and comprehensive data on the road network was, at the start of the study period, provided by the OS Landline format, but roads were not represented in a way that allowed easy conversion into a form suitable for GIS-based network analysis.

Spatial data on the boundaries of output areas from the 2001 Census was downloaded from EDINA UKBorders for the whole area of England and Wales, allowing demographic data to be displayed and analysed spatially. Other boundary data is also available from this source, and an outline of the coastline of England and Wales was obtained to provide a boundary for the interpolated grid-based mapping required for the models developed which incorporate spatially-varying parameters.

3.6.4.2 OS Mastermap data

It was hoped that the Integrated Transport Networks layers of the new OS Mastermap data, which became available in 2007 and replaced Landline, would make it possible to include a more realistic representation of the road network and potentially also the rail network in the modelling process. However, no rail network layer has so far been made available, and while the road network data appears to be of very good quality, the file sizes involved are extremely large making it impractical to store and convert sufficient data to be of use for the case studies in this project.

3.6.4.3 Google Earth

High resolution aerial and satellite imagery is available via Google Earth, and was used here to check the feasibility of new station construction and to provide data on station car park capacity.

3.6.5 Demographic data

Demographic data was obtained from the web interface to census aggregate outputs and digital boundary data (CASWEB). A huge range of data is available from this source, with the following variables used for the models developed here, downloaded at output area level for all of England and Wales:

- resident population (KS001)
- socio-economic class (KS014)
- mode of travel to work (KS015)
- car/van ownership (KS017)

Similar data is available from the Neighbourhood Statistics website, which also supplies information on the average distance travelled to work and the mode used for both the workplace and resident population, along with the size of the workplace population. Data on these variables was again downloaded for the whole of England and Wales.

3.7 Geographical Information Systems (GIS)

3.7.1 Overview of GIS

Geographical Information Systems (GIS) have risen enormously in prominence over the last two decades, so that 'the integration of transport models and technologies such as GIS has become a major requirement in any process of transport planning' (Dueker & Ton, 2000). A GIS can be defined as a system incorporating the following three basic components for handling spatial data (Dueker & Ton, 2000):

- A Graphical User Interface (GUI) by which the user interacts with the system.
- A database management system (DBMS) which stores and manages the data required by the system.
- Spatial modelling tools enabling users to manipulate the data stored by the GIS. This capability to perform spatial operations is the key difference between GIS and other computer programs .

A range of proprietary GIS systems are available, which can be used to fulfil the following four key generic functions (Dueker & Ton, 2000):

- GIS digital mapping, which allows users to communicate interactively with geographic features. Information is structured into geocoded spatial data and attribute data, with spatial objects represented as either raster or vector data
- GIS data management, which organises map features into layers and links non-spatial attribute data to spatial features.
- GIS data analysis, which allows users to set up spatial, attribute or combined queries for analysing relationships between information from databases.
- GIS data presentation, which allows users to extract and or present data using thematic mapping concepts. Wide ranging visualisation capabilities allow detailed exploratory spatial analysis (Hearnshaw & Unwin, 1994). Internet-based GIS can allow a wide range of users to visualise, share and manage geographic information (Tang & Waters, 2005).

3.7.2 GIS data formats

Spatial data can be represented in a GIS in one of two formats. The first is the vector format, where features are represented using a coordinate based structure as points, lines or polygons. The resulting maps closely resemble conventional paper maps, with each possible position in the map having a unique coordinate value. Alternatively, data may be represented in raster format, where attributes are associated with a particular set of grid cells. The location of geographic objects is defined by the row and column position of the cells they occupy, with the spatial resolution (and accuracy) of the map dependent on cell size. Raster data has a simple structure (although less compact than vector data), and makes overlay operations straightforward. However, its graphical output may be unattractive (the cell-based structure can give a blocky appearance), and network analysis (crucial in many transport applications) using such data is extremely difficult. In contrast, network analysis is straightforward using vector data, due to its efficient encoding of topology (Ochieng, 2006), and the spatial resolution of the resulting maps is limited only by coordinate accuracy. Fortunately, most GIS can process and display vector and raster data in parallel, allowing the most appropriate format to be used for each application.

3.7.3 Use of GIS in rail demand modelling

The use of GIS in rail demand modelling so far has been rather limited, but Preston (2001) suggests a number of potential areas where they could be useful, such as in improving aggregate models by adopting a more flexible approach to catchment areas definition (Lythgoe, 2004) or by providing a wider range of geo-demographic data for trip rate models. Some GIS-based enhancements have already been realised, with Whelan & Wardman (1999b) using a GIS to provide information on populations and socio-economic characteristics for their summation model, and Lane et al. (2006) using TransCad to collect and process data for their trip end models. Similarly Wardman et al. (2007) used the CACI InSite system to provide information on population, average household income, car ownership, employment status, socio-economic group, age, gender and drive time to station for station catchments. Another example is the use by Lythgoe (2004) of a GIS representation of the road network based on OS 'Strategi' tiles to determine access times and distances from zonal centres of population to stations.

At the disaggregate level Buliung & Kanaroglou (2006) developed extensions to ArcGIS to investigate household level activity/travel behaviour in Portland, Oregon, which while not specifically related to rail demand indicates the potential for work in this area. Hsiao et al. (1997) used GIS to conduct a system-wide analysis bus stop access in Orange County, California, comparing catchments defined using a buffer area (area within a certain straight line distance from station) and using a maximum road-network based distance.

GIS are also used in commercial applications such as Experian's MOSAIC segmentation systems, which provide a wide range of geodemographic classifications enabling businesses to improve their marketing strategies (Experian, 2006).

3.7.4 Use of GIS in this study

The sheer volume of data available for use in rail demand models can seem overwhelming, and some means of processing and integrating these datasets for this study was obviously needed. GIS undertook this role, being able to integrate data from different sources, to process this data to provide further information, and to present the results of data processing and modelling in a form allowing easy understanding and interpretation. For example, GIS were used to combine numerical data on station usage with spatial data on the geographical location of stations from Digimap, demographic data on population characteristics and digital boundary data to allow trip end models to be calibrated. The visual display capabilities of GIS were crucial when using local analysis techniques, as they allowed local statistics to be mapped, enabling the identification of spatial variations in these statistics and allowing models and datasets to be corrected and enhanced.

The main GIS used in the work reported here was MapInfo Professional 8.5, with the RouteFinder 3.41 add-on used to provide network analysis functionality. This package does have some shortcomings, particularly with regard to network analysis, but was the best available at the time. The use of an alternative system, ArcGIS, was also investigated, as it was hoped that this could improve the quality of network analysis as it had more extensive functionality in this are. However, while ArcGIS was used for some aspects of the site search procedure, problems with the University license for ArcGIS meant that its availability was sporadic. It was also found to be slightly less stable in operation than MapInfo, with a tendency to occasionally crash and fill the hard disk with temporary files. A third GIS, Google Earth, was used to collect data and in the final stage of the site search procedure, but was not sophisticated enough for more extensive use.

3.8 Choice of Case Studies

Because the data required for the calibration of trip end models (see Chapter 4) was readily available for all areas of Great Britain, there were no restrictions on the choice of case study area. An area centred on Southampton was chosen for initial analysis, because it was local to Southampton University and its characteristics were therefore familiar to those involved in the project. Consideration was given to generalising these models across the whole of the UK, but collation of the required transport network data proved to be extremely time consuming. Such data had already been obtained for the whole of England

and Wales for other purposes, and it was felt that the time required to add Scotland to the case study area would be better spent on other aspects of the project. England and Wales were therefore chosen as the extended case study area for calibrating the trip end models, and because of this the site search procedure was applied to the same area (see Chapter 6).

The LENNON data required for flow level models (see Chapter 5) is not freely available, and it was therefore necessary to consider carefully the choice of case study area before approaching train operating companies about the supply of such data. Problems with the allocation of trips from PTE travelcards meant that a non-PTE area was preferred, and eight such areas were identified. The characteristics of these areas are compared in Table 3.12 with their boundaries being set visually to give the most discrete network possible. The availability of a Network Rail Route Utilisation Strategy (RUS) for the area was seen as an advantage because additional survey data will often have been collected during consultation for these documents and could prove useful for demand modelling. Table 3.12 reflects the availability of RUSs at the time the choice of case study for flow level models was made (early 2008).

Table 3.12: Characteristics of possible case study areas for flow level models

Area	Stations	Advantages	Disadvantages
South Hampshire	40	Centred around Southampton In SWML RUS (Network Rail, 2006)	Most lines shared by long-distance services
Bristol	25	Fairly close to Southampton	Most lines served by non-local services Some stations have very low service levels
South-East Wales	87	Reasonably discrete network Ebbw Vale line reopened in 2008 and Vale of Glamorgan line reopened in 2005	High station density on most lines - perhaps limited potential for new stations
Edinburgh	47	New stations and lines recently opened and proposed All routes contained in Scotland RUS (Network Rail, 2007a)	A long way from Southampton Some routes served by longer distance services
East Midlands	67	Variety of service types – suburban and more rural Perhaps the most potential for new station sites	Several routes form part of longer distance services Not a very well-defined network
Norwich	33	Reasonably discrete network All routes contained in Greater Anglia RUS (Network Rail, 2007c)	Some stations have very low service levels
North Lancashire	34	Some routes contained in North-West RUS (Network Rail, 2007b)	Not a very discrete network Some routes served by longer distance services
Teesside	41	Reasonably discrete network All routes contained in East Coast Main Line RUS (Network Rail, 2008b)	Some stations have very low service levels

After weighing up the advantages and disadvantages described in Table 3.12, South-East Wales was selected as the preferred case study area, with the East Midlands as a reserve option if data could not be obtained for South-East Wales. Agreement was reached with Arriva Trains Wales for the supply of LENNON data, and South-East Wales was therefore used as the case study for flow level models. The availability of this data was not certain until after the majority of the work on trip end models had been completed. It was therefore necessary to recalibrate the trip end models for South-East Wales to allow the results of the different modelling approaches to be compared (see Section 7.1.2).

3.9 Conclusions

This chapter has described some methodologies which have the potential to enhance rail demand models but which have not previously been applied in this field, including local analysis techniques, such as Geographically Weighted Regression, and cluster analysis. A brief review of appraisal methods has been carried out highlighting the strengths and weaknesses of different approaches. This was followed by a description of ways in which the various costs and benefits associated with new stations can be estimated. The choice of methodology for this study has been explained and its application briefly outlined. The range of data sources available for the study has been described, with any inherent advantages or disadvantages highlighted. An overview of GIS was provided and the use of GIS in this study described. The reasoning behind the choice of case studies for the various elements of the work was explained. The following chapters will show how the selected data sources were used in these case study areas to investigate the two main model types and to develop the site search procedure.

Chapter Four: Trip Rate and Trip End Models

4.1 Introduction

This chapter details the development of generalised trip rate and trip end models of local rail demand. These are the simplest types of cross-sectional rail demand models, and are therefore an obvious place to start when developing an easy-to-use procedure for assessing the case for new stations. Previous applications of trip rate models in demand forecasting for new stations (for example Preston, 1987), while achieving a reasonable level of accuracy, have lacked spatial transferability and by their nature have not accounted for a number of important variables. The use of trip end models in the UK has been fairly limited, and has again tended to focus only on small areas and to account for the effects of only a small number of independent variables (Preston, 1991b). The data collection and aggregation methods have also tended to be rather crude (inevitably, given the technological limitations when they were developed), and this has limited the usefulness of such models. However, promising results have been obtained using trip end models with a sizeable calibration dataset of US commuter and light rail stations (Lane et al., 2006), and this suggests that there is potential for similar models to be effective at forecasting local rail demand in the UK.

The chapter first describes the calibration of trip rate models for the South Hampshire area, with a variety of catchment population definitions tested in Section 4.2.1. These models are then extended by the addition of other independent variables to produce a range of trip end models, described in Section 4.2.2. This small case study area was used to test whether or not such models were likely to give accurate results, and once this had been confirmed the case study area was then extended to cover all local stations in England and Wales (to give a high level of spatial transferability) with a further range of trip end models tested for this area in Section 4.3.1. The use of Geographically Weighted Regression (GWR) to enhance model performance is investigated in Section 4.3.2 with the results compared to those from the global regression models. Section 4.3.3 details the use of cluster analysis to partition the calibration dataset and identify sub-classes of local station, with model results again compared to those obtained using the whole dataset. Sensitivity analysis of the results from the best trip end models is then carried out in Section 4.3.4. The best models are incorporated in a spreadsheet-based demand forecasting tool for new local railway stations in England and Wales (see Section 4.4). Finally, the results of the trip rate and trip end modelling are summarised in Section 4.5.

4.2 South Hampshire case study

4.2.1 Trip rate models

The simplest form of local rail demand model is the trip rate model, which aims to predict the total number of trips from a station based on a single independent variable, the station catchment population. This can either be a single figure or divided into two or more zones. While the ultimate aim of this section of the research was to develop a generalised model capable of forecasting demand at any new station, regardless of location, the modelling techniques were first tested on a relatively small case study area, made up of 40 local stations in and around South Hampshire and the surrounding area, as shown in Figure 4.1. This case study area was used initially to confirm whether or not trip rate and trip end models were likely to give accurate forecasts before undertaking the major data collection and processing exercise which would be required for calibration at a national level.

Figure 4.1: Stations in South Hampshire case study area

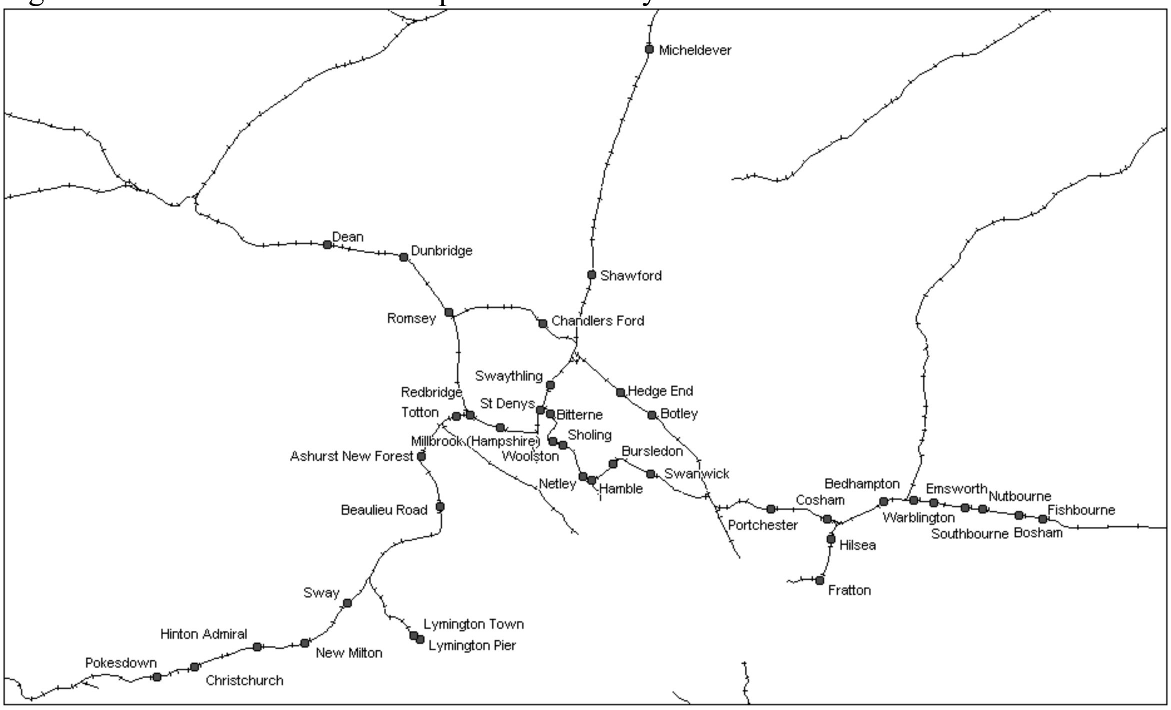

Population data from the 2001 census on the resident population in all output areas in the study area was loaded into MapInfo, and the Distance Calculator tool was used to calculate the straight line distance from the geographical centroid of each output area to its nearest station. This data was then sorted by station, allowing output areas within a set distance of the stations to be isolated. In reality catchments for different stations are likely to overlap, but it is difficult to model this without having information on actual station access/egress trip patterns and it was therefore assumed that population units would only use the station closest to them. Data is available from the Neighbourhood Statistics website on the number of people travelling to work by different modes from each output area. This data was used to calculate the percentage of people travelling to work by train from each output area in Hampshire, Southampton and Portsmouth (5446 in total). These percentages were plotted against the straight line distance from each output area to its nearest railway station, giving Figure 4.2. This appears to show, somewhat surprisingly, that there is no significant correlation between distance from a railway station and the propensity to commute by train. This conclusion was backed up by an attempt to fit a best fit line to the data using SPSS, which gave results which were very insignificant.

Figure 4.2: Scatter plot showing % of workers in output areas commuting by train plotted against distance from nearest railway station

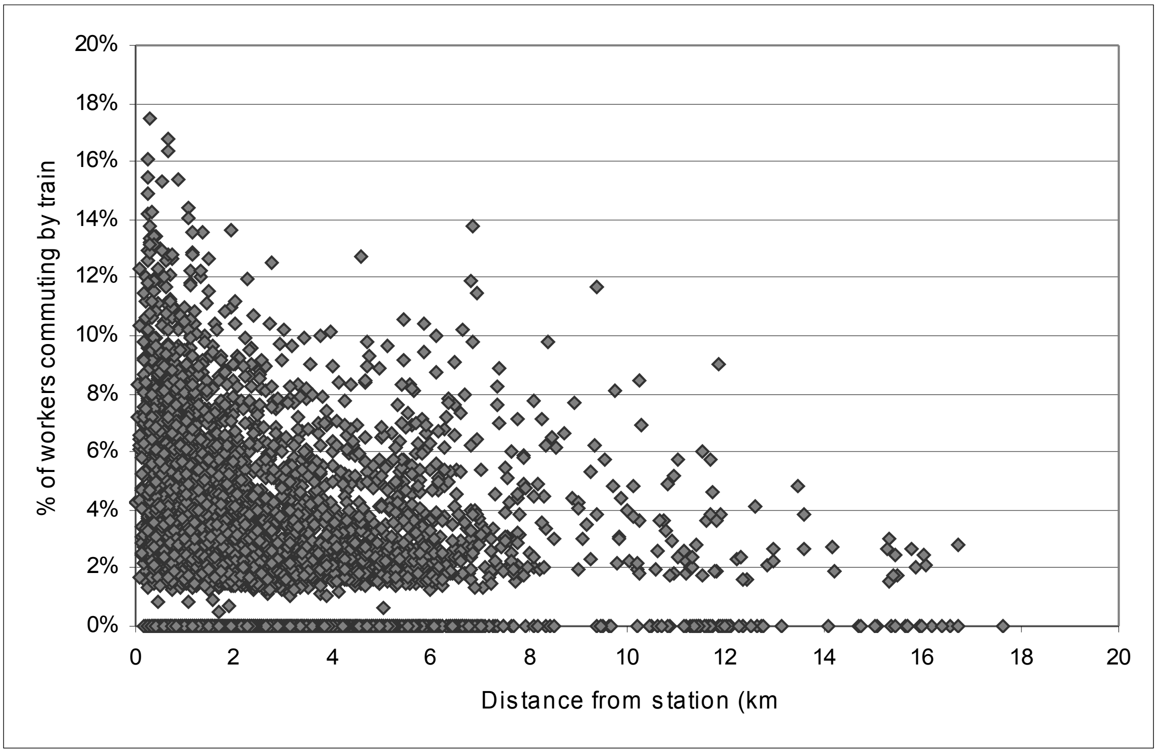

While the PDFH (ATOC, 2002) suggests that catchments for minor stations should have a 2 km boundary, as Figure 4.2 shows no obvious station access distance cut-off point a range of distance values were tested. Observed data on total station origins and destinations was obtained from the ORR station usage spreadsheets. SPSS was used to estimate parameter values for the simple linear regression models 4.1 and 4.2 (the latter model containing two population zones) using these values, giving the results in Table 4.1.

$$\hat{V}_i = \alpha + \beta P_{id} \tag{4.1}$$

$$\hat{V}_i = \alpha + \beta P_{id} + \gamma P_{ide} \tag{4.2}$$

Where:

$\hat{V}_i$ is the predicted total number of passenger entries and exits per year at station i

P_{id} is the total resident population within d km of station i for whom station i is their nearest station

P_{ide} is the total resident population between d and e km of station i for whom station i is their nearest station

Model fit is reported here using both the conventional R_{adj}^2 values and also mean absolute deviation (AD) values, defined by (4.3). AD can be used to compare results of different model forms where the variance of the dependent variable will be different, but is not a perfect measure of model fit, because it gives more weight to errors at stations with a low level of demand than to errors of the same magnitude at stations with a high level of demand.

$$AD = \left(\sum_i^n \left| \frac{\hat{V}_i - V_i}{V_i} \right| \right) / n \tag{4.3}$$

Where:

V_i is the observed total of passenger entries and exits per year at station i

n is the number of flows in the calibration dataset

Table 4.1: Summarised results from calibration of Models 4.1 and 4.2

Model	d	e	Intercept		β parameter		γ parameter		R_{adj}^2	AD
			Value	t stat	Value	t stat	Value	t stat		
4.1	1.5	n/a	-4783.251	-0.125	16.691	6.271	n/a	n/a	0.496	1.494
	2	n/a	-5786.506	-0.158	12.379	6.659	n/a	n/a	0.526	1.504
	3	n/a	-8346.501	-0.199	9.795	5.645	n/a	n/a	0.442	1.634
	5	n/a	-15188.1	-0.328	9.077	5.122	n/a	n/a	0.393	1.679
	10	n/a	-49570.1	-1.036	9.653	5.645	n/a	n/a	0.442	2.018
4.2	1.5	5	583.022	0.014	17.269	5.101	-1.139	-0.282	0.483	1.454
	2	5	25179.701	0.602	13.187	6.890	-6.745	-1.458	0.540	1.465
	2	10	-1685.639	-0.035	12.433	6.453	-0.564	-0.132	0.514	1.474

Table 4.1 shows that using a single 2 km population zone to represent the catchments does in fact give the best results, as while adding a second 2-5 km population zone improves model fit, the additional parameter is both insignificant and of the wrong sign. Figure 4.3 shows the catchments which result for the case study dataset when all output area centroids

within 2 km of a station are allocated to their nearest station. An alternative allocation method would be to consider the spatial extent of output areas rather than their centroids, allocating the whole population to a station if the whole output area was within 2 km (if this was the distance cut-off point) but only a proportion of the population if only part of the output area was within 2 km. However, implementing this for even a small number of stations made the computation of catchment populations extremely complex, and this was not therefore pursued here. If such a method were to be used, then it would be desirable to incorporate information on the distribution of populations within output areas, as particularly for larger output areas the population is likely to be concentrated in a small part of the total area.

Figure 4.3: Station catchments for South Hampshire case study based on 2 km radius

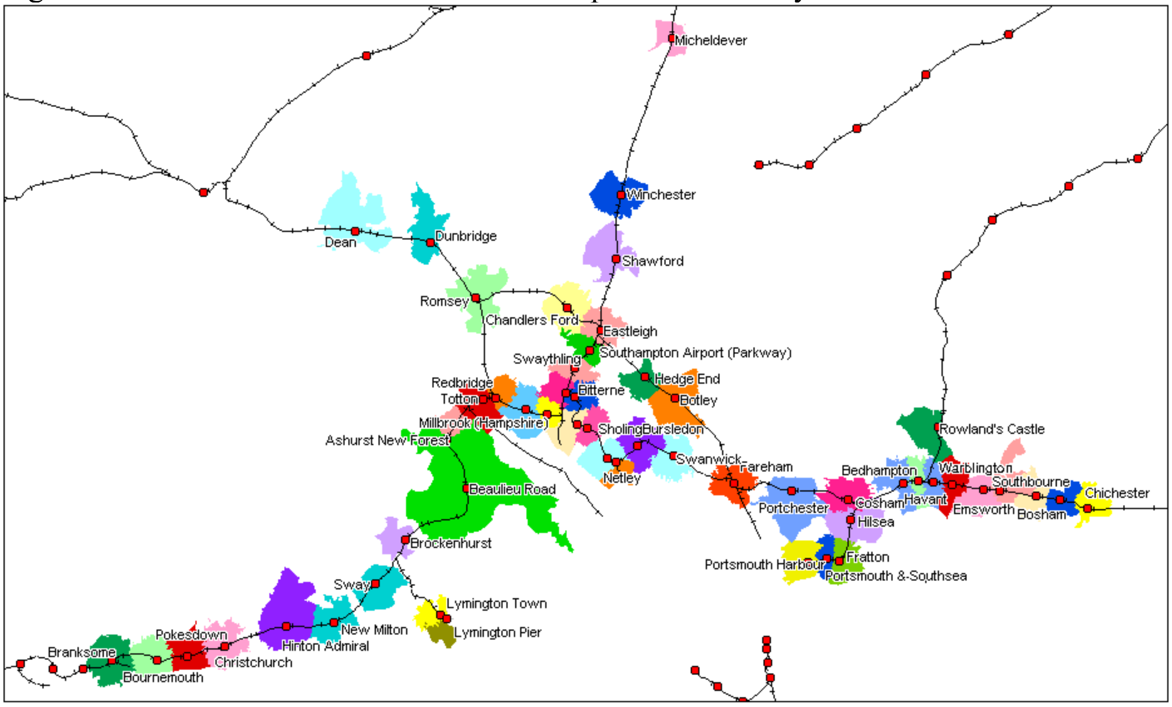

The choice of catchment boundary definition method can lead to a great deal of variation in the shapes and sizes of catchment areas. This is partly because output areas in more rural areas can be very large, and while their centroid may be within 2 km of a station their boundaries are often much further from the station meaning that, for example, in Figure 4.3 Beaulieu Road has a very large catchment area made up of a single output area. Equally parts of output areas may be very close to stations, but because their centroid is more than 2 km from a station they are not included in its catchment area. In Figure 4.3 Shawford is a good example of this where the area immediately to the south-east of the station does not form part of its catchment. In addition to the geographically-weighted centroids used in the calibrations described above, the boundary data used to map the catchments includes the coordinates of population-weighted centroids. Using these to allocate output areas to stations should increase the realism of catchments, particularly in sparsely-populated areas. The 2 km boundary catchments defined using population-weighted centroids are shown in Figure 4.4, and the 2 km variant of Model 4.1 was recalibrated using population values based on these catchments giving the results summarised in Table 4.2.

Table 4.2: Summarised results from best variant of Model 4.1 using population-weighted centroids to define catchments

Model	d	Intercept		β parameter		R_{adj}^2	AD
		Value	t stat	Value	t stat		
4.1	2	-8259.752	-0.231	12.559	6.928	0.546	1.521

Figure 4.4: Station catchments for South Hampshire case study based on 2 km radius using population-weighted centroids

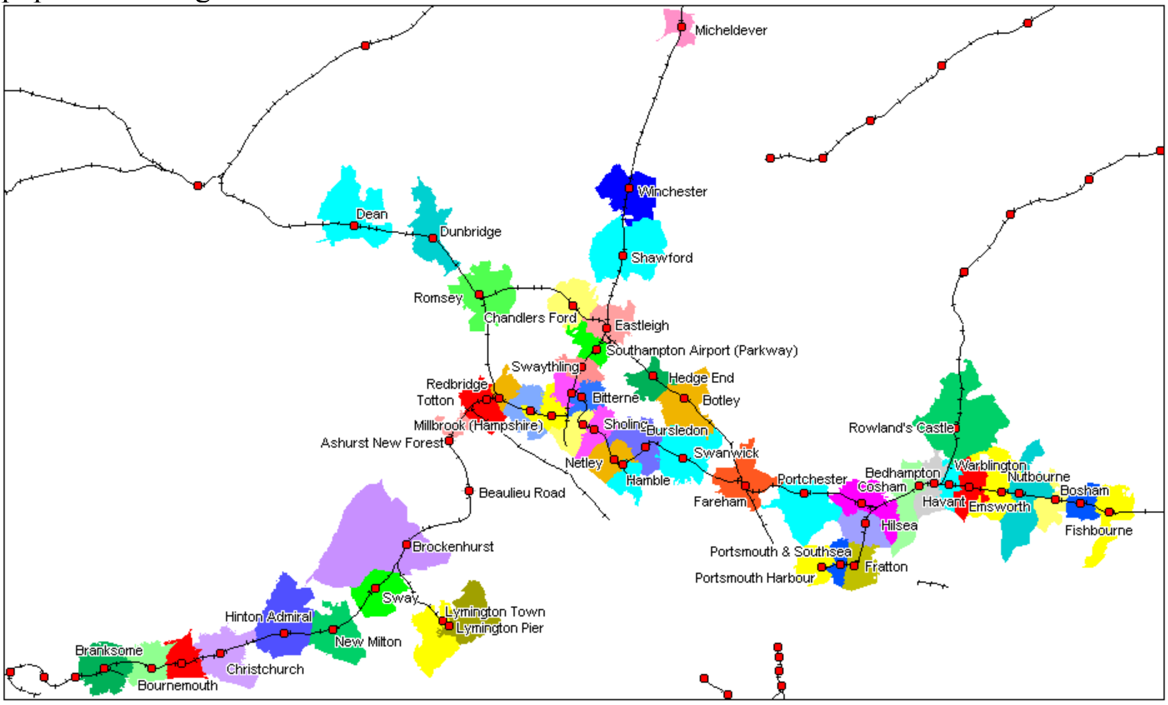

Table 4.2 shows that using population-weighted centroids to define catchments gives a better model fit than using geographically-weighted centroids, and the former method was therefore used in later models. However, it was not possible to verify whether this method gives a better representation of actual travel behaviour as data on observed trip origins and destinations was not available.

Allocating output areas by straight line distance is rather unrealistic, as it takes no account of geographical boundaries to access. It can lead to obvious errors in catchment definition, with for example output areas in Gosport being allocated to Portsmouth Harbour in Figures 4.3 and 4.4 despite being separated from the station by the harbour itself. Routefinder for

MapInfo was therefore used together with Ordnance Survey Meridian data on the road network to calculate road distances from output areas to stations. As detailed in Section 3.6.4.1, this data is far from perfect, but was the best option available when this work took place. All output areas within the case study area were therefore allocated to their nearest station based on road network distances and population-weighted centroids. The resulting catchments based on a 2 km road network distance boundary are mapped in Figure 4.5.

Figure 4.5: Station catchments for South Hampshire case study based on 2 km road network distance

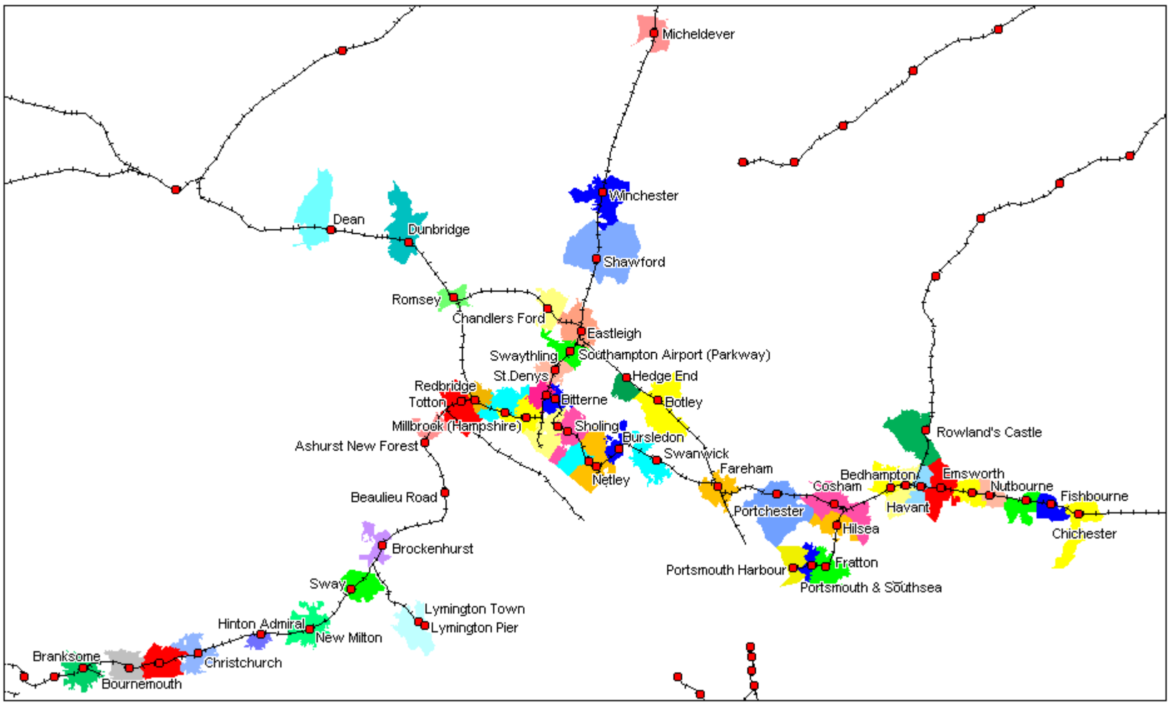

In general the network-based catchment areas and populations were smaller than the straight-line based populations for the same distance cut-off, although there were a few cases where catchment populations increased. This is because some output areas are within 2 km of more than one station and the nearest station in a straight line was not the same as the nearest station via the road network. The best variants of Models 4.1 and 4.2 were recalibrated using these network-based catchment populations, giving the results summarised in Table 4.3.

Table 4.3: Summarised results from best variants of Model 4.1 and 4.2 using road distance to define catchments

Model	d	e	Intercept		β parameter		γ parameter		R_{adj}^2	AD
			Value	t stat	Value	t stat	Value	t stat		
4.1	2	n/a	10897.282	0.325	14.853	6.986	n/a	n/a	0.551	1.391
4.2	2	5	9870.076	0.270	14.742	5.646	0.244	0.076	0.539	1.385

Table 4.3 shows that using road distances to define catchments gives a better model fit than using straight line distances. Using a single zone still gives superior results, with only one population parameter significant in Model 4.2. A further alternative catchment definition method is to use a maximum access time to set catchment boundaries rather than a maximum distance. Ideally the estimated access times would have accounted for the effects of mandatory speed limits and congestion, but no data was available on these, meaning that speeds were only differentiated by road class. This distinction between access distance and access time seemed likely to be more important for inter-urban journeys than for local journeys because distances between stations and ultimate origins/destinations would be expected to be shorter for the latter, and therefore there would be less variation in the types of road used. It is also likely that a higher proportion of local passengers use non-motorised access/egress modes, where road type will have little effect on travel time. Road speeds were initially based on the defaults from the sample road data supplied with RouteFinder, as shown in Table 4.4.

Table 4.4: Initial road speeds used in defining access-time based catchments

Road type	Speed
Motorway	80 kph
A road	65 kph
B road	40 kph
Other road	25 kph

Model 4.4 was calibrated using a range of maximum access time cut-off points to define catchments, giving the results summarised in Table 4.5

$$\hat{V}_i = \alpha + \beta P_{it} \tag{4.4}$$

Where:

P_{it} is the total resident population within t minutes drive time of station i for whom station i is their nearest station

Table 4.5: Summarised results from calibration of Model 4.4

t	Intercept		β parameter		R_{adj}^2	AD
	Value	t stat	Value	t stat		
2	42670.553	1.575	29.169	8.172	0.628	1.565
3	36937.062	1.574	14.442	9.984	0.717	1.377
3.5	28863.726	1.201	12.227	9.945	0.715	1.348
4	25343.913	1.003	11.008	9.463	0.694	1.383
5	17102.334	0.608	9.614	8.515	0.647	1.402
6	11058.853	0.371	9.009	8.083	0.623	1.368
7	11430.01	0.362	8.377	7.502	0.586	1.389
8	13215.857	0.406	7.998	7.175	0.564	1.402

This shows that catchments defined based on access times give a much better model fit than catchments based on distance. The choice of access time cut-off was not obvious, as while 3.5 minutes gave a clearly superior AD value, 3 minutes had a marginally better fit in terms of R^2. The former value was adopted, but the road speeds used to calculate access times are not necessarily realistic. There is little data available on actual average speeds, and while the DfT supplies limited data on average speeds in uncongested conditions, these are unlikely to be sustained continuously. The effect on model fit of using a variety of different road speeds was therefore investigated. It was necessary to vary the relative speeds of different road types instead of merely altering all speeds by the same proportion, as the latter had already been indirectly investigated by the comparison of different drive time cut-off points. The speed for 'other roads' was therefore held constant and the level of differentiation with other classes altered. This also meant that it was not necessary to test the different sets of road speeds using different maximum travel times, so a 3.5 minute maximum was used in all cases. Table 4.6 details the sets of speeds used, with the DfT speeds based on the average speeds in uncongested conditions given in the Transport Statistics Bulletin (DfT, 2006a). Some adaptation of these was necessary because the DfT classification of road types was not the same as that used in the road network data.

Table 4.6: Road speeds

Road type	Low speeds	Original speeds	Speed limits	DfT 1	DfT 2
Motorway	80 kph	80 kph	115 kph	115 kph	115 kph
A road	50 kph	65 kph	105 kph	95 kph	65 kph
B road	30 kph	40 kph	80 kph	80 kph	50 kph
Other road	25 kph	25 kph	25 kph	25 kph	25 kph

Model 4.4 was recalibrated using these speed sets (with t = 3.5) giving the results summarised in Table 4.7. This shows that the original set of speeds used gave the best model fit and were therefore the most accurate in defining station catchments.

Table 4.7: Summarised results from calibration of Model 4.4 using different speed sets

Speed set	Intercept		β parameter		R_{adj}^2	AD
	Value	t stat	Value	t stat		
Low	23144.287	0.869	15.347	8.935	0.669	1.320
Original	28863.73	1.201	12.227	9.945	0.715	1.348
Speed limits	46277.203	1.422	9.306	6.168	0.487	1.776
DfT 1	43283.96	1.278	10.034	6.013	0.474	1.744
DfT 2	31806.469	0.935	13.503	6.28	0.496	1.621

Despite the good results obtained above, using an arbitrary cut off point to define station catchment populations is unlikely to give an accurate representation of actual travel behaviour, because logically those people living close to a station are, all other things being equal, more likely to use that station than people living further away. This can be represented in the models by weighting population units by distance from their nearest station and then combining them to give catchment populations. Several weighting functions were tested (4.5-4.10), similar to those used by Wardman et al. (2007) to enhance direct demand models. It was important that the weight given by the functions was a real number where $0 \leq t < \infty$, and thus (t + 1) was used rather than t in the power and logit functions. The comparative 'shapes' of these functions are shown in Figure 4.6.

Figure 4.6: Weights given by functions 4.5-4.10 when ψ=1

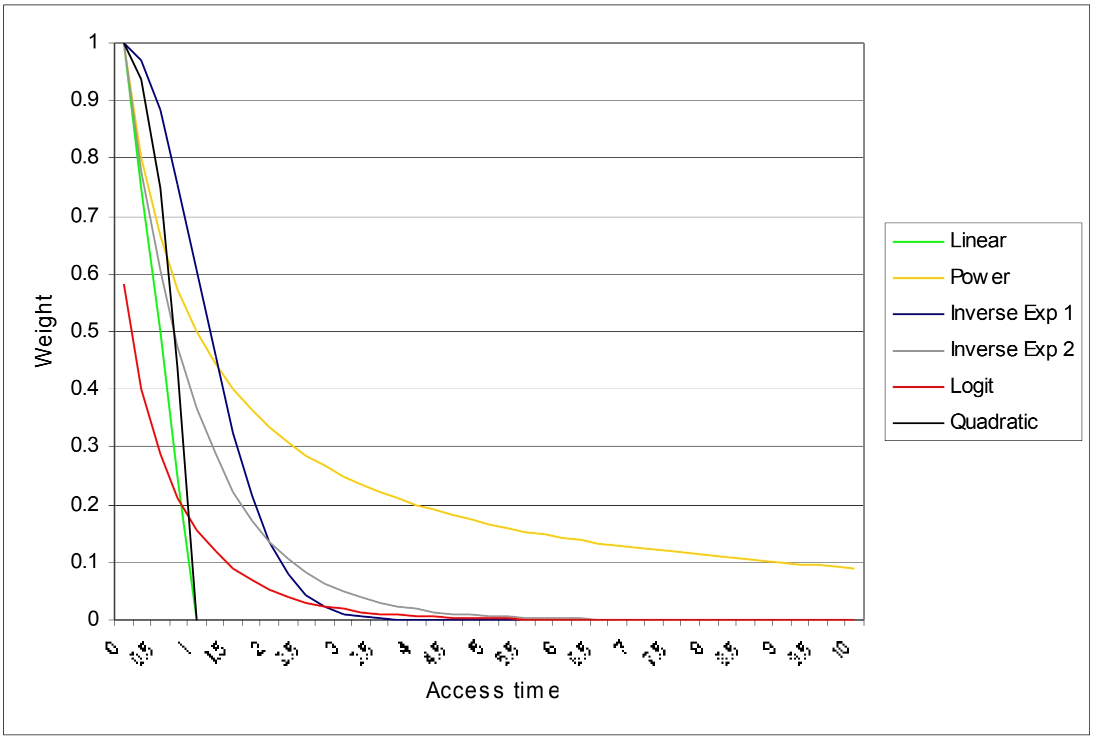

Linear function: $w_a = 1 - \frac{t}{\psi}$ where $\frac{t}{\psi} \leq 1$ or $w_a = 0$ otherwise (4.5)

Power function: $w_a = (t + 1)^{-\psi}$ (4.6)

Inverse exponential function 1: $w_a = \exp\left[-\frac{1}{2}\left(\frac{t}{\psi}\right)^2 \right]$ (4.7)

Inverse exponential function 2: $w_a = e^{-\psi t}$ (4.8)

Logit function: $w_a = \frac{\left(1/e^{\psi (t+1)}\right)}{1 - \left(1/e^{\psi (t+1)}\right)}$ (4.9)

Quadratic function: $w_a = 1 - \psi t^2$ where $\psi t^2 \leq 1$ or $w_a = 0$ otherwise (4.10)

Where:

w_a is the weight attached to population unit *a*

t is the travel time from population unit a to its closest station defined by road speed
ψ is a parameter which can either be determined by calibration or set at a predefined level

Applying the weighting function to the trip rate model gives Model 4.11. To simplify calibration a range of arbitrary values were assigned to the ψ parameter, allowing linear regression to be used. All output areas were first allocated to their nearest railway station by access time. The weighting function was then applied to each output area, with the resulting weights multiplied by the population and the results summed over the set of output areas allocated to each station to give an 'effective population' value equivalent to $\sum_a P_a w_a$. The results of calibrating Model 4.11 are summarised in Table 4.8.

$$\hat{V}_i = \alpha + \beta \sum_a P_a w_a \qquad (4.11)$$

Table 4.8: Summarised results from calibration of Model 4.11

Function 4.5							
ψ		2	3	4	4.5	5	6
Intercept	Value	39285.460	34639.334	30882.786	27866.330	24864.596	19438.198
	t stat	1.279	1.329	1.255	1.122	0.985	0.746
β parameter	Value	89.558	41.817	26.526	22.896	20.380	17.224
	t stat	7.004	8.831	9.597	9.586	9.489	9.289
R_{adj}^2		0.552	0.664	0.700	0.700	0.695	0.686
Mean AD		1.662	1.446	1.377	1.362	1.355	1.334
Function 4.6							
Ψ		0.5	1	1.5	1.75	2	3
Intercept	Value	14015.096	8783.612	9636.070	10630.845	11506.734	11466.664
	t stat	0.427	0.296	0.341	0.379	0.409	0.372
β parameter	Value	14.138	28.789	52.605	69.124	89.462	220.259
	t stat	7.083	8.193	8.687	8.733	8.675	7.758
R_{adj}^2		0.558	0.629	0.656	0.659	0.656	0.603
Mean AD		1.409	1.309	1.288	1.298	1.312	1.412
Function 4.7							
Ψ		1	2	2.5	3	4	5
Intercept	Value	35876.676	25722.976	20564.633	26469.490	11868.932	10382.842
	t stat	1.262	1.023	0.801	0.624	0.421	0.349
β parameter	Value	62.105	21.091	16.661	14.244	11.678	10.306
	t stat	7.866	9.508	9.426	9.195	8.640	8.138
R_{adj}^2		0.610	0.696	0.693	0.682	0.654	0.626
Mean AD		1.541	1.348	1.325	1.486	1.324	1.346
Function 4.8							
Ψ		0.25	0.5	0.625	0.75	1	1.5
Intercept	Value	28327.868	26380.153	20496.723	22619.184	24961.117	24695.567
	t stat	0.372	0.668	0.781	0.857	0.916	0.821
β parameter	Value	1.976	30.962	39.837	50.253	76.278	151.645
	t stat	8.632	9.200	9.169	9.041	8.631	7.619
R_{adj}^2		0.653	0.682	0.681	0.674	0.662	0.594
Mean AD		0.867	1.463	1.336	1.360	1.409	1.506

Table continued on next page

Function 4.9							
Ψ		0.25	0.375	0.5	0.75	1	1.5
Intercept	Value	12246.306	15115.991	17656.519	21230.825	22920.508	22536.035
	t stat	0.446	0.560	0.657	0.781	0.817	0.732
β parameter	Value	11.856	22.475	37.729	89.033	184.632	640.244
	t stat	8.924	9.012	8.996	8.764	8.381	7.444
R_{adj}^2		0.668	0.673	0.672	0.660	0.640	0.582
Mean AD		1.294	1.311	1.331	1.376	1.420	1.517
Function 4.10							
Ψ		0.01	0.05	0.075	0.1	0.25	1
Intercept	Value	10866.862	27626.153	33084.716	35292.195	41571.700	42607.738
	t stat	0.355	1.122	1.369	1.421	1.400	1.063
β parameter	Value	9.377	15.954	20.412	25.425	59.167	228.407
	t stat	7.838	9.693	9.743	9.362	7.265	4.718
R_{adj}^2		0.608	0.704	0.707	0.690	0.570	0.353
Mean AD		1.363	1.360	1.379	1.411	1.635	2.069

Table 4.8 shows that none of the weighting functions gave as a good a fit as the R_{adj}^2 value of 0.715 given using unweighted catchment populations limited by a 3.5 minute drive time, although some of the AD values were slightly superior. While weighting populations by distance from the station seemed intuitively to be more realistic, this was not reflected in the model results, and unweighted catchment populations were therefore retained, albeit using an access time cut-off point rather than the distance cut-off suggested by the PDFH (ATOC, 2002). Allowing stations to 'compete' for patronage from output areas based on the quality of service provided (as in Lythgoe's (2004) models of interurban travel) might give better model results, but this would be much more computationally complex, and it is questionable whether the increased modelling effort would be worthwhile given the quality of the results obtained from these simpler models. The best of the more complex models developed by Lythgoe (2004) had an R_{adj}^2 value of 0.59, which is inferior to the model fit obtained here.

4.2.2 Trip End Models

As described in Section 2.5.3, trip rate models can be extended by incorporating additional explanatory variables to produce linear regression 'trip end' models. This was tested for the South Hampshire area by gradually introducing further variables which seemed likely to affect rail demand levels to the basic model and examining their impact on model fit and parameter significance. The first additional variable tested was train frequency, giving Model 4.12. This was calibrated using data from the DB journey planner website (Deutsche Bahn AG, 2007), giving the results summarised in Table 4.9

$$\hat{V}_i = \alpha + \beta P_{it3.5} + \delta F_i \qquad (4.12)$$

Where:

F_i is the number of trains calling at station i on a normal weekday

Table 4.9: Summarised results from calibration of Model 4.12

Intercept		β parameter		δ parameter		R_{adj}^2	AD
Value	t stat	Value	t stat	Value	t stat		
-53702.112	-1.880	6.034	3.291	2383.386	4.090	0.798	1.096

This shows that the inclusion of train frequency in the model gave a major improvement in model fit, and that both parameters were significant and of the correct sign. Examination of the model residuals suggested that those stations where demand was overpredicted had no direct service to London. A number of representations of rail service levels, ease of access and distance to London were therefore tested, giving Models 4.13-4.17. Distances were calculated using MapInfo, with 'London' defined as the centre of London Bridge when measuring straight line distances. For rail journey time, where an interchange was involved each minute spent waiting for a connection was taken to contribute two minutes to the overall journey time, in line with PDFH recommendations (ATOC, 2002). The results of calibrating these models are summarised in Table 4.10.

$$\hat{V}_i = \alpha + \beta P_{it3.5} + \delta F_i + v L_i \qquad (4.13)$$

$$\hat{V}_i = \alpha + \beta P_{it3.5} + \delta F_i + v L_{iD} \qquad (4.14)$$

$$\hat{V}_i = \alpha + \beta P_{it3.5} + \delta F_i + v L_{iS} \qquad (4.15)$$

$$\hat{V}_i = \alpha + \beta P_{it3.5} + \delta F_i + v L_{iR} \qquad (4.16)$$

$$\hat{V}_i = \alpha + \beta P_{it3.5} + \delta F_i + v L_{iT} \qquad (4.17)$$

Where:

L_i is a dummy variable which takes the value 1 if 10 or more trains to a London terminal station call at station *i* on a normal weekday, and 0 otherwise

L_{iD} is the number of direct trains to London from station *i* on a normal weekday

L_{iS} is the straight line distance in km from station *i* to the centre of London Bridge

L_{iR} is the actual distance in km by rail from station *i* to London Waterloo

L_{iT} is the journey time in minutes by rail from station *i* to the nearest London terminal station

Table 4.10: Summarised results from calibration of Models 4.13-4.17

Parameter		Model 4.13	Model 4.14	Model 4.15	Model 4.16	Model 4.17
Intercept	Value	-59775.918	-45183.406	-301939.966	-323222.270	-15024.919
	t stat	-2.106	-1.468	-2.297	-2.496	-0.134
β parameter	Value	6.606	6.007	4.972	5.130	6.060
	t stat	3.584	3.258	2.683	2.846	3.263
δ parameter	Value	2005.187	2013.883	2783.444	2794.373	2334.214
	t stat	3.203	2.661	4.644	4.742	3.853
ν parameter	Value	53485.997	1632.626	2098.634	1953.101	-303.787
	t stat	1.501	0.772	1.932	2.129	-0.356
R_{adj}^2		0.805	0.796	0.812	0.816	0.794
AD		1.191	1.012	0.930	0.930	1.059

While Models 4.13, 4.15 and 4.17 gave a slight improvement in fit, the London parameters were only significant in Models 4.15 and 4.16, and both of these parameter values suggested that rail demand increases with distance from London, when the reverse was expected to be the case. Despite these concerns, the rail distance to London variable was retained for further consideration.

The next variable to be tested was based on the distance people living in the vicinity of a local station travel to work. Data was therefore obtained from the Neighbourhood Statistics website on commuting distances at output area level, enabling Model 4.12 to be extended to give Model 4.18. A minimum 5 km commuting distance was used as rail commutes shorter than this length are unlikely to occur. The summarised results from calibrating Model 4.18 are given in Table 4.11, which shows that Model 4.18 gave a marked improvement in model fit. The commuting variable was therefore retained in the model, although the rail distance to London variable still seemed to be of the wrong sign.

$$\hat{V}_i = \alpha + \beta P_{it3.5} + \delta F_i + \nu L_{iR} + \xi C_i \qquad (4.18)$$

Where:

C_i is the proportion of the total population living within 2 km of station i who travel more than 5 km to work

Table 4.11: Summarised results from calibration of Model 4.18

Parameter	Value	t stat
Intercept	-701415.885	-4.723
β parameter	6.171	3.952
δ parameter	2961.554	5.871
ν parameter	3277.169	3.828
ξ parameter	820868.076	3.808
R_{adj}^2	0.866	
AD	0.928	

It seemed possible that the proportion of the catchment population which fell within the higher socio-economic classes might affect rail demand, and a variable representing this was added to the model, giving Model 4.19. A summary of the results from calibrating this model is given by Table 4.12.

$$\hat{V}_i = \alpha + \beta P_{it3.5} + \delta F_i + \nu L_{iR} + \xi C_i + \omega S_i \qquad (4.19)$$

Where:

S_i is the percentage of the catchment population in NS-SEC 1 and 2

Table 4.12: Summarised results from calibration of Model 4.19

Parameter	Value	t stat
Intercept	-700265.622	-4.588
β parameter	6.166	3.884
δ parameter	2962.015	5.786
ν parameter	3279.836	3.769
ξ parameter	826518.040	3.315
ω parameter	-11168.822	-0.047
R_{adj}^2	0.862	
AD	0.931	

This modification was obviously not successful, as the socio-economic parameter was very insignificant and reduced model fit slightly. This suggests that the socio-economic class of the catchment population has little impact on rail demand.

The size of station car parks will partly determine the ease with which passengers can access stations and is therefore likely to affect the number of trips they generate. Model 4.20 was therefore calibrated, giving the results summarised in Table 4.13. This gave a major improvement in model fit, although the C_i (commuting) variable was no longer significant. The Spearman's rank correlation coefficient did not indicate the presence of any significant correlation between this and the car park variable. However, Model 4.21 which omitted the C_i variable was also calibrated and the results from this are also shown in Table 4.13 to allow comparison.

$$\hat{V}_i = \alpha + \beta P_{it3.5} + \delta F_i + \nu L_{iR} + \xi C_i + \rho Pk_i \qquad (4.20)$$

$$\hat{V}_i = \alpha + \beta P_{it3.5} + \delta F_i + \nu L_{iR} + \rho Pk_i \qquad (4.21)$$

Where:

Pk_i is the number of car parking spaces at station *i*

Table 4.13: Summarised results from calibration of Models 4.20-4.21

	Model 4.20		Model 4.21	
Parameter	Value	t stat	Value	t stat
Intercept	-447093.840	-3.424	-303271.132	-3.387
β parameter	5.454	4.351	5.055	4.058
δ parameter	2842.403	7.058	2777.635	6.820
ν parameter	1996.344	2.710	1423.664	2.227
ξ parameter	305651.025	1.491	n/a	n/a
ρ parameter	2380.150	4.599	2801.709	6.354
R_{adj}^2	0.915		0.912	
AD	1.266		1.427	

Removing the C_i (commuting) variable reduced model fit slightly but meant that all parameters were significant. The final variable tested with this case study was the level of car ownership within the catchment population, which is the demand side equivalent of the size of the station car park. This variable was incorporated in Model 4.22, and the results of calibrating this model are summarised in Table 4.14.

$$\hat{V}_i = \alpha + \beta P_{it3.5} + \delta F_i + \varphi L_{iR} + \rho Pk_i + \kappa Ca_i \qquad (4.22)$$

Where:

Ca_i is the average number of cars per household in the station catchment

Table 4.14: Summarised results from calibration of Model 4.22

Parameter	Value	t stat
Intercept	-414798.072	-2.912
β parameter	5.821	3.988
δ parameter	2719.764	6.613
φ parameter	1602.167	2.415
ρ parameter	2617.960	5.487
κ parameter	66600.167	1.007
R_{adj}^2	0.912	
AD	1.331	

The car ownership (κ) parameter is insignificant and gives no improvement in model fit, so either Model 4.20 or 4.21 should be taken forward as the preferred model form, with the former having the best fit but all parameters significant in the latter. These models are extremely effective at predicting rail demand at local stations in this case study area, with the model fit much better than that of most rail demand models reviewed in Chapter 2.

The residuals from Model 4.20 are mapped in Figure 4.7, which gives an indication of where the model performs well and where it is less accurate. There are no obvious spatial patterns in the residuals, with stations where demand is underpredicted and overpredicted distributed seemingly at random across the study area. However, Figure 4.7 only shows the absolute error, and it is arguably more useful to examine the percentage error in the prediction relative to the observed trip totals. These errors are therefore mapped in Figure 4.8, which shows that the largest percentage errors are generally associated with the stations with the smallest observed demand, with Beaulieu Road attracting very few passengers, and Dean, Dunbridge, Redbridge and Millbrook also little used compared to other stations in the dataset. Such an error pattern is perhaps to be expected from regression models of this type.

Figure 4.7: Residuals from Model 4.20

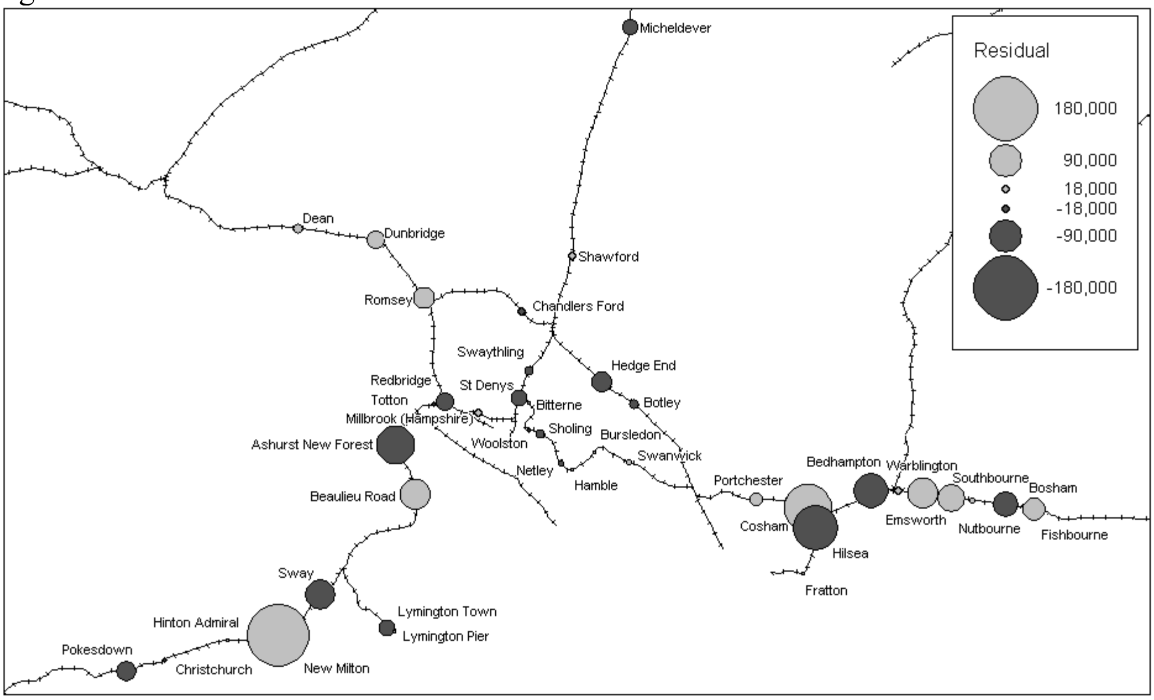

Figure 4.8: Percentage errors in predictions from Model 4.20

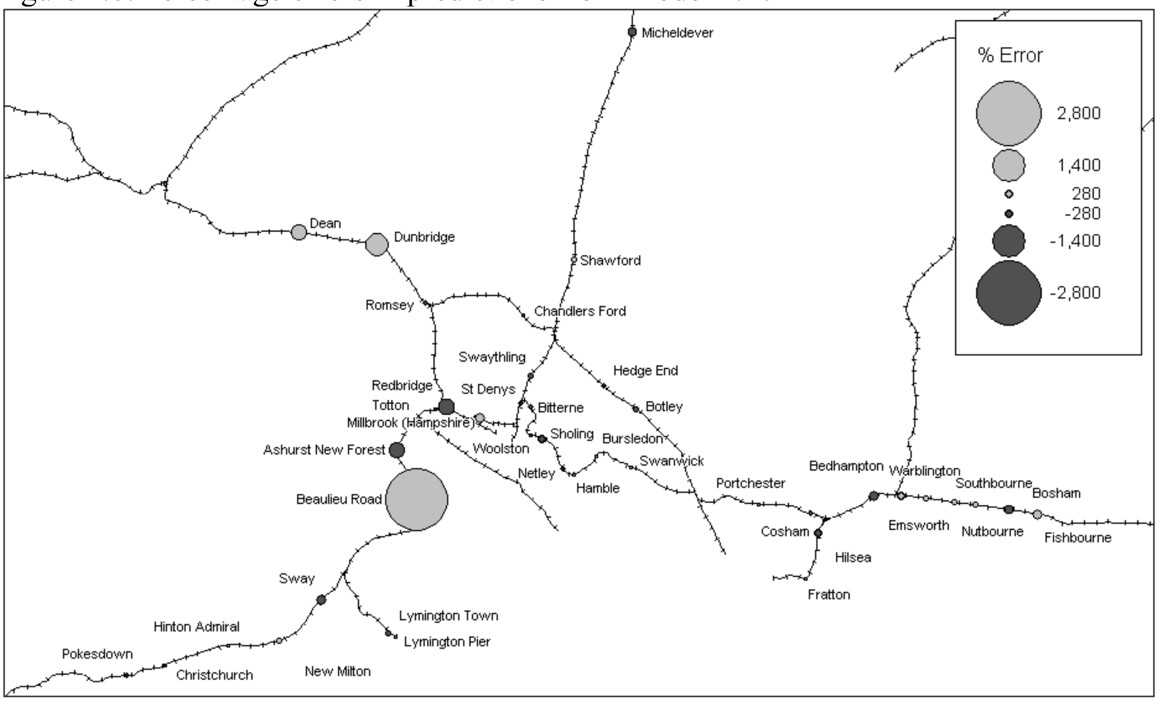

The models described in this section showed that it was possible to use trip rate and trip end models to accurately forecast demand at local railway stations and they were therefore applied to a much larger dataset covering the whole of England and Wales, as described in Section 4.3.

4.3 England and Wales case study

4.3.1 Global trip end models

The ultimate aim of this section of the work was to develop a generalised trip end model for England and Wales, and as good results had been obtained from the models calibrated on the initial South Hampshire case study area, similar models were tested over this much larger area, with a range of independent variables again progressively added to the model form. The calibration dataset for England and Wales was made up of all stations in Network Rail's station categories E and F ('small staffed' and 'small unstaffed' respectively) in these two countries. While the station categorisation carried out by Network Rail is not without its anomalies, this seemed the best way of identifying local stations over a large area without having to resort to the use of discretion in considering the 'localness' of individual stations. As discussed in Section 3.8, while the inclusion of Scottish stations in the dataset was considered, this would have greatly increased the time required for data processing, as spatial and demographic data for England and Wales had already been converted into a suitable form for analysis.

In practice problems with the data meant that it was not possible to include all local stations in England and Wales in the calibration dataset. A number of stations were included in local 'groups' for ticketing purposes, which meant that trips could not be reliably allocated to particular stations within the group. Several other stations were not served by regular trains on weekdays, or have had their train services replaced by buses, and as this would be expected to suppress demand these stations were also removed from the dataset. Finally, a catchment definition problem affecting two stations on Merseyside made it impossible to obtain realistic population allocations for these stations. A total of 1510 stations were therefore available for model calibration, with details of the excluded stations given in Appendix 2.

Population data was downloaded for the whole of England and Wales at output area level, giving a total of 174,198 units, and MapInfo was used to allocate each population unit to its nearest station based on road travel time. Light rail and underground stations were not included in demand forecasts, but were considered to be equivalent to heavy rail stations when allocating population units to their nearest station, as it seems likely that they would be considered as a local railway station by prospective travellers. As before, arbitrary time cut-off points were initially used to define the boundaries of station catchments, with a number of different values tested for these points to establish which gave the best fit for Model 4.4. This gave the results summarised in Table 4.15.

Table 4.15: Summarised results from calibration of Model 4.4 on England and Wales dataset

t	Intercept		β parameter		R_{adj}^2	AD
	Value	t stat	Value	t stat		
1.5	76717.077	13.501	23.321	13.646	0.109	16.483
2	72628.700	12.444	15.745	13.846	0.112	16.087
2.5	73259.725	12.286	11.278	13.143	0.102	16.212
3	73991.715	12.257	8.815	12.680	0.096	16.393
4	77430.991	13.260	5.389	10.932	0.073	16.646

This indicates that using a 2 minute access time cut-off gives the best results, although model fit is extremely poor. Although they gave inferior results for the South Hampshire area, the weighting functions 4.5-4.10 were tested for this larger dataset. As before, to simplify calibration a range of arbitrary values were assigned to the parameter for the weighting function, allowing linear regression to be used. The results from calibrating Model 4.11 using these various values and functions are summarised in Table 4.16. This shows that all the weighting functions gave a better model fit than the arbitrary cut-off points, with the best fit obtained for Model 4.11 using function 4.6 with a ψ value of 3.25, although the R^2_{adj} value of 0.123 was still very low.

Table 4.16: Estimated parameter values and significance of Model 4.11 calibrated on England and Wales dataset using various ψ values with weighting functions 4.5-4.10

Function 4.5							
ψ		1	1.75	2	2.25	2.5	3
Intercept	Value	84011.494	76416.238	74281.097	72753.225	71779.555	70604.370
	t stat	14.940	13.243	12.743	12.356	12.072	11.687
β parameter	Value	131.906	51.909	42.508	35.728	30.565	23.496
	t stat	12.597	13.690	13.893	13.961	13.911	13.748
Adjusted R^2		0.096	0.111	0.114	0.115	0.114	0.112
Function 4.6							
Ψ		2	2.75	3	3.25	3.5	4
Intercept	Value	66745.298	64080.677	64267.268	64818.143	65647.384	67894.199
	t stat	10.557	10.300	10.395	10.551	10.753	11.255
β parameter	Value	48.854	98.894	119.818	142.666	167.269	220.849
	t stat	13.427	14.353	14.476	14.527	14.519	14.362
Adjusted R^2		0.107	0.121	0.123	0.123	0.123	0.121

Table continued on next page

Function 4.7							
Ψ		0.5	0.75	1	1.25	2	3
Intercept	Value	80761.169	74986.191	71422.785	70003.274	72928.599	80187.007
	t stat	14.206	12.907	12.023	11.579	11.785	12.970
β parameter	Value	95.601	49.699	32.224	23.255	11.814	6.809
	t stat	13.084	13.827	14.015	13.858	12.667	11.145
Adjusted R^2		0.102	0.113	0.116	0.114	0.097	0.076
Function 4.8							
Ψ		0.75	1	1.25	1.5	1.75	2
Intercept	Value	68394.824	67305.928	67796.927	68961.518	70353.294	71773.781
	t stat	11.024	10.972	11.182	11.493	11.828	12.152
β parameter	Value	28.457	42.866	59.411	78.008	98.626	121.223
	t stat	13.545	14.052	14.250	14.278	14.212	14.094
Adjusted R^2		0.109	0.116	0.119	0.120	0.119	0.117
Function 4.9							
Ψ		0.75	1	1.25	1.5	1.75	2
Intercept	Value	66775.905	66710.222	67587.743	68886.634	70323.429	71762.929
	t stat	10.781	10.893	11.159	11.486	11.823	12.149
β parameter	Value	50.415	102.486	188.703	325.887	538.244	860.104
	t stat	13.919	14.219	14.321	14.306	14.220	14.092
Adjusted R^2		0.114	0.119	0.120	0.120	0.119	0.117
Function 4.10							
Ψ		0.05	0.1	0.15	0.2	0.25	1
Intercept	Value	75892.279	71694.384	72006.054	72943.568	74244.722	84429.105
	t stat	12.293	11.787	12.041	12.368	12.716	15.090
β parameter	Value	8.870	14.731	19.725	24.375	28.672	87.881
	t stat	12.097	13.343	13.711	13.874	13.858	12.638
Adjusted R^2		0.089	0.106	0.111	0.114	0.114	0.096

Including train frequencies in the model was expected to improve model fit, and the relationship between train frequency and trips made is shown in Figure 4.9.

Figure 4.9: Train frequency and trips made for 1510 stations in England and Wales

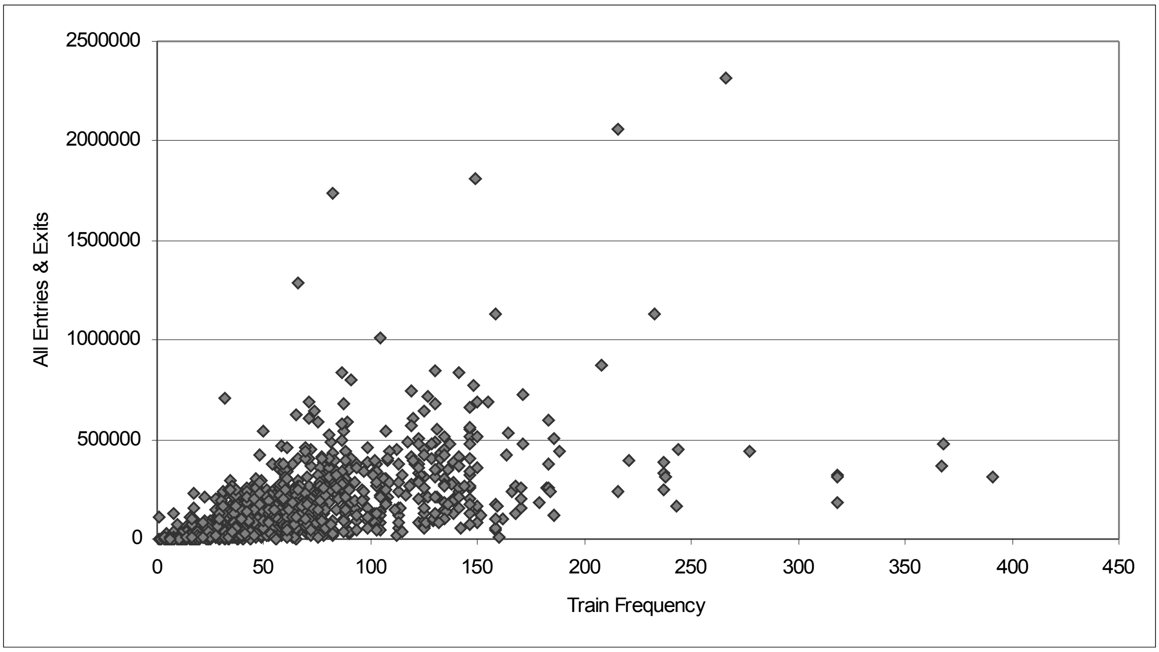

Figure 4.9 indicates that there is an element of heteroskedasticity within the dataset, with greater variation in the number of trips made at higher train frequencies. Before including

train frequencies in the model it seemed sensible to check for a correlation with catchment populations, and the relationship between these two variables is illustrated in Figure 4.10.

Figure 4.10: Catchment population and train frequency for 1510 stations in England and Wales

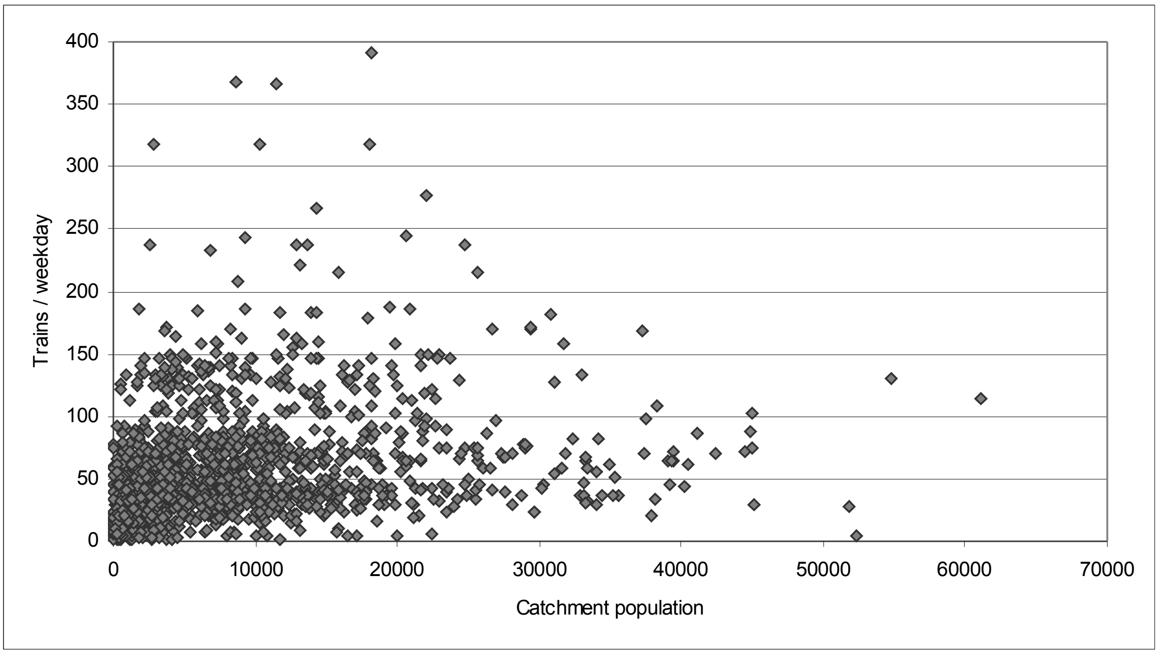

The graph showed no obvious correlation between the two variables, but when the relationship was tested further by calculating the Spearman's rank correlation coefficient, this gave a value of 0.450. This suggested that there was a significant but weak positive correlation between train frequency and catchment population, which meant that calibrating Model 4.23 which includes both variables might be expected to give counterintuitive results. The results of this calibration are given in Table 4.17.

$$\hat{V}_i = \alpha + \beta \sum_a P_a w_a + \delta F_i \qquad (4.23)$$

Table 4.17: Summarised results from calibration of Model 4.23

Intercept		β parameter		δ parameter		R_{adj}^2	AD
Value	t stat	Value	t stat	Value	t stat		
-14264.4	-2.294	49.885	5.392	2135.258	23.638	0.362	3.173

As expected including frequency greatly improved model fit and the parameter values indicate that it had a greater influence in determining demand than catchment population, although the latter parameter was still significant and of the correct sign. It is questionable whether frequency has such a large impact in reality, as the mean population elasticity implied by Table 4.17 (0.172) is lower than the range recommended by ATOC (2002).

However, these recommendations are for actual catchment populations rather than the weighted populations used here, so direct comparison is not possible. The frequency parameter may have been inflated by the simultaneity inherent in this approach to rail demand modelling, and if such simultaneity is present then global regression models may not be appropriate as the estimators will be both biased and inconsistent. Potentially the primary direction of causality could be determined by studying time series data where train frequencies have changed to establish whether this is in response to, or a trigger for, increased demand, but such data was not available for this study. It was however also possible that the unexpected magnitude of the frequency parameter resulted from it acting as a proxy for variables omitted from the model, and therefore the inclusion of additional independent variables was investigated.

While competition between stations was not considered in defining station catchments, it is likely that demand at local stations would be affected by the proximity of larger stations. Models 4.24-4.25 were therefore calibrated for the England and Wales case study, incorporating the distance to the nearest station in Network Rail categories A-D as a single variable (Model 4.24), and the distance to the nearest station in each of these four categories as separate variables (Model 4.25), giving the results summarised in Table 4.18. This shows that the improvement in fit over previous models was only marginal. None of the category specific λ parameters were significant, and several were of the wrong sign, which meant that Model 4.24 was taken forward as the preferred form.

$$\hat{V}_i = \alpha + \beta \sum_a P_a w_a + \delta F_i + \lambda_S T_S \tag{4.24}$$

$$\hat{V}_i = \alpha + \beta \sum_a P_a w_a + \delta F_i + \lambda_A T_A + \lambda_B T_B + \lambda_C T_C + \lambda_D T_D \tag{4.25}$$

Where:

T_S is the distance in km from station i to the nearest category A-D station

T_A is the distance in km from station i to the nearest category A station

T_B is the distance in km from station i to the nearest category B station

T_C is the distance in km from station i to the nearest category C station

T_D is the distance in km from station i to the nearest category D station

Table 4.18: Summarised results from calibration of Models 4.24-4.25

Model		4.24		4.25	
Parameter		Value	t stat	Value	t stat
Intercept		-32529.9	-3.487	-22740.7	-2.287
β parameter		55.362	5.848	51.602	5.456
δ parameter		2216.949	23.241	2182.466	22.128
λ parameters	S	754.057	2.622	n/a	n/a
	A	n/a	n/a	-55.548	-0.475
	B	n/a	n/a	75.253	0.594
	C	n/a	n/a	140.285	0.830
	D	n/a	n/a	81.058	0.359
R_{adj}^2		0.365		0.361	
AD		4.855		3.672	

As the level of employment around a station was found to have a significant effect on rail demand in South Hampshire, Models 4.26 and 4.27 which incorporate this variable were calibrated, giving the results summarised in Table 4.19. Both variants of Model 4.26 gave an improvement in model fit, with an *m* value of 4 minutes giving slightly superior results, and while the lower *m* value was trialled in some later models 4 minutes was retained as the preferred value. The form of Models 4.26 and 4.27 is perhaps questionable, as it combines a weighted population value with a jobs value defined using an arbitrary cut-off point. Using a weighted jobs value might be more realistic, but given the significant amount of time this would have required and the marginal improvement in model fit given by introducing a weighted population parameter, this variant was not pursued.

$$\hat{V}_i = \alpha + \beta \sum_a P_a w_a + \delta F_i + \lambda T_S + \tau J_{idm} \quad (4.26)$$

$$\hat{V}_i = \alpha + \beta \sum_a P_a w_a + \delta F_i + \lambda T_S + \tau Jp_{idm} \quad (4.27)$$

Where:

J_{idm} is the number of jobs located within *m* minutes drive of station *i*

Jp_{idm} is the number of jobs located within *m* minutes drive of station *i* divided by the population within *m* minutes drive time of station *i*

Table 4.19: Summarised results from calibration of Models 4.26-4.27

Model	4.26				4.27			
m	2		4		2		4	
Parameter	Value	t stat	Value	t stat	Value	t stat	Value	t stat
Intercept	-36462.3	-3.986	-43067.5	-4.624	-41339.0	-4.348	-38970.7	-3.830
β parameter	28.343	2.505	26.924	2.458	57.017	6.086	56.373	5.846
δ parameter	2218.500	23.567	2260.176	23.957	2210.679	23.418	2221.515	23.025
λ parameter	822.224	2.975	926.238	3.338	823.627	2.968	870.307	2.852
τ parameter	8.434	4.302	5.317	5.016	14233.967	3.211	10141.164	1.735
R_{adj}^2	0.374		0.376		0.370		0.361	
AD	5.199		5.307		4.973		5.180	

The next variable tested was the level of car ownership within the station catchment, even though this did not give a significant parameter for the South Hampshire dataset. Model 4.28 was therefore calibrated, giving the results summarised in Table 4.20. The car ownership parameter was significant in both model variants, with a 4 minute 'm' (catchment boundary) value again giving superior results. The parameter was positive in both cases, suggesting that as car ownership rises, so does rail use. This may be because the car ownership variable is acting as a proxy for income, which would be expected to have a positive correlation with rail demand. Despite the significance of the parameter, the improvement in fit given by this model over previous variants is marginal.

$$\hat{V}_i = \alpha + \beta \sum_a P_a w_a + \delta F_i + \lambda T_S + \tau J_{idm} + \zeta M_{ihdm} \tag{4.28}$$

Where:

M_{ihdm} is the mean number of motor vehicles owned per household within m minutes drive time of station i

Table 4.20: Summarised results from calibration of Model 4.28

m	2		4	
Parameter	Value	t stat	Value	t stat
Intercept	-65974.5	-4.722	-90113.0	-4.967
β parameter	30.552	2.700	34.973	3.110
δ parameter	2234.426	23.746	2297.784	24.210
λ parameter	996.013	3.523	1104.410	3.904
τ parameter	9.035	4.591	5.678	5.337
ζ parameter	23240.117	2.789	32767.596	3.019
R_{adj}^2	0.377		0.380	
AD	5.767		5.603	

Incorporating a variable representing the size of station car parks improved the fit of the South Hampshire trip end models, and a similar variable was therefore tested for the larger dataset. For eight stations no car park data was given on the National Rail website and the resolution of images of the stations on Google Earth was too low for car park size to be discerned, and these stations therefore had to be removed from the calibration dataset. The results of calibrating Model 4.29 on the reduced dataset are summarised in Table 4.21, which shows that the inclusion of a car park variable in the model gave a clear improvement in model fit. However, the car ownership variable became insignificant, and Model 4.30, which omits this variable, was therefore calibrated giving results which are also shown in Table 4.21. All parameters in Model 4.30 were significant and of the expected sign, with no significant reduction in model fit, and this was therefore taken forward as the preferred form.

$$\hat{V}_i = \alpha + \beta \sum_a P_a w_a + \delta F_i + \lambda_S T_S + \tau J_{id4} + \zeta M_{ihd4} + \rho P k_i \qquad (4.29)$$

$$\hat{V}_i = \alpha + \beta \sum_a P_a w_a + \delta F_i + \lambda_S T_S + \tau J_{id4} + \rho P k_i \qquad (4.30)$$

Table 4.21: Summarised results from calibration of Models 4.29-4.30

	Model 4.29		Model 4.30	
Parameter	Value	t stat	Value	t stat
Intercept	-68909.3	-4.111	-46857.1	-5.401
β parameter	37.312	3.604	33.869	3.349
δ parameter	1998.027	22.526	1976.220	22.561
λ parameter	839.343	3.083	746.175	2.810
τ parameter	4.678	4.753	4.473	4.585
ζ parameter	15433.533	1.538	n/a	n/a
ρ parameter	855.517	11.801	872.838	12.182
R_{adj}^2	0.430		0.430	
AD	4.357		4.398	

All trip end models described so far aimed to predict the absolute number of trips made from a local station over a one year period. However, it is often the case in rail demand modelling that better results are obtained using a logarithmic transformation of the dependent variable. To test whether this applied here, Model 4.31 was calibrated on the reduced England and Wales dataset, giving the results summarised in Table 4.22.

$$\ln \hat{V}_i = \alpha + \beta \sum_a P_a w_a + \delta F_i + \lambda_S T_S + \tau J_{id4} + \rho P k_i \qquad (4.31)$$

Table 4.22: Summarised results from calibration of Model 4.31

Parameter	Value	t stat
Intercept	9.474	118.193
β parameter	0.0004	4.704
δ parameter	0.017	21.598
λ parameter	-0.012	-4.849
τ parameter	0.00004	4.046
ρ parameter	0.009	14.194
R_{adj}^2	0.495	
AD	2.637	

While Model 4.31 appears to give a major improvement in fit over Model 4.30, it is not mathematically valid to compare the R_{adj}^2 values of the linear and loglinear forms, because R^2 is the ratio of explained variance to total variance and the variances of V_i and lnV_i are different (Maddala, 2001). It was therefore necessary to carry out a formal test to compare the performance of the two models. Maddala (2001) suggests the use of three possible tests, with the first two (the MacKinnon, White and Davidson test and the Bera and

MacAleer test) both based on artificial regressions. However, these tests are unsuitable when, as in this case, the predicted values of the dependent variable can be negative. This means that the third test, a special case of the Box-Cox test, must be used. The first stage in this procedure is to divide each V_i value by the geometric mean of the V_is, where the geometric mean is defined by (4.32).

$$\left(\prod_{i=1}^{n} V_i \right)^{1/n} = \sqrt[n]{V_1 \cdot V_2 \cdot \cdots \cdot V_n} \tag{4.32}$$

The geometric mean of the V_is was calculated to be 49872.872, allowing transformed values of the dependent variable to be calculated for Models 4.30 and 4.31. The regressions were then reestimated for both models, giving the results shown in Table 4.23.

Table 4.23: Summarised results from transformed versions of Models 4.30-4.31 calibrated on England and Wales dataset

Model	4.30		4.31	
Parameter	Value	t stat	Value	t stat
Intercept	-0.940	-5.401	-1.343	-16.755
β parameter	0.001	3.349	0.0004	4.704
δ parameter	0.040	22.561	0.017	21.598
λ parameter	0.015	2.810	-0.012	-4.849
τ parameter	0.00009	4.585	0.00004	4.046
ρ parameter	0.018	12.182	0.009	14.194
R_{adj}^2	0.430		0.495	
AD	4.398		2.637	
Residual sum of squares	9916.114		9183.060	

The Box-Cox procedure says that the functional form with the lowest residual sum of squares (RSS) should be chosen as the preferred model (Maddala, 2001). The results in Table 4.23 appears to show that the log-linear form of Model 4.31 provides a better fit than the linear form of Model 4.30. However, it is necessary to carry out a significance test of the difference between the two RSS values using (4.33) to confirm that Model 4.31 gives an improvement. This gave a value of 24.999, which is much larger than the critical value of 3.84, confirming that Model 4.31 should be taken forward as the preferred form.

$$BoxCox = \frac{n}{2}\log\left(\frac{RSS_{l\arg est}}{RSS_{smallest}} \right) \sim \chi^2 \tag{4.33}$$

Taking logarithms of some or all of the independent variables has also been found to improve model fit in previous demand modelling work, and therefore all 32 possible variants of logarithmic and untransformed variables in Model 4.31 were tested, giving the results summarised in Table 4.24. This shows that the best results are given by using natural logarithms of all the independent variables, giving the double log Model 4.34 (displayed here in both additive and multiplicative forms). More detailed results from the calibration of this model are summarised in Table 4.25, which shows that all the model parameters were highly significant and of the correct sign.

Table 4.24: Summarised results from calibrating semilog variants of Model 4.31

R_{adj}^2				$\sum_a P_a w_a$	$\ln \sum_a P_a w_a$
F_i	T_S	J_{id4}	Pk_i	0.495	0.571
			$LnPk_i$	0.552	0.615
		LnJ_{id4}	Pk_i	0.571	0.584
			$LnPk_i$	0.611	0.624
	LnT_S	J_{id4}	Pk_i	0.488	0.575
			$LnPk_i$	0.544	0.615
		LnJ_{id4}	Pk_i	0.574	0.589
			$LnPk_i$	0.611	0.625
LnF_i	T_S	J_{id4}	Pk_i	0.684	0.721
			$LnPk_i$	0.714	0.746
		LnJ_{id4}	Pk_i	0.729	0.732
			$LnPk_i$	0.750	0.754
	LnT_S	J_{id4}	Pk_i	0.693	0.730
			$LnPk_i$	0.719	0.750
		LnJ_{id4}	Pk_i	0.735	0.739
			$LnPk_i$	0.753	0.757

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \rho \ln Pk_i \qquad (4.34)$$

$$\hat{V}_i = \alpha \left(\sum_a P_a w_a \right)^{\beta} F_i^{\delta} T_S^{\lambda} J_{id4}^{\tau} Pk_i^{\rho} \qquad (4.34)$$

Table 4.25: Summarised results from calibration of Model 4.34

Parameter	Value	t stat
Intercept	2.834	15.418
β parameter	0.142	5.862
δ parameter	1.375	42.766
λ parameter	0.284	9.229
τ parameter	0.156	7.622
ρ parameter	0.196	15.435
R_{adj}^2	0.757	
AD	0.892	

A transcendental logarithmic (translog) model form (Model 4.35) was also tested, as such functions allow a more flexible relationship between the variables (Greene, 1993) and are often used in cost functions. This gave the results summarised in Table 4.26.

$$\hat{V}_i = \alpha + \beta_1 \ln \sum_a P_a W_a + \delta_1 \ln F_i + \lambda_1 \ln T_S + \tau_1 \ln J_{id4} + \rho_1 \ln Pk_i + \sum_N \gamma_N \ln^2 X_N + \ldots$$
$$\ldots \sum_{MN} \eta_{MN} \ln X_M \ln X_N \tag{4.35}$$

Where:

X_M and X_N denote the set of independent variables included in the model

Table 4.26: Summarised results from calibration of Model 4.35

Parameter	Value	t stat
Intercept	-382130.666	3.078
β parameter	-33384.373	-1.120
δ parameter	-208321.348	-5.247
λ parameter	-42544.229	-1.008
τ parameter	-24359.548	-1.118
ρ parameter	-52474.421	-2.728
γ_{PaWa} parameter	-241.514	-0.111
γ_F parameter	32955.791	8.245
γ_T parameter	7525.360	1.693
γ_J parameter	-1097.718	-2.278
γ_{Pk} parameter	6654.681	4.338
η_{PaWaF} parameter	8178.756	1.579
η_{PaWaT} parameter	111.242	0.020
η_{PaWaJ} parameter	3642.368	2.027
η_{PaWaPk} parameter	-2669.633	-1.156
η_{FT} parameter	708.054	0.101
η_{FJ} parameter	4435.922	1.078
η_{FPk} parameter	5532.685	1.631
η_{TJ} parameter	3745.371	0.899
η_{TPk} parameter	278.526	0.096
η_{JPk} parameter	4752.242	2.403
R_{adj}^2	0.472	
AD	12.105	

A large number of the variables in Model 4.35 were insignificant, as was expected given the addition of so many additional variables to the model. The model was therefore recalibrated using backwards stepwise calibration, to select the optimal combination of variables from the original set. This gave Model 4.36, and the results from calibrating this model are summarised in Table 4.27. This shows that both the jobs and the squared jobs variables were of the wrong sign, implying that as employment around a station increases, rail use decreases. Because of this, the double log Model 4.34 was retained as the preferred model form and the residuals from this model were mapped to enable any spatial patterns in model performance to be identified, as shown in Figure 4.11.

$$\hat{V}_i = \alpha + \delta \ln F_i + \tau \ln J_{id4} + \rho \ln Pk_i + \gamma_F \ln^2 F_i + \gamma_J \ln^2 J_{id4} + \gamma_{Pk} \ln^2 Pk_i + \ldots$$
$$..\eta_{P_aW_aT_S} \ln \sum_a P_a W_a \ln T_S + \eta_{P_aW_aJ} \ln \sum_a P_a W_a \ln J_{id4} + \eta_{F_iJ} \ln F_i \ln J_{id4} + \ldots \tag{4.36}$$
$$..\eta_{T_SJ} \ln T_S \ln J_{id4} + \eta_{JPk} \ln J_{id4} \ln Pk_i$$

Table 4.27: Summarised results from calibration of Model 4.36

Parameter	Value	t stat
Intercept	338252.877	7.533
δ parameter	-217041.121	-10.384
τ parameter	-47474.277	-5.974
ρ parameter	-39.774.425	-4.068
γ_F parameter	36271.774	11.455
γ_J parameter	-1088.765	-2.318
γ_{Pk} parameter	7543.815	5.297
η_{PaWaTs} parameter	-4578.591	-2.345
η_{PaWaJ} parameter	3728.052	4.441
η_{FJ} parameter	9559.188	4.189
η_{TJ} parameter	6354.222	3.724
η_{JPk} parameter	3511.889	2.952
R_{adj}^2	0.472	
AD	11.635	

Figure 4.11: Residuals from Model 4.34

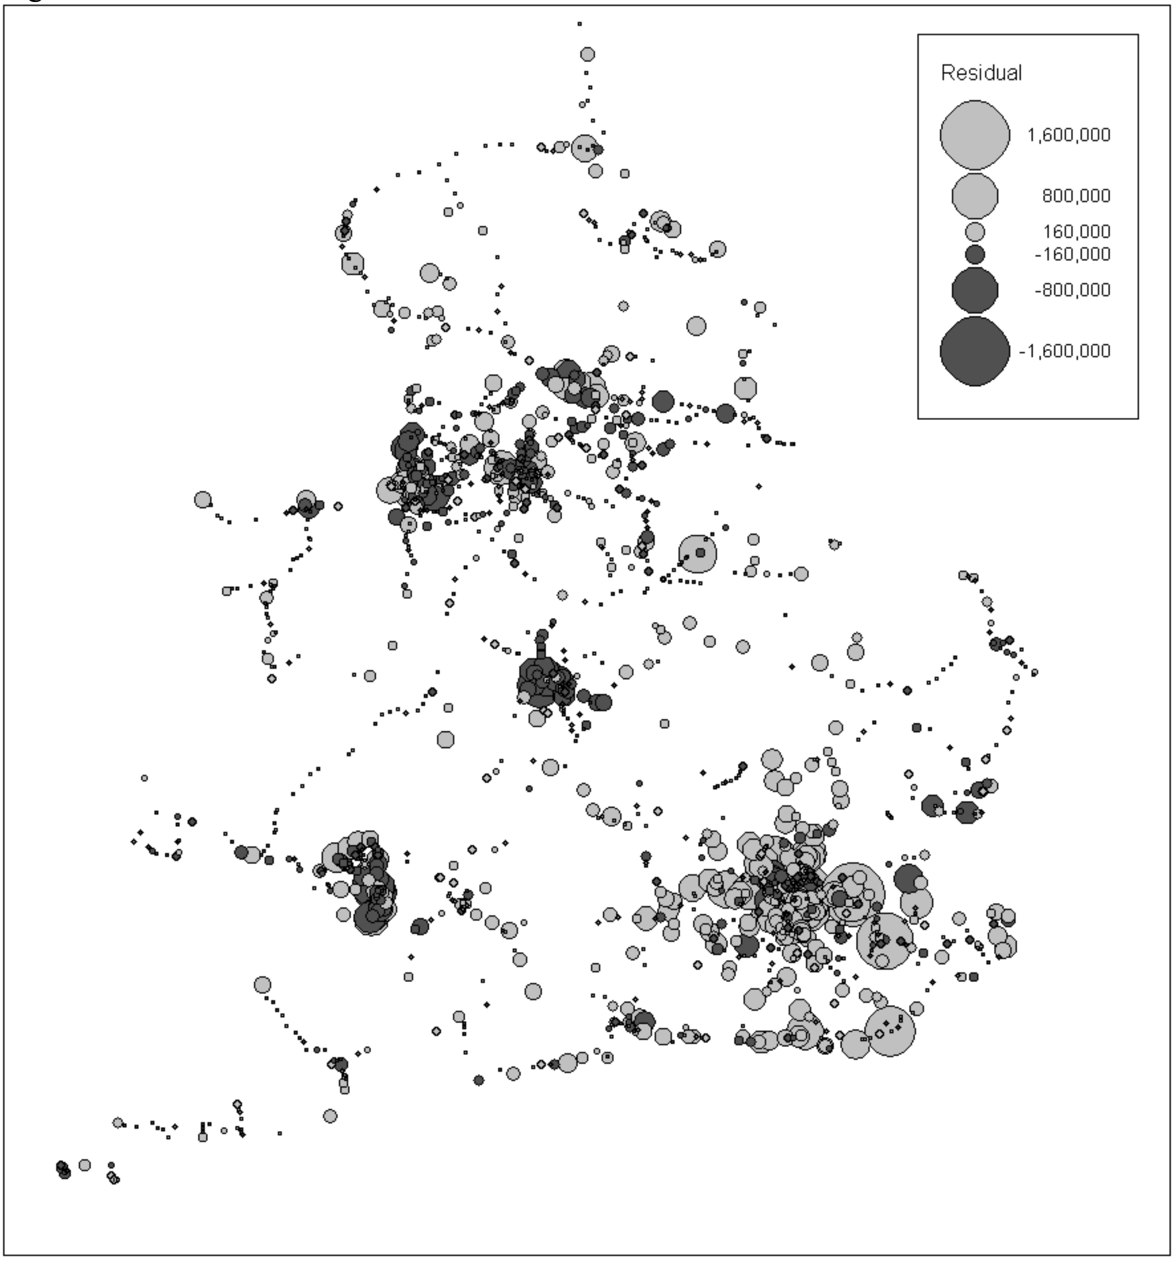

There was a clear spatial pattern in the residuals around London, with demand tending to be overpredicted at stations close to the city centre, and underpredicted at stations in the outer suburbs. Dummy variables were therefore added to the model to represent stations within certain straight line distance bands of London, as illustrated in Figure 4.12, giving Model 4.37. The results of calibrating this model are summarised in Table 4.28.

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \rho \ln Pk_i + \upsilon_1 L_1 + \upsilon_2 L_2 + \upsilon_3 L_3 + \upsilon_4 L_4$$

(4.37)

Where:

L_1 is a dummy variable which takes the value 1 if Station i is within 12.5 km straight line distance from Marble Arch, and 0 otherwise

L_2 is a dummy variable which takes the value 1 if Station i is between 12.5 km and 50 km straight line distance from Marble Arch, and 0 otherwise

L_3 is a dummy variable which takes the value 1 if Station i is between 50 km and 100 km straight line distance from Marble Arch, and 0 otherwise

L_4 is a dummy variable which takes the value 1 if Station i is between 100 km and 150 km straight line distance from Marble Arch, and 0 otherwise

Figure 4.12: 12.5 km, 50 km, 100 km and 150 km buffer zones around Marble Arch

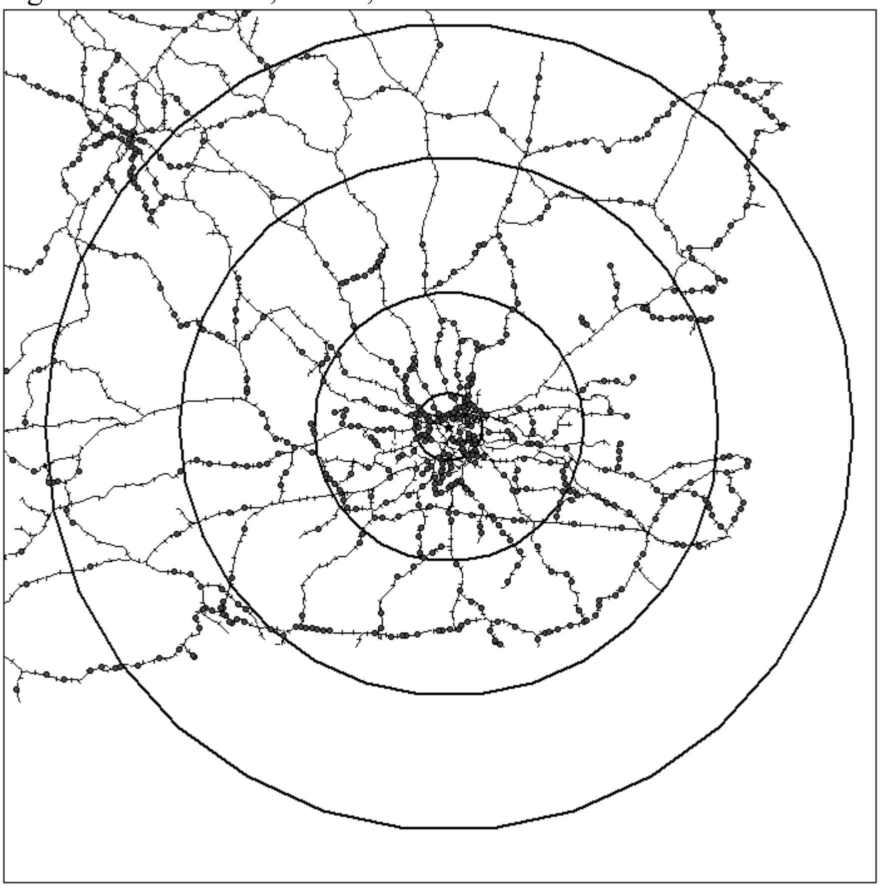

Table 4.28: Summarised results from calibration of Model 4.37

Parameter	Value	t stat
Intercept	2.649	14.481
β parameter	0.153	6.305
δ parameter	1.315	39.883
λ parameter	0.336	10.477
τ parameter	0.172	8.352
ρ parameter	0.192	15.029
υ_1 parameter	0.286	2.588
υ_2 parameter	0.411	5.699
υ_3 parameter	0.395	5.681
υ_4 parameter	0.171	2.146
R_{adj}^2	0.765	
AD	0.864	

All four distance from London parameters were positive and significant, with the 12.5-50 km parameter being most strongly positive. The model fit was though only slightly improved from Model 4.34, which may indicate the presence of multicollinearity given that three significant parameters have been added to the model, although it is not obvious why these dummies would be correlated with any other parameters. A single continuous distance from London variable might be preferable to these dummies, although it is not immediately clear what form this should take. The relationship between the size of the residuals from Model 4.34 and distance from Central London is plotted in Figure 4.13.

Figure 4.13: Residuals from Model 4.34 plotted by distance from Marble Arch

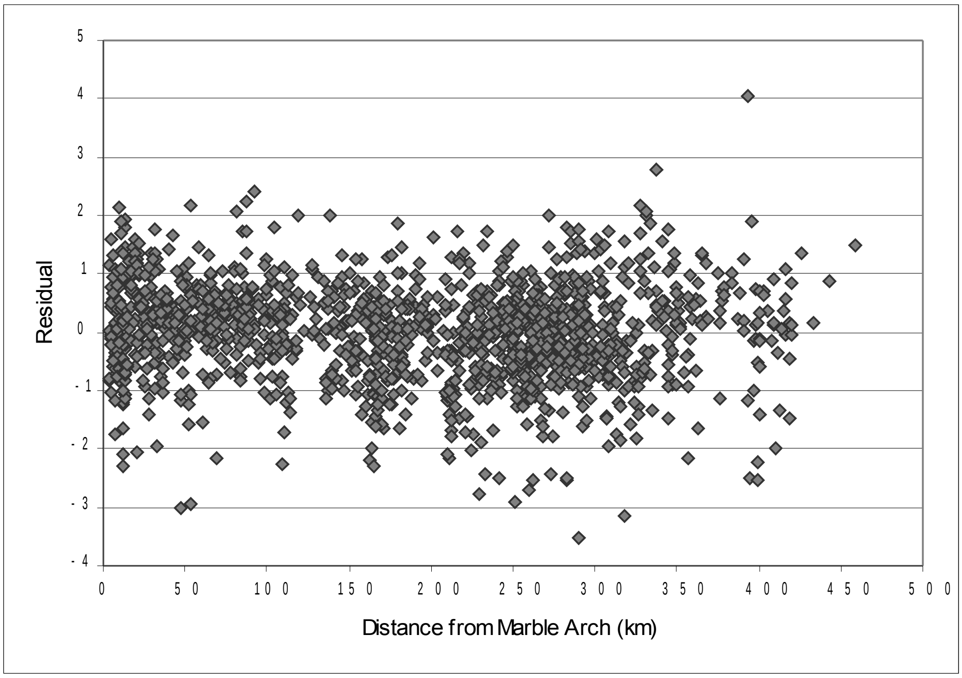

Figure 4.13 does not show any obvious linear relationship between distance from Central London and residual size or sign. Concentrating attention on stations within 150 km of Central London (Figure 4.14) did not make it any easier to spot a pattern, although the

residuals tend to be slightly positive, which reflects the positive values of the distance to London parameters in Model 4.37. This model was therefore adopted as the preferred global model at the end of this phase of the analysis.

Figure 4.14: Residuals from Model 4.34 for stations within 150 km of Marble Arch plotted by distance from Marble Arch

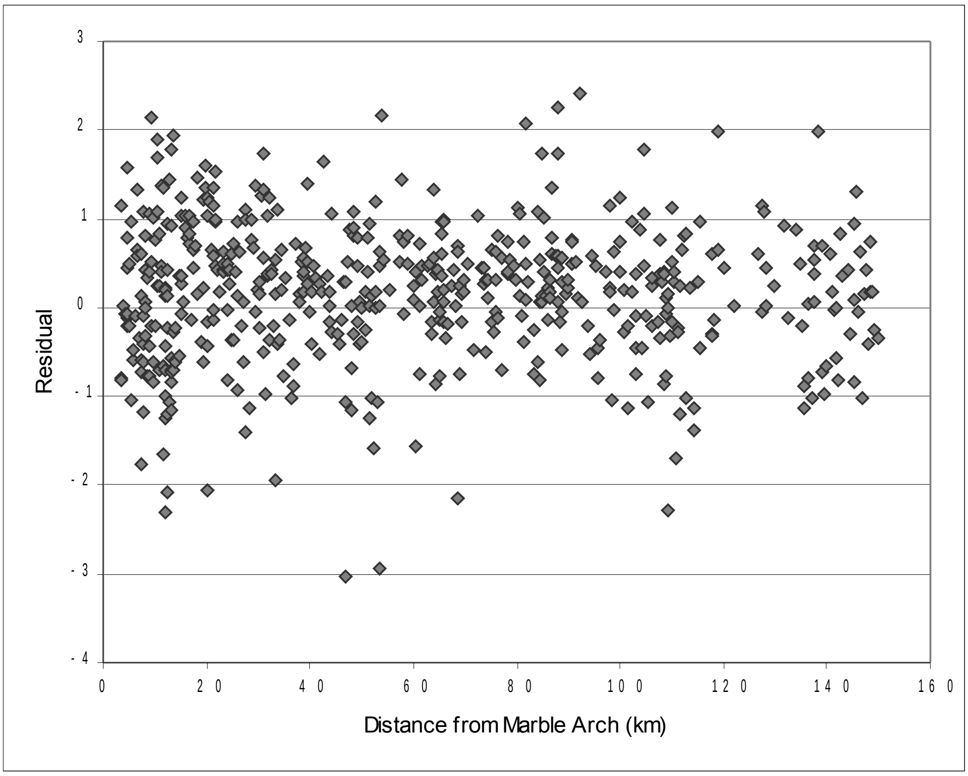

4.3.2 GWR-based trip end models

While Model 4.37 explained around three-quarters of the variation in the logarithm of rail demand across local stations in England and Wales, this still left a quarter of the variation unaccounted for. It seemed possible that this resulted from spatial variations in the propensity to travel by rail, and geographically weighted regression (GWR) was therefore applied to these models to test whether this would give an improvement in model fit and allow the spatial variation to be quantified. A detailed review of GWR and other local analysis techniques was given in Section 3.2, where GWR and the spatial expansion method were identified as the techniques with the greatest potential for enhancing local rail demand models. The use of the spatial expansion method was briefly investigated, but it rapidly became clear that, because it could only deal with unidirectional spatial variations and because the number of model parameters escalates rapidly with the number of independent variables (with a deleterious effect on parameter significance), it did not give any improvement over simpler global regression models. This is a general problem with this method, and is not restricted to its application in this instance, meaning that it is unlikely to be suitable for use in rail demand modelling.

GWR appeared more promising, but it could not be applied to Model 4.37 because it contained distance-based dummy variables and logically there should not be spatial variation in spatially-defined parameters such as these. GWR was therefore applied to the next best model (4.34), giving the results summarised in Table 4.29. A Gaussian model was used with adaptive kernel bandwidths determined by Akaike Information Criterion (AIC) minimisation. Adaptive bandwidths were used because the density of local railway stations varies widely across the study area and therefore varying bandwidths were necessary to avoid there being a high level of variation in the number of data points used to make local parameter estimates.

Table 4.29: Summary of results from global and GWR calibration of Model 4.34

Calibration method		Global	GWR					
m		4	Minimum	Lower quartile	Median	Upper quartile	Maximum	Monte Carlo P-value
Intercept	Value	2.834	-0.454	1.959	2.676	3.333	5.329	0.060
	t stat	15.418						
β parameter	Value	0.142	-0.235	0.058	0.160	0.243	0.363	0.010
	t stat	5.862						
δ parameter	Value	1.375	0.829	1.230	1.360	1.436	1.738	0.150
	t stat	42.766						
λ parameter	Value	0.284	-0.055	0.213	0.334	0.453	0.759	0.000
	t stat	9.229						
τ parameter	Value	0.156	-0.068	0.106	0.165	0.256	0.498	0.000
	t stat	7.622						
ρ parameter	Value	0.196	0.084	0.155	0.188	0.214	0.314	0.410
	t stat	15.435						
Adj. R^2		0.757	0.802					
ANOVA		SS	DF	MS	F			
OLS Residuals		10111.7	6.00					
GWR Improvement		248.0	111.41	2.2265				
GWR Residuals		763.6	1384.59	0.5515	4.0370			

The GWR model had an R_{adj}^2 value of 0.802, which was a major improvement on the value of 0.765 from the global model. The AIC value of 3503.706 was lower than that of 3682.997 for the global model, which also indicated that the GWR model was superior. The significant F statistic from the ANOVA testing the null hypothesis that the GWR model represented no improvement over the global model confirms that the GWR model gave an improvement in fit over the global model. Table 4.29 also gives a summary of the local parameter variation in Model 4.34, along with the results of a Monte Carlo test of the significance of the spatial variation in the parameters. These p values indicate that the spatial variation in the population, distance to higher category station and jobs parameters was statistically significant. However, there was a reasonable probability that the spatial variation in the other parameters occurred by chance.

Figure 4.15 shows the residuals from the GWR version of Model 4.34. The most noticeable difference between this map and Figure 4.11 is that the overprediction of demand at many stations in the West Midlands has been reduced. However, much of the other variation in the residuals from the global model shown in Figure 4.11 appeared to be reproduced in Figure 4.15.

Figure 4.15: Mapped residuals from GWR calibration of Model 4.34

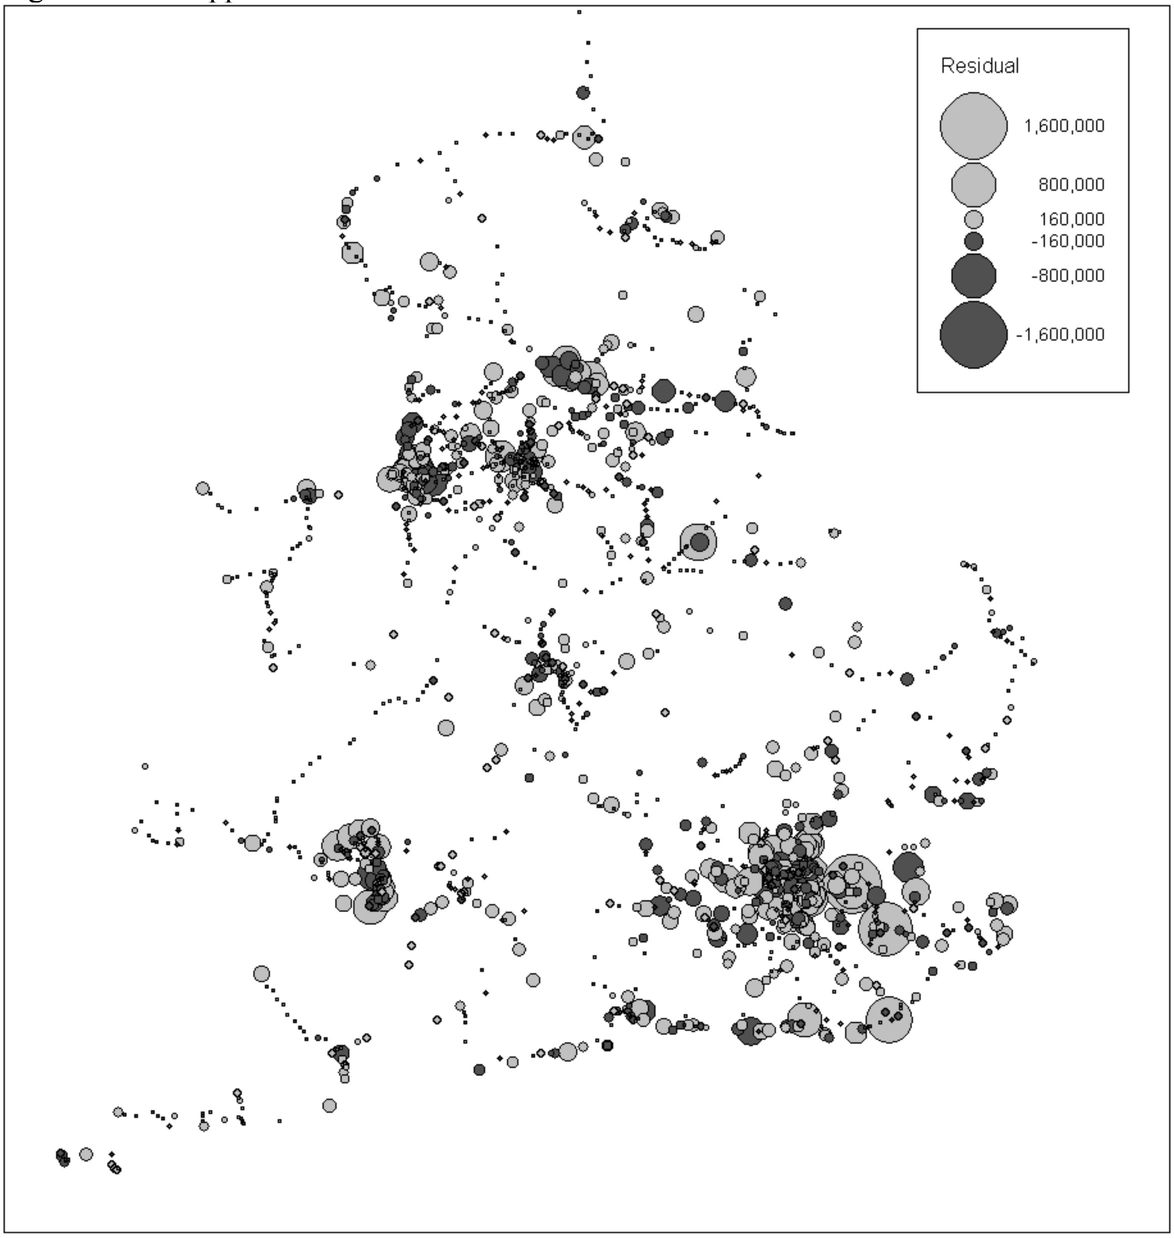

Comparing residuals from GWR calibrations is not straightforward, because they each have a different standard error. However, the GWR software provides a solution by producing standardised (internally Studentised) residuals which are mapped in Figure 4.16.

Figure 4.16: Standardised residuals from GWR calibration of Model 4.34

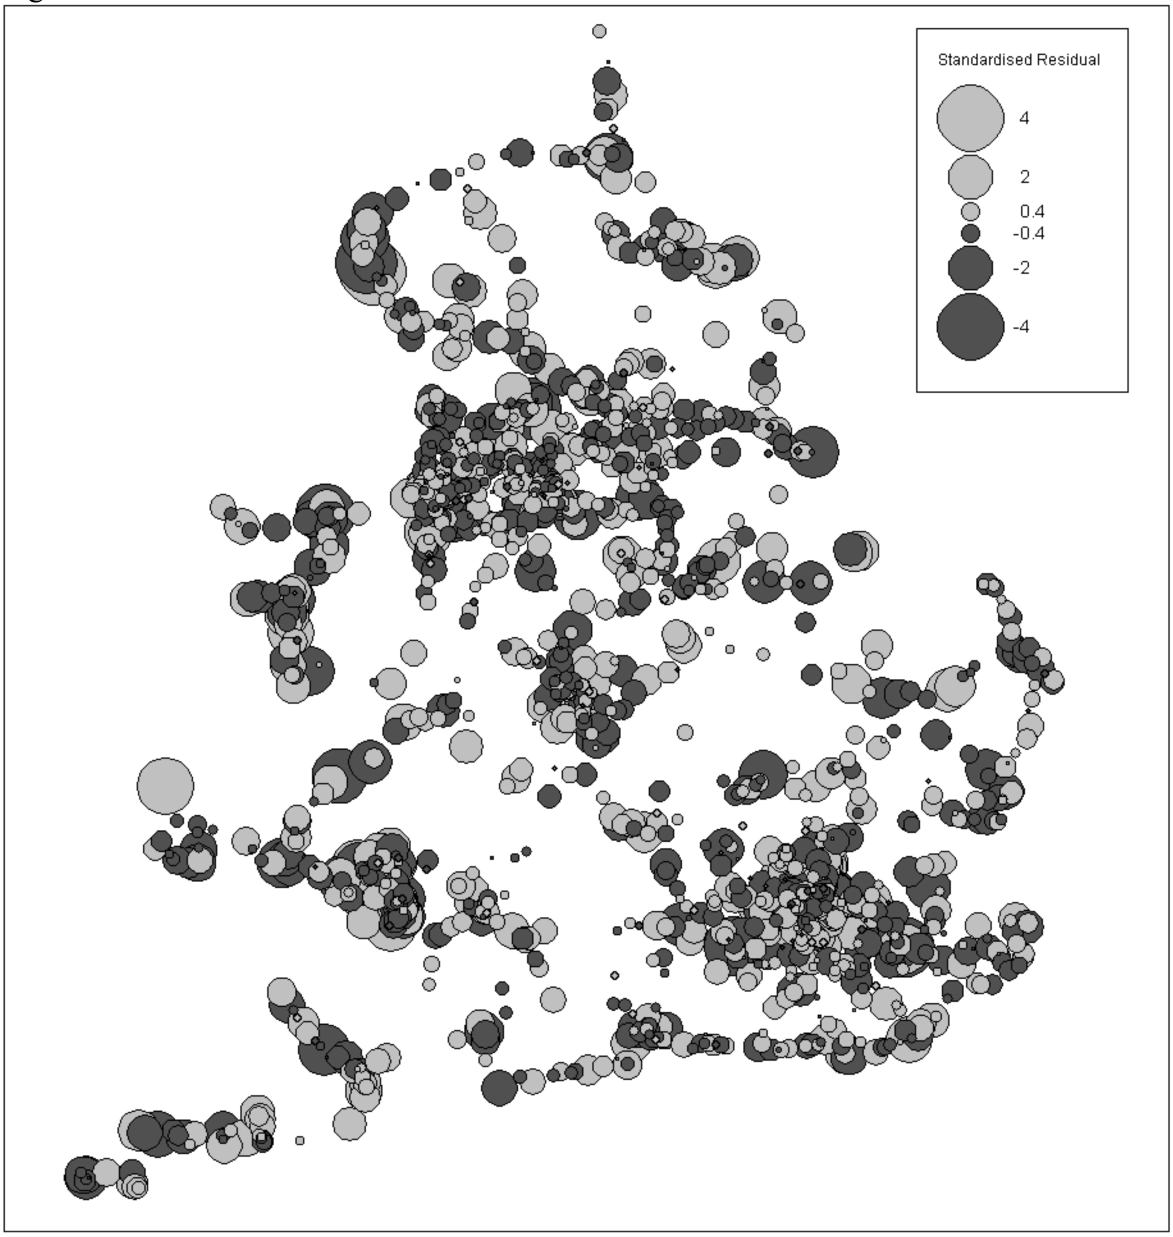

The most obvious difference between Figures 4.16 and 4.15 is that the residuals for stations with low usage levels have greater prominence than before. Fotheringham et al. (2003) suggest that standardised residuals larger than ±3 should be considered as unusual and investigated further, and therefore Figure 4.16 was amended to only show only the residuals which are classified as 'unusual' using these criteria, giving the results shown in Figure 4.17. The reasons why the stations highlighted in Figure 4.17 have these extreme residual values are considered in more detail below, but further exploratory analysis of the results from the GWR calibration was undertaken first.

Figure 4.17: Standardised residuals with values ±3 from GWR calibration of Model 4.34

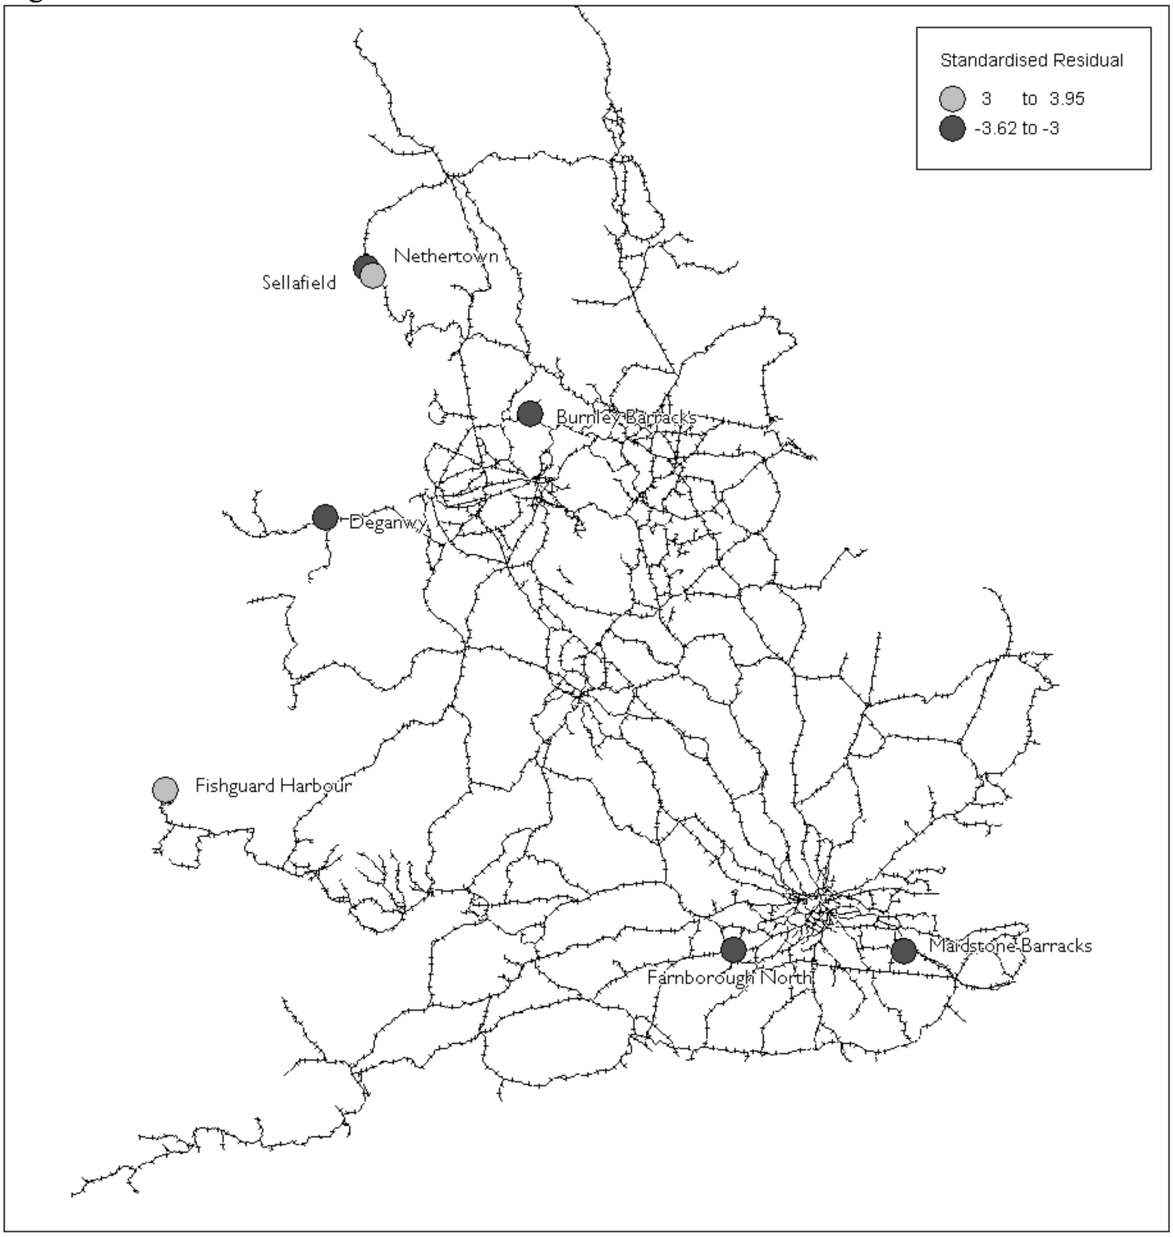

The GWR software also produces two other diagnostic measures which indicate the influence particular observations have on model calibration. The first of these is the leading diagonal of the hat matrix ('influence'), which is a leverage value. Data points with high leverage values are likely to be outliers (Rousseeuw & van Zomeren, 1999), and have the potential to dominate a regression analysis (Fotheringham et al., 2002). The leverage values from the GWR calibration of Model 4.34 are mapped in Figure 4.18. In general these are small, but for a few stations they are unusually large. In some areas these stations are clustered together, with the Cumbrian Coast line, the East Coast Main Line north of Newcastle, the Oxford-Worcester line and the Teesside area being particularly noticeable examples. While points with a high leverage value may be influential, this is not necessarily always the case. The Cook's distance and residual for such points should also be checked, and if all these values are high then there is reason to believe that the point in question may have biased the estimation of the regression coefficients.

Figure 4.18: Leverage values from GWR calibration of Model 4.34

The Cook's distance is the second measure of influence provided by the GWR software, and is calculated using equation 4.38.

$$CD_i = \frac{r_i^2 s_{ii}}{p(1 - s_{ii})} \tag{4.38}$$

Where:

CD_i is Cook's distance

r_i is the internally studentised residual

s_{ii} is the leading diagonal of the hat matrix (the matrix which maps $\hat{y}$ onto y)

p is the number of parameters

Influential observations will have values greater than 1, and may be unusual either in terms of the dependent variable (high residual, low leverage) or the independent variables (low residual, high leverage) (Fotheringham et al., 2002). The Cook's distances for the GWR version of Model 4.34 are mapped in Figure 4.19.

Figure 4.19: Values of Cook's distance from GWR version of Model 4.34

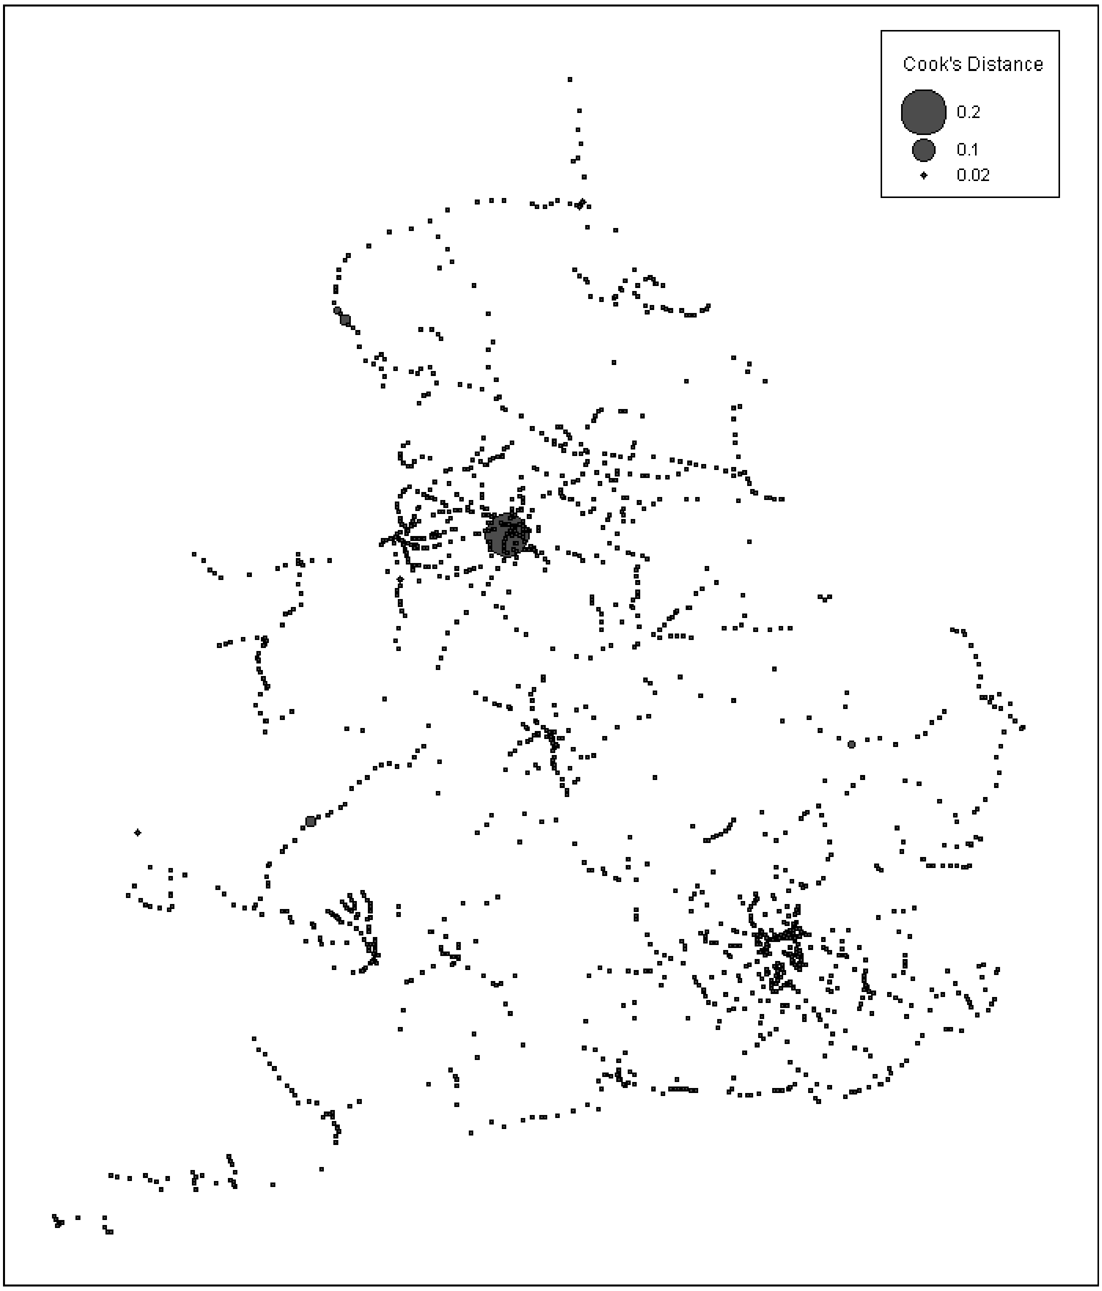

Figure 4.19 shows that the Cook's distance does not exceed Fotheringham et al.'s critical value of 1 at any of the stations in the dataset, but there are still a few stations where the

values are extreme compared to the majority of the dataset, notably Burnage. Examination of the calibration data showed that the recorded train frequency value of 1 was likely to be incorrect given the number of trips made from the station. The source of the error was not immediately apparent, as the TIPLOC code used by the Perl script to calculate the frequency appeared to be correct. The National Rail Timetable suggested that the correct frequency was in fact 72 trains per day (Network Rail, 2008a), and the calibration dataset was therefore amended accordingly.

Several other stations had unusually high standardised residuals, high leverage values and/or large Cook's distances. As a result of the discovery of the data error at Burnage the values of the independent variables for these other stations were investigated in case they were incorrect, or skewing the model results. These values are summarised in Table 4.30.

Table 4.30: Observed values of model variables for stations with extreme Cook's distance from GWR version of Model 4.34

Station	Trips 2005-6	Weighted pop	Frequency	Category A-D time	Catchment jobs	Parking spaces
Burnley Barracks	964	807.520	29	12.825	3549	0
Deganwy	3697.1	362.251	40	23.789	1343	20
Dunston	131	744.675	4	4.148	4400	0
Farnborough North	1647.11	171.712	44	2.874	823	5
Fishguard Harbour	23873	21.900	4	48.537	758	0
Longcross	345	3.068	23	5.474	0	0
Maidstone Barracks	3002.15	360.998	65	2.582	3477	0
Nethertown	297	6.767	8	75.985	87	6
New Clee	246.3	578.814	8	3.287	6143	0
Sellafield	232449.35	0	17	75.908	0	6
Shippea Hill	26	0	2	65.835	0	6
Sugar Loaf	122	0	8	99.244	0	0
Mean values	125697.69	437.377	55.93	14.849	3349.18	24.231

Sellafield appeared to have an unexpectedly high number of trips given that it has a catchment population of zero, but in fact the station is immediately adjacent to the nuclear reprocessing facilities and is thus likely to attract a large number of commuters. These jobs are not recorded as being within Sellafield's catchment because the population weighted centroids of the relevant output areas are further than 4 minutes drive time from the station, as shown in Figure 4.20. This illustrates a more general problem, which is that while station catchment areas are based on population-weighted output area centroids, these may not reflect the centres of employment within output areas, and if a large number of jobs are located within an output area with a low population density (which will be geographically extensive as a result) they may not be assigned to the station which is closest to the job locations. It was not obvious how this could be resolved, as no better data on employment location was available, and if the data for one location (such as Sellafield) was corrected this would only have led to biases in other locations where the misallocation of employment to stations was less apparent.

Figure 4.20: Output area boundaries and centroids around Sellafield station

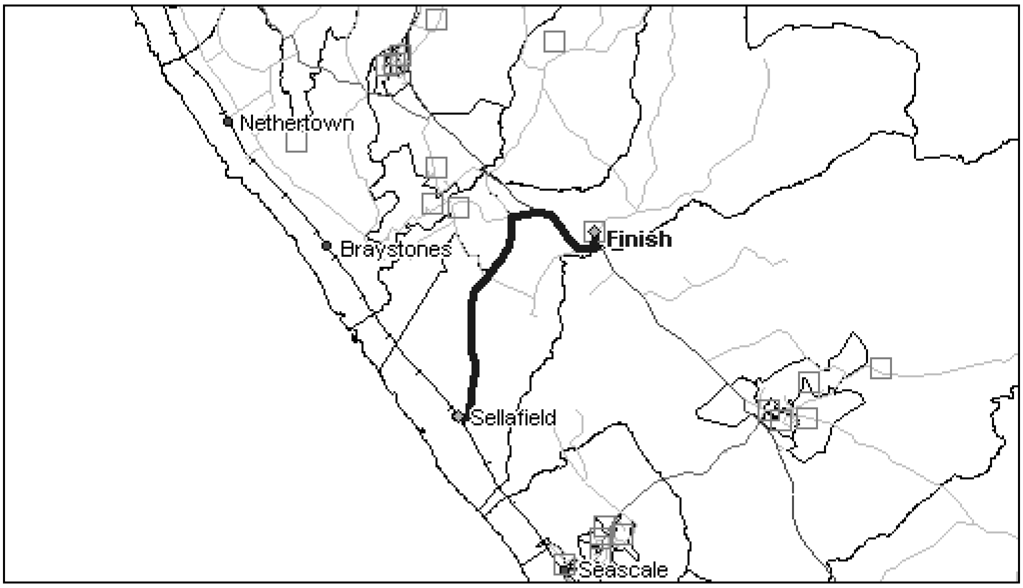

The reasons for the relatively extreme Cook's distances at Nethertown, Shippea Hill and Sugar Loaf were less obvious, but probably resulted from the very low number of trips recorded despite there being more than a minimal number of trains (Nethertown and Sugar Loaf) and/or parking spaces being provided (Nethertown and Shippea Hill). Again it was not obvious how this problem could be resolved in the model, but as the Cook's distances were much smaller than those at Burnage the discrepancies were likely to have much less effect on the model results than the data error at the latter station.

Burnley Barracks, Farnborough North and Maidstone Barracks were all recorded as having significantly fewer trips than would be expected given the values of the independent variables at these locations. The latter two stations were both part of station groups, and the low recorded trip frequency probably resulted from the misallocation of trips within the group. As a result all stations within the Farnborough and Maidstone groups were deleted from the calibration dataset, as they were likely to be biasing the model results. Burnley Barracks was not part of a station group, despite the presence of other stations in Burnley, and the low number of trips recorded may therefore result from a combination of factors, including the presence of parking and staff at Burnley Central (the next station along the line), and possibly a tendency among ticket-issuing staff to automatically sell tickets to one of the other Burnley stations, although this is pure conjecture. Because of this uncertainty

Burnley Barracks was retained in the calibration dataset.

New Clee, Longcross, Dunston and Deganwy also had lower usage levels than expected given the characteristics of their catchments and of the services provided. New Clee has a large number of residents and jobs within its estimated catchment, but is only served by eight trains a day and is located close to both Grimsby Town and Cleethorpes which enjoy a vastly superior service. This is likely to be the cause of the low number of trips, although this effect should have been captured by the 'time to nearest Category A-D station' and 'frequency' variables. A similar situation existed at Dunston where trips were likely to be lost to the nearby MetroCentre station, which has a much better train service, and also to local buses given the proximity of the station to central Newcastle. The trip prediction at Longcross was probably distorted by the surprisingly high train frequency provided at the station, which is not in the vicinity of any significant centres of population. While the frequency is high given the minimal population it is lower than at adjacent stations and is therefore unlikely to attract passengers from further afield, and it was not clear how the model could be adjusted to reflect this. The situation at Deganwy was less obvious, as the station has a good-sized car park, a reasonably frequent train service and a moderately large number of residents and jobs within its catchment, yet the number of trips was extremely low. This may be because Deganwy is located quite close to Llandudno Junction, which has a much higher service frequency and a larger car park, but is still a Category E station and is not therefore included in the 'time to nearest Category A-D station' variable. There could also have been some 'loss' of trips as a result of passengers being sold tickets to/from Llandudno on board trains.

The situation at Fishguard Harbour was the reverse of that at the other stations considered here, as it has a much higher number of trips than would be expected given its extremely low train frequency, lack of car park and small catchment population. This was however easily explained by the presence of the ferry terminal, and the consequent flow of passengers to and from Ireland via the station, and including a dummy variable to represent 'harbour' stations might improve model fit.

In addition to a harbour dummy variable, the pattern of model residuals suggested that dummy variables representing three other station types should be tested, for stations on travelcard boundaries, stations on electrified lines, and terminus stations. The travelcard boundary dummy was applied to stations which are the first/last station on a particular line to fall within the travelcard zones in areas where such schemes exist (London, West Midlands, Merseyside, Greater Manchester, South Yorkshire, West Yorkshire and Tyne & Wear). The terminus dummy variable was applied to stations which form the limit of passenger services on a particular line. Stations which mark the limit of local services but where long distance services continue on the same line, such as Bedwyn, were not considered to be termini. The electrification dummy variable was applied to stations where services are provided by electric trains. It was not applied to stations on electrified lines where services are provided by diesel trains, such as Acklington. The ferry interchange dummy was applied to stations where direct interchange with a regular sea-going ferry service is provided, such as Fishguard Harbour and Holyhead. These dummy variables were added to both Models 4.34 and 4.37, as it was not obvious whether they would be compatible with the distance dummy variables, giving Models 4.39-4.40. The results from globally calibrating these models on the corrected England and Wales dataset are shown in Table 4.31. The additional dummy variables gave positive results, with a noticeable improvement in fit over the global calibration of previous models. However, the harbour dummy variable was insignificant in both models, possibly because of correlation with the terminus dummy variable, and the outermost London distance dummy variable is also insignificant. The value and significance of the electrification parameter were noticeably lower in Model 4.40 than Model 4.39, probably because services near to London are more likely to be electrified than those further away. The insignificant variables (the harbour and outermost London distance band dummies) were removed to give Models 4.41 and 4.42, and the results from calibrating these are also summarised in Table 4.31.

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \rho \ln Pk_i + \eta B_i + \kappa Te_i + \ldots$$
$$\nu El_i + \chi H_i \tag{4.39}$$

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \rho \ln Pk_i + \upsilon_1 L_1 + \upsilon_2 L_2 + \upsilon_3 L_3 \ldots$$
$$\upsilon_4 L_4 + \eta B_i + \kappa Te_i + \nu El_i + \chi H_i \tag{4.40}$$

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \rho \ln Pk_i + \eta B_i + \kappa Te_i + \nu El_i \tag{4.41}$$

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \rho \ln Pk_i + \upsilon_1 L_1 + \upsilon_2 L_2 \ldots$$
$$+ \upsilon_3 L_3 + \eta B_i + \kappa Te_i + \nu El_i \tag{4.42}$$

Where:

B_i is a dummy variable which takes the value 1 if Station i is a Travelcard boundary station, and 0 otherwise

Te_i is a dummy variable which takes the value 1 if Station *i* is a terminus, and 0 otherwise

El_i is a dummy variable which takes the value 1 if Station *i* is served by electric trains, and 0 otherwise

H_i is a dummy variable which takes the value 1 if Station *i* provides interchange with a passenger ferry service, and 0 otherwise

Table 4.31: Summarised results from calibration of Models 4.39-4.42

Model	4.39		4.39		4.41		4.42	
Parameter	Value	t stat	Value	t stat	Value	t stat	Value	t stat
Intercept	2.976	16.890	2.837	16.020	2.979	16.916	2.841	16.039
β parameter	0.145	6.280	0.152	6.496	0.145	6.259	0.150	6.437
δ parameter	1.293	38.690	1.281	38.185	1.293	38.694	1.278	38.109
λ parameter	0.293	9.629	0.315	9.991	0.293	9.637	0.320	10.219
τ parameter	0.151	7.725	0.159	8.045	0.152	7.740	0.161	8.135
ρ parameter	0.194	15.994	0.191	15.515	0.194	15.988	0.192	15.602
υ_1 parameter	n/a	n/a	0.222	2.045	n/a	n/a	0.206	1.907
υ_2 parameter	n/a	n/a	0.311	4.192	n/a	n/a	0.288	3.927
υ_3 parameter	n/a	n/a	0.326	4.576	n/a	n/a	0.301	4.299
υ_4 parameter	n/a	n/a	0.144	1.834	n/a	n/a	n/a	n/a
η parameter	0.433	3.470	0.448	3.609	0.433	3.473	0.438	3.526
κ parameter	0.877	7.726	0.862	7.641	0.898	8.275	0.879	8.165
ν parameter	0.328	6.346	0.191	3.345	0.329	6.378	0.214	3.833
χ parameter	0.187	0.601	0.210	0.678	n/a	n/a	n/a	n/a
R_{adj}^2	0.778		0.782		0.778		0.782	
AD	0.816		0.802		0.824		0.804	

The removal of the insignificant parameters did not reduce model fit, and gave a more parsimonious model, so these reduced forms were taken forward. A final modification of the global model was to test the inclusion of dummy variables representing single zones around four further cities (Manchester, Liverpool, Birmingham and Cardiff), as Figure 4.11 suggested that these might improve model fit. This map suggested that zones of 10 km around Manchester, 15 km around Liverpool, 15 km around Birmingham and 20 km around Cardiff would be most suitable. Incorporating these variables in the model gave Model 4.43 and the results from calibrating this model are summarised in Table 4.32. This shows that it gave a small improvement in fit over Model 4.42, and the Manchester, Birmingham, and Cardiff dummies were all negative and significant, suggesting that these city areas have lower rail use than would be expected elsewhere, all else being equal. It seemed possible that this reflected competition with local bus services in these large urban areas, but while it would be desirable to include a variable to represent this no suitable data was available. The Liverpool dummy variable was insignificant, as was the dummy for the innermost London zone, and these variables were therefore removed from the model, giving Model 4.44. However, it also seemed sensible to test the inclusion of variables earlier discarded, in case their insignificance had been caused by the omission of variables added later, and this modification gave Model 4.45. Some variables were still found to be insignificant in Model 4.45, and the omission of these gave Model 4.46. The results of calibrating these models are also given in Table 4.32.

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \rho \ln Pk_i + \upsilon_1 L_1 + \upsilon_2 L_2 + \upsilon_3 L_3 + \ldots$$
$$\gamma Ma_i + \psi Li_i + \xi Bi_i + \omega Ca_i + \eta B_i + \kappa Te_i + \nu El_i \tag{4.43}$$

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \rho \ln Pk_i + \upsilon_2 L_2 + \upsilon_3 L_3 + \gamma Ma_i + \ldots$$
$$\xi Bi_i + \omega Ca_i + \eta B_i + \kappa Te_i + \nu El_i \tag{4.44}$$

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \zeta \ln M_{ihd4} + \rho \ln Pk_i + \upsilon_1 L_1 + \ldots$$
$$\upsilon_2 L_2 + \upsilon_3 L_3 + \upsilon_4 L_4 + \gamma Ma_i + \psi Li_i + \xi Bi_i + \omega Ca_i + \eta B_i + \kappa Te_i + \nu El_i + \chi H_i \tag{4.45}$$

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \zeta \ln M_{ihd4} + \rho \ln Pk_i + \upsilon_1 L_1 + \ldots$$
$$\upsilon_2 L_2 + \upsilon_3 L_3 + \gamma Ma_i + \xi Bi_i + \omega Ca_i + \eta B_i + \kappa Te_i + \nu El_i \tag{4.46}$$

Where:

M_{ihd4} is the mean number of motor vehicles owned per household within 4 minutes drive time of station i (ζ parameter)

Ma_i is a dummy variable which takes the value 1 if station i is within 10 km straight line distance of central Manchester, and 0 otherwise

Li_i is a dummy variable which takes the value 1 if station i is within 15 km straight line distance of central Liverpool, and 0 otherwise

Bi_i is a dummy variable which takes the value 1 if station i is within 15 km straight line distance of central Birmingham, and 0 otherwise

Ca_i is a dummy variable which takes the value 1 if station i is within 20 km straight line distance of central Cardiff, and 0 otherwise

H_i is a dummy variable which takes the value 1 if Station i provides interchange with a passenger ferry service, and 0 otherwise (χ)

Table 4.32: Summarised results from calibration of Models 4.43-4.46

	Model 4.43		Model 4.44		Model 4.45		Model 4.46	
Parameter	Value	t stat	Value	t stat	Value	t stat	Value	t stat
Intercept	2.817	15.840	2.821	15.963	2.037	9.657	2.021	9.596
β parameter	0.144	6.202	0.147	6.405	0.182	7.909	0.182	7.905
δ parameter	1.311	37.282	1.321	38.667	1.292	37.487	1.300	38.203
λ parameter	0.303	9.619	0.296	9.690	0.340	10.787	0.342	10.906
τ parameter	0.166	8.447	0.163	8.413	0.235	10.734	0.235	10.805
ζ parameter	n/a	n/a	n/a	n/a	0.325	3.366	0.320	3.391
ρ parameter	0.186	15.136	0.184	15.228	0.171	14.252	0.171	14.202
υ_1 parameter	0.112	0.986	n/a	n/a	0.276	2.431	0.216	1.989
υ_2 parameter	0.229	2.973	0.194	2.755	0.267	3.415	0.215	2.910
υ_3 parameter	0.260	3.607	0.236	3.430	0.290	3.964	0.245	3.510
υ_4 parameter	n/a	n/a	n/a	n/a	0.108	1.410	n/a	n/a
γ parameter	-0.376	-2.679	-0.402	-2.905	-0.313	-2.292	-0.349	-2.583
ψ parameter	0.084	0.767	n/a	n/a	0.173	1.625	n/a	n/a
ξ parameter	-0.454	-3.352	-0.483	-3.635	-0.392	-3.014	-0.428	-3.327
ω parameter	-0.466	-3.380	-0.486	-3.551	-0.425	-3.229	-0.453	-3.461
η parameter	0.405	3.285	0.400	3.247	0.414	3.522	0.403	3.438
κ parameter	0.888	8.291	0.888	8.293	0.826	7.687	0.833	8.025
ν parameter	0.188	3.250	0.206	3.706	0.175	3.071	0.212	3.921
χ parameter	n/a	n/a	n/a	n/a	0.120	0.382	n/a	n/a
R_{adj}^2	0.785		0.786		0.783		0.783	
AD	0.792		0.795		0.745		0.746	

The removal of the insignificant variables gave a marginal improvement in model fit with Model 4.44 compared to Model 4.43. Several parameters were insignificant in Model 4.45, and while all parameters were significant in Model 4.46 the model fit was slightly inferior to that of Model 4.44. Because it is desirable to have as parsimonious a model as possible, Model 4.44 was therefore adopted as the preferred global model form. The residuals from this model are mapped in Figure 4.21, which shows that the size of the residuals around the major cities was in general reduced by the modifications, although there were still a number of stations where the prediction error was extremely large.

Figure 4.21: Mapped residuals from calibration of Model 4.44

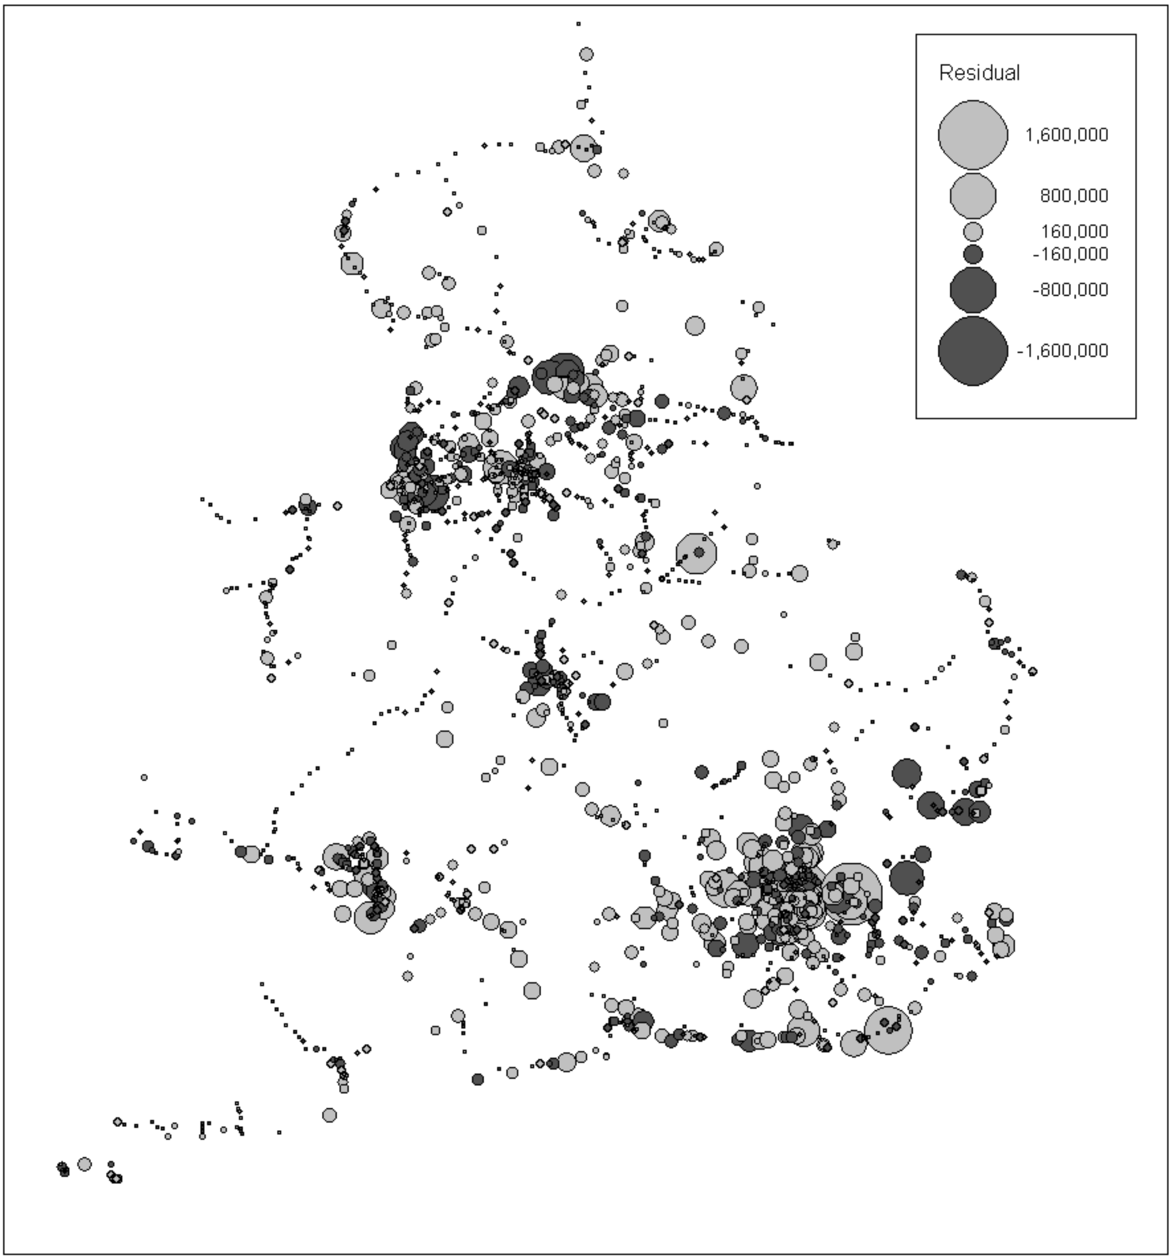

GWR was applied to the best model variant without distance-based dummy variables, Model 4.41, giving the results summarised in Table 4.33. This gave an obvious improvement over the global calibration, with an improvement in R_{adj}^2, a significant F statistic and a reduction in the AIC (3346.45 compared to 3534.021 for the global calibration). It also gave an improvement over the previous GWR calibration (of Model 4.34) with the R_{adj}^2 value increasing from 0.802 to 0.824. A noteworthy feature of the calibration results is the electrification parameter of 0.147. Taking the exponential of this implies that the 'sparks effect' increases patronage by 15.7% which, while smaller than the uplift suggested by Evans (1969), could still provide a strong case for electrification.

Table 4.33: Summary of results from GWR calibration of Model 4.41

Calibration method		GWR					
m		Minimum	Lower quartile	Median	Upper quartile	Maximum	Monte Carlo P-value
Intercept	Value	0.48825	2.24917	2.91773	3.58144	5.58828	0.000
β parameter	Value	-0.20824	0.07928	0.18352	0.23621	0.37086	0.020
δ parameter	Value	0.80878	1.14774	1.27900	1.41187	1.67249	0.040
λ parameter	Value	0.02518	0.15924	0.29441	0.43710	0.78191	0.000
τ parameter	Value	-0.03966	0.09157	0.14741	0.24103	0.45176	0.010
ρ parameter	Value	0.07484	0.15508	0.18583	0.21304	0.30038	0.250
η parameter	Value	-0.86311	0.19069	0.36069	0.54649	0.93960	0.990
κ parameter	Value	-0.25348	0.45314	0.70179	1.03740	1.61809	0.070
ν parameter	Value	-0.85591	0.00000	0.14669	0.41677	1.41956	0.000
Adj. R^2		0.824					
		SS	DF	MS	F		
OLS Residuals		914.9	9.00				
GWR Improvement		256.0	136.38	1.8773			
GWR Residuals		658.9	1353.62	0.4867	3.8569		

Significant spatial variation at the 5% level was found in the population, frequency, distance to higher category station, jobs and electrification parameters and in the intercept. The spatial variation for these parameters was mapped to allow it to be examined more closely as shown in Figures 4.23-4.28. MapInfo was used to create a continuous raster grid by interpolating the values of the parameter in question at each station in the dataset across the spaces between these stations, allowing shaded contours to be produced. This was carried out using inverse distance weighting (IDW) interpolation, which uses a distance weighted average of data points to calculate grid cell values. Inevitably the local values shown on the resulting maps will be more reliable in areas where there is a high density of data points (in this case urban areas where stations are closer together), but the general trends shown should still be reasonably reliable. The maps also show the local t statistics for the parameters, which illustrate where they are statistically significant at the 95% level. It should be noted that, using a 95% significance level, it would be expected that 5% of the t statistics will generate false positives; in other words that at least 5% of the points would be significant, as a result of the multiple hypothesis testing problem (Fotheringham et al., 2002).

Figure 4.22 suggests that variations in catchment population have the greatest impact on rail demand in the South East, Eastern England, North Wales and the Liverpool and Manchester areas. These areas coincide with the locations where the population parameter was statistically significant. The reasons for this variation were not obvious, and therefore it was not possible to identify any further variables which could be added to the model to account for this variation.

Figure 4.22: Local significance of β parameter from GWR calibration of Model 4.41

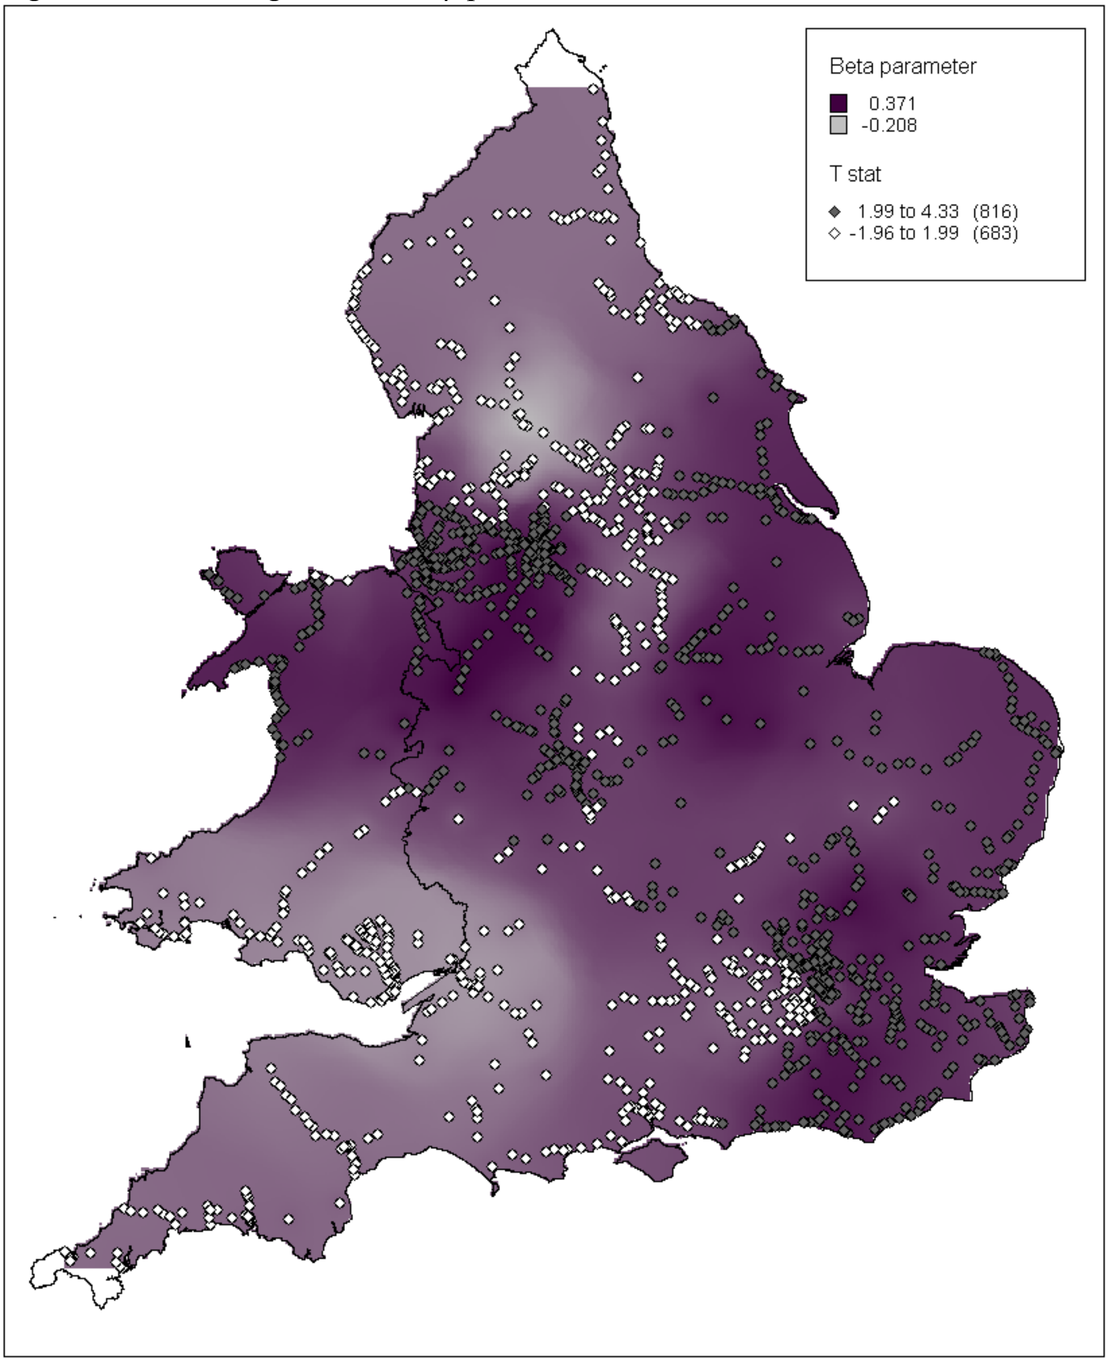

Figure 4.23 indicates that train frequency has the greatest impact on rail demand in the North of England, the East Midlands, West Sussex and Cornwall, but has less of an impact in London, the northern Home Counties and Wessex. Once again, the reasons for this variation were not apparent. The train frequency variable is significant in all areas, suggesting that it has a major impact on rail demand regardless of location.

Figure 4.23: Local significance of δ parameter from GWR calibration of Model 4.41

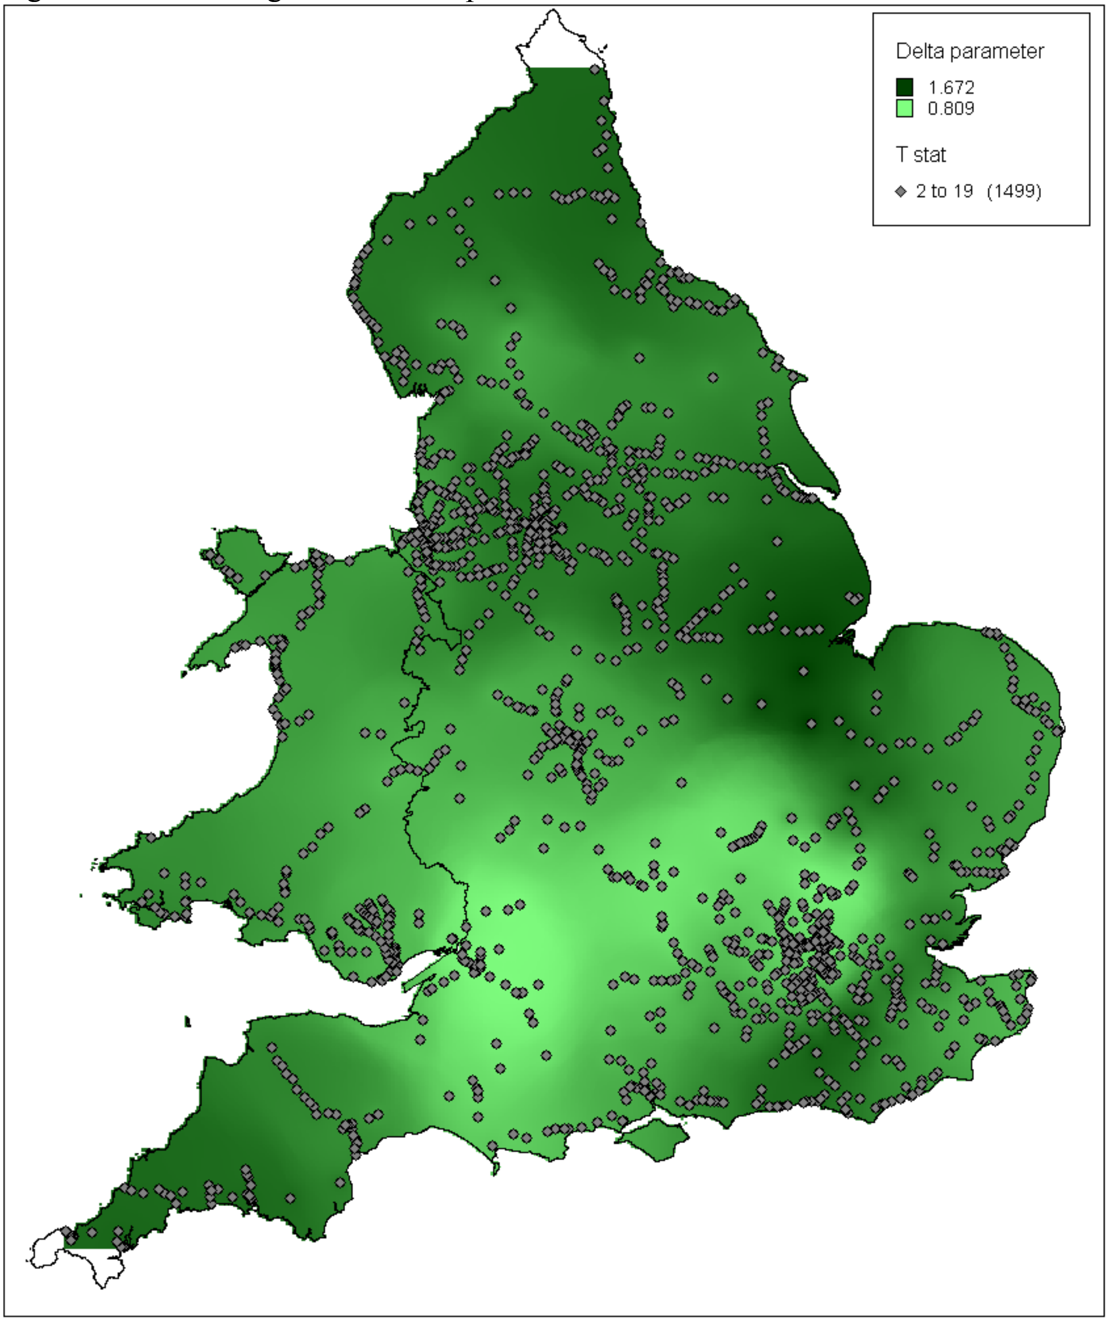

Figure 4.24 suggests that distance from the nearest category A-D station has the most positive impact on demand in Central England and along the East Coast, whereas it is relatively unimportant in Central London, South-West England and the South Wales Valleys. The reasons for this variation are not obvious, although the high density of category A-D stations in Central London may mean that there is little variation in distances from local stations to their nearest larger station. The local t statistics show that the parameter is in general significant where it is most positive.

Figure 4.24: Local significance of λ parameter from GWR calibration of Model 4.41

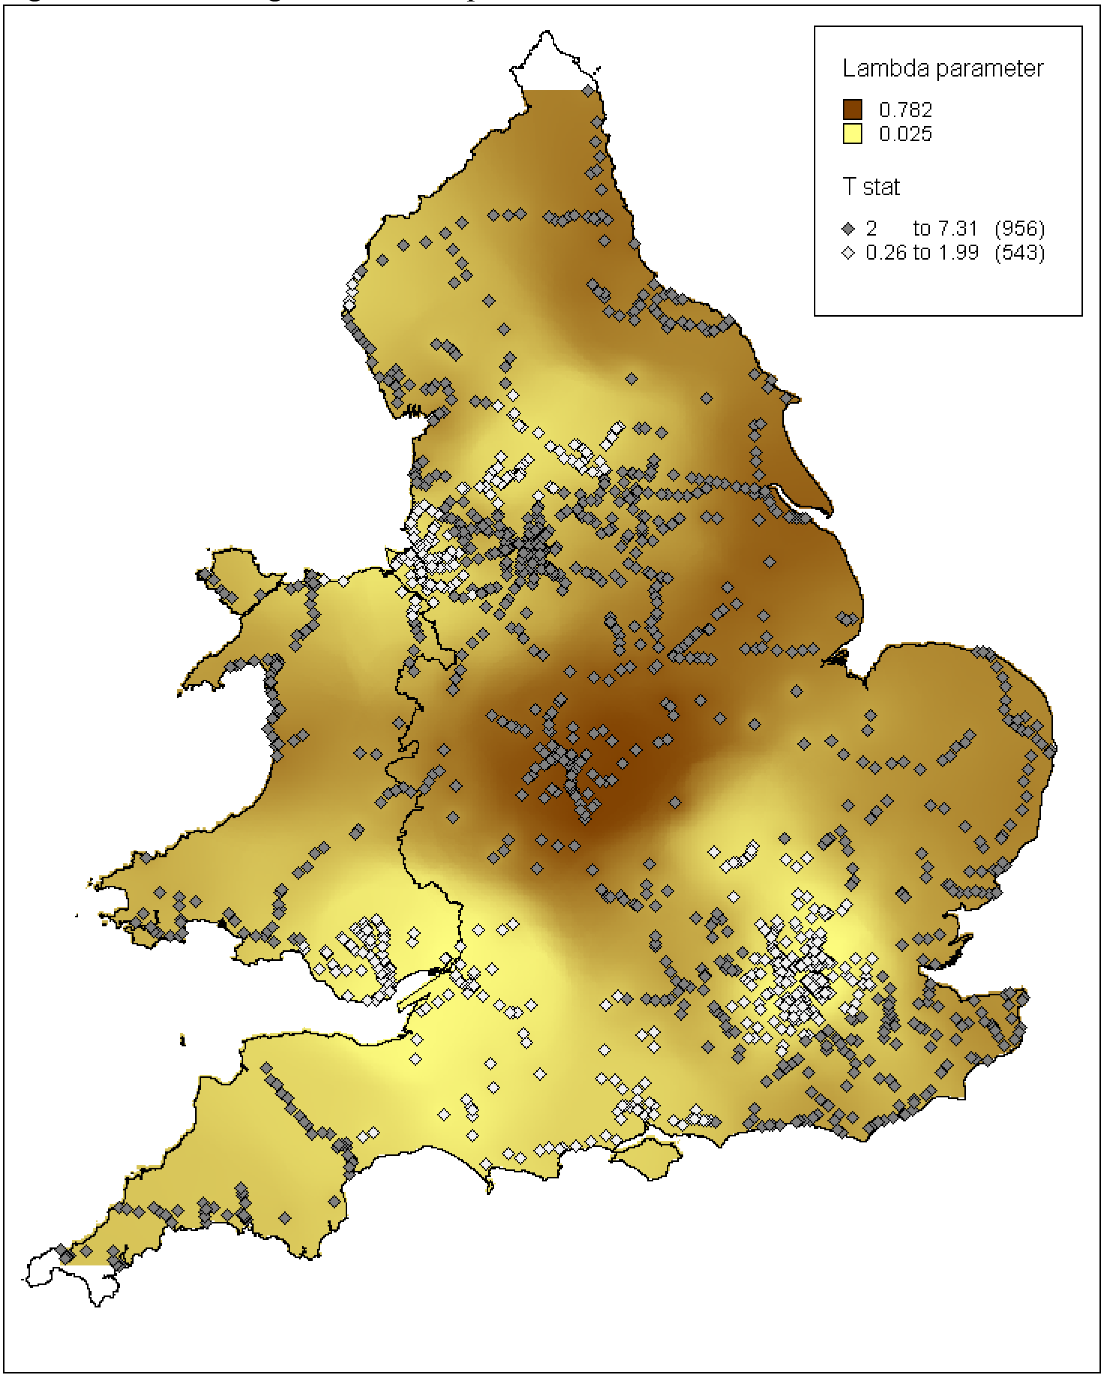

Figure 4.25 shows that the number of jobs located close to a station has the strongest impact on rail demand to the west of London and in Norfolk, and also has a significant impact in South Wales and the South-West. In contrast, this factor has no significant impact on rail demand in the North West, North Wales, and to the South and East of London. The latter pattern may result from the prevalence of rail commuting in this area, which means that the majority of stations are predominantly trip generators, and that only a small proportion of trips are attracted by jobs around the stations. The other variations are harder to explain.

Figure 4.25: Local significance of τ parameter from GWR calibration of Model 4.41

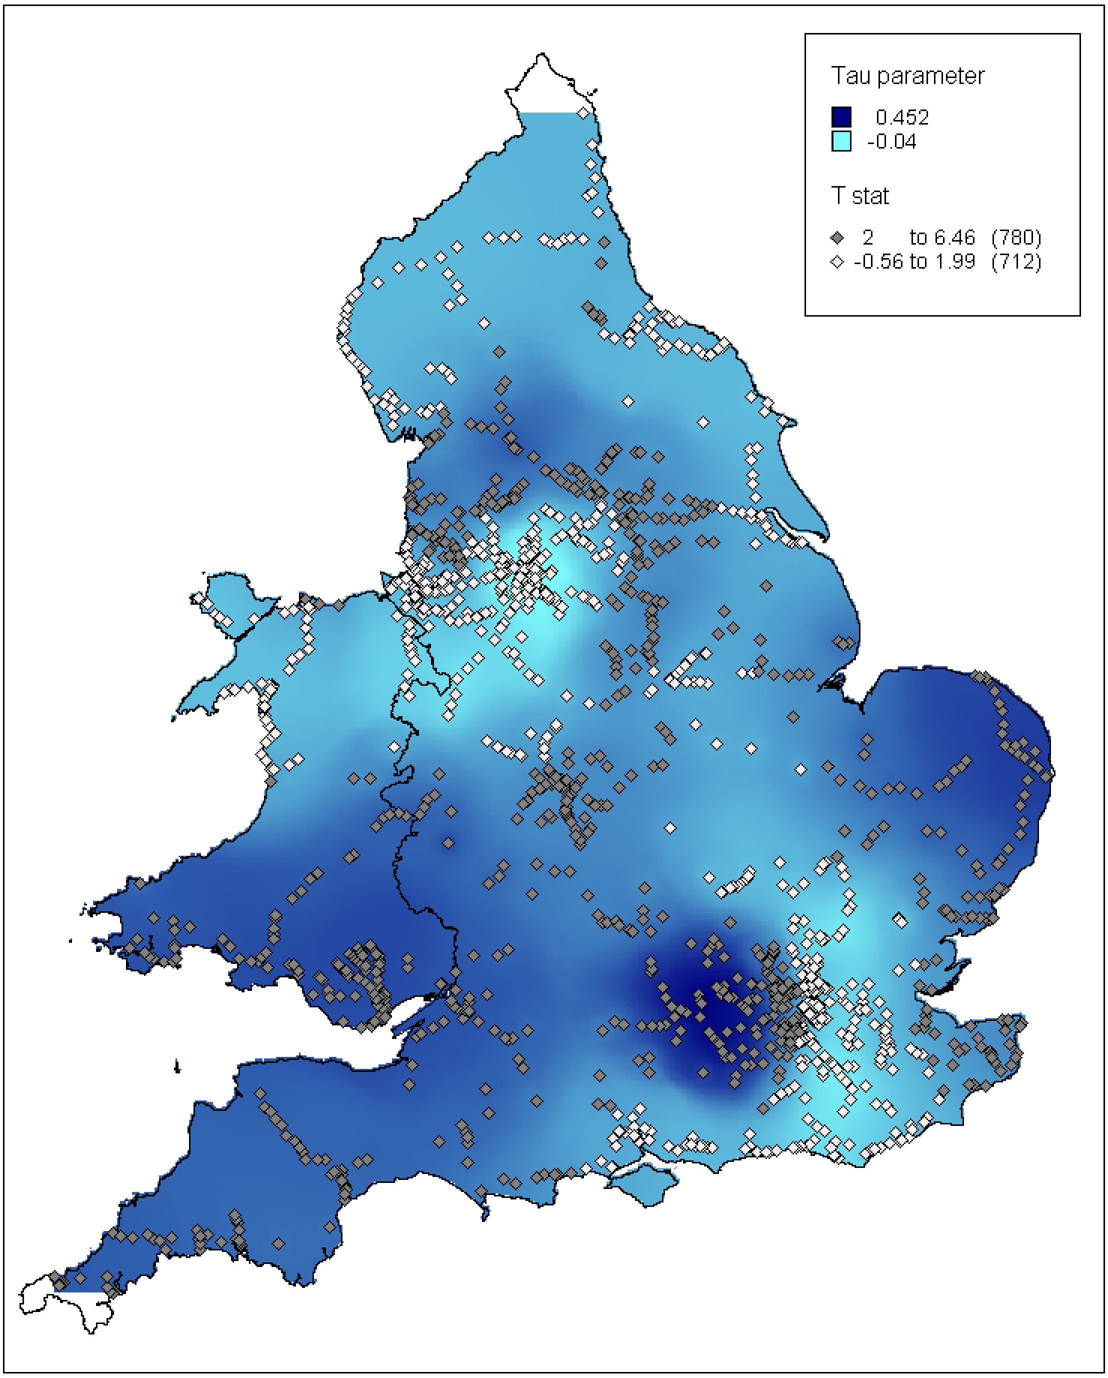

Figure 4.26 shows that electrification has the greatest impact on local rail demand to the north of London, and also around Merseyside. This was to be expected, as in these areas some but not all routes are electrified, allowing competition between electric- and diesel-hauled services. In contrast, in areas where either all or no routes are electrified the parameter is insignificant, as there is no such choice available to passengers. It was perhaps surprising though that the parameter was not significant in the West Midlands, Greater Manchester and West Yorkshire where only certain routes are electrified.

Figure 4.26: Local significance of ν parameter from GWR calibration of Model 4.41

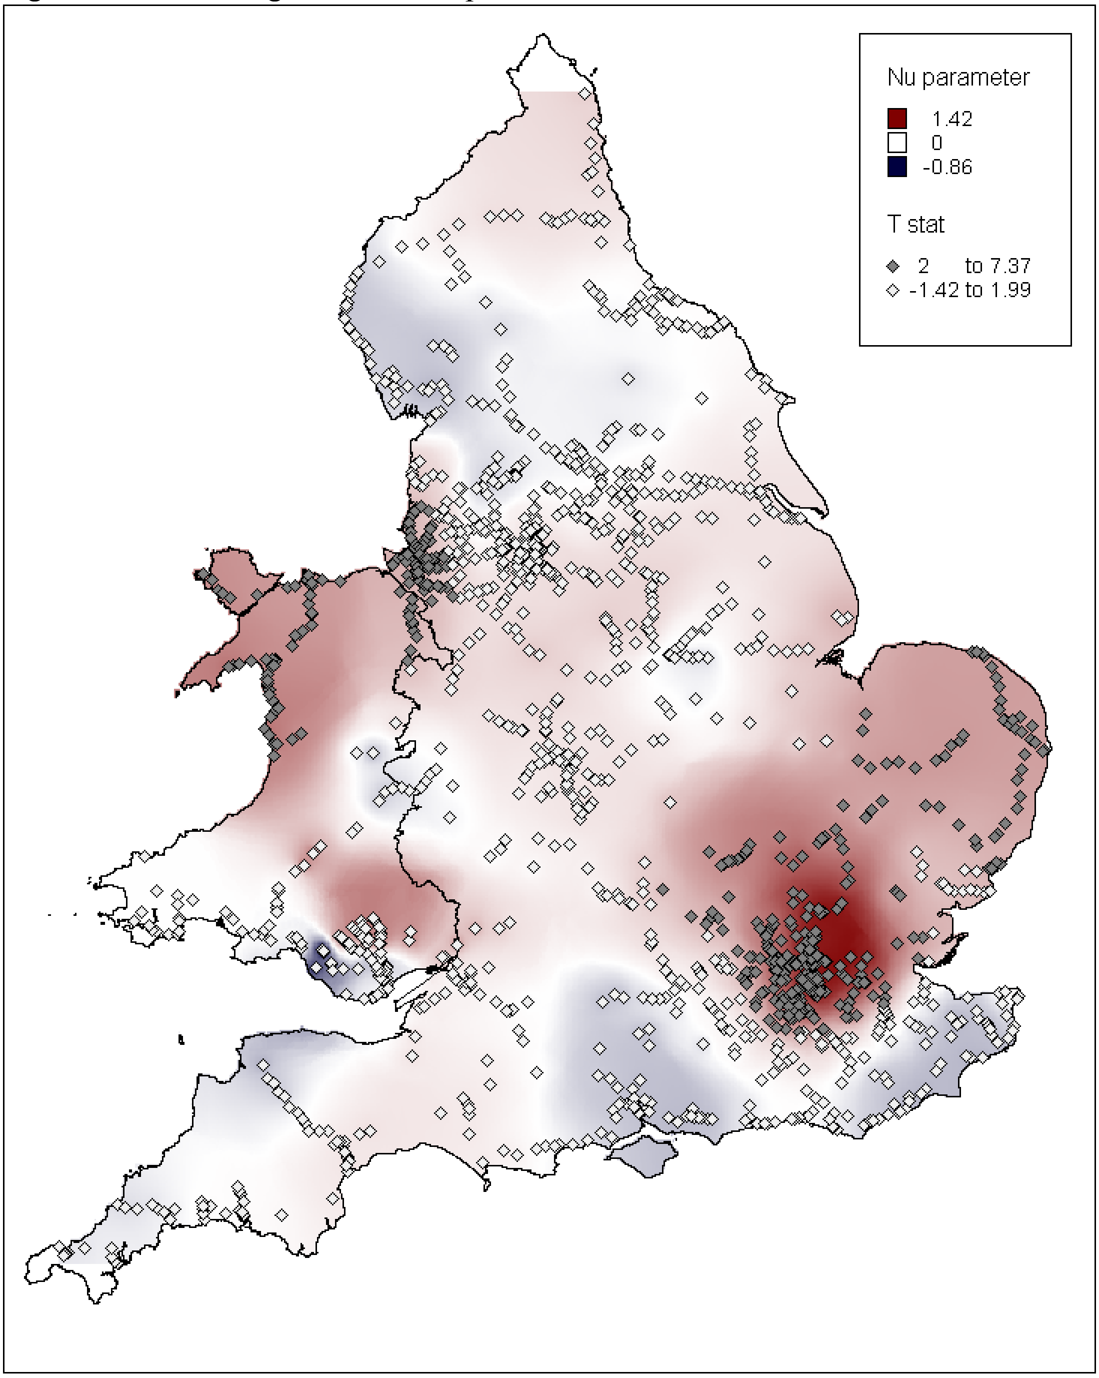

Figure 4.27 shows that the intercept term has the greatest effect on model predictions in the Wessex area, to the north of London and in northern England, indicating that factors not included in the model have a significant impact on rail demand in these areas. Despite lengthy consideration it was not possible to identify what these omitted factors might be. Conversely, the factors included in the model explain the greatest proportion of rail demand in Norfolk and Lincolnshire.

Figure 4.27: Local significance of intercept from GWR calibration of Model 4.41

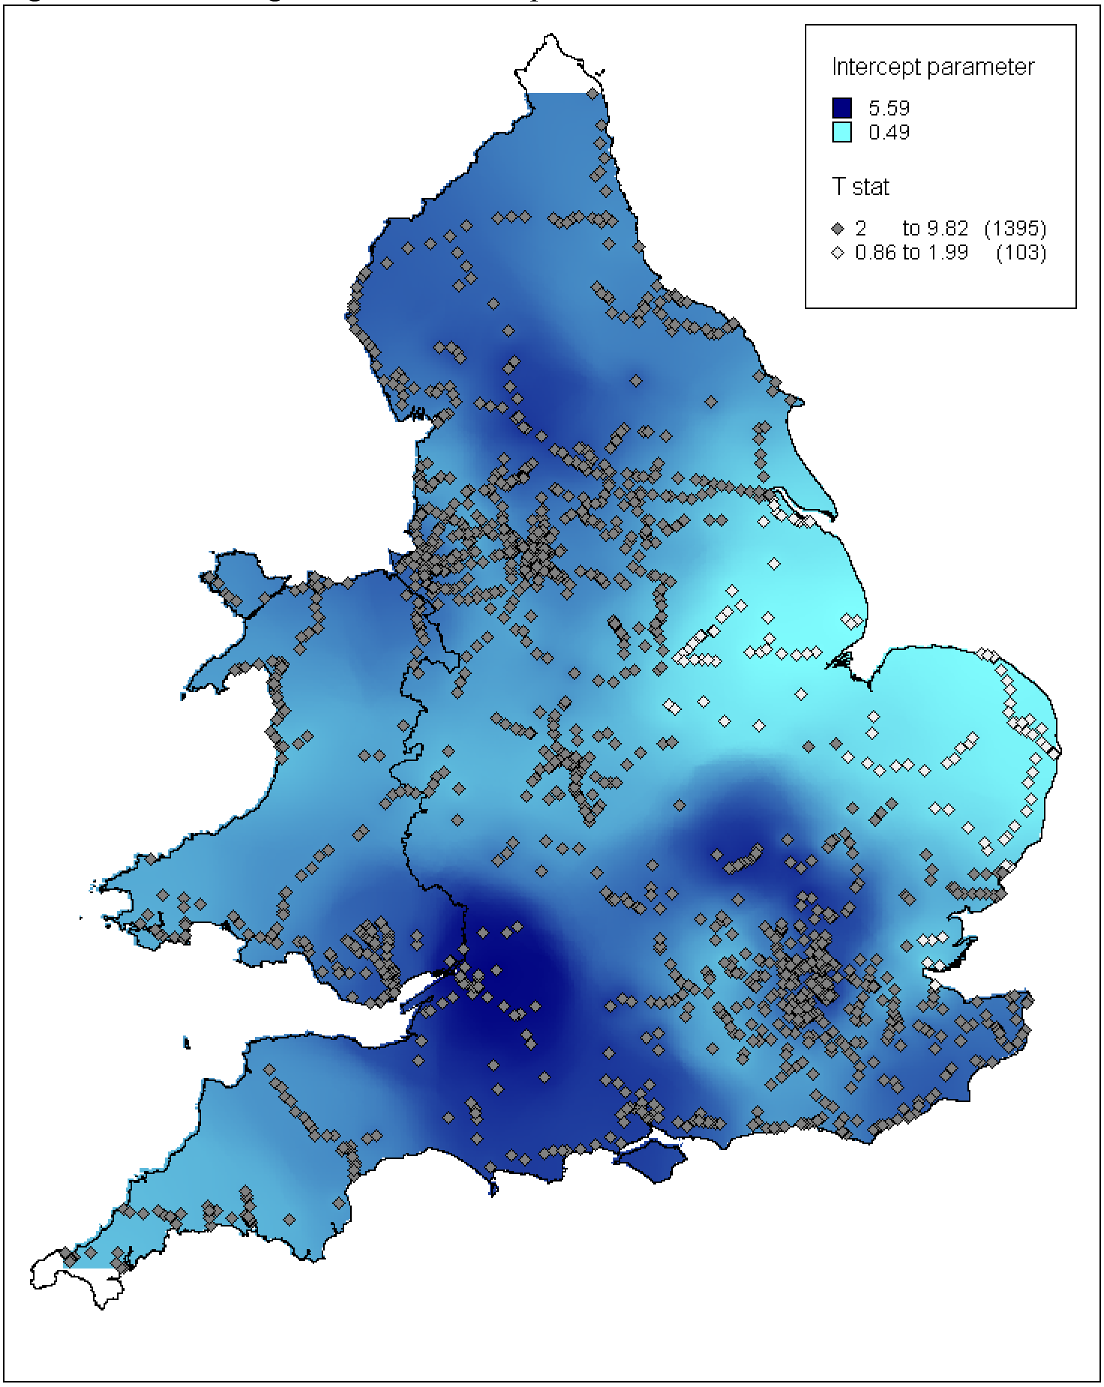

The prediction errors from the GWR calibration of Model 4.41 were mapped in Figure 4.28, showing that as before the largest residuals were concentrated around the major cities. This calibration had a superior fit to the best global model (4.44), and the form of the GWR model seems inherently more realistic because it allows the effect of parameters on rail demand to vary gradually across space. This contrasts with the inclusion of several arbitrarily defined zones as variables in the global model, where the sudden 'step change' in demand across zonal boundaries is unlikely to be representative of reality. While if intermodal and interstation competition could be represented in the latter model this might account for some of the spatial parameter variation identified by GWR, the main conclusion drawn from this section of the study was that the GWR model 4.41 gave the best results of all the trip end models developed and should therefore be the preferred method for forecasting total demand at new local railway stations.

Figure 4.28: Mapped residuals from GWR calibration of Model 4.41

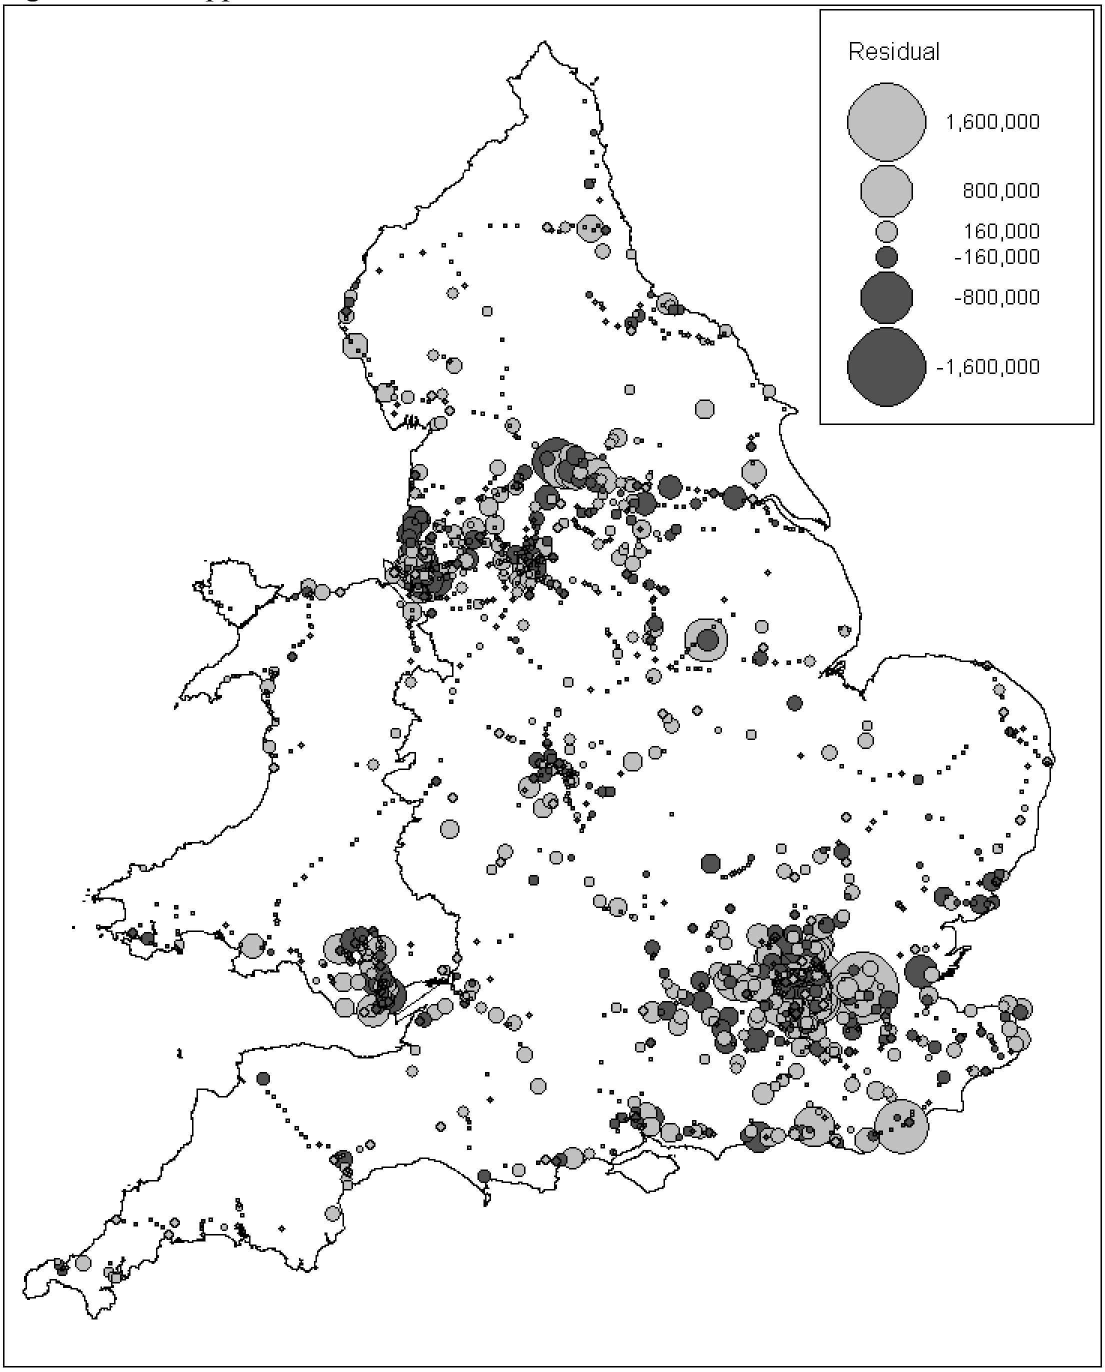

4.3.3 Cluster Analysis

While the best trip end models gave good results, their performance may have been impeded by the disparate range of station types contained within the calibration dataset. The 1499 stations included can all be described as local, but they are not all similar, and cluster analysis was therefore used to investigate whether the dataset should be partitioned. Hierarchical cluster analysis was identified as being the most appropriate approach, and several methods for defining the distance between clusters were investigated. Table 3.3 indicated that the squared Euclidean distance was the most appropriate measure of similarity and this was initially used along with the between-groups linkage clustering method. Variables were standardised based on their Z-scores (to a mean of 0 and to unit variance), as while this can affect correlation levels and relative distances between cases (Aldenderfer & Blashfield, 1984) it was necessary here to prevent variable size having an undue effect on clustering. All independent variables included in the best global trip end model (4.44) were included in the cluster analysis, with the exception of the distance to major city dummy variables. The dependent variable was not included, as for forecasting purposes it was necessary that a new station site should be able to be assigned to a cluster. The variables used were therefore population (β), train frequency (δ), distance to a higher category station (λ), employment (τ), car park size (ρ), travelcard boundary dummy (η), terminus dummy (κ) and electrification dummy (ν).

Hierarchical clustering results are usually represented in a dendrogram, but for such a large dataset this is too unwieldy to display in its entirety. No consensus exists on the rules which should be applied to select the optimal number of clusters from the dendrogram, and Everitt et al. (2001) therefore suggest that using informal and subjective criteria based on subject expertise is perhaps the best approach. SPSS provides a numerical agglomeration schedule which summarises the cluster solution, and the SPSS help system suggested that the best way to select the number of clusters to use was to scan the coefficients column of this schedule for large gaps. This is because a good cluster solution would see a sudden jump in the distance coefficient. The first cluster analysis indicated that using either two or four clusters gave the optimal results. The number of stations included in these clusters are summarised in Table 4.34, which suggests that the clustering was unsuccessful, as the vast majority of stations are contained within a single cluster in both cases.

Table 4.34: Clusters produced using between-groups linkage

Clusters used	2	4
Cluster	Number of stations	Number of stations
1	5	5
2	1494	41
3	n/a	21
4	n/a	1432

A further test of the success of the clustering was provided by recalibrating Model 4.44 separately for each cluster, and comparing the overall AD value with the value given by a single overall calibration. The results from this calibration are summarised in Table 4.35.

Table 4.35: Summarised results from calibration of Model 4.44 using between-groups linkage clusters

Clusters		Overall	2		4			
Cluster			1	2	1	2	3	4
Intercept	Value	2.821	-0.564	2.821	-0.564	3.056	9.509	2.752
	t stat	15.963	0	15.928	0	1.967	1.723	15.235
β parameter	Value	0.147	0.500	0.147	0.500	0.245	0.102	0.145
	t stat	6.405	0	6.362	0	1.600	0.266	6.142
δ parameter	Value	1.321	2.530	1.320	2.530	1.290	0.551	1.334
	t stat	38.667	0	38.560	0	5.880	1.033	37.599
λ parameter	Value	0.296	0	0.296	0	0.392	-0.266	0.305
	t stat	9.690	0	9.659	0	1.287	-0.707	9.754
τ parameter	Value	0.163	0	0.164	0	0.134	-0.023	0.165
	t stat	8.413	0	8.416	0	1.143	-0.082	8.314
ρ parameter	Value	0.184	-0.072	0.183	-0.072	0.072	-0.091	0.184
	t stat	15.228	0	15.076	0	1.015	-0.342	14.722
υ_2 parameter	Value	0.194	0	0.191	0	0.336	-0.099	0.196
	t stat	2.755	0	2.702	0	0.945	-0.111	2.653
υ_3 parameter	Value	0.236	0	0.235	0	0	0	0.228
	t stat	3.430	0	3.406	0	0	0	3.267
γ parameter	Value	-0.402	0	-0.403	0	0	0	-0.402
	t stat	-2.905	0	-2.906	0	0	0	-2.884
ξ parameter	Value	-0.483	0	-0.484	0	0.737	0	-0.514
	t stat	-3.635	0	-3.633	0	0.830	0	-3.781
ω parameter	Value	-0.486	0	-0.485	0	0	0	-0.498
	t stat	-3.551	0	-3.545	0	0	0	-3.609
η parameter	Value	0.400	0	0.401	0	0	0	0
	t stat	3.247	0	3.246	0	0	0	0
κ parameter	Value	0.888	0	0.888	0	0	0	0.885
	t stat	8.293	0	8.280	0	0	0	8.221
ν parameter	Value	0.206	0.258	0.206	0.258	0.020	0	0.211
	t stat	3.706	0	3.693	0	0.065	0	3.684
R_{adj}^2		0.786	n/a	0.785	n/a	0.664	-0.204	0.783
Overall R^2		0.787	0.772		0.774			
Overall AD		0.795	0.988		0.984			

Table 4.35 shows that clustering the data in this way does not improve model results, as model fit as measured by R^2 was reduced and overall AD values were markedly increased. Furthermore, because the smallest clusters in each case only contained five stations calibration on these clusters did not give sensible results. In an attempt to solve these problems the clustering was repeated using within-groups linkage, centroid clustering, median clustering and Ward's method, giving the clusters shown in Table 4.36.

Table 4.36: Clusters produced using other clustering methods

Clustering method	Within-groups linkage		Centroid	Median		Ward's method	
Clusters used	4	5	2	2	5	2	5
Cluster	Stations	Stations	Stations	Stations	Stations	Stations	Stations
1	948	877	1493	1494	1439	615	474
2	49	49	6	5	41	884	828
3	465	71	n/a	n/a	8	n/a	56
4	37	465	n/a	n/a	5	n/a	100
5	n/a	37	n/a	n/a	6	n/a	41

Table 4.36 suggests that only the within-groups linkage and Ward's clustering methods produced clusters which might give useful results from the model, as the other methods produced clusters with very few members. Model 4.44 was therefore calibrated separately for each of the clusters produced by these methods, giving the results in Table 4.37.

Table 4.37: Results from calibration of Model 4.44 using between-groups linkage clusters

Clustering method		Within-groups linkage								
Clusters		4				5				
Cluster		1	2	3	4	1	2	3	4	5
Intercept	Value	2.683	4.910	3.656	4.051	2.606	4.910	1.792	3.656	4.051
	t stat	11.986	6.054	6.571	2.529	11.217	6.054	0.438	6.571	2.529
β parameter	Value	0.104	0.055	0.255	0.280	0.106	0.055	-0.028	0.255	0.280
	t stat	3.604	0.442	6.109	1.894	3.602	0.442	-0.110	6.109	1.894
δ parameter	Value	1.310	1.175	1.218	0.986	1.326	1.175	1.137	1.218	0.986
	t stat	29.888	8.120	18.824	3.504	29.113	8.120	5.920	18.824	3.504
λ parameter	Value	0.329	0.445	0.205	0.275	0.346	0.445	0.205	0.205	0.275
	t stat	8.330	3.999	3.821	0.948	8.355	3.999	1.281	3.821	0.948
τ parameter	Value	0.195	0.134	0.094	0.186	0.188	0.134	0.494	0.094	0.186
	t stat	7.940	1.174	2.741	1.537	7.441	1.174	1.418	2.741	1.537
ρ parameter	Value	0.225	0.009	0.156	0.025	0.233	0.009	0.188	0.156	0.025
	t stat	13.252	0.168	8.517	0.354	12.906	0.168	3.023	8.517	0.354
υ_2 parameter	Value	0.266	0.126	0.147	0.154	0.252	0.126	0.273	0.147	0.154
	t stat	1.844	0.351	1.855	0.383	1.659	0.351	0.576	1.855	0.383
υ_3 parameter	Value	0.347	0.399	0.205	0	0.341	0.399	0	0.205	0
	t stat	3.048	1.167	2.344	0	2.960	1.167	0	2.344	0
γ parameter	Value	-0.285	0	-0.777	0	-0.181	0	-0.423	-0.777	0
	t stat	-1.683	0	-3.029	0	-0.808	0	-1.699	-3.029	0
ξ parameter	Value	-0.447	0	-0.634	0.862	-0.505	0	-0.500	-0.634	0.862
	t stat	-2.448	0	-3.199	1.017	-2.235	0	-1.726	-3.199	1.017
ω parameter	Value	-0.623	0.198	-0.351	0	-0.605	0.198	-1.132	-0.351	0
	t stat	-3.853	0.514	-0.655	0	-3.622	0.514	-1.520	-0.655	0
η parameter	Value	-0.197	0	0.363	0	-0.004	0	0	0.363	0
	t stat	-0.338	0	0.739	0	-0.006	0	0	0.739	0
κ parameter	Value	0.463	0	0.248	0	0.559	0	0	0.248	0
	t stat	1.365	0	0.360	0	1.586	0	0	0.360	0
ν parameter	Value	0.304	0.037	-0.044	0.321	0.032	0.037	0.404	-0.044	0.321
	t stat	1.879	0.133	-0.106	0.908	0.109	0.133	1.761	-0.106	0.908
R_{adj}^2		0.769	0.716	0.650	0.538	0.767	0.716	0.543	0.650	0.538
Overall R^2		0.781				0.782				
Overall mean AD		0.940				0.954				

Table continued on next page

Clustering method		Ward's method						
Clusters		2		5				
Cluster		1	2	1	2	3	4	5
Intercept	Value	3.270	2.595	3.597	2.636	4.977	4.816	3.056
	t stat	10.395	11.282	8.995	10.832	6.959	1.767	1.967
β parameter	Value	0.234	0.097	0.260	0.100	0.033	-0.026	0.245
	t stat	6.209	3.301	6.388	3.333	0.293	-0.139	1.600
δ parameter	Value	1.251	1.339	1.211	1.326	1.184	1.290	1.290
	t stat	23.903	29.531	18.965	27.931	8.722	9.201	5.880
λ parameter	Value	0.242	0.347	0.203	0.331	0.386	0.415	0.392
	t stat	5.173	8.433	3.890	7.533	4.161	3.395	1.287
τ parameter	Value	0.103	0.193	0.094	0.192	0.156	0.085	0.134
	t stat	3.391	7.552	2.804	7.343	1.486	0.370	1.143
ρ parameter	Value	0.161	0.219	0.162	0.237	0.027	0.124	0.072
	t stat	10.286	11.932	8.902	12.195	0.575	2.668	1.015
υ_2 parameter	Value	0.172	0.205	0.149	0.212	0.231	0.456	0.336
	t stat	2.339	1.357	1.926	1.301	0.734	1.112	0.945
υ_3 parameter	Value	0.202	0.341	0.223	0.353	0.438	-0.280	0
	t stat	2.436	2.937	2.621	2.863	1.657	-0.531	0
γ parameter	Value	-0.553	-0.174	-0.815	-0.164	0	-0.426	0
	t stat	-3.385	-0.773	-3.074	-0.719	0	-1.848	0
ξ parameter	Value	-0.577	-0.259	-0.681	-0.272	0	-0.607	0.737
	t stat	-4.051	-1.022	-3.498	-1.061	0	-2.434	0.830
ω parameter	Value	-0.700	-0.399	-0.658	-0.501	0.180	-1.536	0
	t stat	-2.392	-2.455	-1.886	-2.884	0.492	-2.074	0
η parameter	Value	0.379	0	0	0	0	0	0
	t stat	3.137	0	0	0	0	0	0
κ parameter	Value	0	0.940	0	0	0	0	0
	t stat	0	7.043	0	0	0	0	0
ν parameter	Value	0.169	-0.014	0.007	0.839	-0.058	0.303	0.020
	t stat	1.912	-0.058	0.039	1.009	-0.236	1.621	0.065
R_{adj}^2		0.654	0.759	0.651	0.749	0.709	0.648	0.664
Overall R^2		0.775		0.797				
Overall mean AD		1.041		0.765				

Table 4.37 shows that only Ward's clustering method gave an improvement in model fit measured by mean AD over the overall calibration, and only when 5 clusters were used. However, the parameters were not all significant in all calibrations, perhaps because some of the clusters were too small. Once clusters have been defined, it is necessary to establish the features of the observations within each cluster that make them similar.

The mean values of each variable within each cluster from the set of 5 clusters defined using Ward's method are shown in Table 4.38 along with the overall mean values for the dataset. Some of the clusters have very obvious distinguishing features, with all stations in cluster 3 being termini, and all stations in cluster 5 being on travelcard boundaries. However, these differences can easily be accounted for in an unclustered model by dummy variables. All stations in cluster 4 have a large catchment population and a large number of jobs within their catchments, although there are stations in other clusters which have larger catchment populations or larger number of jobs than some stations in cluster 4. The majority of stations in cluster 1 are electrified, and stations in this cluster tend to be fairly close to a higher category station. Finally, stations in cluster 2 tend to have low service frequencies and small catchment populations and car parks.

Table 4.38: Within-cluster variable variation for 5 clusters defined by Ward's method

Cluster		1	2	3	4	5	Overall
Weighted population	5%ile	20.97	2.68	40.75	514.72	26.05	6.18
	Mean	521.14	297.95	520.86	1151.44	435.26	437.54
	95%ile	1481.08	796.53	1074.34	2091.95	839.98	1280.62
Frequency	5%ile	36.00	7.00	8.75	22.85	24.00	8.00
	Mean	90.14	36.14	43.66	62.29	62.93	55.97
	95%ile	180.05	82.00	94.75	115.65	137.00	141.00
Distance to higher category station	5%ile	1.74	3.73	2.54	2.08	4.10	2.13
	Mean	7.83	19.68	21.79	6.83	9.44	14.87
	95%ile	19.47	56.22	44.78	18.71	18.09	44.49
Jobs	5%ile	206.15	27.05	567.75	8270.65	237.00	66.90
	Mean	2979.07	2156.68	4506.64	14072.31	3896.00	3347.00
	95%ile	8462.25	7855.65	10809.25	24713.05	12631.00	11868.70
Parking	5%ile	0.00	0.00	0.00	0.00	0.00	0.00
	Mean	36.29	16.36	25.05	28.29	32.54	24.23
	95%ile	162.00	70.00	80.00	101.00	136.00	110.00
Travelcard boundary	5%ile	0.00	0.00	0.00	0.00	1.00	0.00
	Mean	0.00	0.00	0.00	0.00	1.00	0.03
	95%ile	0.00	0.00	0.00	0.00	1.00	0.00
Terminus	5%ile	0.00	0.00	1.00	0.00	0.00	0.00
	Mean	0.00	0.00	1.00	0.00	0.00	0.04
	95%ile	0.00	0.00	1.00	0.00	0.00	0.00
Electrification	5%ile	1.00	0.00	0.00	0.00	0.00	0.00
	Mean	0.95	0.00	0.27	0.29	0.29	0.34
	95%ile	1.00	0.00	1.00	1.00	1.00	1.00

While partitioning the dataset based on these clusters for model calibration slightly improved model fit, there is no straightforward way to allocate new stations to a cluster, as while all the clusters have distinctive features these tend not to be entirely exclusive to single clusters. This means that, while the partitioned calibration does have some attractive features, the single overall calibration (described in Section 4.3.2) was retained as the preferred method. The problems associated with identifying the defining characteristics of different clusters are not unique to this application of cluster analysis, and are a known limitation of this methodology. In some cases this may not be a major issue, but in applications such as this where it is necessary to allocate observations which are not included in the clustering (such as new stations) to the clusters it will form a major barrier to the use of cluster analysis.

4.3.4 Sensitivity Analysis

Validation of the trip end models calibrated on the England and Wales dataset presents . some problems because the entire dataset available was used for calibration. The best global model (4.44) was therefore validated by removing ten stations at random from the dataset, recalibrating the model, and using this calibration to predict demand at the omitted stations. This procedure was then repeated ten times, and the variations in parameter values and fit between calibrations were compared along with the prediction errors. Z statistics were also calculated in each case to test for a significant difference between the parameter estimates from the full calibration and from calibration using the reduced dataset. The range of the parameter estimates from the ten recalibrations is summarised in Table 4.39, with full details of parameter values and significance given in Blainey (2009c) (see Appendix 1).

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \rho \ln Pk_i + \upsilon_2 L_2 + \upsilon_3 L_3 + \gamma Ma_i + \ldots$$
$$\xi Bi_i + \omega Ca_i + \eta B_i + \kappa Te_i + \nu El_i \qquad (4.44)$$

Table 4.39: Summary of parameter variation from sensitivity analysis of Model 4.44

	Min	Max	Range	% Variation
Intercept	2.772	2.845	0.073	2.59%
β parameter	0.145	0.149	0.004	2.72%
δ parameter	1.318	1.331	0.013	0.98%
λ parameter	0.293	0.3	0.007	2.36%
τ parameter	0.161	0.166	0.005	3.07%
ρ parameter	0.182	0.185	0.003	1.63%
υ_2 parameter	0.186	0.2	0.014	7.28%
υ_3 parameter	0.219	0.242	0.023	9.79%
γ parameter	-0.408	-0.396	0.012	-2.98%
ξ parameter	-0.489	-0.452	0.037	-7.70%
ω parameter	-0.521	-0.48	0.041	-8.38%
η parameter	0.398	0.415	0.017	4.23%
κ parameter	0.888	0.897	0.009	1.01%
ν parameter	0.202	0.215	0.013	6.28%

None of the parameter values from the calibrations on the reduced dataset were found to be significantly different from those given by calibration on the full dataset (Z statistic >1.96 or <-1.96) and Table 4.39 shows that in general there is relatively little variation in the parameter values between calibrations. However, there are some exceptions to this, in particular in the parameters for the two distance to London dummy variables (υ_1 and υ_2) and in the Birmingham (ξ) and Cardiff (ω) dummy variables. This perhaps emphasises the point that the GWR calibration should be preferred to this global calibration, as it dispenses with the need for such zonal dummy variables. However, even for these variables the variation is not so large as to cast serious doubts on the validity of the model form, and most parameters seem to be reasonably stable.

The actual and predicted numbers of trip ends for the stations omitted from each calibration (full details given in Blainey (2009c), see Appendix 1) were compared, giving an overall AD value of 0.753, which was superior to the AD value of 0.795 for Model 4.44 calibrated on the entire dataset. The AD value was inflated by the poor predictions (when measured by % error) at three stations with very low usage, and removing these gave an improved overall AD value of 0.567 for the remaining 97 stations. A similar inflation of the AD value is likely to have occurred when the model was calibrated on the entire dataset, and therefore the model is likely in most cases to give more accurate predictions that the overall AD value would suggest.

While the model parameters appeared to be reasonably stable, no test had yet been made of the model's transferability over time, with the model so far calibrated only on usage data from 2005-06. Such a test was not straightforward, as the data on different variables did not all come from the same time period, and it was not possible to obtain data from different years on demographic variables because the census is only conducted at ten year intervals. Similarly, no data was available on change over time in car parking provision, and only a single set of CIF timetable data was available. Some measurement error may result in all cases from the fact that not all the calibration data refers to the same time period, but because the data collection periods are consistent over space this should not seriously affect model performance.

Data from a number of years on trips made was available from the ORR website, and therefore the best global model (4.44) and GWR model (4.41) were recalibrated for several years using the same data for the independent variables (Blainey, 2009c; see Appendix 1). Table 4.40 summarises the model fit from these recalibrations along with the original 2005-6 calibrations, and Table 4.41 gives the results of a significance test of the difference between the parameter values obtained from the original global calibration and the values obtained from calibrations on data from the other years. It was not possible to carry out a similar significance test of the parameters from the GWR calibrations because the GWR procedure does not produce suitable standard error values. This comparison showed that there was relatively little difference in model fit for both the global model and the GWR model between the four years for which usage data is available. There was more variation in the parameter estimates, as would perhaps be expected given the inevitable year-on-year variation in rail use from particular stations, and the Z statistics in Table 4.41 indicate that a minority of these variations are statistically significant. However the overall stability in model fit gave confidence that the model form was accurate and that it would be valid to use the GWR calibration of Model 4.41 to forecast demand at new stations.

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \rho \ln Pk_i + \eta B_i + \kappa Te_i + \nu El_i \quad (4.41)$$

Table 4.40: Summarised model fit from global calibration of Model 4.44 and GWR calibration of Model 4.41 for four different periods

Year	2002-03	2004-05	2005-06	2006-07
Global R_{adj}^2	0.778	0.778	0.786	0.792
GWR R^2	0.803	0.817	0.824	0.836

Table 4.41: Z statistics from significance test of difference between global parameter estimates from 2005-6 and from other years

Year	2002-3	2004-5	2006-7
Intercept	1.023	0.141	1.684
β parameter	2.870	1.000	1.565
δ parameter	-0.559	-0.588	1.088
λ parameter	-0.484	0.226	-3.194
τ parameter	-3.632	-1.053	-1.789
ρ parameter	-2.000	-0.333	-2.167
υ_2 parameter	0.814	0.429	0.471
υ_3 parameter	-0.014	0.246	-1.565
γ parameter	-1.399	-0.094	-1.130
ξ parameter	0.030	0.271	-1.158
ω parameter	1.000	0.672	-0.190
η parameter	-1.244	-0.154	-0.146
κ parameter	-0.364	0.252	0.187
ν parameter	0.161	0.036	0.732

4.4 Spreadsheet demand forecasting tool

Once effective models had been developed, it was necessary to incorporate them into a tool which could easily and quickly produce forecasts of demand for new stations. An Excel spreadsheet was therefore created, based on GWR Model 4.41, which can provide location-specific demand forecasts for any site in England and Wales. The user has to enter values for the independent variables for the proposed station site, along with its coordinates, and the spreadsheet will then almost instantly provide a forecast (see Figure 4.29). The demand impact of changing supply side variables such as train frequency or car park size at the site in question can be easily checked by altering the relevant values.

The demand forecasts are calculated using the model form defined above, based on matrices of values for each parameter defined during the GWR calibration and stored in the forecasting spreadsheet. The matrices allow location-specific demand forecasts to be made which are accurate to the nearest kilometre. The sheets and cells containing the parameter values and estimation formulae are locked to prevent inadvertent editing by users. As far as the author is aware, no such demand forecasting tool capable of producing instant location-specific forecasts over such a wide area has ever been produced in the past. The use of this spreadsheet to estimate demand at a large number of potential station sites across England and Wales is described in Section 7.3.

Figure 4.29: Spreadsheet tool for trip end forecasts

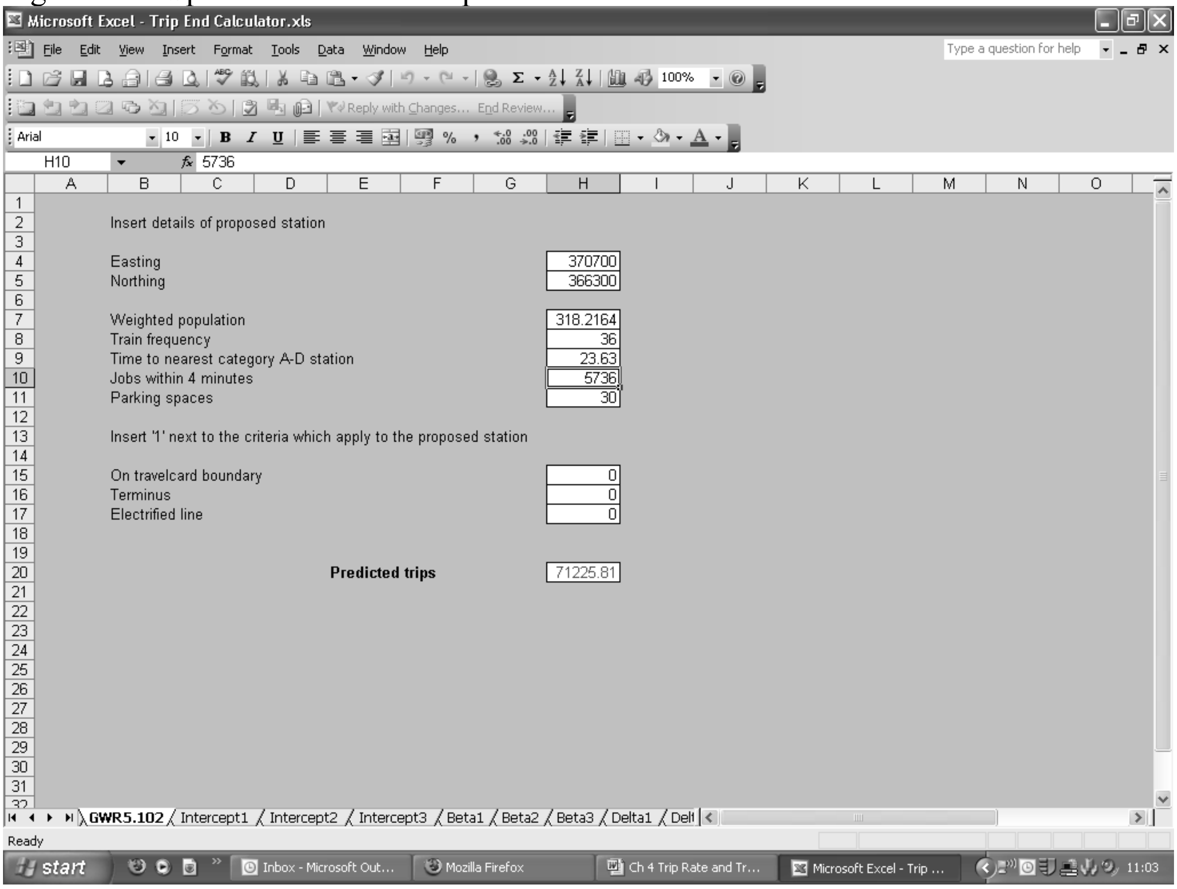

4.5 Conclusions

This chapter has described the development of a range of generalised trip rate and trip end models initially for a case study area centred on South Hampshire and then for a larger case study dataset containing nearly all local stations in England and Wales. Unweighted catchment populations limited by a 3.5 minute drive time were found to give the best results in trip rate models for South Hampshire, while a power weighting function was found to be most effective in defining the population variable with the larger dataset. The best trip end model for the South Hampshire area was Model 4.21, which had an extremely good fit ($R_{adj}^2 = 0.912$) and incorporated population, train frequency, rail distance to London and car park size variables.

To increase the spatial applicability of the models they were developed further with the much larger England and Wales dataset. A wide range of independent variables were tested in the models to establish which had a significant impact on the level of rail demand. The mapping of model residuals allowed the identification of additional independent variables which were then also tested in the models. The ability to depict model residuals in this way is a major advantage of using GIS in the demand modelling process, and one that has not previously been exploited in this field. Linear and loglinear forms for the dependent variable were compared using the Box-Cox test, with a loglinear dependent variable giving superior results. Linear, semilog and double log combinations of the explanatory variables were then compared, and a double log model form was found to give the best model performance. Trip end models have never before been calibrated on a dataset as large as that used here, and the range of explanatory variables tested is much greater than that used in previous models developed in the UK.

The use of GWR to enhance trip end models has been investigated, with the best global models recalibrated using this technique. The diagnostic measures included allowed some errors in the calibration dataset to be identified and also highlighted stations which might be having an undue influence on model results. Significant spatial variation was identified in several parameters and this variation was mapped. GWR was found to give superior results to the global regression models, with the best GWR model (4.41) incorporating population, train frequency, distance to larger station, employment, car park size, travelcard boundary, terminus and electrification variables. The best global model (4.44) included dummy variables representing distance to London, Manchester, Birmingham and Cardiff in addition to the variables from Model 4.41, but had an inferior fit compared to the GWR model. The successful implementation of a GWR-based modelling approach is a significant development in rail demand modelling as such local analysis techniques have not previously been applied in this field (Blainey 2009a). The fit of both these models ($R_2^{adj} = 0.824$ for Model 4.41 and $R_2^{adj} = 0.786$ for Model 4.44) compares favourably to those from the largest previous application of trip end models, which achieved R_{adj}^2 values of 0.760 for light rail and 0.571 for commuter rail in the USA (Lane et al., 2006). There are no comparable UK applications of this type of model, but the level of fit obtained here is superior to that from most other types of rail demand model used in this country, with the exception of dummy variable direct demand models which are not suitable for forecasting demand from new stations (see for example Wardman et al., 2007).

The use of hierarchical cluster analysis to partition the dataset was investigated but, while distinct clusters were identified and this clustering gave a marginal improvement in model predictions, no straightforward method to allocate new stations to clusters could be developed. This meant that, while the partitioned calibration had some attractive features, the single overall calibration was retained for use in demand forecasting for new stations. Nonetheless, cluster analysis may still have some potential for categorising railway stations, and as it has not (as far as the author is aware) ever been used before in this context it could be worthy of further investigations.

Sensitivity analysis of the best global regression model (4.44) showed that the parameter values were relatively insensitive to the removal of stations from the calibration dataset. This model, along with GWR Model 4.41, was also recalibrated on trip data from three other years and, while there was some variation in parameter values, the overall fit remained stable giving confidence that the model form was accurate.

The best model, GWR Model 4.41, was used to create a spreadsheet-based demand forecasting tool which will automatically produces a forecast of demand for any location in England and Wales once values of the independent variables have been input. No local rail demand model with this level of spatial transferability has ever previously been developed, and this represents a major advance on previous trip end models such as those contained in the PDFH, showing that such models can be generalised over a large area.

Chapter Five: Flow Level Models

5.1 Introduction

Most recent work on rail demand modelling in the UK has focused on flow level models, with increasingly complex cross-sectional direct demand models developed for interurban travel and for parkway stations (Lythgoe, 2004; Wardman et al., 2007), some of which also include a station choice element. However, the fit of these models is only slightly superior to that of much simpler direct demand models developed by Preston (1987, 1991), even though model specification, calibration and application is much more complex. Given that the data and computing techniques now available for rail demand modelling are far superior to those available when the simpler models were developed (see Sections 3.6-3.7), and given that data and computing deficiencies were one of the main limitations to the effectiveness of these models (Preston 1987, 1991), it seemed sensible in this project to concentrate (at least initially) on the development of relatively simple flow level models. Simplicity of calibration is particularly important for this type of model, as the sheer volume of data required for flow level models meant that calibration of a single model for local rail services throughout the UK was not feasible. To allow spatial transferability it must therefore be relatively straightforward to recalibrate the preferred model forms in areas other than those used for initial model development.

This chapter therefore details the development of location-specific flow level models for sections of the local rail network in South-East Wales. Firstly, in Section 5.2, cross-sectional direct demand models based on the general form derived from demand theory (see Section 2.5.4) are calibrated for several different subsets of flows from this area. Linear and loglinear models are compared, before the issue of incorporating intermodal competition in direct demand models is investigated in some detail. Catchment definition methods are then considered, with flow-specific catchments defined and tested in model calibration. Attempts are then made to constrain the total number of trips predicted by flow level models to equal the total number of trip origins predicted or observed at each origin station. Several different representations of intervening and competing opportunities are tested with models calibrated in the same area.

Section 5.3 details a survey of ultimate passenger origins and destinations carried out on the Cardiff to Rhymney line in South-East Wales. The survey methodology is first outlined, before the results from the survey are described and analysed. The observed trip ends are converted into estimated station catchments using a GIS interpolation procedure. Various issues affecting catchment shape and size are then discussed with particular attention paid to the relationship between access distance and access mode.

In Section 5.4 investigations into an alternative form of flow level model, the intervening opportunity trip distribution (IOTD) model, are described, with both linear and non-linear variants tested. The issue of destination selection for flow level models is then discussed in Section 5.5, and the findings of this chapter are drawn together in Section 5.6.

5.2 Direct demand models

5.2.1 Data processing

A range of origin trip end models were described in Chapter 4, and the best of these models gave extremely good results. However, these models did not consider trip destinations. Being able to predict the number of trips made to particular destinations is important, because it helps identify where services should be provided to and allows evaluation of the demand impacts of changing service patterns.

The local rail network in South-East Wales was chosen as the case study area for the development of flow level models, primarily because it is one of the very few large urban areas in the UK where a travelcard ticketing scheme does not operate. Where such schemes do exist LENNON electronic ticket sales data does not accurately reflect the total number of trips made, as it is impossible to establish what trips were made using travelcards without specially commissioned OD surveys. This is a major problem given that LENNON data is the best source of data on the size of rail passenger flows. While several other local rail networks where no travelcard scheme operates were identified, South-East Wales appeared to be the most suitable area for this work (see Section 3.8). Arriva Trains Wales (the local train operator) supplied LENNON data for the 2007 fiscal year of settlement for all flows to and from all 85 Network Rail category E (small staffed) and F (small unstaffed) stations in the case study area. While Network Rail's allocation of stations to these categories is questionable in some cases, this categorisation still provided the most straightforward way to define local stations. The stations for which data was obtained are mapped in Figure 5.1.

Figure 5.1: Stations included in South Wales case study

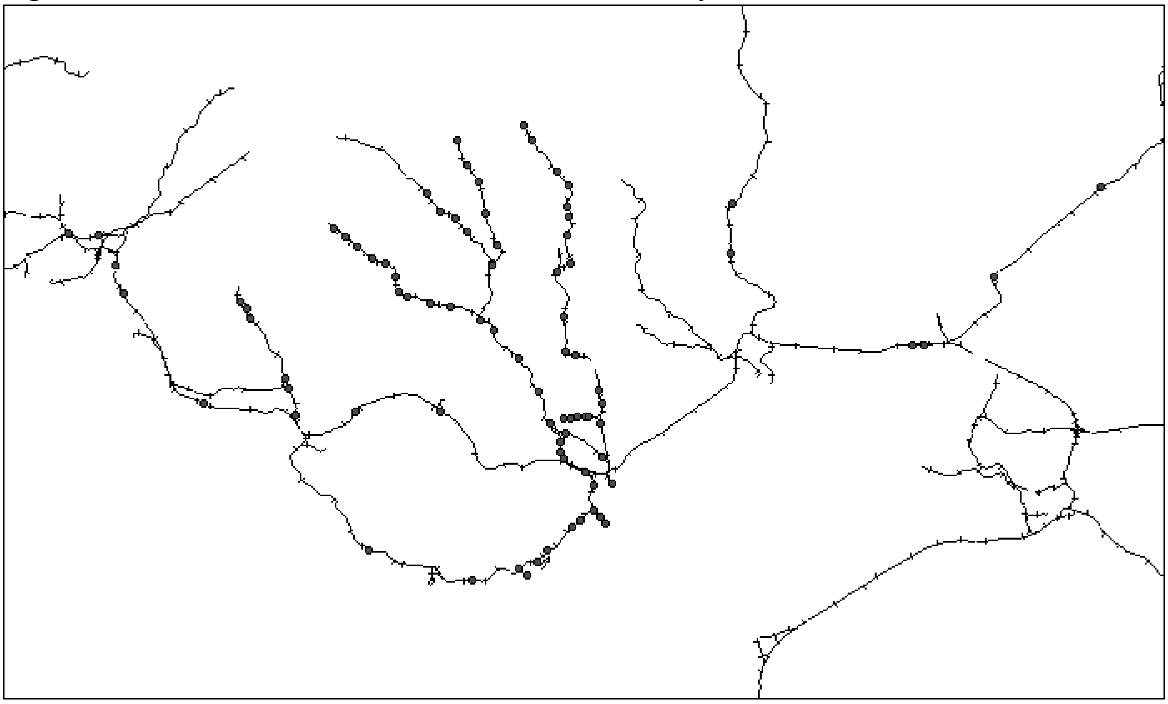

The LENNON data was supplied in a form which gave the total number of trips from each station in the dataset to and from every origin and destination to or from which trips were made, giving a total of 41,089 flows. This included flows to a number of origins and destinations which were not railway stations, for example ‘Tredegar Bus’ and ‘Belfast NI-M190’. These flows tended to be small, and as such origins and destinations are likely to have very different characteristics to railway stations they were removed from the dataset. As supplied the trips in the dataset were disaggregated by ticketing route, but for this analysis these were aggregated together to give the total number of trips between each station pair regardless of route. Flows to London posed a problem, as flows could exist from the same station to ‘London Terminals’, specific London termini and various combinations of London Underground zones. These destinations were combined to give a single flow of trips from each origin to ‘Central London’. Rail zones (destinations such as ‘Zone 1256 London’) were not included in these composite flows, as they are likely to represent travel to a specific non-terminal station in the London area.

Flows to station groups were also problematic, as it was obvious that it would not be possible to calculate journey times or frequencies to such groups. In most cases though it was apparent that one particular station in the group would almost certainly be used for all trips in that flow, and this station was thus set as the destination. Where trips were recorded to this individual station in the group as well as to the station group these trips were aggregated into a single flow. Some flows were recorded as being from a station to the same station and all such flows were removed on the grounds of realism, along with all flows where zero trips were recorded. Finally, flows between stations which were both in the case study area appeared twice, and this duplication was also removed to avoid bias in model calibration, giving a dataset of 28,071 usable flows.

Data on train frequencies and rail journey times were obtained by using Perl scripts developed by Dr John Armstrong (Armstrong et al., 2007) which interrogated Common Interface Format (CIF) timetable data files to produce mean journey times and service frequencies for particular flows. The scripts were only capable of calculating train frequencies and journey times for flows with direct services, which reduced the size of the dataset to 2,439 flows. This greatly reduced the model degrees of freedom, but the 2,439 flows selected contained 88.4% of the total trips in the dataset, so the majority of travel from the case study stations was still considered during model calibration. Spatial, demographic and socio-economic data and data on station car parks were obtained from the same sources as for the trip end models.

5.2.2 Basic models

Initial models calibrated on the dataset of 2,439 flows aimed to forecast flow level demand as a function of rail journey time, rail service frequency, the generating potential of the origin and the attractiveness of the destination. Three different representations of the origin and destination were tested, with the first model (5.1) based on the simplest model form used by Whelan & Wardman (1999b) and using dummy variables to represent demand at origin and destination stations. While this type of model would not be suitable for forecasting flows to/from new stations as no dummy variable parameter for these stations can be estimated from the calibration dataset, the model fit obtained should provide a benchmark for more transferable model forms. A total of 188 origin dummies and 181 destination dummies were required to represent the calibration flows in the model with the default origin and destination stations being set as Aber and Abercynon North respectively. The choice of default stations would affect the parameter values for the origin and destination dummy variables, but not the overall model fit, and was therefore not important. In Model 5.2 the dummy variables were replaced with continuous variables representing the total number of trips made from origin stations and to destination stations. This type of model is potentially more useful than Model 5.1, as it could be used to provide predictions for flows to/from new stations by using total trip predictions for such stations from trip end models. Values for the En_i and Ex_j variables were obtained from the ORR station usage data as complete LENNON data was not available for all stations within the dataset. Model 5.3 is similar to Model 5.2, but uses the total number of trip origins and destinations at the origin and destination stations, rather than only origins at the former and destinations at the latter. The results of calibrating Models 5.1-5.3 are summarised in Table 5.1.

$$\hat{T}_{IJ} = \alpha + \sum_{i}^{n} \beta_i O_i + \sum_{j}^{n} \gamma_j D_j + \delta R_{ij} + \eta F_{ij} \tag{5.1}$$

$$\hat{T}_{IJ} = \alpha + \beta En_i + \gamma Ex_j + \delta R_{ij} + \eta F_{ij} \tag{5.2}$$

$$\hat{T}_{IJ} = \alpha + \beta T_i + \gamma T_j + \delta R_{ij} + \eta F_{ij} \tag{5.3}$$

Where:

T_{IJ} is the predicted number of trips made from Station *I* to Station *J* in the 2007 fiscal year of settlement

O_i is a dummy variable which takes the value 1 if *i* is Station *I*, and 0 otherwise

D_j is a dummy variable which takes the value 1 if *j* is Station *J*, and 0 otherwise

R_{ij} is the average journey time for direct trains from Station *I* to Station *J*

F_{ij} is the number of direct trains from Station *I* to Station *J* on a normal weekday

En_i is the total number of trips originating at station *I* in 2006-7

Ex_j is the total number of trips terminating at station *J* in 2006-7

T_i is the total number of trips originating or terminating at station *I* in 2006-7

T_j is the total number of trips originating or terminating at station *J* in 2006-7

Table 5.1: Summarised parameter values and significance from calibration of Models 5.1-5.3 on all flows with direct services

	Model 5.1		Model 5.2		Model 5.3	
	Value	t stat	Value	t stat	Value	t stat
Intercept	-120.809	-0.025	-2603.589	-3.590	-2601.259	-3.587
β parameter	n/a	n/a	0.001	4.490	0.001	4.475
γ parameter	n/a	n/a	0.007	21.758	0.004	21.763
δ parameter	-27.447	-1.704	-13.656	-1.895	-13.726	-1.904
η parameter	220.732	7.508	167.263	9.440	167.000	9.425
R_{adj}^2	0.212		0.220		0.220	

The fit of Model 5.1 was poor, and while both the frequency (δ) and journey time (η) parameters were of the correct sign, the journey time parameter was not significant. It was expected that the inclusion of station dummy variables would account for most of the variation in the data, but few of these parameters were significant, possibly because of the large number of dummy variables relative to the size of the dataset. The results from

Models 5.2 and 5.3 were very similar, but while they gave a marginal improvement in fit over Model 5.1, model fit was still extremely poor and the rail journey time parameter was insignificant.

The poor fit of the models might partly result from the fact that a significant proportion of the flows in the calibration dataset were only served by a very low number of direct services, but that a much greater service frequency was possible by changing trains. The best way to resolve this problem would be to include such journey opportunities in the model, but a more immediate solution was to progressively exclude all flows where the number of direct services available was less than a set number, and then recalibrate the model. This gave the results summarised in Table 5.2 for Model 5.2. This shows that excluding all flows with less than ten direct services per day from the calibration dataset gave the best model fit, although the journey time (η) parameter was still insignificant and the model fit was still poor.

Table 5.2: Parameter values and significance from calibration of Model 5.2 with flows excluded based on minimum direct service level

Min direct services		2	3	4	5	6
Dataset size		2196	2087	2045	2011	1987
Intercept	Value	-3205.735	-3743.404	-4228.060	-4279.247	-4492.791
	t stat	-3.781	-4.117	-4.597	-4.561	-4.751
β parameter	Value	0.001	0.001	0.001	0.001	0.001
	t stat	3.958	4.003	4.260	4.248	4.369
γ parameter	Value	0.004	0.004	0.004	0.004	0.005
	t stat	19.345	19.532	20.781	20.653	21.393
δ parameter	Value	-14.211	-12.710	-7.698	-9.341	-7.366
	t stat	-1.612	-1.329	-0.766	-0.890	-0.679
η parameter	Value	189.569	196.113	193.796	195.717	193.473
	t stat	9.489	9.296	9.143	9.093	8.960
R_{adj}^2		0.203	0.210	0.228	0.228	0.239

Min direct services		7	8	9	10
Dataset size		1959	1931	1909	1877
Intercept	Value	-4376.607	-4042.602	-3741.005	-3716.831
	t stat	-4.502	-3.995	-3.554	-3.417
β parameter	Value	0.001	0.001	0.001	0.001
	t stat	4.37	4.408	4.452	4.464
γ parameter	Value	0.005	0.005	0.005	0.005
	t stat	21.317	21.228	21.165	21.483
δ parameter	Value	-11.452	-21.202	-30.704	-29.606
	t stat	-0.961	-1.516	-1.942	-1.832
η parameter	Value	192.692	190.307	188.468	183.853
	t stat	8.815	8.609	8.428	8.041
R_{adj}^2		0.240	0.240	0.240	0.247

While the inclusion of flows with few direct services may have caused some problems, the main reason for the poor model fit seemed likely to be an inappropriate model form. The linear form used so far does not allow for any interrelationship between the variables, when in fact the number of flows between two stations would be better represented as a fraction of the overall demand at the origin, multiplied by variables representing the attraction factor of the destination and the quality of the service between them. Using dummy variables or indeed continuous variables for the total number of trips at the origin and destination stations should then provide an effective constraint on the trips predicted for flows to and from those stations. A multiplicative form is also more consistent with the original derivation of the direct demand model from demand theory than an additive form. Using a multiplicative form gave the dummy variable Model 5.4 and the continuous origin and destination variable Model 5.5.

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} R_{ij}^{\delta} F_{ij}^{\eta} \tag{5.4}$$

$$\hat{T}_{IJ} = \alpha\, En_i^{\beta}\, Ex_j^{\gamma} R_{ij}^{\delta} F_{ij}^{\eta} \tag{5.5}$$

These models could still be calibrated using linear regression by taking natural logarithms of each side of the equations, as shown below. The results of calibrating these models on both the dataset of 1,877 flows with ten or more direct services per day and on the dataset of 2,439 flows with one or more direct services per day are summarised in Table 5.3.

$$\ln \hat{T}_{IJ} = \ln\alpha + \sum_{i}^{n} \beta_i \ln O_i + \sum_{j}^{n} \gamma_j \ln D_j + \delta \ln R_{ij} + \eta \ln F_{ij} \tag{5.4}$$

$$\ln \hat{T}_{IJ} = \ln\alpha + \beta \ln En_i + \gamma \ln Ex_j + \delta \ln R_{ij} + \eta \ln F_{ij} \tag{5.5}$$

Table 5.3: Summarised parameter values and significance from calibration of Models 5.4-5.5

Model		5.4		5.5	
Min direct services		1	10	1	10
Intercept	Value	8.741	5.015	-8.448	-9.247
	t stat	20.980	7.394	-26.228	-24.295
β parameter	Value	n/a	n/a	0.686	0.766
	t stat	n/a	n/a	36.287	34.732
γ parameter	Value	n/a	n/a	0.859	0.968
	t stat	n/a	n/a	43.645	42.088
δ parameter	Value	-1.263	-1.066	-1.321	-1.343
	t stat	-34.697	-24.447	-36.992	-32.407
η parameter	Value	0.335	1.245	0.347	-0.036
	t stat	8.641	9.035	11.376	-0.482
R_{adj}^2		0.787	0.749	0.689	0.648

Using this multiplicative form appears to give a huge improvement fit over the previous additive models, but the R_{adj}^2 values are not directly comparable because the variances of T_{ij} and LnT_{ij} are different. A Box-Cox test was therefore carried out (as in Section 4.3) which confirmed that the multiplicative form gave a much better fit, with the residual sum of squares based on the transformed independent variables much lower for these models and the χ^2 statistic highly significant. This was reassuring given the findings of previous studies of inter-urban demand, such as that by Wardman et al. (2007), where a similar dummy variable model achieved an R_2^{adj} value of 0.932. The inferior model fit obtained here is likely to reflect an inferior representation of the generalised cost of travel, as for example information on the monetary cost of travel between station pairs was not available for inclusion in these models. A better model fit was also obtained when all flows with direct services were included in the calibration dataset than when only flows with ten or more direct services were included, although this may largely result from the additional model degrees of freedom. Including the flows with only a few direct services seemed likely to distort model results, as the direct frequency data did not give an accurate representation of the rail service provided for such flows, and therefore the smaller dataset was used for further model developments.

A translog model form (5.6) was also tested on this smaller dataset (as in Section 4.3.1), and this gave the results summarised in Table 5.4. As with the trip end models, this form had several insignificant variables, so the model was recalibrated using backward stepwise calibration to establish the optimal model form. This was found to be given by Model 5.7, the results from which are also summarised in Table 5.4. While the results from Model 5.7 were more promising, the train frequency and rail time parameters were of the wrong sign, indicating that rail demand increases as train frequency reduces and journey time increases. The values of the En_iR_{ij} and Ex_iR_{ij} parameters were also counterintuitive, suggesting that the negative impact of rail journey time was magnified at trips to and from larger stations, when the reverse would be expected to be the case. Model 5.4 was therefore retained as the 'benchmark' dummy variable model, with Model 5.5 taken forward as the initial preferred model for use with new stations.

$$\hat{T}_{IJ} = \alpha + \beta \ln En_i + \gamma \ln Ex_j + \delta \ln R_{ij} + \eta \ln F_{ij} + \sum_N \kappa_N \ln^2 X_N + \sum_{MN} \upsilon_{MN} \ln X_M \ln X_N \quad (5.6)$$

$$\begin{aligned} \hat{T}_{IJ} &= \alpha + \beta \ln En_i + \gamma \ln Ex_j + \delta \ln R_{ij} + \eta \ln F_{ij} + \kappa_{En} \ln^2 En_i + \kappa_{Ex} \ln^2 Ex_i + \ldots \\ &\upsilon_{EnEx} \ln En_i \ln Ex_j + \upsilon_{EnR} \ln En_i \ln R_{ij} + \upsilon_{ExR} \ln Ex_i \ln R_{ij} + \upsilon_{ExF} \ln Ex_i \ln F_{ij} + \ldots \quad (5.7) \\ &\ldots \upsilon_{RF} \ln R_{ij} \ln F_{ij} \end{aligned}$$

Where:

X_M and X_N denote the set of independent variables included in the model

Table 5.4: Summarised results from calibration of Models 5.6-5.7

Parameter	Value	t stat	Value	t stat
Intercept	565818.355	17.302	569505.928	19.313
β parameter	-30790.996	-10.316	-29893.396	-10.185
γ parameter	-68696.082	-22.466	-68213.695	-22.918
δ parameter	12658.116	2.613	11560.430	2.671
η parameter	-40153.086	-3.784	-46064.790	-7.631
κ_{En} parameter	323.826	3.150	373.329	4.013
κ_{Ex} parameter	1654.313	15.427	1688.080	16.078
κ_R parameter	-129.178	-0.287	n/a	n/a
κ_F parameter	-2630.963	-1.743	n/a	n/a
υ_{EnEx} parameter	2388.780	13.860	2461.485	15.788
υ_{EnR} parameter	-823.675	-2.700	-981.034	-3.919
υ_{EnF} parameter	740.998	1.319	n/a	n/a
υ_{ExR} parameter	-881.118	-2.693	-988.218	-3.558
υ_{ExF} parameter	3772.201	6.457	3228.564	6.445
υ_{RF} parameter	1677.759	1.565	2690.93	3.628
R_{adj}^2	0.393		0.393	

It would be desirable to replace the station dummy and total trip variables in Models 5.4 and 5.5 with a set of variables which represent specific characteristics of stations and their catchment areas. The origin variable was therefore replaced with several variables from the trip end models described in Chapter 4, giving Model 5.8. The weighted population term used was that found to be most effective in the trip end models for England and Wales. The distance to the nearest category A-D station variable could not be included in this model because some flows in the calibration dataset originate at stations in categories A-D, and the total train frequency variable was obviously not suitable for a flow level model. The results of calibrating Model 5.8 on the 1,877 flow dataset are summarised in Table 5.5.

$$\hat{T}_{ij} = \alpha \left(\sum_a P_a w_a \right)^{\beta} J_{i4}^{\tau} Pk_i^{\rho} \prod_j^n D_j^{\gamma_j} R_{ij}^{\delta} F_{ij}^{\eta} \tag{5.8}$$

Where:

P_a is the population in output area a, for which station i is the closest station

$w_a = (d+1)^{-3.25}$

d is the travel time by road from output area a to station i

J_{i4} is the number of jobs located within four minutes drive of station i

Pk_i is the number of parking spaces at station i

Table 5.5: Summarised parameter values and significance from calibration of Model 5.8 on all flows with more than ten direct services per day

	Value	t stat
Intercept	1.863	3.725
β parameter	0.185	4.300
τ parameter	0.534	16.714
ρ parameter	0.155	7.878
δ parameter	-1.064	-22.098
η parameter	0.646	5.815
R_{adj}^2	0.599	

While as expected this modification reduced model fit, Model 5.8 still captured more than half of the variation in the observed data, similar to the earlier direct demand models developed by Preston (1991b). This suggested that the three origin-specific variables were fairly effective at capturing the variations in trip generation rates. All parameters were also significant and of the expected sign. The fit of Model 5.8 is good enough to indicate that this general form may have potential for forecasting demand at new stations, and it was therefore retained for further consideration.

All models tested so far represent trip destinations using a set of dummy variables, and give no consideration to the actual characteristics of the destination station or the area around it. This is not unique to this study, as previous work on flow level models (for example Lythgoe (2004)) has tended to treat destinations in the same way. However, there is no reason why destinations should not be represented by continuous variables which describe their characteristics in the same way as origin stations. Model 5.8 was therefore modified with the destination dummy variables replaced by variables representing the population and number of jobs located around the destination station, giving Model 5.9. A significant correlation was found to exist between the destination population and destination jobs variables (Spearman's ρ = 0.384), and a model was also tested (Model 5.10) with the form of the latter variable altered to give the number of jobs per resident. Calibrating Models 5.9-5.10 on the dataset of 1,877 flows with ten or more direct services per day gave the results summarised in Table 5.6.

$$\hat{T}_{ij} = \alpha \left(\sum_a P_a w_a \right)^{\beta} J_{i4}^{\tau} Pk_i^{\rho} \left(\sum_b P_b w_b \right)^{\gamma} J_{j4}^{\chi} R_{ij}^{\delta} F_{ij}^{\eta} \qquad (5.9)$$

$$\hat{T}_{ij} = \alpha \left(\sum_a P_a w_a \right)^{\beta} J_{i4}^{\tau} Pk_i^{\rho} \left(\sum_b P_b w_b \right)^{\gamma} Jp_{j4}^{\chi} R_{ij}^{\delta} F_{ij}^{\eta} \qquad (5.10)$$

Where:

P_b is the population in output area b, for which station j is the closest station

$w_b = (d + 1)^{-3.25}$

d is the travel time by road from output area b to station i

J_{j4} is the number of jobs located within four minutes drive of station j

Jp_{j4} is the number of jobs located within four minutes drive of station j divided by the resident population with four minutes drive of station j

Table 5.6: Summarised results from calibration of Models 5.9-5.10

	Model 5.9		Model 5.10	
	Value	t stat	Value	t stat
Intercept	-4.712	-8.566	-2.086	-3.956
β parameter	0.143	3.088	0.141	3.041
τ parameter	0.397	11.587	0.395	11.581
ρ parameter	0.142	6.492	0.149	6.827
γ parameter	-0.089	-1.883	0.608	12.372
χ parameter	0.746	22.843	0.976	23.111
δ parameter	-0.818	-16.558	-0.794	-16.157
η parameter	1.129	13.103	1.112	12.925
R_{adj}^2	0.436		0.438	

The fit of both models is worse than that of Model 5.8, although this was expected given the replacement of the destination dummy variables. While Model 5.10 only has a marginally better fit than Model 5.9, the former model should be preferred because all parameters are significant and of the correct sign. However, both these models explain less than half of the variation in the data, and thus can not be used with confidence to predict flow sizes from new stations.

5.2.3 Intermodal competition

The extent of intermodal competition on particular flows will almost certainly affect rail demand, and a road journey time variable was therefore added to the best model developed so far (5.4), giving Model 5.11. RouteFinder for MapInfo was used to calculate road journey times between station pairs, and the difference between this and the rail journey time was incorporated in the model as a proportion. A proportional measure was used because there was a strong correlation between rail and road journey times, and using the absolute time difference would have made logarithmic calibration impossible because some of the time differences were negative. The road journey times were calculated using the same set of road speeds used for catchment definition in Models 5.8-5.10 and in the trip end models developed previously (see Table 4.4).

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} R_{ij}^{\delta} F_{ij}^{\eta} Cp_{ij}^{\kappa} \tag{5.11}$$

Where:

Cp_{ij} is the car journey time from station i to station j divided by the rail journey time from station i to station j

To allow detailed investigation of intermodal competition, analysis focused on flows to the main stations in Cardiff (Queen Street and Central), Newport and Swansea. While this reduced the size of the dataset and therefore the model degrees of freedom, it made it feasible to manually interrogate online journey planners to collect the required data on bus journey times and fares and rail fares for later models. Flows to Central and Queen Street stations in Cardiff were kept separate rather than being aggregated together, because both stations form major trip attractors, and this increased the number of flows which could be included in the model, giving a calibration dataset of 174 flows with direct services. The results of calibrating Model 5.11 on this dataset are summarised in Table 5.7 along with details of the recalibration of Models 5.4, 5.8 and 5.10 on this smaller dataset to provide a benchmark level of model fit

Table 5.7: Summarised results from calibration of Models 5.4, 5.8, 5.10 and 5.11 on all flows with direct services to Cardiff Central, Cardiff Queen Street, Newport and Swansea

Model		5.4	5.8	5.1	5.11
Intercept	Value	9.497	3.530	-13.400	14.574
	t stat	3.573	2.171	-1.339	5.266
β parameter	Value	n/a	0.198	0.186	n/a
	t stat	n/a	1.488	1.385	n/a
τ parameter	Value	n/a	0.192	0.148	n/a
	t stat	n/a	1.748	1.347	n/a
ρ parameter	Value	n/a	0.338	0.354	n/a
	t stat	n/a	5.092	5.158	n/a
γ parameter	Value	n/a	n/a	2.572	n/a
	t stat	n/a	n/a	1.741	n/a
χ parameter	Value	n/a	n/a	1.930	n/a
	t stat	n/a	n/a	2.245	n/a
δ parameter	Value	-1.196	-0.573	-0.580	-1.950
	t stat	-3.208	-2.743	-2.736	-4.960
η parameter	Value	0.888	1.002	1.176	0.279
	t stat	1.998	4.439	5.813	0.636
κ parameter	Value	n/a	n/a	n/a	-3.159
	t stat	n/a	n/a	n/a	-3.961
R_{adj}^2		0.776	0.510	0.506	0.809

The fit of Models 5.4 and 5.10 was better when calibrated on this small dataset than when calibrated on all direct flows, but the fit of Model 5.8 was worse, and limited confidence can be placed in the values of the two destination-specific variables (γ and χ) in Model

5.10 because only four destinations were included in the calibration dataset. Model 5.11 gave a better fit than any of the earlier flow level models, although the frequency parameter (η) was insignificant.

Several model formulations which gave a more complete representation of intermodal competition were tested on the 174 flow dataset. Models 5.12-5.13 incorporated bus journey time as an absolute figure, whereas Models 5.14-5.15 included it as a value relative to rail journey time. Models 5.13 and 5.15 included the relative car journey time variable in addition to the bus journey time and rail fare variables. Data on rail fares are readily available online, albeit in a disaggregate form, and in this case were obtained from the National Express East Coast ticket sales website (www.nationalexpresseastcoast.com) as this returned results faster than comparable sites. The standard day return fare was used for nearly all flows, with the saver return fare used in the few cases where no standard day return was available. Data on bus journey times was obtained from the Transport Direct website (www.transportdirect.info), with the times retrieved being for a journey by bus from the origin railway station to the destination station. However, the accuracy of some of this data appeared questionable, and it is unlikely that in reality people would use railway stations as the starting point for their bus journeys. Many of the bus journeys involved interchange and as it was not obvious whether the time spent in interchanging should be treated the same as travel time, or as double the equivalent travel time as suggested in the PDFH (ATOC, 2002), both methods were tested in calibration. Access and egress walk times at each end of the journey were not included in the bus journey time, as it was unlikely that passengers would actually be travelling to and from the railway stations when using the bus, and also because experience had shown that the times quoted on Transport Direct for this access and egress sometimes bear little relationship to reality. While it would have been desirable to include bus fares in the model to give a full representation of bus generalised cost, most operators do not give detailed information on fares, making it impossible to collect sufficient data for calibration. The cost of car travel can be assumed to be directly proportional to journey time, removing the need for a car cost variable. The results of calibrating Models 5.12-5.15 are summarised in Table 5.8.

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} R_{ij}^{\delta} F_{ij}^{\eta} Rf_{ij}^{\lambda} Bt_{ij}^{\nu} \qquad (5.12)$$

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} R_{ij}^{\delta} F_{ij}^{\eta} Cp_{ij}^{\kappa} Rf_{ij}^{\lambda} Bt_{ij}^{\nu} \qquad (5.13)$$

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} R_{ij}^{\delta} F_{ij}^{\eta} Rf_{ij}^{\lambda} Bp_{ij}^{\nu} \tag{5.14}$$

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} R_{ij}^{\delta} F_{ij}^{\eta} Cp_{ij}^{\kappa} Rf_{ij}^{\lambda} Bp_{ij}^{\nu} \tag{5.15}$$

Where:

Rf_{ij} is the rail fare from station *i* to station *j*

Bt_{ij} is the journey time by bus from station *i* to station *j*

Bp_{ij} is the journey time by bus from station *i* to station *j* divided by the mean journey time by rail from station *i* to station *j*

Table 5.8: Summarised results from calibration of Models 5.12-5.15

		Actual interchange time				Double interchange time			
Model		5.12	5.13	5.14	5.15	5.11	5.13	5.13	5.15
Intercept	Value	10.880	5.835	10.880	5.835	11.021	5.564	11.021	5.564
	t stat	5.004	2.015	5.004	2.015	5.144	1.965	5.144	1.965
δ parameter	Value	2.299	4.810	1.019	3.159	2.468	5.218	1.053	3.400
	t stat	3.934	4.221	1.890	3.186	4.236	4.624	2.028	3.491
η parameter	Value	0.719	1.201	0.719	1.201	0.647	1.146	0.647	1.146
	t stat	2.015	3.045	2.015	3.045	1.849	3.014	1.849	3.014
κ parameter	Value	n/a	3.130	n/a	3.130	n/a	3.397	n/a	3.397
	t stat	n/a	2.539	n/a	2.539	n/a	2.806	n/a	2.806
λ parameter	Value	-3.268	-5.340	-3.268	-5.340	-3.177	-5.423	-3.177	-5.423
	t stat	-4.781	-5.082	-4.781	-5.082	-4.757	-5.286	-4.757	-5.286
ν parameter	Value	-1.280	-1.651	-1.280	-1.651	-1.415	-1.817	-1.415	-1.817
	t stat	-2.642	-3.360	-2.642	-3.360	-3.119	-3.961	-3.119	-3.961
R_{adj}^2		0.859	0.867	0.859	0.867	0.863	0.874	0.863	0.874

The model fit was in all cases an improvement on previous models, and counting each minute of bus interchange or wait time as two minutes of travel time gave a better fit for all four models. However, while the rail journey time parameters (δ) were expected to be negative, so that as rail journey time increased demand would decrease, in all these models the parameters were positive, suggesting that demand increases with rail journey time, which seemed unlikely to reflect reality. This was confirmed by plotting demand against rail journey time in Figure 5.2, which while showing no clear relationship suggested a slight decrease in demand with journey time. A further problem was that both the absolute and the relative bus journey time parameters (ν) were negative, suggesting that as bus journey time and relative bus journey time increase, rail demand decreases. While the former could be plausible as this indicates a greater distance between the origin and the destination, the latter seemed unlikely, as a relatively higher bus journey time compared to the rail journey time would be expected to increase rail demand. While this might result from imperfections in the model form, it seemed likely that the poor quality of the bus service data was at least partly responsible.

Figure 5.2: Relationship between rail journey time and flow size

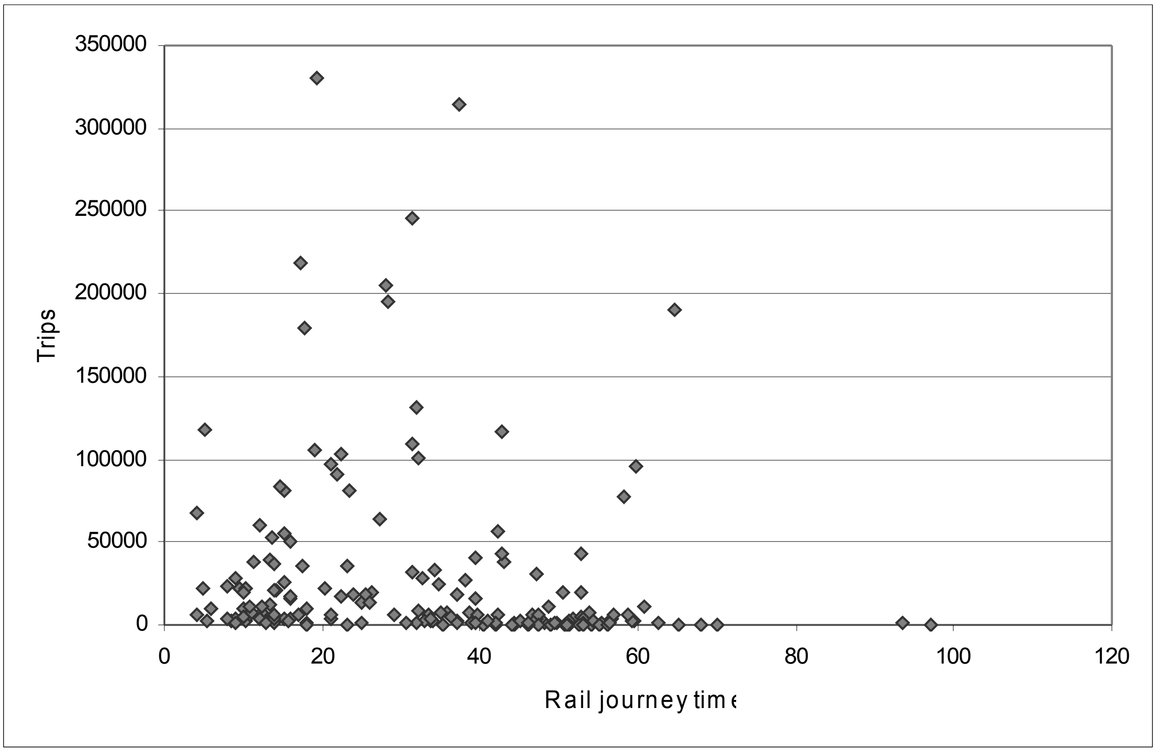

An alternative representation of intermodal competition was tested in Model 5.16, which included a non-mode specific variable for straight-line distance between the station pairs, and represented the three competing modes with speed variables calculated by dividing the mode-specific travel time by either this straight-line distance or by a mode-specific distance (road distance for bus and car, rail distance for train). Such a modification should in theory remove any non-random correlation between the variables and therefore improve model performance. The results of calibrating Model 5.16 (with bus interchange time still counted double) are summarised in Table 5.9.

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} D_{ij}^{\omega} Rs_{ij}^{\delta} Cs_{ij}^{\kappa} Bs_{ij}^{\nu} F_{ij}^{\eta} Rf_{ij}^{\lambda} \qquad (5.16)$$

Where:

D_{ij} is the straight line distance (in km) from station *i* to station *j*

Rs_{ij} is the average rail speed for journeys from station *i* to station *j*

Cs_{ij} is the average car speed for travel from station *i* to station *j*

Bs_{ij} is the average bus speed for travel from station *i* to station *j*

Table 5.9: Summarised results from calibration of Model 5.16

Speed calculation	Straight line distance		Mode-specific distance	
	Value	t stat	Value	t stat
Intercept	5.873	1.925	7.964	3.455
ω parameter	3.433	3.480	2.769	3.041
δ parameter	-1.781	-2.849	-2.939	-3.605
κ parameter	-2.824	-1.195	-1.232	-1.163
ν parameter	1.795	3.833	2.014	4.319
η parameter	1.172	2.980	1.144	3.017
λ parameter	-5.496	-5.167	-4.862	-4.747
R_{adj}^2	0.872		0.881	

Model 5.16 slightly improved model fit, in particular when the speed calculations were based on mode-specific distances, but in other respects the expected improvements did not materialise. While the car speed parameter (κ) was of the expected sign in both models it was also insignificant, and the bus speed (ν), rail speed (δ) and straight line distance (ω) parameters were all of the wrong sign, indicating that there were still problems of collinearity. A check was therefore made for correlations between the variables, giving the results in Table 5.10.

Table 5.10: Spearman's ρ correlation coefficients for relationship between modal speeds and other model variables

Speeds based on mode-specific distance						
	Ln rail speed		Ln road speed		Ln bus speed	
Variable	ρ coefficient	sig.	ρ coefficient	sig.	ρ coefficient	sig.
Ln rail speed	n/a	n/a	-0.420	0.000	0.498	0.000
Ln road speed	-0.420	0.000	n/a	n/a	-0.321	0.000
Ln bus speed	0.498	0.000	-0.321	0.000	n/a	n/a
Ln straight line distance	0.659	0.000	-0.403	0.000	0.355	0.000
Ln rail frequency	-0.594	0.000	0.228	0.000	-0.488	0.000
Ln rail fare	0.651	0.000	-0.430	0.000	0.342	0.000
Ln origin population	0.001	0.985	0.154	0.043	-0.083	0.274
Ln origin jobs	0.063	0.406	0.035	0.649	0.002	0.979
Ln origin parking	0.182	0.016	-0.087	0.252	-0.044	0.563
Speeds based on straight line distance						
	Ln rail speed		Ln road speed		Ln bus speed	
Variable	ρ coefficient	sig.	ρ coefficient	sig.	ρ coefficient	sig.
Ln rail speed	n/a	n/a	0.479	0.000	0.420	0.000
Ln road speed	0.479	0.000	n/a	n/a	0.424	0.000
Ln bus speed	0.420	0.000	0.424	0.000	n/a	n/a
Ln straight line distance	0.633	0.000	0.522	0.000	0.424	0.000
Ln rail frequency	-0.469	0.000	-0.224	0.000	-0.460	0.000
Ln rail fare	0.636	0.000	0.526	0.000	0.402	0.000
Ln origin population	0.005	0.945	-0.082	0.280	-0.066	0.390
Ln origin jobs	0.060	0.434	-0.031	0.687	-0.027	0.720
Ln origin parking	0.178	0.019	0.065	0.392	-0.046	0.545

Table 5.10 shows that there were reasonably strong correlations between the speed variables, but that stronger correlations existed between rail speed and straight line distance, rail frequency and rail fares. Dividing rail fare by rail distance to give a rail fare

per km variable could improve model performance and Model 5.17 was therefore calibrated, giving the results summarised in Table 5.11, along with those from Model 5.18, where the origin dummy variables are replaced with origin characteristic variables. Consideration was given to dividing the fares by straight line distance instead of rail distance, given that some rail distances differed significantly from road distances, but this might have led to biases for flows where both rail and road routes are forced to divert around 'barriers' such as the hills dividing the South Wales valleys. The speed variables were calculated using mode-specific distances as these gave better results in Model 5.16. The models were also calibrated with the bus speed variable omitted because of the problems encountered with Model 5.16, giving Models 5.19 and 5.20, and the results from these models are also shown in Table 5.11.

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} D_{ij}^{\omega} Rs_{ij}^{\delta} Cs_{ij}^{\kappa} Bs_{ij}^{\nu} F_{ij}^{\eta} Rfkm_{ij}^{\lambda} \tag{5.17}$$

$$\hat{T}_{ij} = \alpha \left(\sum_{a} P_a w_a \right)^{\beta} J_{i4}^{\tau} Pk_i^{\rho} \prod_{j}^{n} D_j^{\gamma_j} D_{ij}^{\omega} Rs_{ij}^{\delta} Cs_{ij}^{\kappa} Bs_{ij}^{\nu} F_{ij}^{\eta} Rfkm_{ij}^{\lambda} \tag{5.18}$$

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} D_{ij}^{\omega} Rs_{ij}^{\delta} Cs_{ij}^{\kappa} F_{ij}^{\eta} Rfkm_{ij}^{\lambda} \tag{5.19}$$

$$\hat{T}_{ij} = \alpha \left(\sum_{a} P_a w_a \right)^{\beta} J_{i4}^{\tau} Pk_i^{\rho} \prod_{j}^{n} D_j^{\gamma_j} D_{ij}^{\omega} Rs_{ij}^{\delta} Cs_{ij}^{\kappa} F_{ij}^{\eta} Rfkm_{ij}^{\lambda} \tag{5.20}$$

Where:

$Rfkm_{ij}$ is the rail fare per rail km for travel from station i to station j

Table 5.11: Summarised results from calibration of Models 5.17-5.20

Model	5.17		5.18		5.19		5.2	
	Value	t stat	Value	t stat	Value	t stat	Value	t stat
Intercept	10.129	4.409	5.266	3.355	10.894	4.272	5.894	3.731
β parameter	n/a	n/a	0.220	1.711	n/a	n/a	0.231	1.811
τ parameter	n/a	n/a	0.219	2.075	n/a	n/a	0.202	1.920
ρ parameter	n/a	n/a	0.284	4.254	n/a	n/a	0.319	4.716
ω parameter	-2.124	-4.858	-2.362	-5.605	-2.748	-5.956	-2.360	-5.715
δ parameter	-2.554	-2.842	1.231	1.982	-0.888	-0.974	1.054	1.718
κ parameter	-1.609	-1.448	-1.445	-1.953	-2.070	-1.681	-1.930	-2.669
ν parameter	2.191	4.510	-0.010	-0.026	n/a	n/a	n/a	n/a
η parameter	0.868	2.255	1.218	5.230	0.616	1.453	1.114	4.769
λ parameter	-3.511	-3.646	-3.192	-4.522	-3.762	-3.515	-3.089	-4.490
R_{adj}^2	0.869		0.565		0.838		0.575	

These modifications generally improved model performance although model fit was slightly worse than for Model 5.16. The straight line distance (ω) and fare (λ) parameters were both significant and of the correct sign in all models, although the implied fare elasticities seemed extremely high, and the car speed parameter (κ) was also of the correct sign although not significant. However, the rail (δ) and bus speed (ν) parameters were still both of the wrong sign in Model 5.17, and while they were of the correct sign in Model 5.18 they were both insignificant. Even when the bus speed variable was removed in Model 5.19 the rail speed variable was still of the wrong sign (although it was now insignificant) and the fare elasticity was still high. However, all variables in Model 5.20 were of the correct sign, and model fit was slightly superior to that of Model 5.18. While the origin population (β), origin employment (τ) and rail speed (ρ) variables were only of marginal significance, Model 5.20 was adopted as the preferred model form using generalised origin variables, replacing Model 5.8.

The continued problems with parameter signs and significance in the dummy variable models made it necessary to test a further means of representing intermodal competition, by combining journey times for all three modes in a single total journey time index, giving Model 5.21. As incorporating bus journey times had previously given unexpected results, a total journey time index which only included rail and car journey times was also tested, giving Model 5.22. The form of the journey time index variable in these models may be incorrect, as it was expressed as a continuous variable when in fact it could only take values between 0 and 1, but there was no straightforward way around this problem. The results of calibrating Models 5.21-5.22 are summarised in Table 5.12.

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} D_{ij}^{\omega} JTI_{ijRCB}^{\varphi} F_{ij}^{\eta} Rfkm_{ij}^{\lambda} \qquad (5.21)$$

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} D_{ij}^{\omega} JTI_{ijRC}^{\varphi} F_{ij}^{\eta} Rfkm_{ij}^{\lambda} \qquad (5.22)$$

Where:

JTI_{ij}^{RCB} is the three mode rail journey time index, given by $\dfrac{r_{tij}}{r_{tij} + c_{tij} + b_{tij}}$

JTI_{ij}^{RC} is the two mode rail journey time index, given by $\dfrac{r_{tij}}{r_{tij} + c_{tij}}$

r_{tij} is the rail journey time from station i to station j

c_{tij} is the car journey time from station i to station j

b_{tij} is the bus journey time from station i to station j

Table 5.12: Summarised results from calibration of Models 5.21-5.22

Model	5.21		5.22	
	Value	t stat	Value	t stat
Intercept	14.520	5.178	10.118	3.395
ω parameter	-2.089	-4.703	-2.728	-6.452
φ parameter	2.792	2.817	-0.473	-0.333
η parameter	0.671	1.641	0.700	1.618
λ parameter	-2.609	-2.383	-4.267	-3.747
R_{adj}^2	0.847		0.833	

The fit of both these models was comparable to that of Models 5.17 and 5.19, and the majority of the parameters were of the correct sign. The exception to this was the journey time index (φ) parameter for Model 5.21 which was positive and highly significant, when it was expected to be negative given that rail demand is likely to reduce when the rail journey time increases relative to the journey time by competing modes. This appeared to be a continuation of the problems encountered in previous attempts to incorporate bus journey times in the model, where similar counterintuitive results were obtained. While the value of the journey time index parameter for Model 5.22 was of the correct sign it was insignificant, which meant that this model form could not be entirely trusted.

The prediction errors from the model with the best fit so far (5.17) were mapped to allow any spatial patterns in model accuracy to be identified, despite the problems identified above with some of the model variables. Because only four destinations were included in the calibration of this model, it was possible to display all the errors in four maps, Figures 5.3 to 5.6. The figures mapped here are the prediction errors rather than the residuals, calculated by subtracting the predicted number of trips from the actual number of trips rather than by subtracting the ln of the predicted number of trips from the ln of the actual number of trips and taking the exponent of the result.

Figure 5.3: Prediction errors from Model 5.17 for flows to Cardiff Central

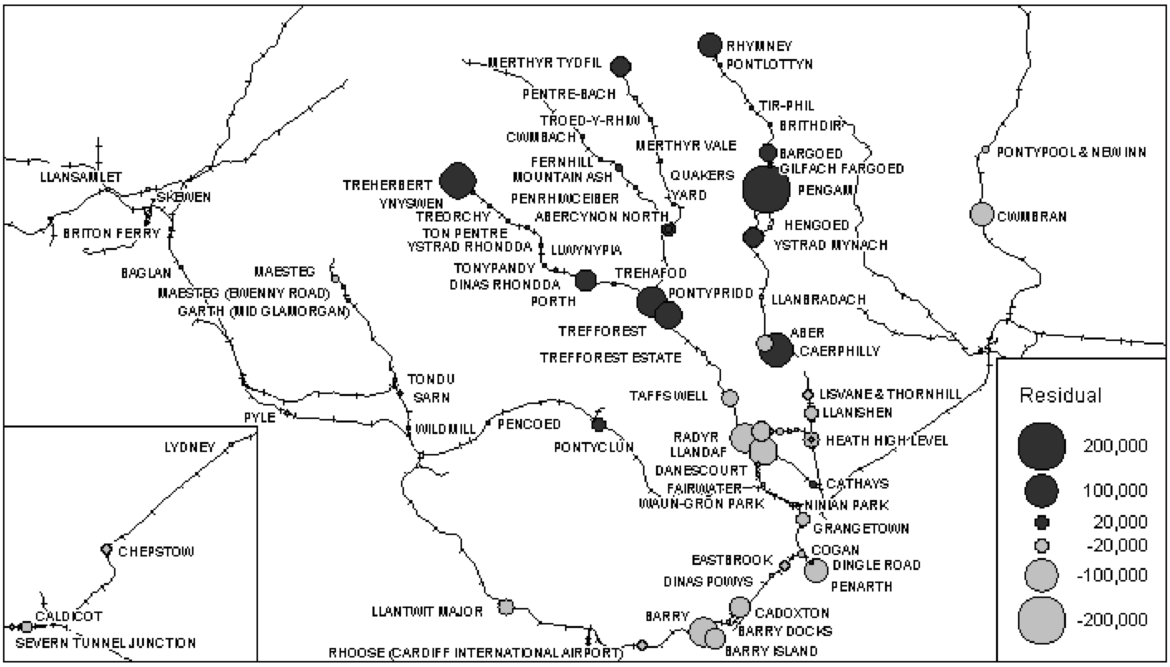

Figure 5.4: Prediction errors from Model 5.17 for flows to Cardiff Queen Street

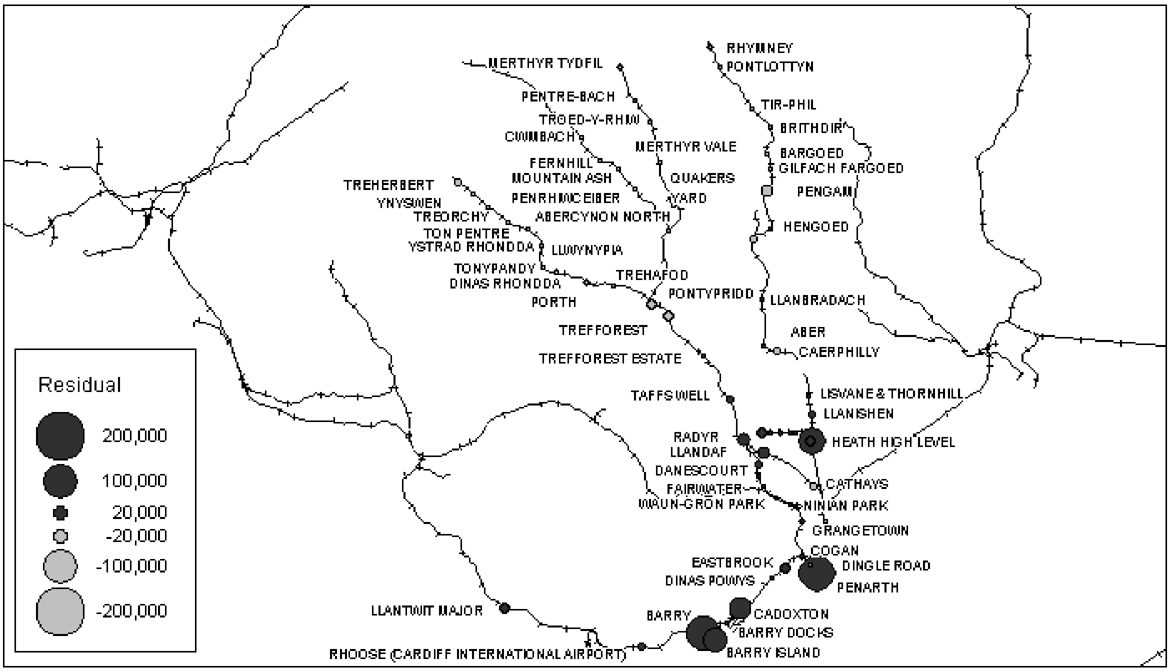

Figure 5.5: Prediction errors from Model 5.17 for flows to Newport

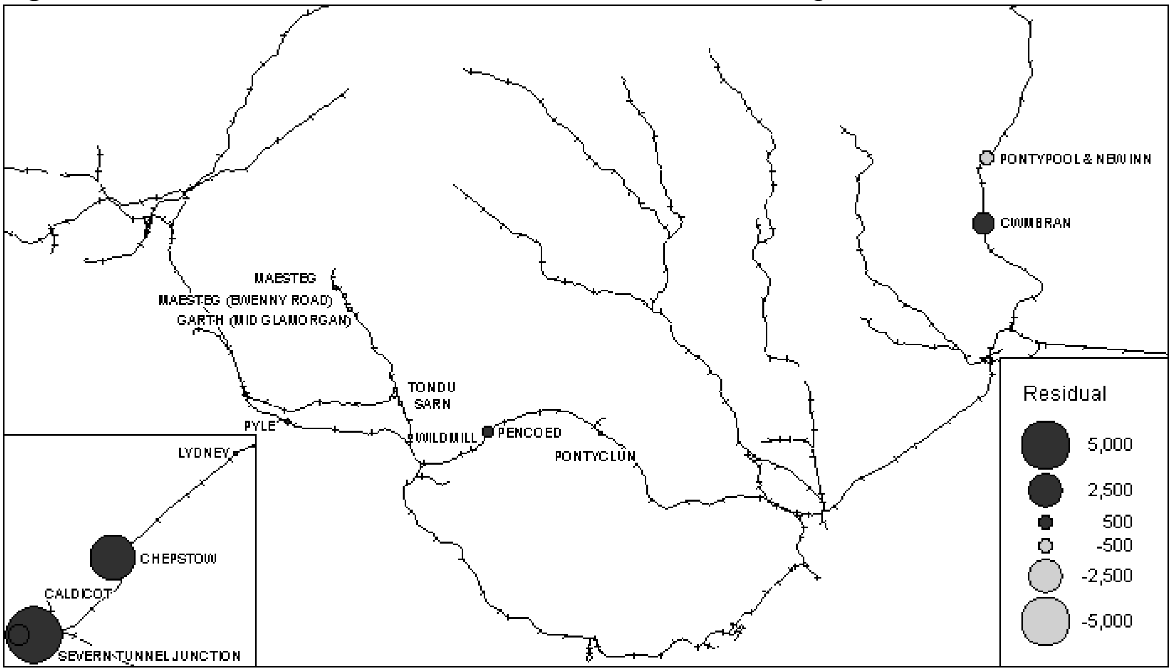

Figure 5.6: Prediction errors from Model 5.17 for flows to Swansea

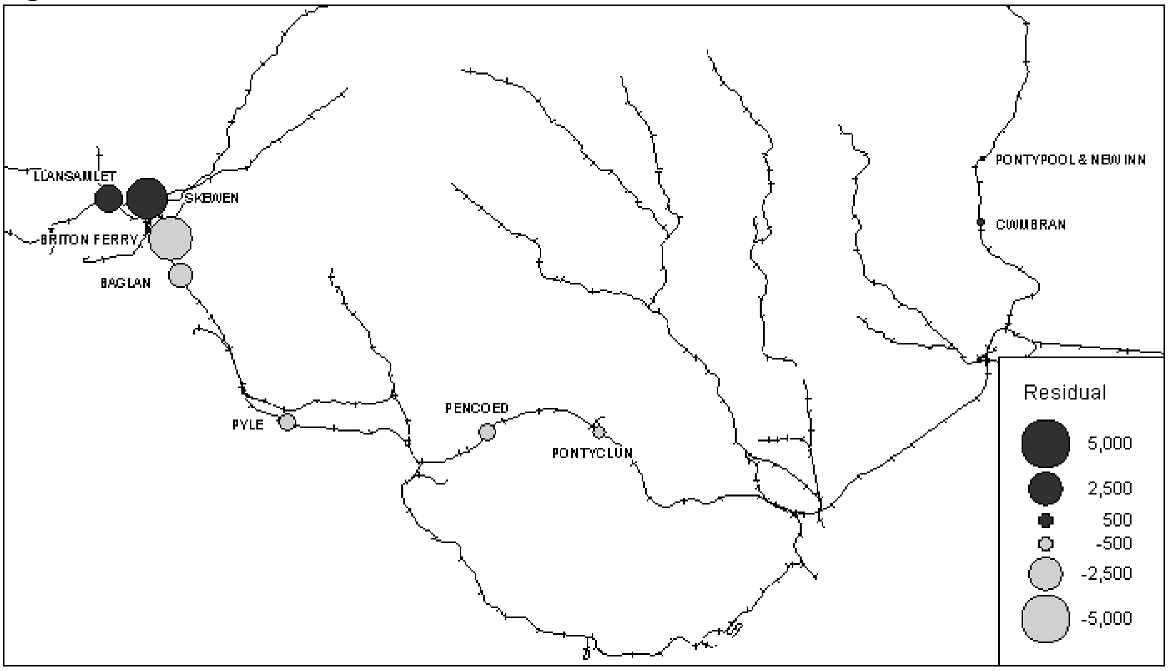

Figures 5.3-5.6 highlight some obvious spatial patterns in the prediction errors. There appears to be a 'north-south' divide in model accuracy for flows to the Cardiff stations, with the model underpredicting demand from Valley Line stations north of Lisvane & Thornhill and Taffs Well to Cardiff Central, and overpredicting demand from all other stations. The reverse seems to apply for flows to Cardiff Queen Street, where the model underpredicts demand south of this imaginary line, and overpredicts demand at stations further up the valleys. This may reflect the effect of intervening opportunities on rail demand, although in this case the 'boundary line' would be expected to fall between the

two major Cardiff stations, but regardless of their cause the resultant errors will have a detrimental effect on model accuracy. The prediction errors for flows to Newport and Swansea are much smaller (note the different scale used for the errors on the maps), and because less flows are involved it is difficult to identify spatial variation. Figure 5.6 indicates that demand is higher than expected on flows to Swansea from stations nearby, and lower than expected for flows from stations slightly further away. There was though no straightforward way to modify the models to account for these error patterns. Because later dummy variable models contained parameters which were insignificant or of the wrong sign, Model 5.4 was retained as the preferred dummy variable model form, although Model 5.20 replaced Model 5.8 as the preferred model with generalised origin variables.

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} R_{ij}^{\delta} F_{ij}^{\eta} \tag{5.4}$$

$$\hat{T}_{ij} = \alpha \left(\sum_{a} P_a w_a \right)^{\beta} J_{i4}^{\tau} Pk_i^{\rho} \prod_{j}^{n} D_j^{\gamma_j} D_{ij}^{\varpi} Rs_{ij}^{\delta} Cs_{ij}^{\kappa} F_{ij}^{\eta} Rfkm_{ij}^{\lambda} \tag{5.20}$$

5.2.4 Catchment definition

The methods used to define the population variable in Models 5.8, 5.18 and 5.20 were unsatisfactory for a direct demand model, because no consideration was given to direction of travel when allocating population units to stations. A more complex but perhaps more realistic method of defining station catchments was therefore tested, and was found to produce catchments which appeared visually more realistic, although several problems with implementation prevented these catchments being adopted as part of the recommended methodology.

This alternative catchment definition method was based on the premise that travellers would choose the railway station which minimised the total journey time to their destination station, rather than simply minimising their origin station access time. To define these flow-specific catchments road travel time was calculated from all population units (census output areas) in the case study area to their nearest four railway stations. While superior results might be obtained by using a larger number of competing stations for each output area, this would increase processing time, and four was felt to be a reasonable compromise. Rail journey times were then calculated from each relevant origin station to each of the destination stations being considered (in this case Cardiff Central, Cardiff Queen Street, Newport and Swansea). While the Perl scripts described previously could be used for station pairs where a direct service exists, because these scripts can not yet cope with interchanging journeys this data had to be supplemented with journey times based on information from the DB online journey planner. Flows with Pilning as an origin were given a journey time of 999 minutes as the absence of trains serving this station on weekdays means that it should not be allocated any population units. Road access and rail journey times were then combined to calculate total travel time from each population unit to each destination via each of its four possible origin stations. These stations were then ranked by total journey time for each output area and destination, and flow-specific catchments were defined by allocating each output area to the origin station which gave the shortest total journey time to the destination in question. Figure 5.7 shows the catchments defined using this method for flows to Cardiff Central.

While this map is difficult to interpret, some problems with the catchments are apparent, with a number of catchments including several stations, while others contain no stations at all. This is almost certainly the result of inaccurate scaling of the road access times to stations. These were based on the same set of road speeds used previously to allocate population units to stations for trip end models, which were found to give optimum results for such models. However, when defining catchments by minimising access times it was not necessary to combine the road access and rail journey times, so while they may be accurate relative to each other, it is likely that they were not accurate in absolute terms.

Figure 5.7: Destination-specific catchments for flows to Cardiff Central

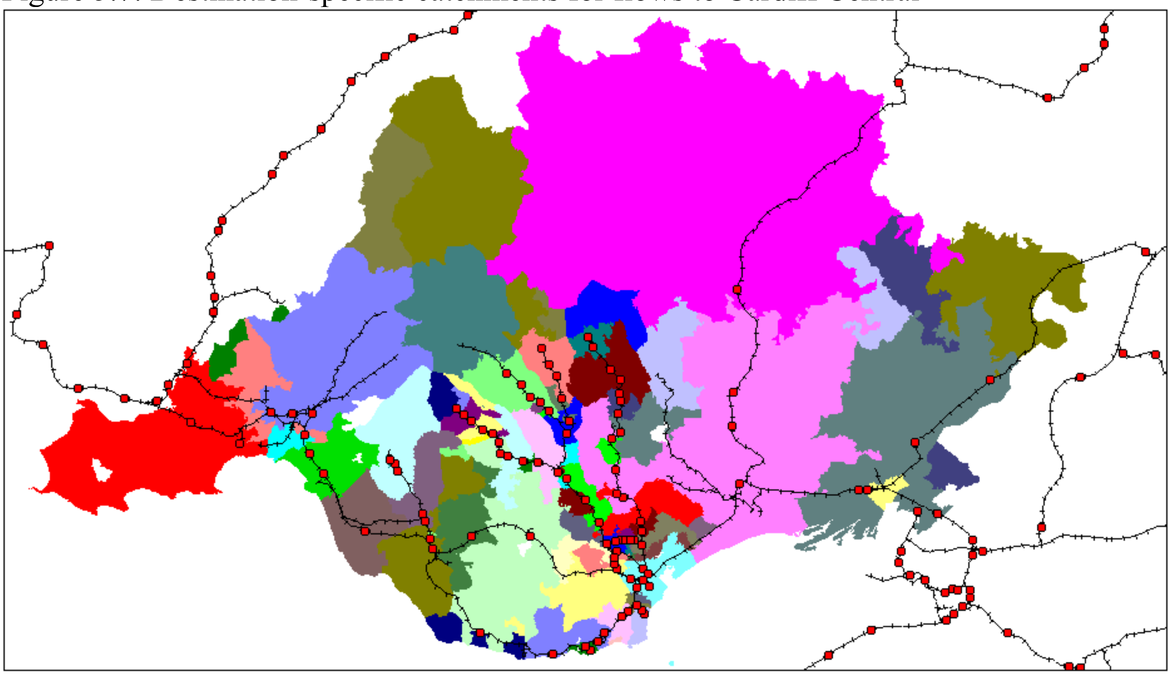

A scaling factor was therefore applied to these access times, and output areas were then reallocated to stations using the revised overall journey times. While ideally an automated procedure would have been used to determine the optimal scaling factor to apply, there was no obvious way to implement such a procedure, and therefore a process of trial and error was used. A scale factor of 1.5 was initially applied, which meant that the car journey times calculated by Routefinder were multiplied by 1.5, giving the effective average speeds shown in Table 5.13.

Table 5.13: Effective road speeds used by Routefinder to calculate scaled car journey times

Road type	Speed
Motorway	53.3 kph
A road	43.3 kph
B road	26.7 kph
Other road	16.7 kph

Reallocating output areas to stations based on the access times given by this scale factor gave some improvements but did not entirely solve the problem, with for example an overallocation of output areas to Newport station in preference to Pontypool & New Inn and Cwmbran. The differentiated speeds by road type were therefore abandoned, and the use of flat average speeds of 25 kph and 15 kph in combination with rail journey times to define catchments was tested. A flat speed of 15 kph gave the best results, and the resulting catchments are shown in Figure 5.8, but it is questionable whether such a low access speed would be mirrored in travel behaviour in the more rural parts of the case study area. This reflects a fundamental problem with this form of catchment definition, which is that catchments are likely to vary depending on access mode, but because access mode split will vary from station to station there is no straightforward way to incorporate this variation in either catchment definition or more generally in demand models.

The destination-specific catchment populations were used to replace the previous weighted catchment populations in Model 5.20, giving Model 5.23. Initially no jobs variable was included in this model, as it was not obvious how such a variable should be defined, and in any case variations in the number of jobs would be expected to have a greater effect on the number of trip destinations at a station than on the number of origins. The results from calibrating Model 5.23 with catchments defined using an average access speed of 15 kph are summarised in Table 5.14.

Figure 5.8: Destination-specific catchments for flows to Cardiff Central with flat average speed of 15 kph used to calculate road journey times

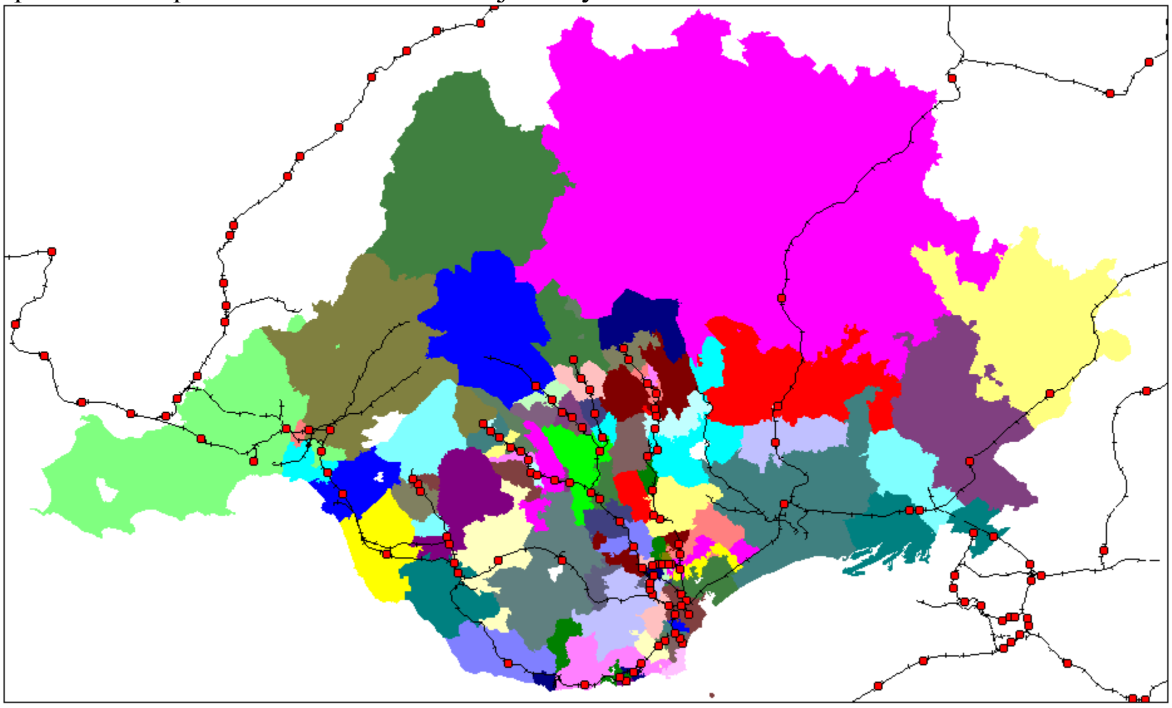

$$\hat{T}_{ij} = \alpha\, P_{iJ}^{\beta}\, Pk_i^{\rho} \prod_{j}^{n} D_j^{\gamma_j}\, D_{ij}^{\omega}\, Rs_{ij}^{\delta}\, Cs_{ij}^{\kappa}\, F_{ij}^{\eta}\, Rfkm_{ij}^{\lambda} \qquad (5.23)$$

Where:

P_{iJ} is the population for whom station i is the station which gives the shortest total journey time when travelling to station j

Table 5.14: Summarised results from calibration of Model 5.23

	Value	t stat
Intercept	6.005	4.305
β parameter	0.207	3.308
ρ parameter	0.281	4.266
ω parameter	-2.435	-5.926
δ parameter	0.998	1.646
κ parameter	-1.194	-1.666
η parameter	1.294	5.675
λ parameter	-3.250	-4.746
R_{adj}^2	0.573	

Despite the incorporation of flow-specific catchments, the fit of Model 5.23 was inferior to that of Model 5.20 calibrated on the same dataset. All parameters in Model 5.23 were of the correct sign, although the fare elasticity (λ) was still high and the car (κ) and rail speed (δ) parameters were not significant. However, this may have resulted from the small size of the calibration dataset rather than from any flaw in the model form. Logically it would be expected that weighting the population units by access distance after they had been allocated to a particular station would improve model fit. Weighting functions 4.5-4.10

were therefore tested in Model 5.23 (using a 15 kph average access speed), with a range of predefined values specified for the weighting parameter (ψ). However, even the best ψ values only equalled the model fit given by the unweighted flow-specific populations, and the weights allocated by the best weighting function (4.6) were a very close approximation to a constant weight regardless of distance from the station. It therefore seemed that the best representation of reality was given by assigning an equal weight to all population units within flow-specific station catchments, in contrast with the generalised catchments where the best fit was obtained by using a distance decay function.

Because the fit of Model 5.23 was inferior to that of Model 5.20 the inclusion of an employment variable in the flow-specific model was tested, both as an absolute figure (Model 5.24) and relative to catchment population (Model 5.25). The results of calibrating these models (with a 15 kph average access speed) are summarised in Table 5.15.

$$\hat{T}_{ij} = \alpha\, P_{iJ}^{\beta}\, J_{iJ}^{\tau} Pk_i^{\rho} \prod_{j}^{n} D_j^{\gamma_j} D_{ij}^{\omega}\, Rs_{ij}^{\delta} Cs_{ij}^{\kappa} F_{ij}^{\eta}\, Rfkm_{ij}^{\lambda} \tag{5.24}$$

$$\hat{T}_{ij} = \alpha\, P_{iJ}^{\beta}\, Jp_{iJ}^{\tau} Pk_i^{\rho} \prod_{j}^{n} D_j^{\gamma_j} D_{ij}^{\omega}\, Rs_{ij}^{\delta} Cs_{ij}^{\kappa} F_{ij}^{\eta}\, Rfkm_{ij}^{\lambda} \tag{5.25}$$

Where:

J_{iJ} is the number of jobs located within the area where station *i* is the station which gives the shortest total journey time when travelling to station *j*

Jp_{iJ} is the number of jobs located within the area where station *i* is the station which gives the shortest total journey time when travelling to station *j*, divided by the resident population in the same area

Table 5.15: Summarised results from calibration of Models 5.24-5.25

	Model 5.24		Model 5.25	
	Value	t stat	Value	t stat
Intercept	6.489	4.603	6.489	4.603
β parameter	0.038	0.225	0.209	3.280
τ parameter	0.171	1.011	0.171	1.011
ρ parameter	0.298	4.426	0.298	4.426
ω parameter	-2.398	-5.858	-2.398	-5.858
δ parameter	0.758	1.220	0.758	1.220
κ parameter	-1.617	-2.304	-1.617	-2.304
η parameter	1.221	5.301	1.221	5.301
λ parameter	-3.185	-4.679	-3.185	-4.679
R_{adj}^2	0.580		0.580	

The fit of both models was slightly superior to that of Model 5.20, but the population (β), employment (τ) and rail speed (δ) parameters were all insignificant in Model 5.24, as were the employment and rail speed parameters in Model 5.25. These problems, together with the time required to define flow-specific catchments for anything other than a very small set of destinations, meant that Model 5.20 was retained as the preferred generalised origin variable model form. Nonetheless, the use of this type of flow specific catchment could still improve the accuracy of rail demand models if sufficient data on road speeds and travel times was available to allow reliable catchment definition using these methods. Their failure to give an improvement in this instance seems likely to result from deficiencies in the data available rather than from fundamental flaws in the methodology.

5.2.5 Constrained models

While the best direct demand models described above gave reasonable results, the forecasts from these models were not constrained by the total number of trips originating or terminating at the stations in question, meaning that it was possible for the sum of distributed trips predicted by the direct demand model to exceed the actual total number of trips. A methodology for constraining the forecasts from direct demand models so that they sum to give the total observed number of trips was therefore tested, and is described here. However, such a constraint was not ultimately found to improve model fit, and this methodology does not therefore form part of the recommended demand modelling procedure, although some of the initial modifications made to allow the previous models to be recalibrated on the Rhymney line were retained.

The simplest way to constrain flow level forecasts so that they sum to give the total observed trip ends is to scale the results from a direct demand model by an appropriate factor. However, the model calibrations described above were not suitable for such scaling, as they did not model the complete set of flows for any particular station. A new subset of flows was therefore selected from the LENNON dataset, comprising all 2,818 flows from the 15 stations on the Rhymney branch, chosen because they encompass a wide variety of station characteristics. However, the majority of these flows had very low demand, and it would be unrealistic for a model to be expected to distribute trips to these flows correctly as their size will be dependent on highly individual factors such as the location of meetings or the residential choice of passengers' relatives. It was therefore necessary to develop a methodology for excluding such flows in a consistent manner across all origin stations in the dataset. This was achieved by ranking flows from each station in descending order of size, and then selecting progressively smaller flows until a certain percentage of the total trip origins at each station were included in the dataset. Table 5.16 shows the total number of trip origins and flows from each station, along with the number of flows required to reach a set percentage of the total trips when flows are ranked in descending order of size.

Table 5.16: Number of flows needed to include set percentages of total trips

Station	Total origins	Total flows	Flows to reach trip cut-off point			
			90%	95%	97.5%	99%
Rhymney	150929	202	8	14	19	31
Pontlottyn	11364	83	10	15	21	34
Tir-Phil	16566	97	11	16	22	38
Brithdir	8170	53	12	15	19	28
Bargoed	103726	233	13	20	28	50
Gilfach Fargoed	1615	34	10	14	18	23
Pengam	388817	256	4	9	14	24
Hengoed	68431	168	11	16	23	38
Ystrad Mynach	161595	274	10	16	24	42
Llanbradach	39070	130	11	17	22	35
Aber	137864	187	7	15	22	34
Caerphilly	430949	496	6	15	27	62
Lisvane & Thornhill	98561	201	8	16	26	43
Llanishen	84770	187	9	18	27	45
Heath High Level	173934	217	7	16	26	41
TOTAL	1876361	2818	137	232	338	568

The 95% cut off point was chosen, as it gave a reasonable number of flows for calibration while reducing the time required for data collection. The simplest of the successful direct demand models (5.4) was recalibrated on this dataset along with the more problematic Model 5.19 and the generalised origin variable Model 5.20, although the train frequency variables had to be replaced with a service headway variable, giving Models 5.26-5.28. This was because 65 of the flows did not have a direct rail service meaning that the Perl scripts could not be used to calculate train frequencies. It would have required a prohibitive amount of time to manually calculate train frequencies for flows where interchange was required, but establishing the approximate service headway was much more straightforward using the Deutsche Bahn online journey planner. Where interchange was necessary, the time spend interchanging was counted double in the rail journey time, as recommended in the PDFH (ATOC, 2002). The results of calibrating Models 5.26-5.28 on the 232 flow Rhymney line dataset are summarised in Table 5.17.

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} R_{ij}^{\delta} H_{ij}^{\eta} \tag{5.26}$$

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} D_{ij}^{\omega} Rs_{ij}^{\delta} Cs_{ij}^{\kappa} H_{ij}^{\eta} Rfkm_{ij}^{\lambda} \tag{5.27}$$

$$\hat{T}_{ij} = \alpha \left(\sum_a P_a w_a \right)^{\beta} J_{i4}^{\tau} Pk_i^{\rho} \prod_j^n D_j^{\gamma_j} D_{ij}^{\omega} Rs_{ij}^{\delta} Cs_{ij}^{\kappa} H_{ij}^{\eta} Rfkm_{ij}^{\lambda} \tag{5.28}$$

Where:

H_{ij} is the service headway in minutes between station i and station j

Table 5.17: Summarised results from calibration of Models 5.26-5.28

	Model 5.26		Model 5.27		Model 5.28	
	Value	t stat	Value	t stat	Value	t stat
Intercept	10.093	9.255	11.597	11.312	9.710	6.348
β parameter	n/a	n/a	n/a	n/a	-0.290	-1.820
τ parameter	n/a	n/a	n/a	n/a	0.388	5.227
ρ parameter	n/a	n/a	n/a	n/a	0.198	4.236
ω parameter	n/a	n/a	-2.081	-6.849	-1.502	-3.710
δ parameter	-0.454	-4.168	0.659	2.775	0.474	1.421
κ parameter	n/a	n/a	4.495	4.905	2.825	2.751
η parameter	-0.252	-0.919	-0.135	-0.527	-0.839	-5.054
λ parameter	n/a	n/a	-1.573	-4.101	-1.545	-3.087
R_{adj}^2	0.819		0.852		0.683	
Mean AD	0.549		0.490		0.870	

The fit of Model 5.26 ($R_{adj}^2 = 0.819$) was superior to that of the previous calibration of Model 5.4 on the dataset of 2439 direct flows ($R_{adj}^2 = 0.787$), although the headway parameter (η) was not significant in Model 5.26. The fit of Model 5.27 ($R_{adj}^2 = 0.852$) is also superior to the previous calibration of Model 5.19 ($R_{adj}^2 = 0.838$), although the headway parameter was not significant and the car speed parameter (κ) was of the wrong sign, suggesting that rail demand increases with car speed. Similarly, the fit of Model 5.28 is superior to that from the previous calibration of Model 5.20, although the population (β) and road speed parameters were of the wrong sign, and the rail speed parameter (δ) was insignificant. Because this was the only model of the three calibrated on this dataset which would be suitable for forecasting new stations, and because the population parameter seemed likely to give counterintuitive results, it was modified with a total entries variable replacing the generalised origin variables, giving Model 5.29. This would still be suitable for forecasting flow level demand from new stations if combined with predictions from a trip end model. The results from calibrating Model 5.29 are summarised in Table 5.18, which shows that the model had a very good fit and that the entries parameter (β) was of the correct sign, although the road speed parameter (κ) was of the wrong sign and the headway parameter (η) was insignificant. The implied rail fare elasticity, while lower than in previous calibrations, is still larger at -1.38 than the PDFH recommended values for local and interurban trips of -0.7 to -1.1 (ATOC, 2002). However, the size of elasticity can perhaps be justified by the extensive bus competition in the case study area, which may mean that rail demand is more elastic with respect to fare changes than would be the case elsewhere. Overall, Model 5.29 seems likely to prove the most accurate of the models developed so far at forecasting flow level demand from new stations.

$$\hat{T}_{ij} = \alpha\, En_i^{\beta} \prod_{j}^{n} D_j^{\gamma_j} D_{ij}^{\omega} Rs_{ij}^{\delta} Cs_{ij}^{\kappa} H_{ij}^{\eta} Rfkm_{ij}^{\lambda} \qquad (5.29)$$

Table 5.18: Summarised results from calibration of Model 5.29

	Value	t stat
Intercept	5.218	5.697
β parameter	0.819	15.886
ω parameter	-1.730	-5.910
δ parameter	0.574	2.337
κ parameter	3.991	5.514
η parameter	-0.097	-0.799
λ parameter	-1.381	-3.797
R_{adj}^2	0.827	
Mean AD	0.567	

The main aim of this section of the analysis was to constrain the total predicted flow level trips to match the total observed level of trip origins from each station. Simple scaling of model results on a station by station basis would not improve model fit as any possible improvement would have been captured during model calibration by the origin dummy variables. The results from the simplest model calibrated for the Rhymney line (5.26) were therefore scaled so that the overall total of trips was consistent with the observed total, using transformation 5.30. The accuracy of the scaled and unscaled predictions are compared in Table 5.19.

$$T_{ijS} = \sum_i T_{iO} \left(\frac{T_{ijP}}{\sum_i \sum_j T_{ijP}} \right) \qquad (5.30)$$

Where:

T_{ijS} is the scaled total of trips predicted between station i and station j

T_{iO} is 95% of the observed total of trip origins at station i (because the model only considers flows making up 95% of the trips)

T_{ijP} is the predicted total of trips between station i and station j given by Model 5.26

Table 5.19: Fit of unscaled and scaled predictions from Model 5.26

Measure of fit	Unscaled predictions	Scaled predictions
R^2	0.858	0.779
AD	0.549	1.049

Both measures of fit indicated that the unscaled predictions were superior to the scaled predictions. While the scaled predictions are still worthy of consideration, because the overall number of trips predicted on the system should be more representative of reality, given that this overall number of trips can be produced by a trip end model this property is of limited use. This methodology does not therefore provide an ideal solution to the problem of constraining flow-level predictions, and this issue will be considered further in Section 5.4.

5.2.6 Intervening and Competing Opportunities

Preston (2001) noted that a particular problem with direct demand models is that while they may predict travel to primary destinations well, they are less accurate in their forecasts of travel to secondary destinations. This is partly because they do not usually address the issue of intervening and competing opportunities, even though the number of flows from an origin to a particular destination will obviously be lower if another destination offering equivalent or superior facilities exists within a shorter journey time of the origin. The 232 flow Rhymney line dataset was used to investigate this issue, as it was necessary for a wide range of destinations to be included in the calibration dataset, with Model 5.26 modified to incorporate several different representations of intervening opportunities. These modifications are described in this section, although they did not prove successful and were therefore not adopted as part of the recommended demand forecasting methodology.

The first means of accounting for the presence of intervening opportunities was by including a dummy variable representing flows where a higher category station was closer to the origin than the destination station in question, giving Model 5.31. However, this definition of an 'intervening opportunity' may be incorrect, as stations in the same category as the destination could form intervening opportunities as well as stations in higher categories. This latter definition was used to define the dummy variable included in Model 5.32. An alternative way to represent intervening opportunities in the model was to rank all destination stations by observed demand, and define an intervening opportunity as being a station of higher rank which is closer to the origin in terms of journey time than the destination station in question. Such a variable was incorporated in Model 5.33. A further possibility was to replace the dummy variable with a continuous variable representing the difference in rank between the destination in question and the highest ranked intervening opportunity, giving Model 5.34. Using differences in rank to define the degree of an intervening opportunity is not an ideal solution, as the ranking system takes no account of the actual difference in patronage from rank to rank. For example, the difference in trip attraction between the stations ranked 1 and 2 was 6.65 million exits, whereas the difference between the stations ranked 2499 and 2500 was only 4 exits, yet Model 5.34 would treat these differences as being equivalent. It may therefore be more realistic to replace the difference in ranks with the difference in the number of exits, giving Model 5.35. The results from calibrating Models 5.31-5.35 are summarised in Table 5.20.

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} R_{ij}^{\delta} H_{ij}^{\eta} IEh_{ik}^{\lambda} \tag{5.31}$$

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} R_{ij}^{\delta} H_{ij}^{\eta} IEsh_{ik}^{\lambda} \tag{5.32}$$

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} R_{ij}^{\delta} H_{ij}^{\eta} IEr_{ik}^{\lambda} \tag{5.33}$$

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} R_{ij}^{\delta} H_{ij}^{\eta} IErd_{ik}^{\lambda} \tag{5.34}$$

$$\hat{T}_{IJ} = \alpha \prod_{i}^{n} O_i^{\beta_i} \prod_{j}^{n} D_j^{\gamma_j} R_{ij}^{\delta} H_{ij}^{\eta} IEexd_{ik}^{\lambda} \tag{5.35}$$

Where:

IEh_{ik} is a dummy variable which takes the value e^1 if a station of higher category than station j is closer in terms of rail journey time to station i than station j, and e^0 otherwise

$IEsh_{ik}$ is a dummy variable which takes the value e^1 if a station of the same or higher category than station j is closer in terms of rail journey time to station i than station j, and e^0 otherwise

IEr_{ik} is a dummy variable which takes the value e^1 if a station of higher rank than station j is closer in terms of rail journey time to station i than station j, and e^0 otherwise

$IErd_{ik}$ is the difference in rank between station j and the highest ranked station which is closer to station i than station j, and takes the value 0.01 if no such higher ranked station exists because a logarithmic transformation is necessary for calibration

$IEexd_{ik}$ is the difference in total trip exits between station j and the station with the highest number of exits which is closer to station i than station j and has a greater number of exits than station j, and takes the value 0.01 if no such station exists

Table 5.20: Summarised results from calibration of Models 5.31-5.35

		Model 5.31	Model 5.32	Model 5.33	Model 5.34	Model 5.35
Intercept	Value	10.210	10.090	9.960	9.754	9.740
	t stat	9.363	9.219	9.085	8.732	8.784
δ parameter	Value	-0.608	-0.458	-0.353	-0.324	-0.289
	t stat	-3.973	-3.614	-2.499	-2.223	-1.924
η parameter	Value	-0.209	-0.251	-0.248	-0.243	-0.249
	t stat	-0.762	-0.914	-0.906	-0.891	-0.914
λ parameter	Value	0.350	0.019	-0.274	-0.031	-0.024
	t stat	1.433	0.075	-1.113	-1.342	-1.580
R_{adj}^2		0.820	0.818	0.819	0.820	0.821

None of the models brought anything more than a marginal improvement in model fit, and in Models 5.31 and 5.32 the intervening opportunity parameter (λ) was insignificant and of the wrong sign, suggesting that demand rose if an intervening opportunity existed. While the intervening opportunity parameter was of the correct sign in Models 5.33-5.35 it was not significant and the headway parameter (η) was also insignificant and of the wrong sign in all models.

Models 5.4 (or 5.26 when train frequency data is not available) and 5.20 (or 5.28 and 5.29 when train frequency data is not available) were therefore retained as the preferred direct demand model forms, with Model 5.29 likely to be the most useful for forecasting demand at new stations. However, the fact that these models did not consider intervening opportunities was a cause for concern. An alternative methodology which allowed the representation of intervening opportunities (the use of probability-based intervening opportunity trip distribution models) was therefore investigated, and this is described in Section 5.4.

$$\hat{T}_{ij} = \alpha \, En_i^{\beta} \prod_{j}^{n} D_j^{\gamma_j} D_{ij}^{\omega} Rs_{ij}^{\delta} Cs_{ij}^{\kappa} H_{ij}^{\eta} Rfkm_{ij}^{\lambda} \qquad (5.29)$$

5.3 Ultimate origin-destination survey

5.3.1 Methodology

While a number of catchment definition methods were tested in Sections 4 and 5.2.4, the preferred methods could only be chosen based on the model fit they provided, as no data was available to show how well these theoretical catchments represented actual travel behaviour. It was not possible to gain access to data from the National Rail Travel Survey (NRTS), and while Greater Manchester PTE supplied data on ultimate trip origins and destinations from the Greater Manchester Area Travel Survey problems with the survey methodology meant that this could not be trusted. A small scale survey of ultimate trip origins and destinations in South Wales was therefore undertaken. This survey took the form of a self-completion questionnaire, based on the relevant sections of the NRTS questionnaire (see Appendix 3). Self-completion questionnaires were chosen to minimise the time and staffing requirements for the survey.

Permission was obtained from Arriva Trains Wales to undertake the survey on board trains on the Cardiff-Rhymney line, as these serve a wide variety of station types and the flows on this line were being investigated as part of the work on flow level models. A pilot study was planned initially, involving 50 questionnaires, to test the design and feasibility of the survey, and these were handed out on four weekday services (Blainey, 2009d; see Appendix 1). This was largely successful, although responses highlighted the need for a few minor changes to the questionnaire before the full study was carried out.

The full-scale survey was carried out over three days on 31 services on the Cardiff Central – Rhymney line (Blainey, 2009d; see Appendix 1), with the aim of obtaining another 450 responses. The researcher (the author of this thesis) attempted to ask every passenger on board these services to fill in a questionnaire. However, in practice this was not possible for several reasons. Many passengers were either asleep, talking on a mobile phone, or carrying out other activities which the researcher felt it would be impolite to interrupt. A number of the services were formed of two diesel multiple units coupled together with no corridor connection, and while the researcher moved from unit to unit at intermediate stations, this inevitably meant that some passengers could not be contacted. Several services were also extremely busy, and this combined with short journey times meant that the researcher did not have time to speak to every passenger on the train. Using multiple researchers might have solved the latter two problems, but on quieter services there would have been insufficient passengers to keep them occupied, and in any case only a single researcher was available for this project. Approximately two-thirds of the passengers who were offered a questionnaire agreed to complete one, and this allowed 464 responses to be collected. Together with the 50 completed questionnaires from the pilot study this gave a total sample of 514 responses available for analysis.

5.3.2 Initial Analysis

The data from the questionnaires was copied into a spreadsheet for analysis. During this process it became evident that some respondents had misunderstood the questions which aimed to elicit origin and destination addresses and given the same address for their trip origin and their trip destination (despite modifications made to the wording of the questions following the pilot survey). Given that neither the wording used in the pilot survey nor that used in the full survey was foolproof, the only way to resolve this problem would be to replace the self-completion questionnaires with an interview-based approach.

Many respondents had provided origin or destination details without giving a postcode, and where possible Transport Direct and/or Google Maps (maps.google.co.uk) were used to obtain these codes. However, a number of respondents gave nonspecific origins or destinations, such as ‘Hengoed’ or ‘Caerphilly’, and as such responses did not allow the distance travelled to/from the station to be calculated to an acceptable level of accuracy these data points were removed from the dataset. The Edina Digimap Postcode Query tool was then used to obtain coordinates for the 409 origins and 404 destinations which remained, allowing catchment analysis to be undertaken using MapInfo.

A few journeys were recorded where either the origin or the destination station was not on the Rhymney line and as not enough trips were recorded to such stations to give a reliable representation of the station catchments these points were also removed from the dataset, reducing it to 753 trip ends. No distinction was made in the analysis between origins and destinations, as in many cases it was difficult to establish which end of the flow was the generator and which was the attractor. The distribution of these trip ends between the stations is given in Table 5.21, together with summary statistics describing the distance of the trip ends from the stations.

Table 5.21: Distance of ultimate trip ends from Rhymney line stations

Station	Trip Ends	Road access distance (km)					
		Mean	St Dev	90th percentile	95th percentile	Max	Min
Aber	27	1.623	1.184	3.172	4.279	4.71	0.12
Bargoed	42	2.479	3.218	6.382	9.7055	15.37	0.17
Brithdir	8	0.673	0.713	1.589	1.8795	2.17	0.18
Caerphilly	108	3.018	3.594	5.873	8.2755	25.83	0.08
Cardiff Central	91	2.222	3.923	4.74	9.97	22.71	0.09
Cardiff Queen Street	186	0.981	1.060	2.25	2.6825	7.88	0.05
Gilfach Fargoed	1	0.830	n/a	0.83	0.83	0.83	0.83
Heath High Level	15	1.777	0.912	2.072	2.836	4.6	0.38
Hengoed	36	1.783	2.038	4.045	4.46	10.85	0.16
Lisvane & Thornhill	25	1.992	3.874	2.632	5.996	19.5	0.14
Llanbradach	19	1.974	2.911	3.782	6.467	12.56	0.16
Llanishen	19	1.354	1.338	2.084	3.552	5.82	0.07
Pengam	52	2.511	2.249	4.011	5.02	13.29	0.02
Pontlottyn	9	0.839	0.439	1.33	1.33	1.33	0.17
Rhymney	20	3.301	2.293	6.603	7.9735	9.56	1.49
Tir-Phil	15	0.793	0.546	1.546	1.574	1.63	0.16
Ystrad Mynach	80	4.112	3.403	7.329	10.7485	18.45	0.07
TOTAL	753	2.143	2.869	4.59	7.15	25.83	0.02
Total excluding Cardiff	476	2.582	2.993	5.69	7.755	25.83	0.02

5.3.3 Fit With Theoretical Catchments

The main reason for the survey was to enable theoretical catchment boundaries to be compared to actual travel behaviour. The mean observed distance by road between stations and ultimate trip ends was (at 2.14 km) greater than the outer catchment boundary used in most previous local rail demand models (2 km). This mean distance was even greater when the large Cardiff stations were removed from the dataset. While the maximum access distance was just over 25 km, 90% of the distances were below 5 km, and 95% below 7.5 km, suggesting that one of these figures should be used if an arbitrary boundary was necessary.

However, not all surveyed trips were undertaken at the same frequencies, and it was possible that there could be a correlation between trip frequency and access distance which would skew these mean distances. The data from the survey was scaled up using the reported frequency of the trips recorded, and compared with the observed annual trip totals from the most recent set of ORR station usage data to establish how complete a picture of rail travel on the Rhymney line had been established. The results from this scaling are shown in Table 5.22.

Table 5.22: Percentage of total trip ends captured by OD survey

Station	Surveyed Trip Ends	Total Trip Ends	% Captured
Aber	7403	169463	4.37%
Bargoed	9856	142857	6.90%
Brithdir	1200	11039	10.87%
Caerphilly	21112	608934	3.47%
Cardiff Central	24840	9126923	0.27%
Cardiff Queen Street	41482	2231784	1.86%
Gilfach Fargoed	480	2284	21.02%
Heath High Level	3699	275582	1.34%
Hengoed	6298	96676	6.51%
Lisvane & Thornhill	4823	129755	3.72%
Llanbradach	5525	57836	9.55%
Llanishen	3850	173289	2.22%
Pengam	13547	426341	3.18%
Pontlottyn	2402	16499	14.56%
Rhymney	5552	176953	3.14%
Tir-Phil	2708	25489	10.62%
Ystrad Mynach	19690	221619	8.88%
Total	174467	13893323	1.26%
Total excluding Cardiff	108145	2534616	4.27%

Excluding the central Cardiff stations (where trips via the Rhymney line form only a small proportion of total travel), the survey captured just over 4% of the total travel on the Rhymney line. In general a larger proportion of trips were captured at quieter stations (Brithdir, Gilfach Fargoed, Pontlottyn and Tir Phil) than at busier stations. Similarly, a greater proportion of trips were captured from stations further up the valley than from those closer to Cardiff. This was probably because the trains were quieter further away from Cardiff, making it easier to question all passengers, whereas passengers travelling to stations close to Cardiff would often have left the train before they could be surveyed. The mean and standard deviation of the road access distances were recalculated for the scaled-up data, with the mean road access/egress distance found to have reduced slightly to 1.97 km from 2.14 km. A similar reduction to 2.39 km from 2.58 km was found for access/egress outside central Cardiff. This suggested that in general access/egress distance reduced slightly with trip frequency, as was expected.

While these results indicate that the catchment boundaries used in previous rail demand models were too small, the catchments used in the trip end models (Section 4) and direct demand models (Section 5.2) developed here do not have arbitrary boundaries, instead allocating output areas either to their nearest station or to the station which minimises overall journey time, with populations weighted using distance-decay functions. The observed trip ends were therefore compared with both the generalised (equivalent to those used in the trip end models) and flow specific station catchments used in the South Wales direct demand models. Table 5.23 details the percentage of observed trip ends which fell within the theoretical catchment boundary for each station. The flow-specific catchments are those based on a 15 kph average access speed, and only flows which are entirely within the Rhymney line were considered for this catchment type to keep the number of catchments which required calculation manageable.

Table 5.23: Percentage of trip ends within catchment boundaries

Catchment type	Generalised			Flow-specific		
Station	Within catchment	Total	Percentage	Within catchment	Total	Percentage
Aber	24	27	88.89%	16	23	69.57%
Bargoed	28	42	66.67%	28	39	71.79%
Brithdir	6	8	75.00%	7	8	87.50%
Caerphilly	63	108	58.33%	62	90	68.89%
Cardiff Central	56	91	61.54%	19	91	20.88%
Cardiff Queen Street	43	186	23.12%	131	183	71.58%
Gilfach Fargoed	0	1	0.00%	2	2	100.00%
Heath High Level	3	15	20.00%	13	15	86.67%
Hengoed	28	36	77.78%	23	31	74.19%
Lisvane & Thornhill	22	25	88.00%	15	21	71.43%
Llanbradach	16	19	84.21%	12	15	80.00%
Llanishen	12	19	63.16%	7	16	43.75%
Pengam	40	52	76.92%	32	45	71.11%
Pontlottyn	9	9	100.00%	9	9	100.00%
Rhymney	17	20	85.00%	11	18	61.11%
Tir-Phil	15	15	100.00%	13	14	92.86%
Ystrad Mynach	19	80	23.75%	32	67	47.76%
Total	401	753	53.25%	432	687	62.88%
Total excluding Cardiff	302	476	63.45%	282	413	68.28%

This suggests that the generalised catchment boundaries are not an accurate representation of reality, given that almost half the observed trip ends fell outside these boundaries. The flow-specific catchments gave a better fit with the observed data but still failed to capture over a third of observed trip ends. There appears to be a great deal of variation in accuracy between stations with some catchments, such as the generalised catchments for Lisvane & Thornhill and Aber, containing the majority of observed trip ends whereas others, such as the generalised catchments for Ystrad Mynach and Cardiff Queen Street and the flow specific catchment for Cardiff Central, contain very few of the observed trip ends. Figures 5.9 and 5.10 show the distribution of observed trip ends around Ystrad Mynach and Cardiff Queen Street together with the generalised catchment boundaries for those stations.

Figure 5.9: Ultimate trip ends to/from Ystrad Mynach station

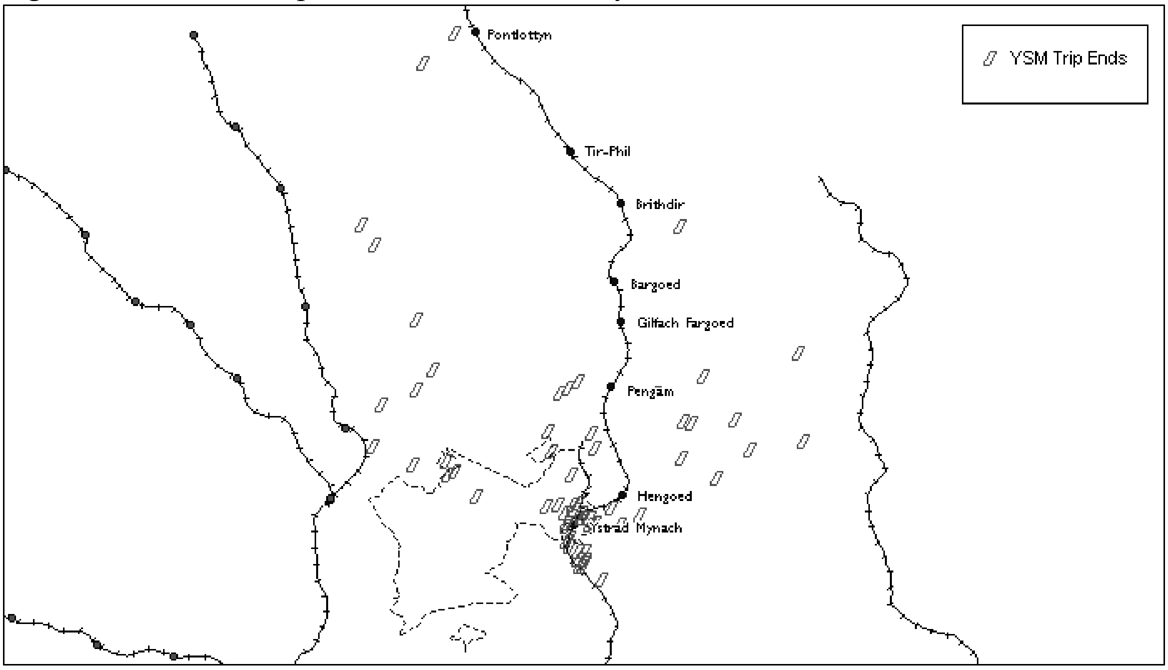

Figure 5.10: Ultimate trip ends to/from Cardiff Queen Street station

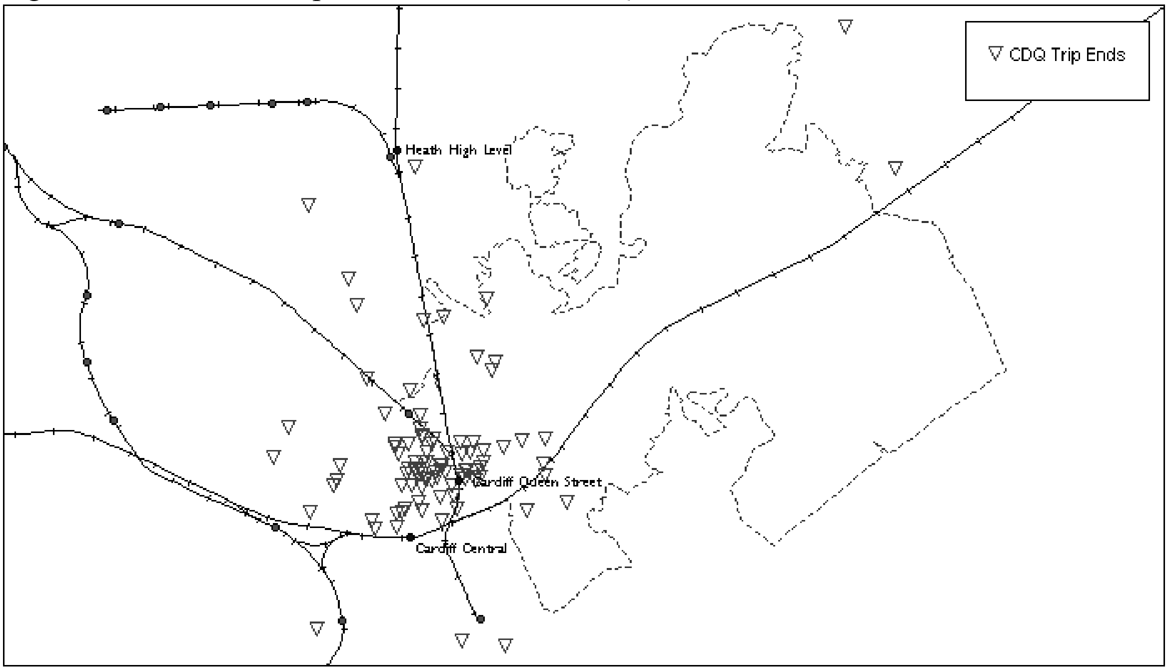

The principal problem at both of these stations was the shape of the catchment areas, which only cover the area on one side of the railway. This resulted from the shape of the output areas around the station used to define the catchment boundaries. For example, the low density of output areas in central Cardiff meant that most of the central shopping area was allocated to Cardiff Central despite being closer to Queen Street. Using catchments based solely on access time and not constrained by output area boundaries might therefore give better results, and such catchments were created using RouteFinder for MapInfo. Table 5.24 gives the percentage of observed trip ends which fall within these catchments.

Table 5.24: Percentage of trip ends within non-output area based catchment boundaries

Station	Within catchment	Total	Percentage
Aber	21	27	77.78%
Bargoed	28	42	66.67%
Brithdir	6	8	75.00%
Caerphilly	73	108	67.59%
Cardiff Central	50	91	54.95%
Cardiff Queen Street	141	186	75.81%
Gilfach Fargoed	1	1	100.00%
Heath High Level	9	15	60.00%
Hengoed	28	36	77.78%
Lisvane & Thornhill	23	25	92.00%
Llanbradach	16	19	84.21%
Llanishen	11	19	57.89%
Pengam	40	52	76.92%
Pontlottyn	9	9	100.00%
Rhymney	19	20	95.00%
Tir-Phil	15	15	100.00%
Ystrad Mynach	17	80	21.25%
Total	507	753	67.33%
Total excluding Cardiff	316	476	66.39%

While these catchments overall gave a better match with the observed data, this was largely due to the improvement at Cardiff Queen Street, where 75% of observed trip ends now fell within the catchment compared to 23% before. Furthermore, such catchments could not easily be used in model calibration, because they do not correspond to census data units meaning that there is no straightforward way to calculate catchment populations. A pragmatic solution to this particular issue might be to treat Cardiff Queen Street and Cardiff Central as a single station for catchment definition and modelling purposes. However, the finding that the theoretical catchments only contain between half and two-thirds of the observed trip ends indicates that there is much room for improvement in catchment definition in rail demand models.

5.3.4 Observed Catchments

In an attempt to identify ways in which such improvements could be achieved, the observed data on ultimate trip ends from the survey was analysed in more detail. The longest mean access distances were recorded at the four busiest stations outside Cardiff, specifically Caerphilly (3.02 km), Ystrad Mynach (4.11 km), Pengam (2.51 km) and Bargoed (2.48 km). A number of services terminate at Bargoed and Ystrad Mynach suggesting that the longer access distances at these stations may result from 'railheading', a theory supported by Figure 5.10 which shows that the majority of passengers at Ystrad Mynach come from the area to the north of the station. This railheading was likely to be a consequence of the higher train frequency at Ystrad Mynach (four services per hour in each direction) compared to stations north to Bargoed (three services per hour in each direction) and further north to Rhymney (one service per hour in each direction), combined with the provision of a medium-sized car park (34 spaces).

The maps of observed trip ends in Figures 5.9-5.10 did not provide a very clear picture of the observed station catchments, as representing trip ends as point locations does not illustrate the frequency of trips from those points. MapInfo was therefore used to create grid-based maps of observed station catchments which show trip density within the catchments based on the scaled trip ends derived from the survey. Two methods of producing these maps were available in MapInfo, inverse distance weighting (IDW) interpolation and triangulated irregular network (TIN) interpolation. IDW interpolation uses a distance weighted average of data points to calculate grid cell values, whereas TIN interpolation draws lines between points, dividing them into triangles and connecting all the points it can, to create a mesh of connectivity so that grid points can be interpolated (MapInfo, 2009). Both methods were tested for Aber station, giving Figures 5.11-5.12, which also show the generalised theoretical catchments for Aber and the local road and rail networks.

Figure 5.11: Observed catchment for Aber using TIN interpolation

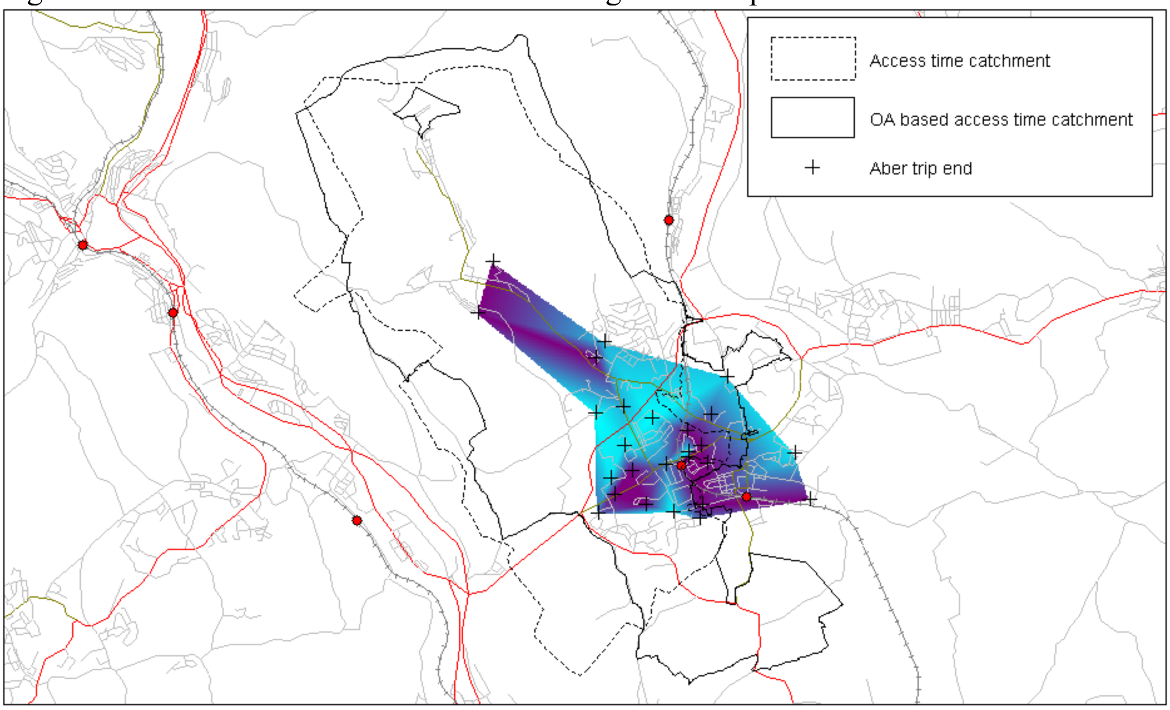

Figure 5.12: Observed catchment for Aber using IDW interpolation

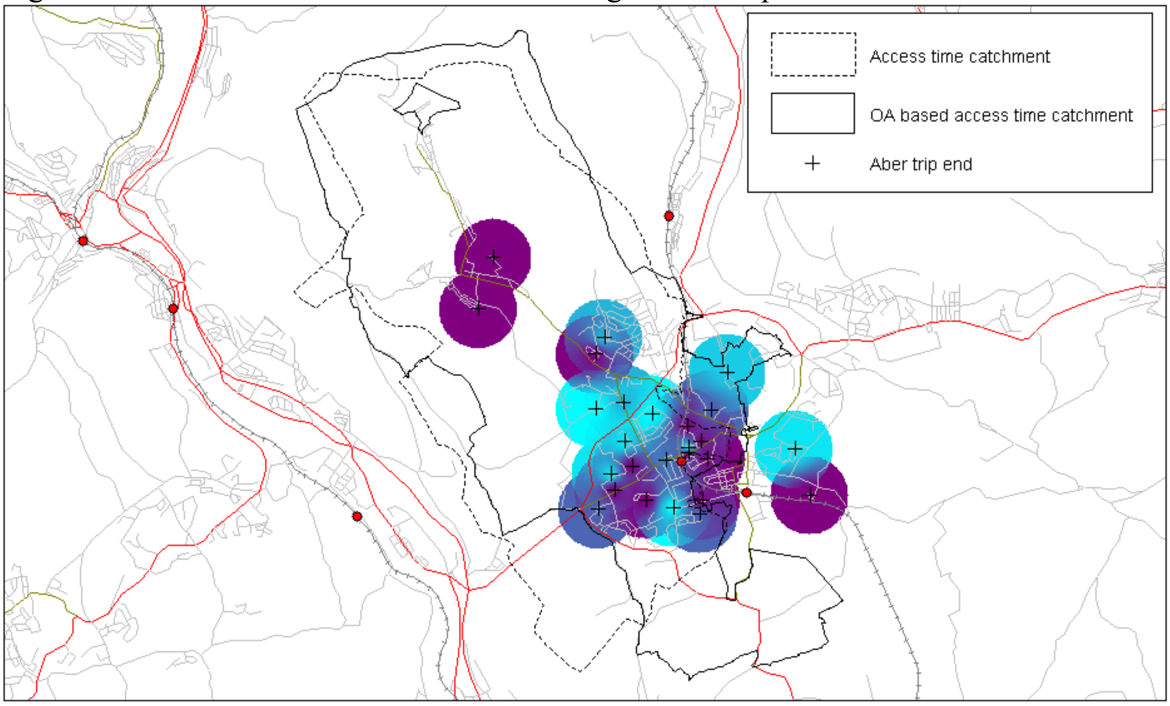

While TIN interpolation produced a more continuous catchment with polygonal boundaries, the observed data points formed the boundaries of the catchment, which was not necessarily realistic. In contrast, IDW interpolation gave a discontinuous catchment with isolated sections, but allowed the catchment to extend slightly beyond the observed points (the extent to which this occurred was dependent on the grid border used in interpolation). Discontinuous catchments may be more realistic than continuous catchments, as the TIN catchment includes Caerphilly station, which would not be expected to fall within Aber's catchment. Furthermore, the MapInfo Knowledge Base suggested that the IDW interpolator was more suitable for data values which do not have any relationship or influence over neighbouring values (MapInfo, 2009), and as rail trip end data falls into this category IDW interpolation was selected as the preferred method. Observed catchments were mapped in this way for the other stations on the Rhymney line, and these are illustrated in full in Blainey (2009d) (see Appendix 1). These interpolated images gave a much clearer picture of observed station catchments, but did not make it obvious how these observed catchments should be represented in demand models. Furthermore, while several catchments appeared to have isolated sections, because the survey only captured around 4% of total travel it is not clear how many of these exist in reality, and it would therefore be unrealistic to draw a continuous catchment boundary containing all observed points. To gain a full understanding of station catchments and suggest suitable generalisation methods it would be necessary to extend this analysis over a much larger number of stations, requiring more data than was available for this study.

5.3.5 Disaggregation by Access Mode

It seemed likely that access mode choice would affect access distance and the shape of catchments, and the results of disaggregating the dataset by access mode are summarised in Table 5.25. Access/egress mode information had not been provided for four of the trip ends, meaning that the dataset described in the table comprises 749 trip ends.

Table 5.25: Mean access/egress distances disaggregated by mode

Station	Walk		Car (driver)		Car (passenger)	
	Trip Ends	Mean dist	Trip Ends	Mean dist	Trip Ends	Mean dist
Aber	16	1.108	4	1.935	4	1.955
Bargoed	28	1.467	2	1.280	5	7.992
Brithdir	7	0.577	1	1.340		
Caerphilly	57	1.818	17	5.878	9	2.583
Cardiff Central	72	0.937	2	0.365	1	6.510
Cardiff Queen St	165	0.854	2	0.805	3	0.583
Gilfach Fargoed	1	0.830				
Heath High Level	13	1.875	2	1.140		
Hengoed	20	1.011	11	3.146	3	2.543
Lisvane & Thornhill	14	0.680	7	4.267	2	4.015
Llanbradach	13	1.369	2	7.145	2	1.355
Llanishen	17	0.977	1	3.300		
Pengam	18	1.497	23	3.512	9	1.926
Pontlottyn	7	0.739			1	1.330
Rhymney	10	2.164	1	9.560	4	4.818
Tir-Phil	10	0.550	2	1.390	2	0.990
Ystrad Mynach	44	2.709	15	6.535	10	6.071
Total	512	1.255	92	4.233	55	3.605
Total excluding Cardiff	275	1.579	88	4.4	51	3.726
% Total	68.36		12.28		7.34	
Station	Bus		Bicycle		Taxi	
	Trip Ends	Mean dist	Trip Ends	Mean dist	Trip Ends	Mean dist
Aber	2	2.905				
Bargoed	4	3.058			3	2.770
Brithdir						
Caerphilly	22	4.195	1	2.190	1	4.220
Cardiff Central	14	8.841			1	3.510
Cardiff Queen St	13	2.549	2	1.595	1	1.860
Gilfach Fargoed						
Heath High Level						
Hengoed	1	1.420	1	0.330		
Lisvane & Thornhill			1	1.340	1	1.040
Llanbradach	1	2.410	1	0.290		
Llanishen	1	5.820				
Pengam					2	2.770
Pontlottyn	1	1.050				
Rhymney	5	3.108				
Tir-Phil	1	1.630				
Ystrad Mynach	8	5.420	2	2.525		
Total	73	4.637	8	1.549	9	2.720
Total excluding Cardiff	46	3.947	6	1.533	7	2.730
% Total	9.75		1.07		1.20	

The table shows, unsurprisingly, that average access/egress distances are much shorter for walk trips than for bus and car trips. Interestingly the average access/egress distance by taxi falls approximately half way between that for walk and that for other motorised modes, although given this is based on only 9 observations it would be dangerous to place too much reliance on this result. Over two-thirds of station access/egress was on foot, indicating that this is the dominant mode used to travel to local railway stations. The frequency distributions of access distance by each mode were plotted in Figure 5.13.

Figure 5.13: Frequency distribution of access distances by mode

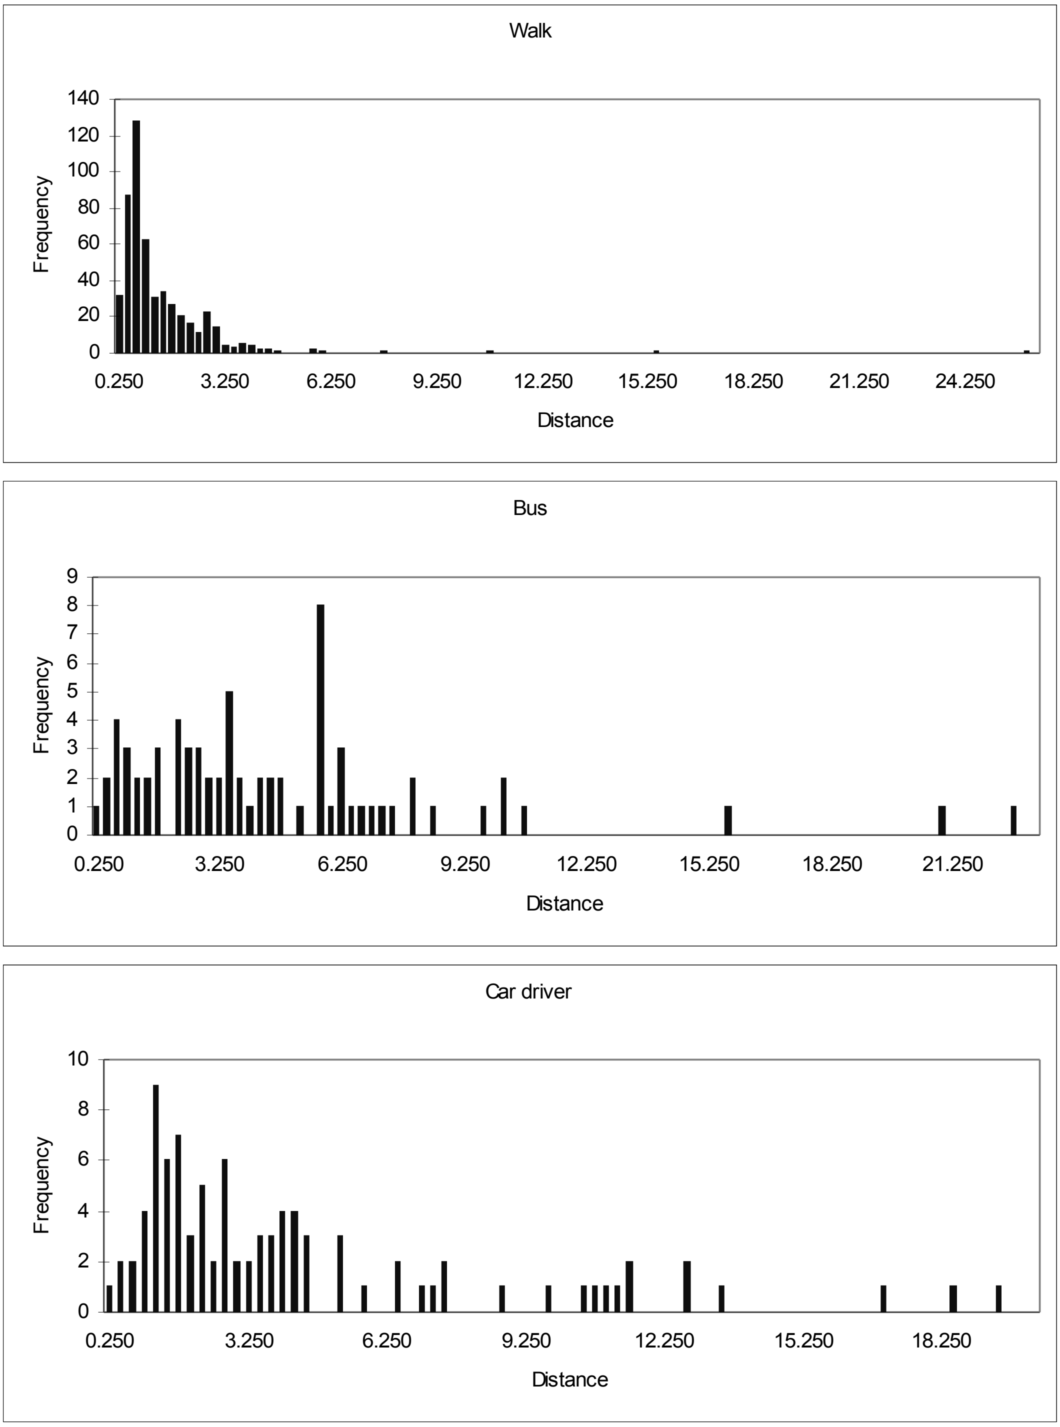

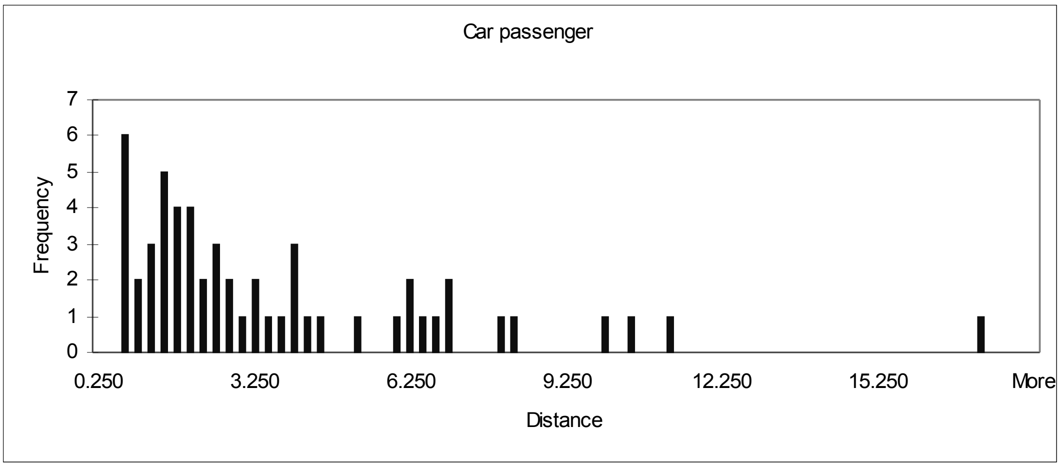

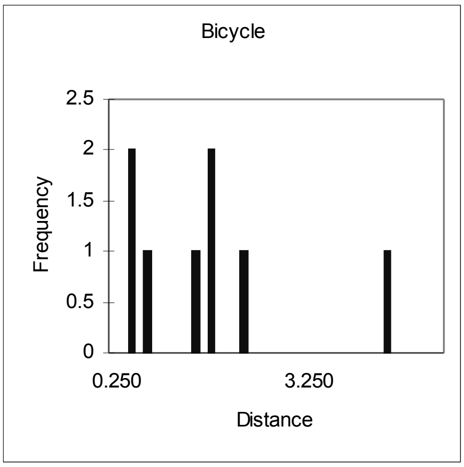

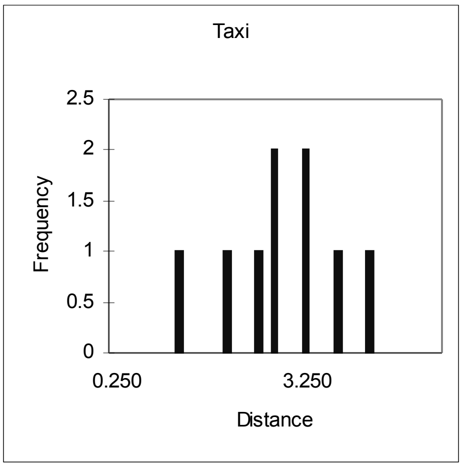

Figure 5.13 shows that in general the number of trips recorded decays with distance from the station, with this effect being particularly marked for access on foot. The 90th and 95th percentile distances were calculated for each mode, and are shown in Table 5.26, which suggests that if arbitrary catchment boundaries are used it would be sensible to assume a 3 km boundary for walk access trips, but that a 10 or 12.5 km boundary would be necessary for motorised trips. However, as access mode split varies from station to station the use of such differential catchments in modelling would not be straightforward.

Table 5.26: 90th and 95th percentile access distances for each mode

Mode	Walk	Bus	Car driver	Car passenger	Bicycle	Taxi
90th percentile	2.6	8.312	10.818	7.47	2.874	3.652
95th percentile	3.087	10.424	12.5645	9.979	3.672	3.936

The data in Table 5.25 were scaled up using the information on trip frequency to give a more accurate picture of access mode split, giving the results in Table 5.27. Scaling the trips by frequency further increased the dominance of walk as the major access mode, and slightly reduced the mean access distance by this mode. The predominance of access/egress trips made on foot cast some doubt on the theoretical methods used to define station catchments in this work, despite the good model fit they gave, as these were based on access by car. Pedestrian speeds will not vary by road class in the same way as car speeds, meaning that relative access times in station catchment definition may have been miscalculated. Furthermore, the routes available to pedestrians are not the same as those for motorists, as while pedestrians would not be able to use motorways a number of additional footpaths might be available. The problems this caused for catchment definition are highlighted by Figure 5.9, which shows that a number of trips originated in the area immediately to the south of Ystrad Mynach station, an area linked to it by road via an extremely circuitous route. During the survey a footpath was observed which provides direct pedestrian access to the station, but this was not included in the road network data used to define catchments, and footpath data was not available in OS Meridian format. While Mastermap data on footpaths is available from EDINA Digimap, the format makes network analysis difficult, and the size of the datasets involved mean that its use is not practical for the case study areas used in this study. Similarly, the raster data available through EDINA Digimap is not suitable for network analysis, despite displaying footpaths clearly. An increasing amount of vector data on footpaths and cycle routes is available from OpenStreetMap (www.openstreetmap.org), but at the time of writing this did not extend to the Rhymney Valley area.

Table 5.27: Mean access/egress distances disaggregated by mode with trip ends scaled by trip frequency

Station	Walk		Car (driver)		Car (passenger)	
	Trip Ends	Mean Dist	Trip Ends	Mean Dist	Trip Ends	Mean Dist
Aber	5379	1.139	772	2.990	483	1.795
Bargoed	7714	1.070	288	1.280	828	9.510
Brithdir	1152	0.564	48	1.340		
Caerphilly	12947	1.557	2028	3.442	1192	3.338
Cardiff Central	20417	0.902	336	0.290	288	6.510
Cardiff Queen Street	37760	0.844	486	1.083	534	0.537
Gilfach Fargoed	480	0.830				
Heath High Level	3603	1.624	96	1.140		
Hengoed	4024	1.204	1278	2.737	36	3.202
Lisvane & Thornhill	3041	0.678	1728	2.782	50	1.434
Llanbradach	4130	1.816	768	5.791	483	0.520
Llanishen	3604	0.887	6	3.300		
Pengam	3873	1.341	7506	3.845	1687	2.054
Pontlottyn	1922	0.882			240	1.330
Rhymney	3366	2.211	480	9.560	840	4.295
Tir-Phil	2072	0.586	54	1.273	576	0.823
Ystrad Mynach	11299	2.865	3902	6.810	2076	3.668
Total	126783	1.240	19776	4.213	9313	3.307
Total excluding Cardiff	68606	1.559	18954	4.363	8491	3.373
% Total	73.09%		11.40%		5.37%	

Table continued on next page

Station	Bus		Bicycle		Taxi	
	Trip Ends	Mean Dist	Trip Ends	Mean Dist	Trip Ends	Mean Dist
Aber	768	3.269				
Bargoed	1014	4.553			12	2.763
Brithdir						
Caerphilly	4393	3.887	480	2.190	24	4.220
Cardiff Central	3318	7.279			1	3.510
Cardiff Queen Street	2123	3.244	576	1.558	3	1.860
Gilfach Fargoed						
Heath High Level						
Hengoed	480	1.420	480	0.330		
Lisvane & Thornhill			3	1.340	1	1.040
Llanbradach	48	2.410	96	0.290		
Llanishen	240	5.820				
Pengam					481	2.332
Pontlottyn	240	1.050				
Rhymney	866	2.703				
Tir-Phil	6	1.630				
Ystrad Mynach	1165	5.441	768	2.039		
Total	14661	4.527	2403	1.542	522	2.426
Total excluding Cardiff	9220	3.833	1827	1.537	518	2.427
% Total	8.45%		1.39%		0.30%	

A related and wider problem results from the automatic allocation of stations to the road network to enable catchment definition. While this is a prerequisite for GIS network analysis, because the data on stations does not contain details of the locations of their entrances/exits there was no guarantee that stations were located on the correct road links. In some cases this makes little difference to catchment shape but if two or more roads run close to a particular station but are not linked in its immediate vicinity, the shape of the catchment will be significantly affected depending on which road the station was allocated to. This problem should in theory be relatively straightforward (if somewhat time-consuming) to solve, as the location of station entrances could be manually checked using GoogleEarth and the National Rail website. If more than one station entrance existed then in this case the main entrance would be used as the network allocation point, although if the entrances were on unconnected roads then the theoretical catchments might be smaller than the actual catchments, as the network analysis procedure does not allow network points (stations) to be located on multiple links of the network. This problem could be overcome by adding road links with zero travel time to connect multiple station entrances. In the work described here time constraints meant that checks were not made of station location on the GIS road network, but such corrections could improve the realism of station catchments.

Table 5.28 compares the number of access/egress trips made by car drivers to the number of parking spaces made at each station. Car passengers are not included in these totals, as while some of them may be travelling with drivers who also use the train, others will be being dropped off or picked up at the station and will not need to use a parking space.

Table 5.28: Car access/egress trips and parking spaces

Station	Total trips	Car (driver) trips	% Car trips	Parking spaces
Aber	27	4	14.81%	0
Bargoed	42	2	4.76%	14
Brithdir	8	1	12.50%	0
Caerphilly	108	17	15.74%	83
Cardiff Central	91	2	2.20%	248
Cardiff Queen Street	186	2	1.08%	0
Gilfach Fargoed	1	0	0.00%	0
Heath High Level	15	2	13.33%	0
Hengoed	36	11	30.56%	5
Lisvane & Thornhill	25	7	28.00%	81
Llanbradach	19	2	10.53%	10
Llanishen	19	1	5.26%	42
Pengam	52	23	44.23%	59
Pontlottyn	9	0	0.00%	10
Rhymney	20	1	5.00%	20
Tir-Phil	15	2	13.33%	10
Ystrad Mynach	80	15	18.75%	34
Spearman's rank correlation coefficient		*0.390*	0.840	

While there is no significant correlation between the number of access/egress trips by car and car park size, such a correlation does exist between the percentage of access/egress trips made by car drivers and car park size. This was expected, but there were still a couple of anomalies, most notably the small proportion of trips from Llanishen made by car drivers given the large car park provided there. Conversely, almost one third of trips to/from Hengoed were made by car drivers, even though only 5 parking spaces exist at the station, suggesting that other parking facilities were available nearby. This highlights the limitations of using station car park size to predict rail demand.

5.3.6 Potential Enhancements to Theoretical Catchments

The Rhymney line survey enabled several possible enhancements to catchment definition methods in rail demand models to be identified, although constraints on time and resources meant that it was not possible to investigate them further during this study. The data from the survey showed that there were extensive overlaps between the catchments of neighbouring stations (see Figure 5.14). Using overlapping catchments similar to those developed by Lythgoe (2004) rather than the 'all or nothing' allocation of population units used here might give a better fit, although this would make model calibration much more complicated. Such catchments would also be unlikely to solve the problems caused by arbitrary output area boundaries. Better data on station access times which gave a more accurate reflection of road speeds would almost certainly improve the accuracy of theoretical catchments, as would using mode-specific catchments, although this would again complicate model calibration and would require prediction of access mode choice decisions when making forecasts for new stations.

Figure 5.14: Observed catchment overlap between Hengoed and Pengam

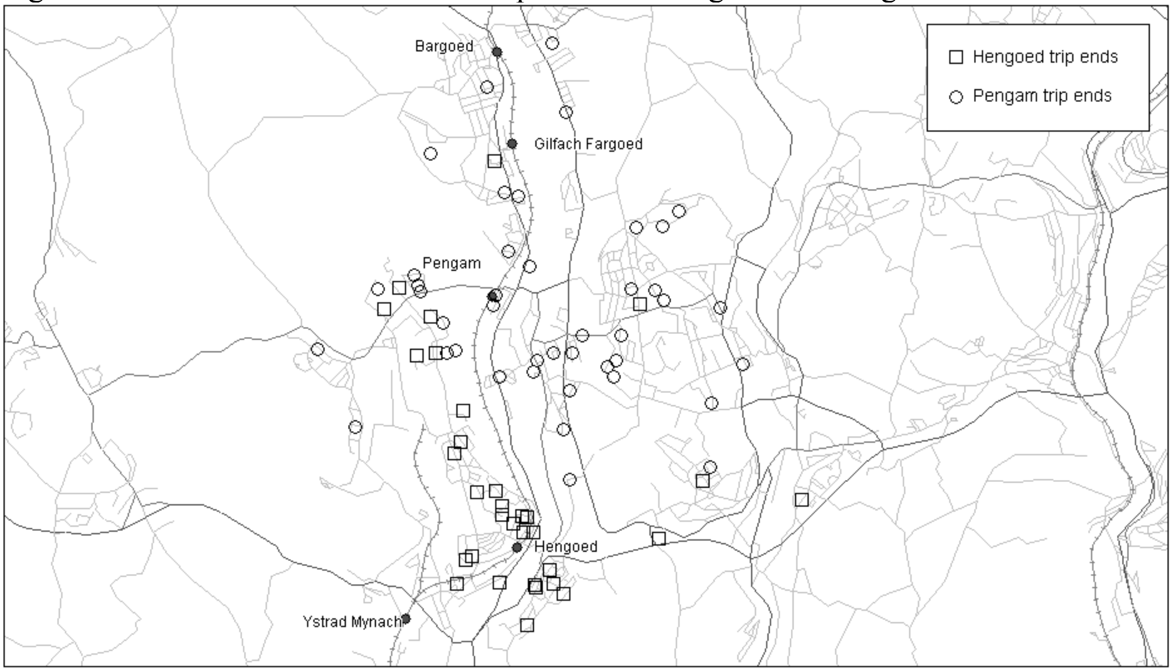

5.4 Intervening opportunity trip distribution (IOTD) models

As the direct demand models described in Section 5.2 could not easily be constrained to reflect the observed total of trip ends and could not give an adequate representation of the effect of intervening or competing opportunities, investigations turned to a different type of model with the potential to resolve both these issues. This was a trip distribution model based on a general methodology dubbed the ‘intervening opportunity model’ by Kanafani (1983), and originally postulated by Stouffer (1940), who proposed “that the number of persons travelling a given distance is directly proportional to the number of opportunities at that distance and inversely proportional to the number of intervening opportunities”. These intervening opportunity models forecast the probability of a particular destination being selected, with flow sizes then calculated by multiplying the probabilities by the total trips from the origin station, obtained either from observed data or from the predictions of a trip end model. Such a methodology has, as far as the author is aware, never before been applied to rail demand modelling, and anecdotal evidence suggests that its success in other areas has been limited.

The 232 flow dataset from the Rhymney line was used for calibration, with the first stage in this process being the definition of a measure of destination opportunities. Ideally for a rail demand model this would be based on the level of employment, population, and retail facilities around the destination station, but a more pragmatic approach was initially adopted here with the total number of observed trip ends at the station used as a proxy for destination attractiveness. For each origin station the destination opportunities were then ranked by a measure of travel impedance. Kanafani (1983) suggests that this can be distance, travel time, or generalised cost, and in this instance rail journey time was used initially. This ordering allows the cumulative function *[V(j)]* to be constructed, showing the accumulated opportunities up to and including the *j*th ranked destination. This means that if *J* is the furthest destination then *V(J)* will give the total destination opportunities for that particular origin. Observed destination choice proportions (π_j) were obtained from the LENNON data, and these were used to construct the empirical probability function of trip attenuation (5.36), which gives the probability that a choice will be made by destination *j*, given that it will definitely be made by the lowest ranked destination *J* (Ruiter, 1967).

$$P[V(j)|V(J)] = \frac{1 - e^{-\beta V(j)}}{1 - e^{-\beta V(J)}} \qquad (5.36)$$

Where:

P[V(j)|V(J)] is the probability that a trip maker will choose a destination within the *V(j)* accumulation of opportunities by the time the *j*th destination has been reached, given that they will definitely choose a destination by the time the lowest ranked destination *J* is reached.

β is a parameter determined by calibration

Linear regression could not be used to calibrate this model, so nonlinear least-squares regression was used instead, once the attraction factors had been scaled by dividing them by 10,000 to prevent the exponent values exceeding the maximum allowed by the SPSS calibration procedure. SPSS was set to use the Levenburg-Marquardt algorithm for calibration, with the parameter value initially set at 0.0001. This gave the results summarised in Table 5.29 when applied to the Rhymney line dataset. Using the total number of rail trips made to destinations describes the general attractiveness of the destination rather than its specific attractiveness to travellers from the Rhymney line, and therefore an alternative attraction factor was tested using the total number of rail trips made to each destination from stations on the Rhymney line (defined as all stations between Rhymney and Heath High Level). Variations in station choice in central Cardiff (whether passengers travel to Queen Street or to Central) are likely to increase the model prediction errors. As this choice is relatively unimportant for the purposes of this work, particularly given that all Rhymney line services call at both stations and that there is no difference in fare, Model 5.36 was therefore calibrated both using separate central Cardiff destinations and with trips to these stations aggregated into a single 'central Cardiff' destination.

Table 5.29: Summarised results from calibration of Model 5.36

Cardiff destinations		Queen Street and Central		Combined	
Attraction factor		Total trips	Rhymney line trips	Total trips	Rhymney line trips
β parameter	Value	0.00351	0.01694	0.00446	0.02443
	Std error	0.00022	0.00143	0.00028	0.00170
	95%CI Lower Bound	0.00308	0.01411	0.00392	0.02107
	95%CI Upper Bound	0.00393	0.01976	0.00501	0.02778
R^2		0.870	0.851	0.902	0.889
AD		1.333	1.164	1.087	1.114

The model appears to give a very good fit to the observed data, and the β parameter is in all cases highly significant and of the expected sign. However, it is not possible to directly compare R^2 values from this model with those from the direct demand models developed in Section 5.2 because the dependent variable is of an entirely different nature, being expressed in terms of probabilities rather than absolute flow sizes. The best comparison method therefore seemed likely to be mean absolute deviation (AD) values, as used in Chapter 4, although as stated there this is not a perfect measure of model fit. Function (5.37) was used to calculate predicted absolute flow sizes based on the probabilities forecast by Model 5.36, allowing AD values to be calculated using (5.38). These could then be used in a like-for-like comparison with the results from the direct demand models.

$$\hat{T}_{ij} = T_i\left(e^{-\beta V(j-1)} - e^{-\beta V(j)}\right) \tag{5.37}$$

$$AD = \left(\sum_{ij}^{n}\left|\frac{\hat{T}_{ij} - T_{ij}}{T_{ij}}\right|\right)/n \tag{5.38}$$

The AD values given in Table 5.29 compare unfavourably to an AD value of 0.490 for the best direct demand model calibrated on this dataset (5.27). However, Model 5.27 is based on dummy variables and is therefore of little use for forecasting flows from new stations, in contrast to Model 5.36 which does not contain station-specific variables. The intervening opportunities trip distribution model therefore seemed worthy of further investigation. Using Rhymney line trips as the attraction factor gave inferior results to using total trips, and therefore the latter method was retained, and as expected prediction errors were reduced by using a single central Cardiff destination.

The form of the generalised cost function in Model 5.36 was unsatisfactory, as ideally rail fares and service frequencies and competition with other modes should be represented alongside rail journey time. However, there was no straightforward way to incorporate them in the cost function because the relative importance of each of these variables was unknown. Parameter values can not be determined during model calibration, as the impedance measure is only used to rank destinations and its absolute values do not appear in the model. However, an improvement might be given by using the PDFH definition of generalised journey time (5.39) as the cost function rather than rail journey time alone.

$$GJT = J + S + I \tag{5.39}$$

Where:

J is the total station-to-station journey time (including interchange time)

S is the service interval penalty

I is the sum of the interchange penalties for any interchange required

The service interval penalty converts the train headway into an equivalent time effect, and recommended penalties for different service intervals are given in the PDFH (ATOC, 2002). The interchange penalty converts the need to interchange into a time effect, with recommended values again given in the PDFH. GJT values were calculated for the 232 flows in the Rhymney line dataset and Model 5.36 was recalibrated using these as the cost function. Total destination trips were used as the attraction factor as these gave better results previously and the results of the recalibration are summarised in Table 5.30.

Table 5.30: Summarised results from calibration of Model 5.36 with destinations ranked using PDFH-based GJT

Cardiff destinations		Separate	Combined
β parameter	Value	0.00360	0.00448
	Std error	0.00022	0.00027
	95%CI Lower Bound	0.00317	0.00394
	95%CI Upper Bound	0.00404	0.00502
R^2		0.872	0.902
AD		1.334	1.097

While the R^2 values were not comparable between the two sets of calibrations because the dependent variables were not identical, using a more sophisticated measure of travel impedance appears to make little difference to model fit. This was because the difference between the two measures was not sufficient to cause many changes in the ranking of destinations, and absolute values of travel impedance are not considered during model calibration. However, the PDFH definition of GJT should be preferred to the simple measure of rail journey time because it incorporates more variables in the model.

Using rail GJT as the travel impedance measure did not take account of intermodal competition and an alternative measure (5.40) was therefore tested, which combined a journey time index with the road distance between the origin and destination. However, using this distance as a multiplier may lead to problems, with some destinations ranked highly for particular origins because they are close together by road, even though the rail journey from the origin to the destination involves a comparatively roundabout route. This could result in overprediction of demand on these flows, and a further impedance measure (5.41) which used rail distance as the multiplier was therefore tested. Table 5.31 summarises the results of recalibrating Model 5.36 using both impedance measures. While using rail distance as the multiplier gave a better model fit than using road distance, the results were still inferior to those obtained when rail GJT was used as the impedance measure, and this was therefore retained as the preferred measure. The inferior results obtained from the journey time index may occur because this model distributes trips made by people who have already decided to use rail, when in reality modal choice and destination choice decisions may be made simultaneously.

$$I_{ij} = D_{ij}\left(\frac{r_{GJTij}}{r_{GJTij} + c_{tij}}\right) \tag{5.40}$$

$$I_{ij} = DRl_{ij}\left(\frac{r_{GJTij}}{r_{GJTij} + c_{tij}}\right) \tag{5.41}$$

Where:

I_{ij} is the travel impedance measure from station i to station j

D_{ij} is the road distance from station i to station j

DRl_{ij} is the rail distance from station i to station j

r_{GJTij} is the rail generalised journey time from station i to station j (defined by 5.39)

c_{tij} is the car journey time from station i to station j

Table 5.31: Summarised results from calibration of Model 5.36 with destinations ranked using (5.40) and (5.41)

Impedance measure		5.40		5.41	
Cardiff destinations		Separate	Combined	Separate	Combined
β parameter	Value	0.00330	0.00403	0.00387	0.00433
	Std error	0.00021	0.00025	0.00023	0.00026
	95%CI Lower Bound	0.00289	0.00353	0.00341	0.00038
	95%CI Upper Bound	0.00371	0.00453	0.00432	0.00048
R^2		0.865	0.897	0.879	0.903
AD		1.534	1.325	1.370	1.187

The probability-based models described above used an attraction factor based on actual travel behaviour, but this introduced an element of circularity into the models. Attraction factors based on demographic variables were therefore tested as an alternative. This is only really necessary where predictions between two new stations are required, as usually data on actual trip totals would be available for the destination. The most obvious demographic attraction factor to use was the total number of jobs within the catchment of the destination station. The destination catchments were defined by allocating output areas to their nearest station in terms of access time. The number of jobs was divided by 1000 to prevent the exponent values exceeding the maximum allowed by the calibration procedure, and as before the parameter value was initially set to 0.0001. It may be necessary to apply a distance cut-off point when allocating jobs to destinations, as otherwise jobs located several kilometres from their nearest station will carry the same weight in the attraction factor as jobs located immediately adjacent to their nearest station. Model 5.36 was therefore recalibrated with the attraction factor defined firstly as the number of jobs for which destination j was the nearest station, and secondly as the number of jobs within 4 minutes estimated drive time of destination j for which destination j was the nearest station. The results of these calibrations are summarised in Table 5.32.

Table 5.32: Summarised results from calibration of Model 5.36 with employment-based attraction factors

Employment cut-off time		None		4 minutes	
Cardiff destinations		Separate	Combined	Separate	Combined
β parameter	Value	0.00962	0.01085	0.01899	0.02295
	Std error	0.00068	0.00063	0.00122	0.00122
	95%CI Lower Bound	0.00828	0.00960	0.01660	0.02055
	95%CI Upper Bound	0.01097	0.01210	0.02139	0.02535
R^2		0.826	0.881	0.866	0.915
AD		2.812	2.612	1.960	1.581

Using a 4 minute cut-off when defining destination catchments gave better results both in terms of R^2 and of AD. However, the AD values were still inferior to those obtained when observed total trip destinations were used as the attraction factor. Given that the results from the OD survey suggested that the majority of passengers would reach their destination on foot (see Section 5.3) the use of a drive time based boundary may have been unrealistic. Model 5.36 was therefore recalibrated with the attraction factor defined as the number of jobs within 3 km of destination *j* (the recommended arbitrary boundary from Section 5.3) for which destination *j* was the nearest station. Given the problems of catchment definition using census output areas for some stations, particularly those in central Cardiff (see Figure 5.10), a further attraction factor was tested using the number of jobs within 1 km of the destination station, regardless of whether these jobs were closer to another station. This meant that some jobs would be allocated to more than one stations, but could give a better indication of the relative attractiveness of different destinations. The results of calibrating Model 5.36 using both these attraction factors are summarised in Table 5.33.

Table 5.33: Summarised results from calibration of Model 5.36 with employment-based attraction factors using distance cut-off points

Employment cut-off distance		3 km		1 km	
Cardiff destinations		Separate	Combined	Separate	Combined
β parameter	Value	0.17549	0.20896	0.34638	0.45205
	Std error	0.01136	0.01117	0.02292	0.02712
	95%CI Lower Bound	0.15310	0.18695	0.30122	0.39859
	95%CI Upper Bound	0.19787	0.23098	0.39153	0.50550
R^2		0.861	0.911	0.877	0.912
AD		2.103	1.965	1.464	1.373

Using 3 km catchment boundaries to select the number of jobs included in the attraction factor gave inferior results to using a 4 minute drive time. The 1 km boundaries gave the best results of any of the employment-based destination attraction measures, but model fit was still inferior to that given when observed total trip destinations were used as the attraction measure. However, the level of employment is not the only factor which determines the attractiveness of a destination. The inclusion of a measure of population alongside a measure of employment in the attraction factor was tested to establish whether this would give a better representation of destination attractiveness. This required the inclusion of two separate attraction factors within the model, giving Model 5.42. The results from calibrating this model are summarised in Table 5.34.

$$P[V(j)|V(J)] = \frac{1 - e^{[-\beta V(j) + \gamma W(j)]}}{1 - e^{[-\beta V(J) + \gamma W(J)]}} \qquad (5.42)$$

Where:

W(j) is the accumulation of population opportunities at destination *j*

W(J) is the accumulation of population opportunities at the lowest ranked destination for the origin in question

Table 5.34: Summarised results from calibration of Model 5.42 with number of jobs and size of population within 4 minutes drive time of destination as attraction factors

Cardiff destinations		Separate	Combined
β parameter	Value	0.01857	0.02273
	Std error	0.00183	0.00177
	95%CI Lower Bound	0.01496	0.01924
	95%CI Upper Bound	0.02218	0.02622
γ parameter	Value	-0.00003	-0.00001
	Std error	-0.00009	-0.00008
	95%CI Lower Bound	-0.00021	-0.00017
	95%CI Upper Bound	0.00015	0.00014
R^2		0.866	0.915
AD		1.857	1.429

Model 5.42 gave an improvement in AD values over the calibration of Model 5.36 (the single attraction factor model) using the number of jobs within 4 minutes drive time as the attraction factor, but the AD values were noticeably inferior to those obtained from Model 5.36 when the number of jobs within 1 km was used as the attraction factor. Model 5.42 was therefore recalibrated with a 1 km cut-off used for job opportunities in combination with both a 4 minute cut-off and a 1 km cut-off for population opportunities, giving the results summarised in Table 5.35.

Table 5.35: Summarised results from calibration of Model 5.42 with number of jobs within 1 km and size of population within 1 km or 4 minutes drive time of destination as attraction factors

Population cut-off		1 km		4 minute	
Cardiff destinations		Separate	Combined	Separate	Combined
β parameter	Value	0.03612	0.05160	0.03455	0.05101
	Std error	0.00332	0.00356	0.00337	0.00366
	95%CI Lower Bound	0.02958	0.04458	0.02791	0.04380
	95%CI Upper Bound	0.04266	0.05863	0.04119	0.05822
γ parameter	Value	0.00012	0.00042	-0.0000005	0.00017
	Std error	0.00020	0.00016	0.00009	0.00008
	95%CI Lower Bound	-0.00027	0.00011	-0.00018	0.00002
	95%CI Upper Bound	0.00051	0.00074	0.00018	0.00032
R^2		0.877	0.914	0.877	0.913
AD		1.309	1.037	1.323	1.029

Model 5.42 gave lower AD values with all these attraction factors than the best variant of Model 5.36 with a job-based attraction factor. The AD values were also lower than those obtained when Model 5.36 was calibrated using the best attraction factor based on observed total trip destinations, meaning that this variant of Model 5.42 should be adopted as the preferred model form. The 1 km population cut-off gave slightly better results with

separate Cardiff destinations, while the 4 minute cut-off gave better results with a single central Cardiff destination. However, the 4 minute cut-off point was more consistent with other models developed in this study and was therefore chosen as the preferred option.

The transferability of this variant of Model 5.42 was tested by using it to predict destination choice for the top 95% of flows from 13 stations on the Merthyr Tydfil line, a dataset of 279 flows. This gave an AD value of 1.786 using separate central Cardiff destinations, and 1.943 using a single Cardiff destination, indicating that the model did not perform as well predicting trip distribution on the Merthyr Tydfil line as it did on the Rhymney line. To test whether this was the result of differences in the demand characteristics of the two lines (requiring model recalibration) or because the model form was not suitable for the Merthyr Tydfil line, the model was recalibrated on the Merthyr Tydfil line flows and also on a combined dataset containing the flows from both lines. The results of these recalibrations are summarised in Table 5.36.

Table 5.36: Summarised results from calibration of Model 5.42 with 4 minute population cut-off on Merthyr Tydfil line and Merthyr and Rhymney lines

Calibration flows		Merthyr Tydfil line		Merthyr Tydfil and Rhymney lines	
Cardiff destinations		Separate	Combined	Separate	Combined
β parameter	Value	0.01877	0.02588	0.02487	0.03690
	Std error	0.00159	0.00197	0.00157	0.00204
	95%CI Lower Bound	0.01565	0.02201	0.02178	0.03289
	95%CI Upper Bound	0.02189	0.02975	0.02796	0.04090
γ parameter	Value	0.00006	0.00018	0.00013	0.00026
	Std error	-0.00005	0.00004	0.00004	0.00003
	95%CI Lower Bound	-0.00003	0.00010	0.00005	0.00020
	95%CI Upper Bound	0.00016	0.00027	0.00021	0.00033
R^2		0.804	0.814	0.828	0.853
AD		1.967	1.717	1.687	1.465

This suggests that the model calibrated on the combined dataset is more effective at predicting flows on the Merthyr Tydfil line than the model calibrated on that line alone. Because the results from the recalibration are reasonable it seems that the model form is transferable, but local calibrations may be less transferable. To be sure of accurate results it would be sensible to recalibrate the model for the area where forecasts are required if suitable data is available. In general Model 5.42 is effective at predicting destination choice and, unlike the earlier loglinear regression models, can account for the effect of intervening opportunities on destination choice and constrains the total number of trips distributed to match the total observed or predicted. This means that it should, when used in conjunction with trip end models, provide a means of accurately predicting travel patterns from new local railway stations. However, its fit (measured by AD) appears inferior to that of the best loglinear regression models, which meant that Model 5.42 could not be adopted as the sole preferred flow level model. Further analysis of the results from the best models of each general form was required, and this is described in Section 7.3.

5.5 Destination Selection

An issue which affects all such flow-level models but which has not so far been addressed is the question of how to select the set of destinations to which travel is forecast. It is obviously not feasible to model travel to all 2519 stations on the rail network, but no procedure is available to identify a suitable subset. For example, stations will exist which are closer to the origin stations included in the calibration dataset for the intervening opportunity models (and therefore of higher rank) than some of the stations which make up the top 95% of flows, but which have very few trips to them from the origin. These stations should arguably be added to the calibration dataset, but this would increase its size significantly given that London Paddington was included as a destination for some of the origins. The inclusion of such additional stations would greatly increase the time necessary to compile the calibration dataset, with no guarantee of an improvement in model fit, and this expansion was not therefore pursued. While the question of which destinations should be modelled when predicting flows from a new station remains unanswered, it should in most cases be straightforward to identify a set of likely destinations for new local railway stations by investigating travel patterns at existing stations in the same area. The issues which arose when the identification of such a set of destinations was attempted will be discussed later in Section 7.4.2.

5.6 Conclusions

This chapter has described the development of a range of flow level models, for various subsets of a case study area centred on Cardiff in South-East Wales, and the results of a survey investigating rail passengers' ultimate trip origins and destinations. Loglinear and linear direct demand models were initially tested, based on the generalised model form described in the literature review. Loglinear models were found to give superior results because of their ability to capture interrelationships between the model variables. The results given by representing origins and destinations using dummy variables, total trip entry and exit variables and generalised variables such as catchment population were compared. Dummy variable models were found to give the best model fit, but are not suitable for modelling demand from new stations.

Attempts were made to incorporate intermodal competition in the direct demand models and variables representing competition with car travel were successfully included. However, it did not prove possible to develop a model form which accounted for competition with bus travel without other variables becoming insignificant, probably because of major imperfections in the available data on bus journey times. Model residuals were again mapped to allow spatial patterns in model accuracy to be identified.

Detailed consideration was given to the issue of station catchment definition for direct demand models which represent trip origins using generalised variables. Flow-specific catchments were developed by allocating census output areas to origin stations by minimising overall journey time to the destination station. This is believed to be the first time that such a methodology has been tested for aggregate local rail demand models at such a detailed level. Weighting of population units within these flow-specific catchments was tested but gave inferior results to the unweighted aggregation of units, and problems were encountered when calculating road journey times. While appearing to be more realistic than generalised catchments these flow-specific catchments did not bring an improvement in model fit and this, together with the significant increase in processing time over generalised catchments, meant that they were not persevered with in this work.

A problem with direct demand models is that the sum of flows forecast from particular origin stations may be much larger (or sometimes smaller) than the total trips either observed or forecast by trip end models at that station. Forecasts from direct demand models for Rhymney line stations were scaled so that they would sum to give the observed total of trip origins at each origin station. However, these scaled predictions had a worse fit than the unscaled predictions, indicating that this methodology did not provide an ideal solution to the problem of constraining flow level predictions.

A number of direct demand model forms were tested in an attempt to incorporate the presence of intervening and competing opportunities in direct demand models. However, it did not prove possible to obtain significant and correct sign intervening opportunities parameters without the majority of other parameters becoming insignificant. The preferred direct demand model forms from this study are therefore the dummy variable Model 5.4

and the generalised origin variable Model 5.20. If full train frequency data is not available then these should be replaced by Model 5.26 and Models 5.28/5.29 respectively.

To allow the accuracy of the various theoretical catchment definition methods to be assessed, a survey of ultimate passenger origins and destinations was carried out on the Rhymney line. This found that even the best theoretical catchments only included between 62% and 69% of observed trip ends, depending on whether central Cardiff stations were included. Observed catchments were produced using GIS spatial interpolation methods and these allowed station access patterns to be visualised in a new and attractive manner. Observations were disaggregated by mode and this showed that unsurprisingly average access/egress distances are much shorter for walk trips than for bus or car trips. The survey indicates that if arbitrary catchment boundaries are used in modelling then a 3 km boundary should be assumed for walk trips, with a 10 or 12.5 km boundary necessary for motorised trips. These boundaries are much larger than those recommended in the PDFH, but there is likely to be significant overlap between such catchments. It is important that such arbitrary boundaries should be based on road network distance rather than straight line distance as geographical barriers to station access will often affect the shape of station catchments, particularly in hilly areas such as the South Wales valleys. Generalising the features of observed catchments to allow theoretical catchments to be enhanced proved difficult, and while several possible enhancements were identified constraints on time and resources meant that they were not investigated further in this study.

Because it had not proved possible to incorporate the effect of intervening opportunities in direct demand models, nonlinear IOTD models were calibrated, a methodology which had not previously been applied to rail demand modelling. Several impedance functions were tested, with rail generalised journey time found to give the best model fit. Both population and employment around the destination station were incorporated in the attraction factor of the best model, Model 5.42. The spatial transferability of this model was tested with reasonable results, although local recalibration is advisable. This model seems effective at predicting destination choice while accounting for the presence of intervening opportunities and constrains predicted flow sizes to match the total trips observed and predicted. However, because it relies on forecasts from trip end models to produce absolute forecasts of flow sizes (rather than probabilities), it is necessary to compare the accuracy of these combined forecasts with those from the best direct demand models before deciding upon a preferred method. Such a comparison is described in Chapter 7.

Blainey & Preston (2009b) give an overview of much of this work on flow level models.

The final section of this chapter briefly considered the issue of selecting a set of suitable destinations for new origin stations. This is not a straightforward process, although in many cases it should be reasonably simple to identify a set of likely destinations by examining flows from nearby stations.

Chapter Six: Site Search Procedure

6.1 Introduction

Chapters 4 and 5 outlined the development of models which can effectively forecast both the total trips generated at new local railway stations, and the destinations of these trips. However, such techniques can only be used once potential locations for stations have been established. A semi-automated search procedure for new station sites was therefore developed, based on that outlined by Preston (1987), and this procedure is described in this chapter. Section 6.2 describes how population units were isolated which are close to railway lines but are not adequately served by existing stations and which therefore form potential targets for new station catchments. In Section 6.3 clusters of such output areas which indicate promising locations for new stations were identified. These locations are then optimised in Section 6.4 so that as many people and jobs as possible fall within the station catchments. Finally, methods for easily assessing the feasibility of station construction are described in Section 6.5, and the procedure is summarised in Section 6.6.

6.2 Isolation of target population units

The first stage in this procedure was to define the geographical area of interest, which in this case was the whole of England and Wales, to correspond with the case study area for the most successful trip end models. All census output areas within this area were then allocated using RouteFinder to their nearest station in terms of access time. Output areas whose centroid was within an acceptable distance of an existing station could then be automatically removed from the dataset using the GIS. The value of this maximum 'acceptable' distance was not easy to determine but 4 minutes drive time was used here, based on the results from the work on catchment areas for trip end and flow level models in Chapters 4 and 5. While the OD survey indicated that many stations draw passengers from a wider catchment than this (see Section 5.3), the spacing of existing stations in many areas suggests that distances greater than this are sub-optimal. Indeed, if the site search procedure was being applied to urban areas where a higher station density and shorter access times/distances were the norm (e.g. London) this maximum acceptable distance would probably need to be reduced further.

A 2 km buffer zone was then drawn around all existing railway lines using ArcMap.

Again, the width of this buffer zone could be varied, but a 2 km cut-off seemed likely to include all communities close enough to a railway line to be effectively served by a station. All output areas which are not within 4 minutes of an existing station but which were outside the buffer zone were then removed from the dataset using a filter. The output areas which remained in the dataset at this point were those which could potentially be served by new stations on existing railway routes. Such output areas in the region around Leeds are mapped in Figure 6.1.

Figure 6.1: Output areas around Leeds within 2 km of a railway line but more than 4 minutes from existing stations

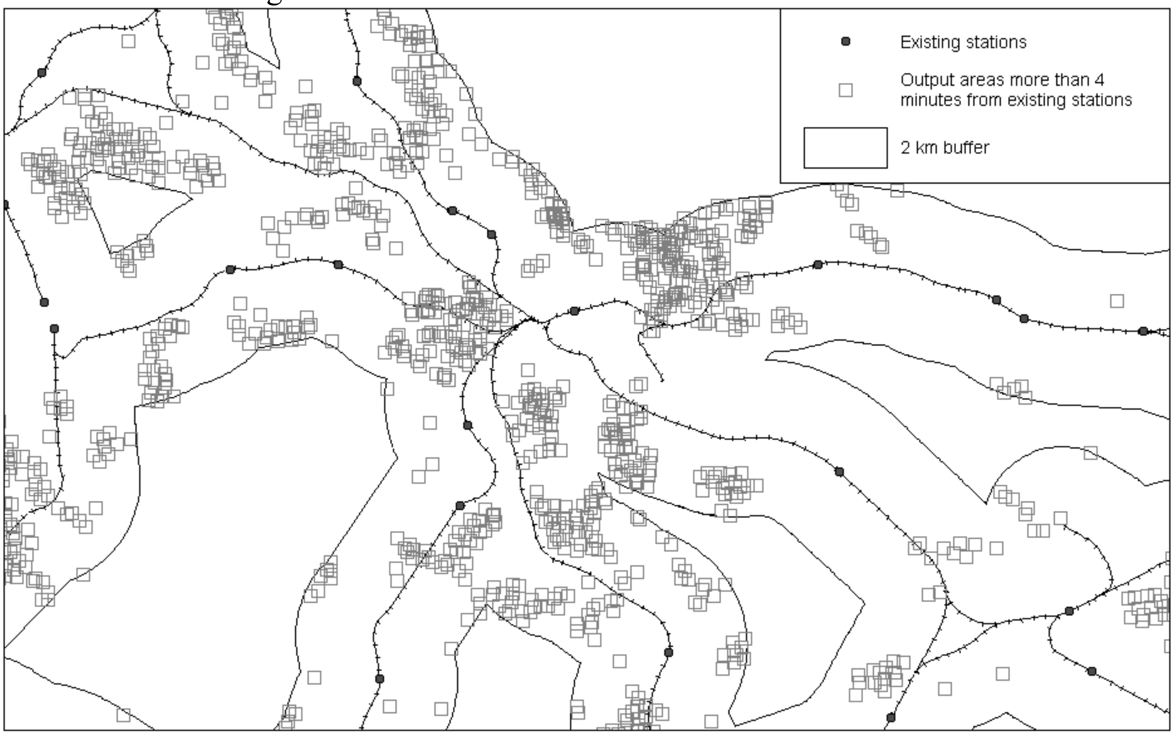

6.3 Cluster Identification

The next step in the procedure was to identify clusters of output areas which might indicate possible sites for a new station. The use of the ArcMap Point Density tool to automatically identify such clusters was investigated, but the computer used did not have enough memory to process the data at a fine enough scale. Clusters were therefore identified manually, but displaying the output areas in the form shown in Figure 6.1 did not give enough information for such a procedure to be reliable, as the output areas represented by the centroids varied greatly in size and also in the populations and employment sites that they contained. Population and employment density for each of the output areas was therefore calculated, and MapInfo was used to represent these in chloropleth maps. An example is given by Figure 6.2 which shows population density for the area around Leeds.

Figure 6.2: Population density for output areas around Leeds within 2 km of a railway line but more than 4 minutes from an existing station

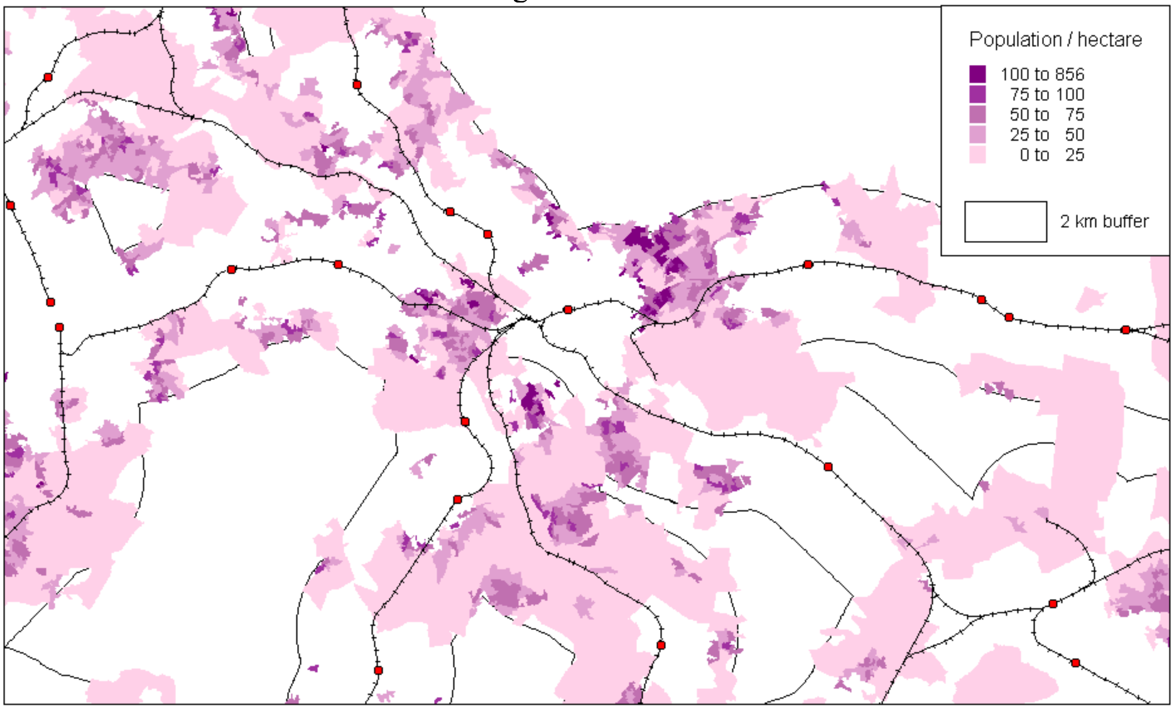

From these chloropleth maps it was easy to visually identify the areas of high population and employment density near existing railway lines which are not served well by existing stations. These show up as darker coloured areas in the maps, and in general the larger these areas were, the greater the potential market. These areas were marked on the maps with circles of 2 km radius, as illustrated in Figure 6.3. While this section of the procedure could not be automated, manual identification of these areas was relatively quick, requiring less than a day's work for the whole of England and Wales.

Figure 6.3: Potential areas for new stations in West Yorkshire

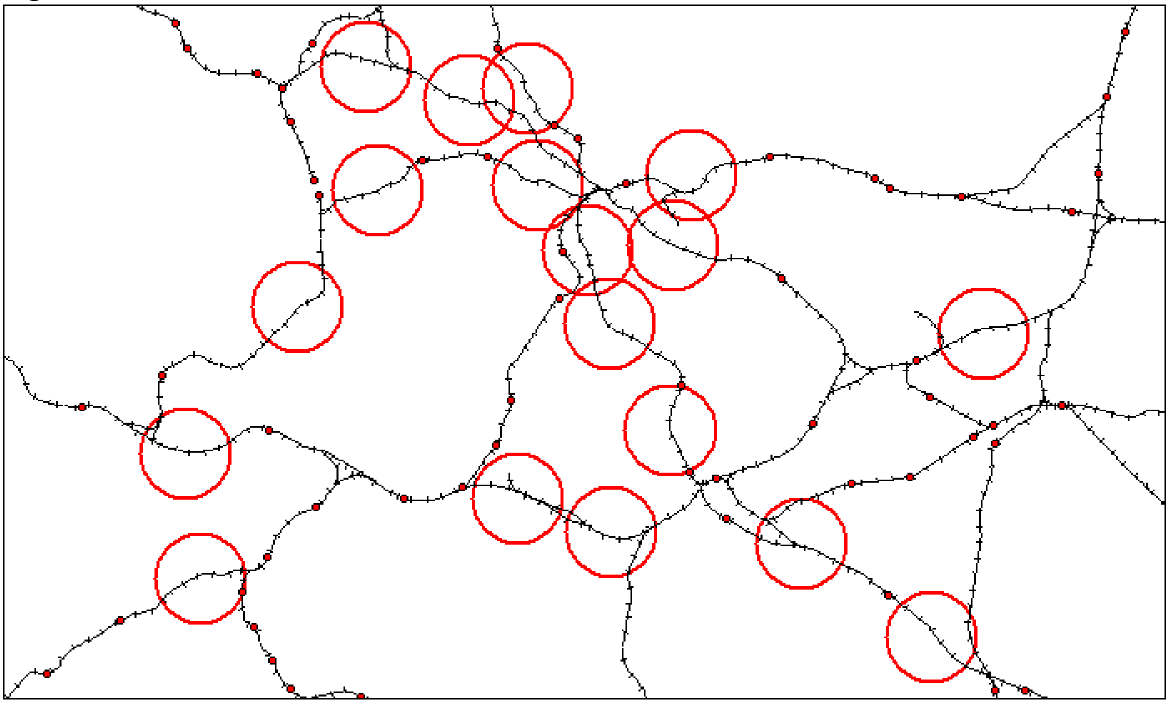

In total these two searches identified 421 possible areas for new station construction in England and Wales. The procedure had not so far considered the feasibility of access to the potential sites, and therefore Google Maps was used together with OS Meridian data in MapInfo to identify more precise point locations for these stations based on the availability of access to the road network. While this element of the procedure again required manual processing, the use of GIS meant that this was a relatively quick process, and was accomplished in two days for this extremely large case study area. The point locations for all 421 potential stations are shown in Figure 6.4, with the red points indicating sites on lines with a passenger service, and the orange points indicating sites on freight only lines.

A small number of the sites shown in Figure 6.4 approximately coincide with sites identified by ATOC (2009). However, the majority do not, as the ATOC report concentrated on new stations away from existing passenger railway routes, and those sites which do coincide are those on existing freight-only rail routes. This means that the results from this search procedure and from the ATOC (2009) report are largely complementary. The procedure outlined here allows more precise sites to be identified, and could be extended to cover the additional locations identified by ATOC. However, it would be desirable to identify potential 'corridors' for new lines before using this procedure for sites away from existing lines, as otherwise the number of possible sites identified would become unmanageable.

Figure 6.4: Possible locations for new stations in England and Wales

6.4 Optimisation of Catchments

Once approximate sites for new stations had been identified the next step was to define the station catchment areas. The ArcMap Network Analyst tool was used to automatically define catchments around the station sites based on drive-time bands. While as described above and in Section 5.2 the exact extent of station catchments is still open to debate, it seems sensible to locate new stations so that the maximum possible population is located within the minimum possible distance. Three access time bands were therefore used, with travel time boundaries of 2, 4 and 6 minutes from the station. Non-overlapping catchments were specified, although there is perhaps an argument for using overlapping catchments in this instance to investigate the level of 'catchment competition' between new stations and existing stations. The catchments given by Network Analyst for both the proposed and existing stations in the South-East Wales area are illustrated in Figure 6.5.

Figure 6.5: Catchments for proposed and existing stations in South-East Wales

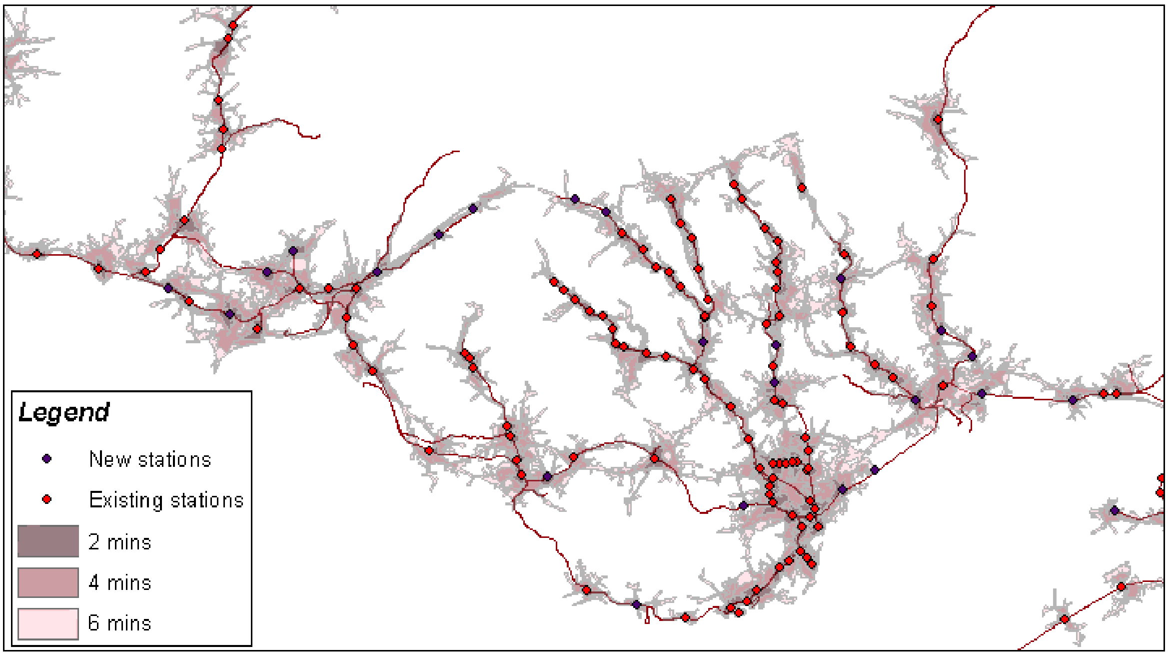

The scale of Figure 6.5 makes it difficult to see the details of individual catchments, but the next stage of the search procedure was to examine the catchments for the proposed station sites to establish the extent to which they incorporate the previously unserved output areas and the extent to which they overlap with the catchments of existing stations. In some cases the stations appeared to have been sited well, with their catchments incorporating the majority of relevant output areas, and complementing rather than competing with neighbouring existing stations, with Glyncoch a good example of this (see Figure 6.6).

Figure 6.6: Catchments for potential site at Glyncoch and neighbouring stations

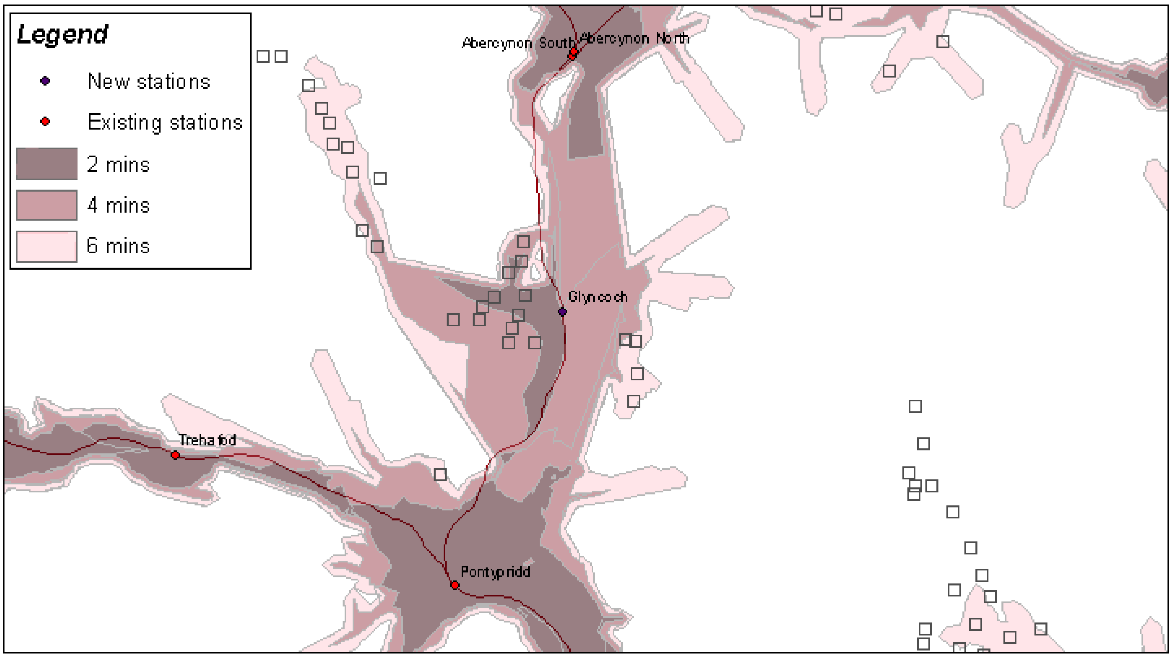

However, for some other sites catchments appeared to overlap excessively with those of other stations, or to be of an unexpected shape, such as at St Mellons (see Figure 6.7).

Figure 6.7: Catchments for potential site at St Mellons and neighbouring stations

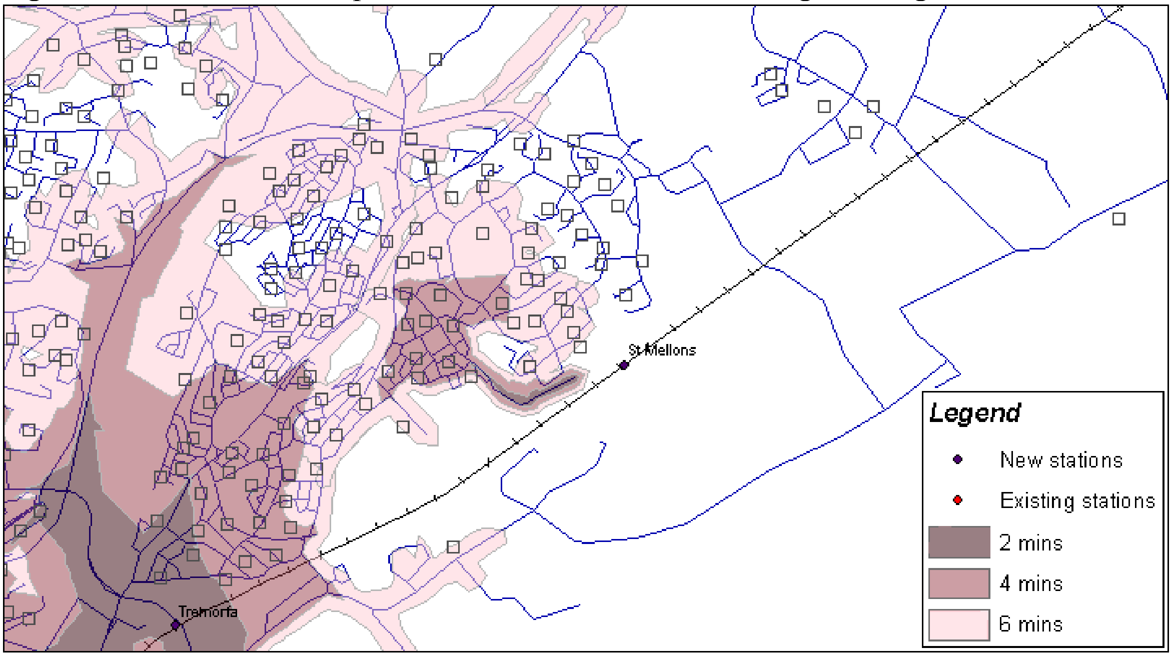

The first problem could usually be resolved by moving the proposed station site to reduce overlap. The latter problem resulted from the automatic allocation of the station to the closest link on the road network, which could mean that sites were allocated to minor roads with a circuitous link to the rest of the road network, even if a road with much greater connectivity was almost as close to the site. When the station was manually linked to this road instead of the closest road the resulting catchments appeared much more realistic, as shown in Figure 6.8. Similar corrections were made where necessary for all of the other potential station sites identified in England and Wales. While this required manual examination of the catchments, this procedure was again speeded by the use of GIS, and took less than a week for the complete set of 421 sites.

Figure 6.8: Catchments for potential site at St Mellons and neighbouring stations with site linked to road with higher connectivity

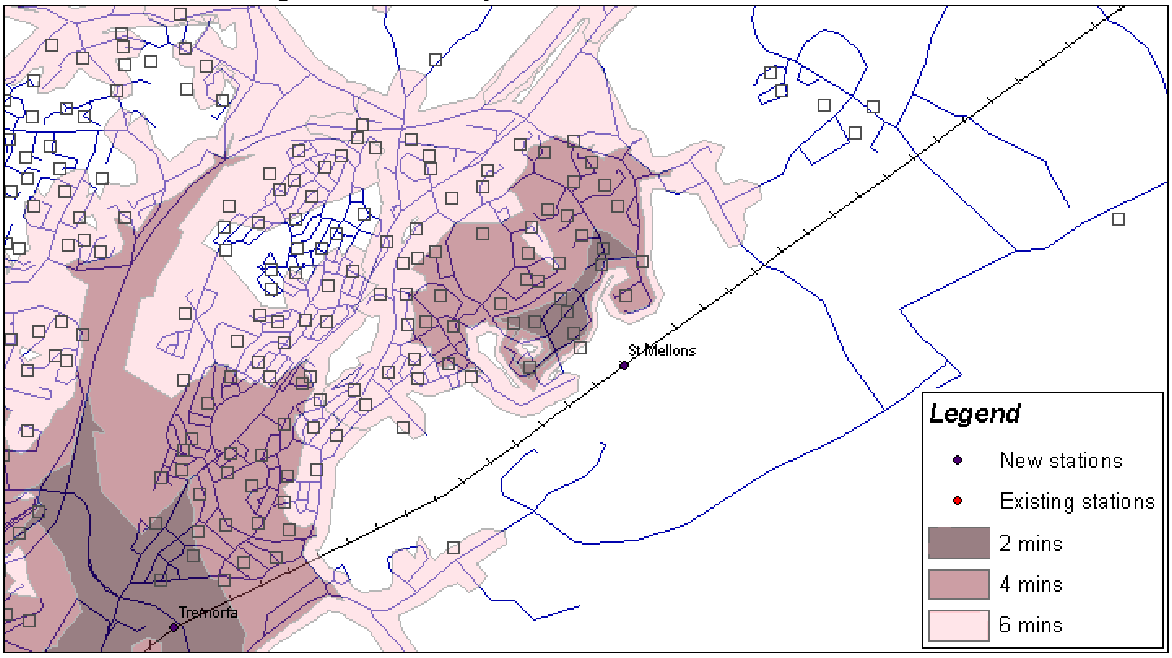

6.5 Feasibility of Construction

Once the potential sites had been relocated to optimise catchment areas, it was necessary to check the feasibility of station construction at these sites. This is an important issue, as engineering constraints could make construction costs prohibitively high, but there was no straightforward way to provide an automated solution. The Ordnance Survey Meridian data used to represent the railway network was not detailed enough to show features such as viaducts or embankments. While OS Mastermap data shows such features in detail, as described in Section 3.6.4.2 the file sizes involved were too large to allow this data to be used for anything other than a small case study area. However, if serious consideration was being given to station construction at a particular site, this data would enable the feasibility of construction to be easily assessed. Figure 6.9 gives an overview of the area around the potential station site at Glyncoch, showing the level of detail available with Mastermap.

Figure 6.9: Mastermap map of area around potential station at Glyncoch

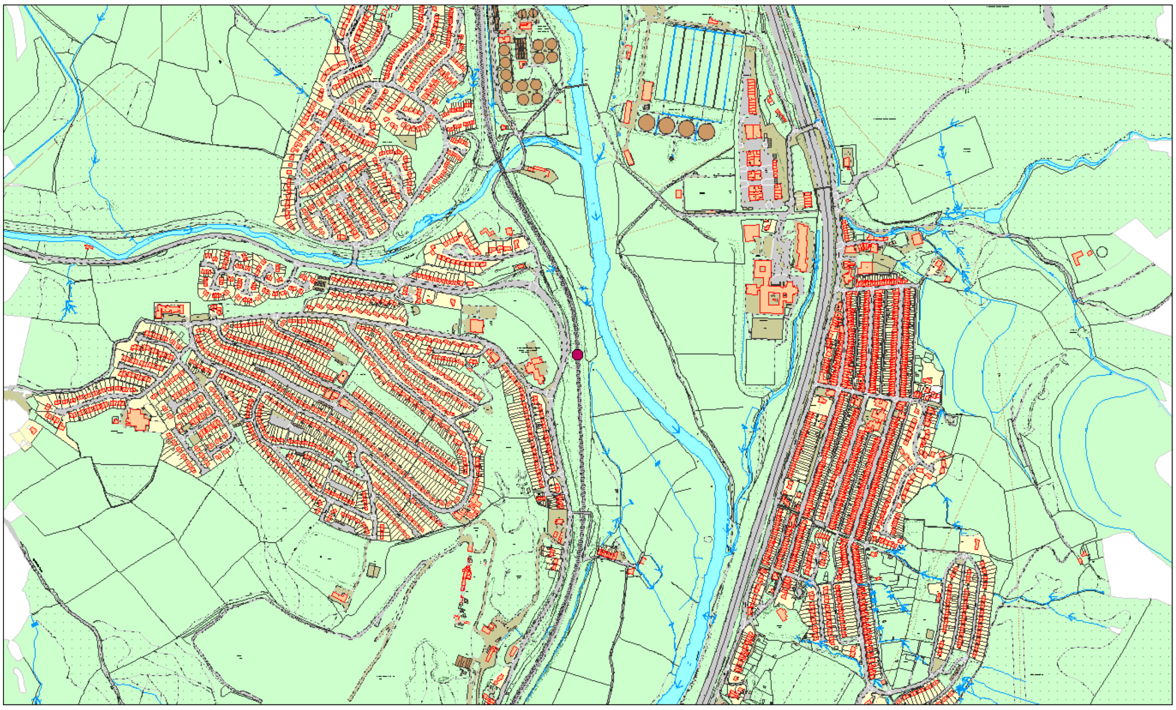

The other possibility for assessing site feasibility was to use the OS raster data which is available in three scales in TIFF format from Digimap. This data is effectively an electronic image of the paper maps available from the Ordnance Survey. It therefore provides a visually clear representation of the area covered, but the raster format means that it is impossible to isolate map objects. This makes it difficult to use for automated analysis in a GIS, but it can still be used to assess construction constraints through visual inspection. This is illustrated by Figure 6.10 which depicts the area around the potential site for Glyncoch station using the largest available scale of raster data (it should be noted that the map is not reproduced at the original scale here). An alternative source of raster data was available in the form of aerial imagery from Google Earth, as shown in Figure 6.11. As the raster format of both the OS and GoogleEarth data makes automated analysis in a GIS difficult, the best solution is therefore to use one or more of these detailed data formats to manually identify any construction constraints once demand models have been used to identify the most promising new station sites.

Figure 6.10: 1:10000 scale raster based map of area around potential station at Glyncoch

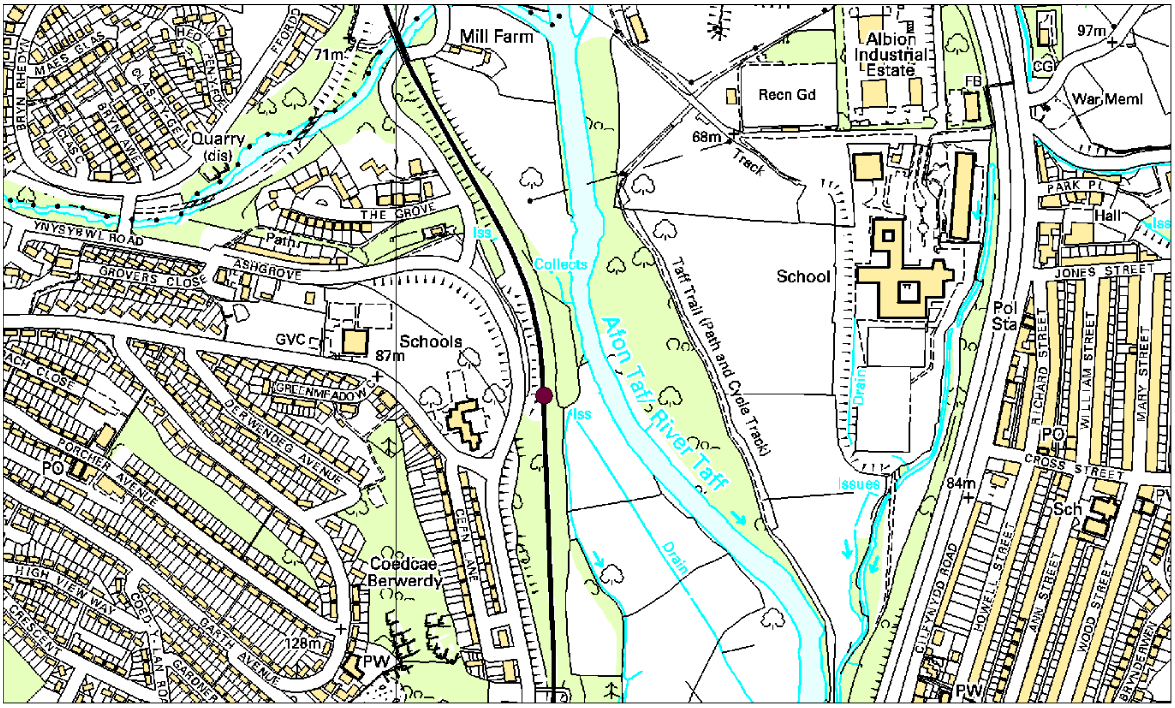

Figure 6.11: Aerial image from GoogleEarth of area around potential station at Glyncoch

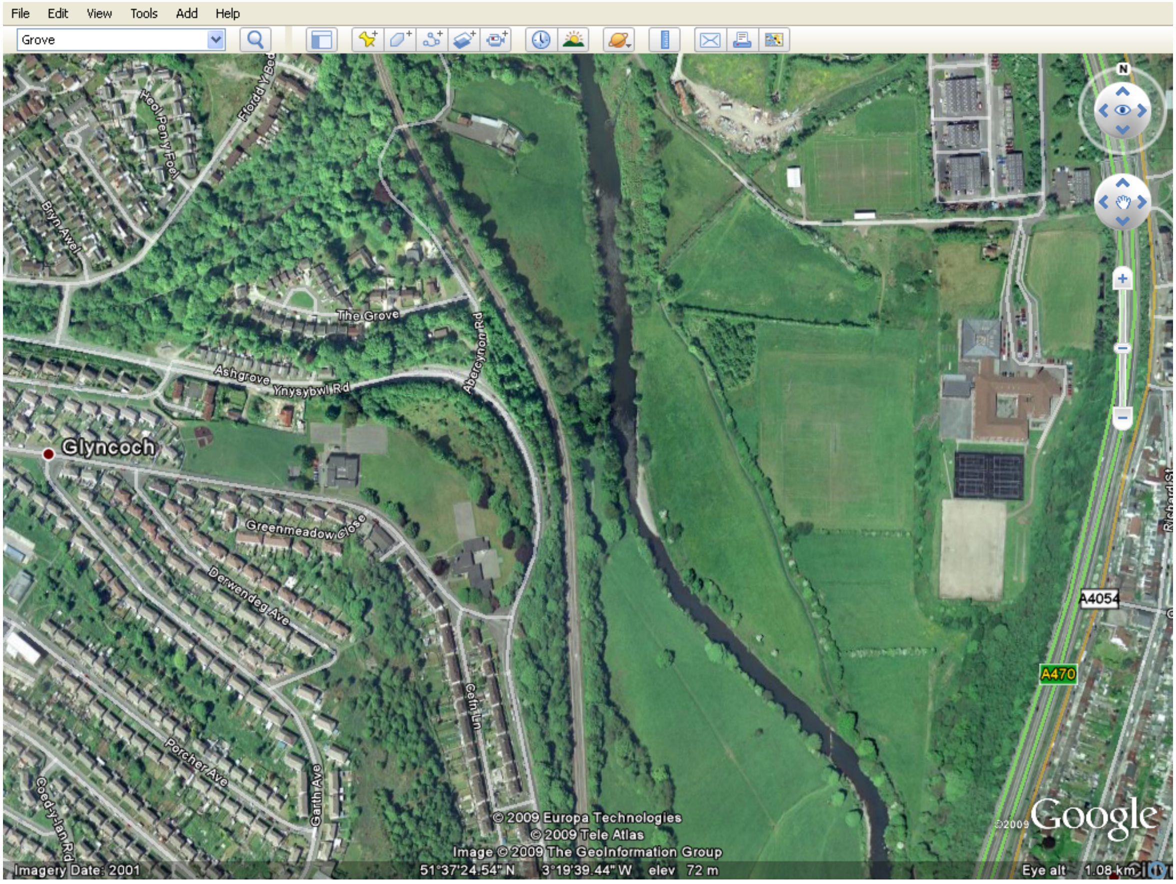

6.6 Conclusions

This chapter has described the development of a semi-automated procedure for establishing the location of new station sites. While the basic steps in such procedures have been outlined in the past (notably by Preston (1987)) this is the first time that GIS have been used to implement a procedure of this type and the first application of a site search procedure over such a large area. Although ATOC have produced a report identifying settlements which could be served by new stations on new or rebuilt lines (ATOC, 2009) since this application of the procedure developed for this study was first presented (to the Passenger Demand Forecasting Council in February 2009), the procedure used does not identify precise sites, unlike the methodology described here.

GIS were used to isolate census output areas which are within 2 km of existing railway lines but which are not within an acceptable access distance of an existing station. These output areas are those which could potentially be served by new stations on existing railway routes. Clusters of output areas with a high population or employment density were then identified and marked on the GIS maps, as these clusters indicated potential sites for new stations. 421 such sites were identified in a search of all railway lines in England and Wales. Next, non-overlapping catchments for these stations were defined using GIS and these were checked to establish the extent to which they overlapped with the catchments of existing stations and to which they included the target output areas. If necessary the proposed stations were relocated to optimise their catchments. Finally, a GIS-based methodology for checking the feasibility of station construction at the sites identified was described.

Demand forecasts for each of the potential sites for new stations identified by this procedure will be produced in Chapter 7 using the demand models developed in earlier chapters, and an appraisal procedure will be applied to a subset of sites to assess the case for station construction.

Chapter Seven: Synthesis and Appraisal

7.1 Introduction

This chapter brings together the best demand models developed in Chapters 4-5 and the site search procedure described in Chapter 6 to create a synthesised procedure for locating and forecasting demand at new station sites. The results from this procedure will then be fed into an appraisal procedure which assesses the business case for constructing these new stations.

Firstly, model transferability over both time and space are investigated in Section 7.2, with trip end models recalibrated for the South-East Wales area. Model results are then compared and the preferred modelling methodology identified in Section 7.3. Section 7.4.1 outlines the use of the preferred trip end model to predict demand at the 421 sites identified in England and Wales by the site search procedure, and these sites are then ranked by predicted demand. In Section 7.4.2 the preferred flow level models are used to make flow level predictions of demand for a subset of these sites in South-East Wales.

Section 7.5 considers the issues of demand build-up over time at new stations, of abstraction of trips from existing stations and from other modes and of the generation of non-user benefits. Section 7.6 then brings estimations of the costs associated with new stations together with the expected revenue generated as a result of the forecast demand at the stations in South-East Wales and associated user and non-user benefits to carry out both financial and social cost-benefit analyses of their construction. The results from these analyses are compared and break-even demand levels are estimated based on mean fares. The procedures described in this chapter are brought together in a spreadsheet-based appraisal tool in Section 7.7, before some conclusions are outlined in Section 7.8.

7.2 Model Transferability

7.2.1 Transferability over time

The best trip end models had previously been recalibrated using observed trip end data from several years (see Section 4.3.4) over time, and while there was some variation in the parameter estimates, the results were found to be stable overall. Therefore while recalibration on the most recent data available would seem sensible before making demand forecasts, this overall stability allowed confidence to be placed in the model form.

It was not possible to investigate the transferability of the flow-level models over time, as the data required for calibration was only available for a single time period. However, some level of temporal transferability could be achieved by scaling the model predictions to reflect the overall change in rail use between the calibration time period and the most recent time period for which trip end data was available.

7.2.2 Transferability over space

The South Wales area was chosen as the case study for developing a synthesised demand modelling and appraisal procedure, as this was the only area for which LENNON data was made available for this study and hence the only area for which flow-level models could be calibrated. The trip end models were therefore recalibrated for this area, allowing investigation of their transferability over space. Recalibration of these models for the South Wales area was to some extent unnecessary, as the area was already covered by the England and Wales models (see Section 4.2). However, as trip end models had been calibrated specifically for the South Hampshire area and shown to give good results, a local calibration of the best global trip end models was carried out for South Wales to allow the results to be compared with those from the England and Wales models. The GWR models were not recalibrated on the South Wales dataset as not enough data points were available to allow confidence to be placed in the model results.

The best global trip end model was Model 4.44, but this included dummy variables representing the distance from several major cities, which were not appropriate for a model calibrated on this more localised dataset. The travelcard boundary and electrification dummy variables were also unnecessary, as the dataset contained no stations in these categories. These variables were therefore removed from the model, giving Model 7.1. The results of calibrating this model on the 83 category E and F stations in South Wales used for investigations into flow level models are summarised in Table 7.1.

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \rho \ln Pk_i + \kappa Te_i \qquad (7.1)$$

Table 7.1: Summarised results from calibration of Model 7.1

	Value	t stat	z stat
Intercept	4.179	3.578	7.672
β parameter	-0.087	-0.877	-10.174
δ parameter	1.097	7.342	-6.588
λ parameter	0.204	1.388	-2.968
τ parameter	0.252	3.110	4.684
ρ parameter	0.211	4.181	2.250
κ parameter	1.718	6.022	7.757
R_{adj}^2	0.630		
AD	0.716		

The model fit was very much inferior to that from the best South Hampshire and England and Wales trip end models. Furthermore, the distance to higher category station (λ) parameter was insignificant, and the population (β) parameter was both insignificant and of the wrong sign. The z statistics also show that all the parameter values were significantly different to those from the calibration on the England and Wales dataset. Despite these problems, the AD value is relatively small, and this was compared to the equivalent values for predictions for these stations from the best global and GWR models calibrated on the England and Wales dataset (reproduced below), as shown in Table 7.2.

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \rho \ln Pk_i + \eta B_i + \kappa Te_i + \nu El_i \qquad (4.41)$$

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \rho \ln Pk_i + \upsilon_2 L_2 + \upsilon_3 L_3 + \gamma Ma_i + \ldots$$
$$\xi Bi_i + \omega Ca_i + \eta B_i + \kappa Te_i + \nu El_i \qquad (4.44)$$

Table 7.2: AD values from trip end models for South Wales dataset

Model	S Wales 7.1	Global 4.44	GWR 4.41
AD	0.716	0.992	0.767

The AD value for Model 7.1 is much smaller than that for the global calibration of Model 4.44 when applied to this subset of the dataset, and also smaller than that for the GWR calibration of Model 4.41. However, the difference with this latter value is relatively small, and given the problems with two parameter values in Model 7.1, Model 4.41 should be the preferred trip end model for forecasting purposes within this area. This is because any anomalies in the dataset at stations in South Wales would have had a much greater impact on parameter values in Model 7.1 than in Model 4.41.

7.3 Comparison of model results and preferred methodology

Identification of the best trip end models was relatively straightforward, as from Sections 4 and 7.2.2 it was clear that the GWR calibration of Model 4.41 gave the best results. However, identifying the best flow level model is more problematic, as while Chapter 5 showed that the IOTD models were better able to deal with the effects of intervening opportunities than the direct demand models, their fit as measured by AD was noticeably inferior. However, the direct demand models with the best AD values were unsuitable for forecasting demand from new stations as they relied on origin and destination dummy variables. A more valid comparison is between the best direct demand model with generalised origin variables (5.28), the best total entry model (5.29), and the best IOTD model (5.42 with 4 minute population cut-off and 1 km employment cut-off). These models are shown below, and their R_{adj}^2 and AD values when calibrated on the 232 flow Rhymney line dataset are shown in Table 7.3.

$$\hat{T}_{ij} = \alpha \left(\sum_{a} P_a w_a \right)^{\beta} J_{i4}^{\tau} Pk_i^{\rho} \prod_{j}^{n} D_j^{\gamma_j} D_{ij}^{\varpi} Rs_{ij}^{\delta} Cs_{ij}^{\kappa} H_{ij}^{\eta} Rfkm_{ij}^{\lambda} \qquad (5.28)$$

$$\hat{T}_{ij} = \alpha\, En_i^{\beta} \prod_{j}^{n} D_j^{\gamma_j} D_{ij}^{\varpi} Rs_{ij}^{\delta} Cs_{ij}^{\kappa} H_{ij}^{\eta} Rfkm_{ij}^{\lambda} \qquad (5.29)$$

$$P[V(j)|V(J)] = \frac{1 - e^{[-\beta V(j) + \gamma W(j)]}}{1 - e^{[-\beta V(J) + \gamma W(J)]}} \qquad (5.42)$$

Table 7.3: Comparison of flow level model fit for 232 flow Rhymney line dataset

Model	Central Cardiff destinations	R_{adj}^2	AD
5.28	Separate	0.678	0.892
	Combined	0.689	0.861
5.29	Separate	0.827	0.567
	Combined	0.836	0.541
5.42	Separate	0.877	1.323
	Combined	0.913	1.029

While Table 7.3 suggests that when measured by R_{adj}^2 Model 5.42 gives superior results, it is not valid to directly compare these values given that the dependent variables are not the same for all models. Model 5.29 appears to give superior results when model fit is measured by AD values, but this is not a perfect measure of model fit either. In any case, Models 5.29 and 5.42 can not be used on their own to forecast absolute flow sizes for new stations, because for Model 5.42 the predicted destination choice probabilities have to be multiplied by the total trips from the origin to give these flow sizes, and for Model 5.29 the total trips from the origin are used as a model variable. For a new station obviously no observed total trip values would be available, meaning that forecasts from trip end models would have to be used instead. The results from combining total trip forecasts from the best trip end model (GWR Model 4.41) with flow-level predictions from Model 5.42 and of replacing the observed total trip origins in Model 5.29 with the predicted total trip origins from Model 4.41 are shown in Table 7.4.

Table 7.4: Results from combining trip end and flow level models to forecast flow level demand for Rhymney line

Model	Central Cardiff destinations	AD
5.29	Separate	0.917
	Combined	0.981
5.42	Separate	1.551
	Combined	1.462

These indicate that Model 5.29 is clearly superior in terms of accuracy compared to the combination of Models 4.41 and 5.42, and while its fit as measured by AD is slightly inferior to that of Model 5.28 more confidence can be placed in the form of Model 5.29 because the population (β) parameter in Model 5.28 is of the wrong sign. However, because AD has shortcomings as a measure of model fit, it seemed sensible to make two other tests of the relative performance of Models 5.29 and 5.42 (in combination with Model 4.41).

If a good forecast can be obtained of the total number of trips made from a new station, then arguably the ranking of particular destinations by passengers would be just as important as actual flow sizes when planning the services to be offered from the station. The accuracy of the modelling procedures in ranking destinations was therefore compared, by calculating for the 232 flow Rhymney line dataset the difference between each destination's predicted rank and observed rank, and averaging this difference over all flows. This gave the results summarised in Table 7.5, which again indicates that the combination of Models 4.41 and 5.42 gives inferior results to that of Models 4.41 and Model 5.29.

Table 7.5: Comparison of ranking accuracy for 232 flow Rhymney line dataset

Model	Central Cardiff destinations	Mean difference in rank
5.28	Separate	2.26
	Combined	2.17
5.41	Separate	3.45
	Combined	3.26

When carrying out appraisal for a new station, the forecasts of revenue generated are arguably more important than the forecasts of demand, because the business case for any new station will depend on the amount of revenue it can generate. The total revenue predicted by each of the models for the Rhymney line was therefore estimated by multiplying the demand forecast for each flow by the Standard Day Return fare for that flow. Actual revenue was also estimated by multiplying the observed number of trips on each flow by the SDR fare, and these figures are compared in Table 7.6.

Table 7.6: Comparison of actual and predicted revenue for 232 flow Rhymney line dataset

Model	Central Cardiff destinations	Predicted revenue	Actual revenue	% difference
5.29	Separate	£4,415,524.96	£8,408,247.60	-47.48%
	Combined	£5,053,608.50	£8,460,847.70	-40.30%
5.42	Separate	£5,118,237.40	£8,408,247.60	-39.13%
	Combined	£4,143,475.60	£8,460,847.70	-51.03%

This analysis suggests that all models significantly underpredict the level of revenue generated by the Rhymney line, but that there is comparatively little difference between the revenue forecasts from the two models. The level of underprediction is a cause for some concern, but comparison of the predictions with observed demand indicated that much of the underprediction occurred because the Model 4.41 forecast of demand at Caerphilly was far below the actual observed demand. It is questionable whether Model 4.41 should be used to model demand from Caerphilly, as it is in Network Rail Category D and therefore not 'local', and if this station is excluded then the cost underprediction reduces to ~25%. The figures in Table 7.6 indicate that for appraisal purposes there may be little to choose between Model 5.29 and Model 5.42 in terms of their accuracy in forecasting revenue. However, given that Model 5.29 was slightly better at ranking destinations than Model 5.28, and that several parameters in Model 5.28 had unexpected values, the combination of Model 4.41 and Model 5.29 was adopted as the preferred method of making flow-level predictions. While the total forecasts for each station will not be equivalent to those from the trip end models, the latter forecasts are still taken into account during the production of flow level forecasts through their use as values for the total entries variable.

7.4 Demand predictions and ranking of new station sites

7.4.1 Trip end predictions

Once a preferred demand forecasting methodology had been identified, only the collection of relevant data remained before forecasts could be made for the station sites identified in Section 7. Trip end forecasts were produced for all 421 sites, as the spreadsheet tool described in Section 4.4 allowed forecasts to be rapidly obtained based on Model 4.41.

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \rho \ln Pk_i + \eta B_i + \kappa Te_i + \nu El_i \quad (4.41)$$

Population and employment totals for the new stations were obtained by reallocating all census output areas within England and Wales to their nearest station site (in terms of journey time) assuming that all proposed stations were constructed. These units were then aggregated into catchments and weighted in the same way as during model calibration. Similarly, distances to the nearest higher category stations were calculated in the same way as before.

Where possible, train frequencies were assumed to be the same as those at local stations on the same route as the proposed station. However, for a number of sites no suitable service currently exists, because the potential stations are located on freight-only routes or routes where only express/intercity passenger services operate. The need to provide a new service would obviously have a negative impact on the viability of these potential stations, and they were therefore placed in a separate group. For demand forecasting purposes, it was assumed that an hourly service would operate, with a daily total of 34 trains.

The exact size of new station car parks would depend on the precise area available for this purpose at each site. However, Google Earth was used to assess approximately how much land was available at each site, with the sites then allocated to one of four car park sizes defined based on the number of spaces provided at recently opened stations, as shown in Table 7.7. Despite its lack of precision, this method should still give a reasonable indication of the demand potential at each station site.

Table 7.7: Approximated car park sizes for new stations

Land available	Parking spaces
None	0
Limited	10
Some restrictions	40
Plentiful	100

Finally, the values of the three dummy variables in Model 4.41 were set to reflect the characteristics of the station sites. Terminus stations were easily identified, as were stations which were located on electrified lines. It was assumed that none of the new sites would become travelcard boundary stations, as even if they were adjacent to existing boundary stations it is likely that the current boundaries would be retained. The only exceptions to this were proposed stations which were located on lines which cross the boundaries of travelcard areas but where no local rail service currently exists. Four such sites were located in South Yorkshire, and two in the West Midlands. Trip end forecasts could then be made for all 421 sites using the GWR calibration of Model 4.41, and Blainey (2009e) (see Appendix 1) gives full details of these forecasts and the values of the independent variables at each of the sites. The magnitude of these forecasts is summarised in Table 7.8, which gives the number of stations where demand is forecast to be at each level, both as totals and disaggregated into sites on passenger and freight routes.

Table 7.8: Summarised magnitude of trip end forecasts at 421 station sites

Forecast trips per year	Number of stations		
	Total	Passenger routes	Freight routes
0-49,999	140	105	35
50,000-99,999	137	102	35
100,000-149,999	69	53	16
150,000-199,999	33	27	6
200,000-249,999	22	18	4
250,000-299,999	8	7	1
300,000-349,999	8	7	1
350,000-399,999	3	2	1
400,000-449,999	1	1	0

Obviously those stations with the highest forecast demand levels are likely to be those with the greatest case for construction, although this will still be dependent on construction and operating costs, the revenue generated per trip and the level of abstraction from existing stations. Sites where over 200,000 trips per year are forecast are detailed in Table 7.9, with those located on freight only lines shown in italics. All these sites are mapped in Figure 7.1 to show their geographical distribution. This shows that the majority of these high demand station sites are located in the South-East of England, although there is also a cluster of such sites in West Yorkshire. They tend to be located on electrified lines offering a high service frequency, although several sites feature which would become termini if freight lines were reopened to passenger traffic.

Table 7.9: Sites where more than 200,000 trips per year are forecast

Station	Model 4.41 Trip Ends	Station	Model 4.41 Trip Ends
Rodley	410160	Woodley	236523
Rustington	386072	Vange	234083
Aldwarke	372509	Hawkwell	233333
Newbiggin-by-the-Sea	*371075*	Culcheth	232256
Burpham	341190	Werrington	230335
Corringham	331184	*Portishead*	*229753*
Mossack Hall	326461	Bewbush	229012
Kildwick	320006	*Ludgershall*	*227077*
Farnham Road (Slough)	316072	Paulsgrove	225523
Great Salterns	315627	Roffey	220558
Lydd	*315207*	South Hildenborough	217836
Shotgate	304947	Armley	215525
Beehive Road (Bracknell)	299810	*Hardway*	*212813*
Surrey Research Park	296359	Calcot	211085
Cliftonville	291765	*Wisbech*	*210977*
Apperley Bridge	289494	Lexden Heath	207481
Kendal Wood	269707	Grange	206566
Norsey Wood	267253	Sheerwater	206465
Turnford	263858	Creekmouth	203086
Leiston	*254285*	Middleton	203040
Laisterdyke	240831	North Stifford	200951

Figure 7.1: Locations of sites where more than 200,000 trips per year are forecast

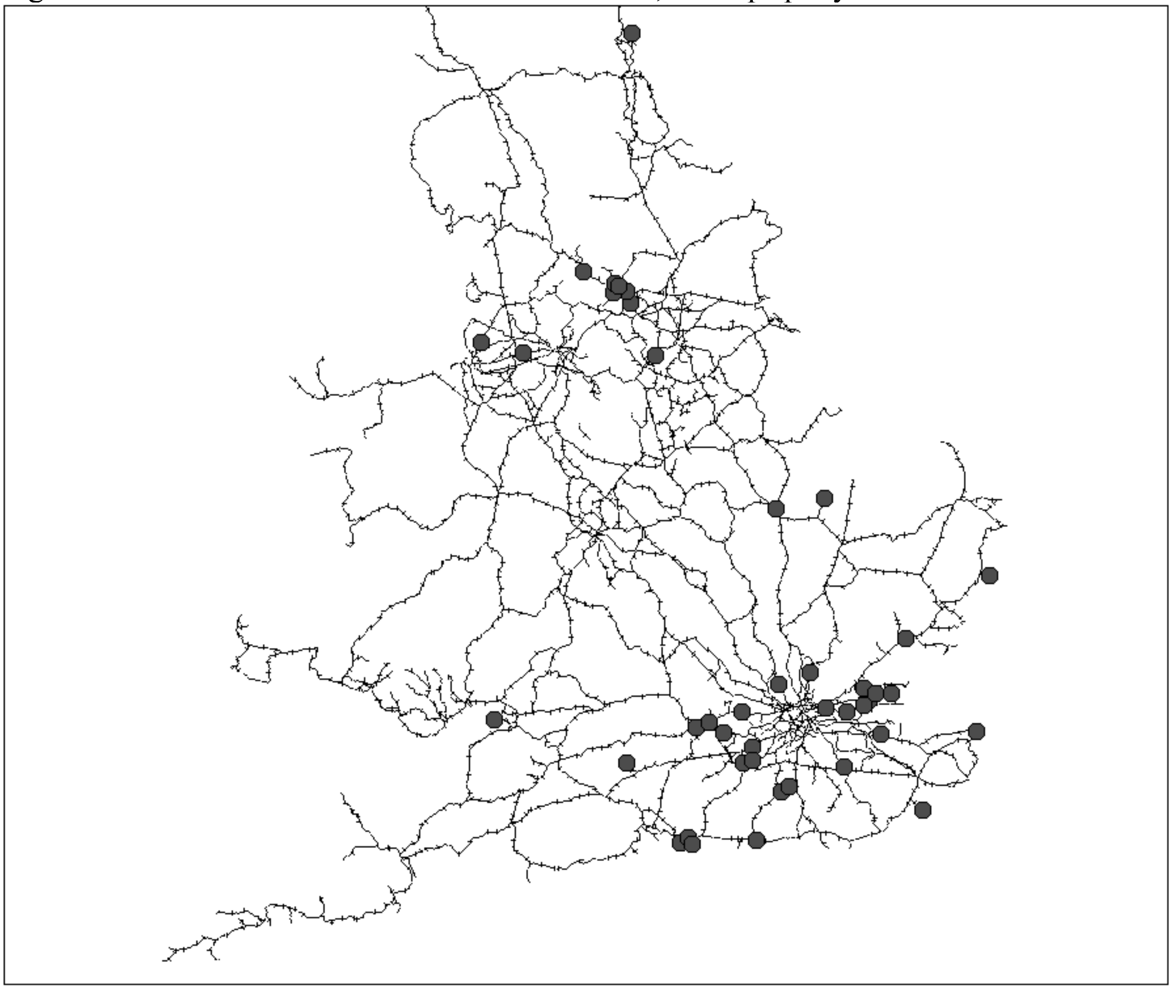

It is possible that sites with demand lower than 200,000 trips per year would still be viable for construction, as approximately 54% of existing stations in the UK have demand levels lower than this, and Table 7.10 shows that less than 200,000 trips were made in 2007-8 at half of the stations opened in the last ten years (see also Section 7.6.4). More detailed appraisal of station viability would be needed to establish which stations have the best case for construction, and a procedure to establish this is described below. However, the trip end forecasts produced here provide a quick means of checking the likely viability of a large number of potential sites for new stations, and to the author's knowledge this is the first time that such a large scale analysis of possible sites for new stations has been undertaken.

Table 7.10: Trips made in 2007-8 from stations opened since 1999

Station	Opening Date	2007-8 Trips
Horwich Parkway	01/05/1999	303858
Braintree Freeport	10/11/1999	37038
Dunfermline Queen Margaret	25/01/2000	202477
Brighouse	28/05/2000	89309
Wavertree Technology Park	13/08/2000	205232
Warwick Parkway	13/08/2000	438722
Lea Green	17/09/2000	181660
Beauly	15/04/2002	41878
Chandlers Ford	18/05/2003	212987
Edinburgh Park	04/12/2003	382644
Glasshoughton	21/02/2005	122178
Gartcosh	09/05/2005	110967

(Stations opened since April 2007 excluded as full year's demand data not available)

7.4.2 Flow level predictions

To assess the revenue likely to be generated by new stations it is necessary to model the distribution of trips to destinations in addition to the total number of trip origins. The flow level models developed in this study have only been calibrated for the South-East Wales area, and therefore flow-level predictions were attempted for the 16 station sites identified in this area, which are shown in Figure 7.2.

Figure 7.2: Potential sites for new stations in South-East Wales

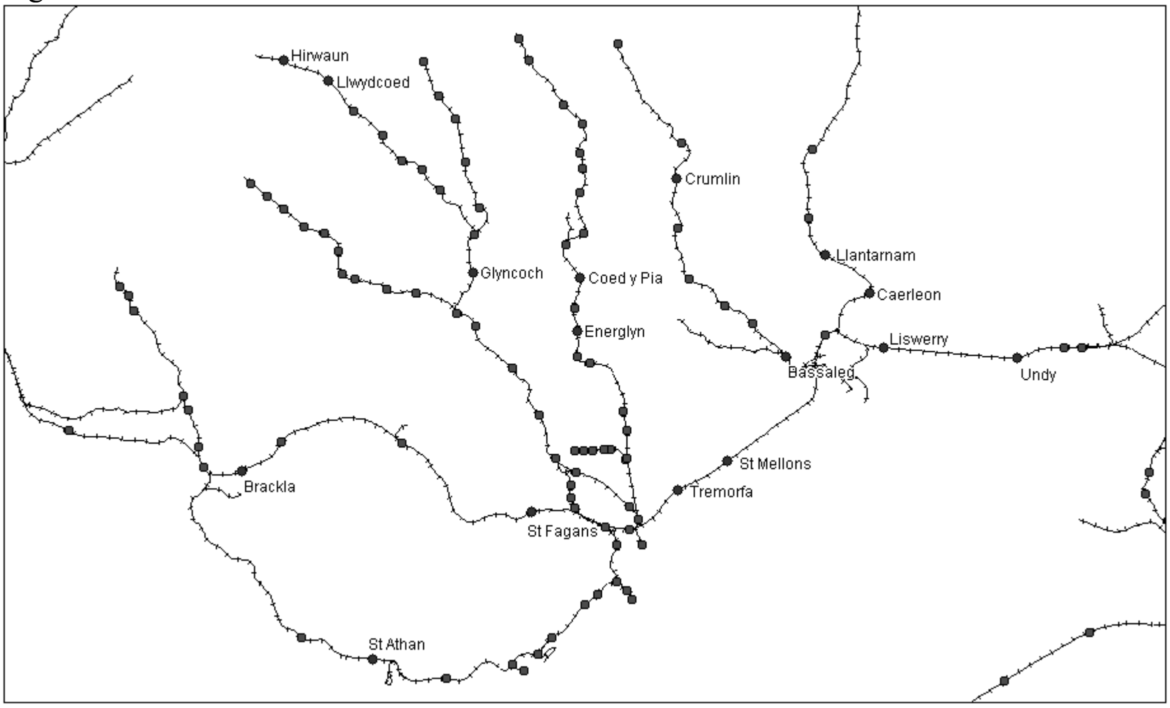

Model 5.29 was used together with the predictions from Model 4.41 to forecast flow-level demand from these 16 stations. The first element of this forecasting is the identification of the destinations which are expected to account for 95% of the trips from the new stations. As discussed in Section 5.4 this destination selection is not straightforward, with the only obvious solution being to base the set of destination stations for a new station on travel patterns at adjacent stations. The numbers of destinations required to account for the top 95% of trips from the stations adjacent to each of the 16 potential new stations were therefore calculated, and these are shown in Table 7.11. For some of the station sites the identification of adjacent stations was obvious, but in other cases there were either no adjacent existing stations in one direction or the adjacent station was much larger (for example Cardiff Central was adjacent to Tremorfa) meaning that it could not be used as a comparator. In such cases the nearest two comparable stations on the same route were used as comparators. It was not possible to define suitable comparator stations for the two sites on the Ebbw Vale line (Bassaleg and Crumlin) as because the line has only recently been reopened flow level usage data was not available for the existing stations on the line. These two sites were therefore removed from the dataset for which flow level predictions were to be made.

$$\hat{T}_{ij} = \alpha \, En_i^{\beta} \prod_{j}^{n} \; D_j^{\gamma_j} D_{ij}^{\omega} Rs_{ij}^{\delta} Cs_{ij}^{\kappa} H_{ij}^{\eta} Rfkm_{ij}^{\lambda} \qquad (5.29)$$

Table 7.11: Minimum number of destinations required to account for 95% of trips from comparator stations for new station sites in South-East Wales

New Station	Comparator station 1	Number of destinations	Comparator station 2	Number of destinations	Average destinations	Stations in common
Bassaleg	Rogerstone	n/a	Risca & Pontymister	n/a	n/a	n/a
Brackla	Pencoed	10	Pontyclun	14	12	7
Caerleon	Cwmbran	34	Pontypool & New Inn	22	28	20
Coed y Pia	Llanbradach	17	Ystrad Mynach	16	16.5	14
Crumlin	Newbridge	n/a	Llanhilleth	n/a	n/a	n/a
Energlyn	Aber	15	Llanbradach	17	16	11
Glyncoch	Pontypridd	33	Abercynon	18	25.5	13
Hirwaun	Aberdare	18	Cwmbach	15	16.5	14
Liswerry	Severn Tunnel Junction	12	Caldicot	9	10.5	7
Llantarnam	Cwmbran	34	Pontypool & New Inn	22	28	20
Llwydcoed	Aberdare	18	Cwmbach	15	16.5	14
St Athan	Rhoose	17	Llantwit Major	13	15	12
St Fagans	Pencoed	10	Pontyclun	14	12	7
St Mellons	Cwmbran	34	Pontypool & New Inn	22	28	20
Tremorfa	Cwmbran	34	Pontypool & New Inn	22	28	20
Undy	Severn Tunnel Junction	12	Caldicot	9	10.5	7

If using an IOTD model (or other constrained model form) it would be necessary to restrict the destinations selected for each origin to the minimum required to account for 95% of trips from that origin. However, because Model 5.29 is an unconstrained model the only restriction on the number of destinations modelled is the time required to assemble the data necessary to make forecasts. To avoid further difficulty in selecting destination stations, forecasts were therefore made from each origin to all destinations required to make up the top 95% of trips from both of that origin's comparator stations.

Data was required on the rail journey time for each flow, and this was estimated based on the journey time from adjacent stations to the destination using equation 7.2. The addition of an extra minute is necessary to account for the time taken for trains to accelerate away from the additional stop at the new station.

$$\hat{R}_{ij} = R_{kj} + \left(R_{lj} - R_{kj}\right)\frac{D_{ik}}{D_{kl}} + 1 \qquad (7.2)$$

Where:

R_{ij} is the estimated rail journey time in minutes from new station i to station j

R_{kj} is the observed rail journey time from existing station k to station j, where existing station k is closer to station j than new station i but is on the same route

R_{lj} is the observed rail journey time from existing station l to station j, where existing station l is further from station j than new station i but is on the same route

D_{ik} is the distance from new station i to existing station k

D_{kl} is the distance from existing station k to existing station l

Rail fares were estimated in a similar manner based on the Standard Day Return fare at adjacent stations, with figures rounded up to the nearest higher multiple of 5 pence, in line with railway fare policy.

Once all the required data had been collated, Model 5.29 could be used to predict flow sizes for the 14 proposed stations. The sum of trips forecast by flow for each station is summarised in Table 7.12, along with the total forecast trip origins at each station, with the forecasts detailed in full in Blainey (2009f) (see Appendix 1). The flow level forecasts can be illustrated geographically and schematically, and Figures 7.3 and 7.4 show such illustrations for flows from Energlyn.

Table 7.12: Comparison of sum of flow level forecasts and total forecast trip origins

Station	Sum of flows	Total trip origins	% Difference
Brackla	46978	29234	60.70%
Caerleon	33172	30932	7.24%
Coed y Pia	55024	70685	-22.16%
Energlyn	105462	79727	32.28%
Glyncoch	33324	39338	-15.29%
Hirwaun	60618	87012	-30.33%
Liswerry	114835	73244	56.78%
Llantarnam	85259	61131	39.47%
Llwydcoed	21414	28802	-25.65%
St Athan	10097	34119	-70.41%
St Fagans	36401	16504	120.56%
St Mellons	37031	37940	-2.40%
Tremorfa	282848	58281	385.32%
Undy	49951	39958	25.01%

Figure 7.3: Geographic representation of predicted flow sizes from Energlyn

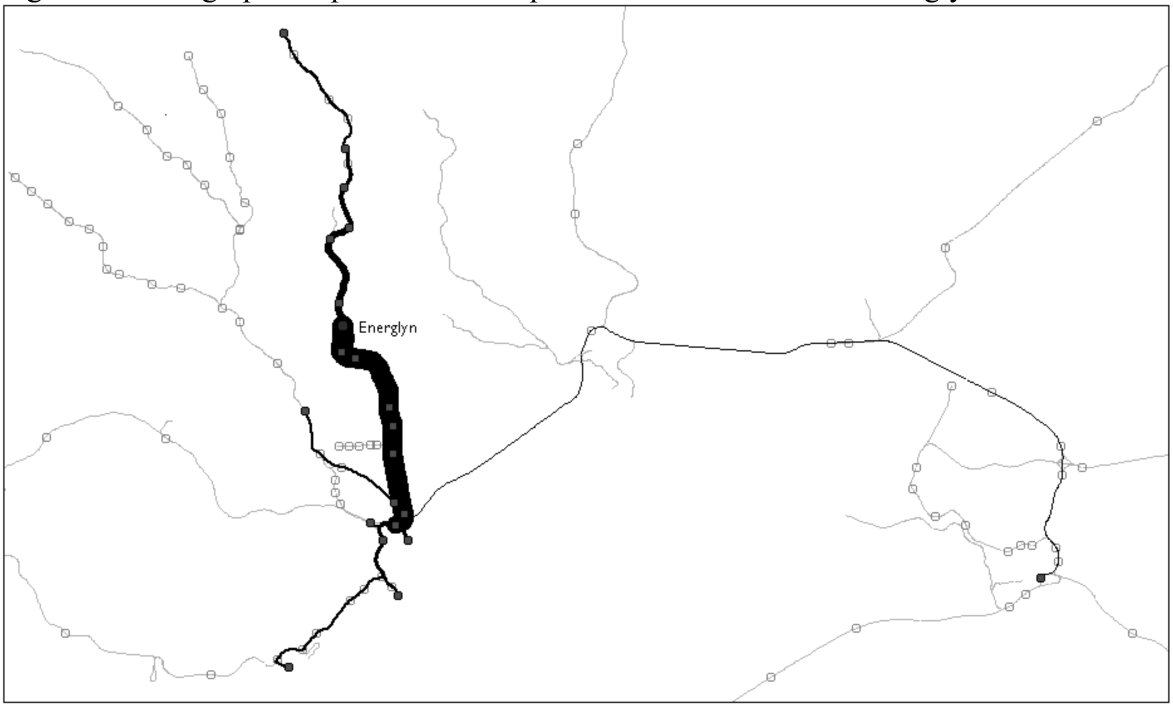

Figure 7.4: Schematic representation of predicted flow sizes from Energlyn

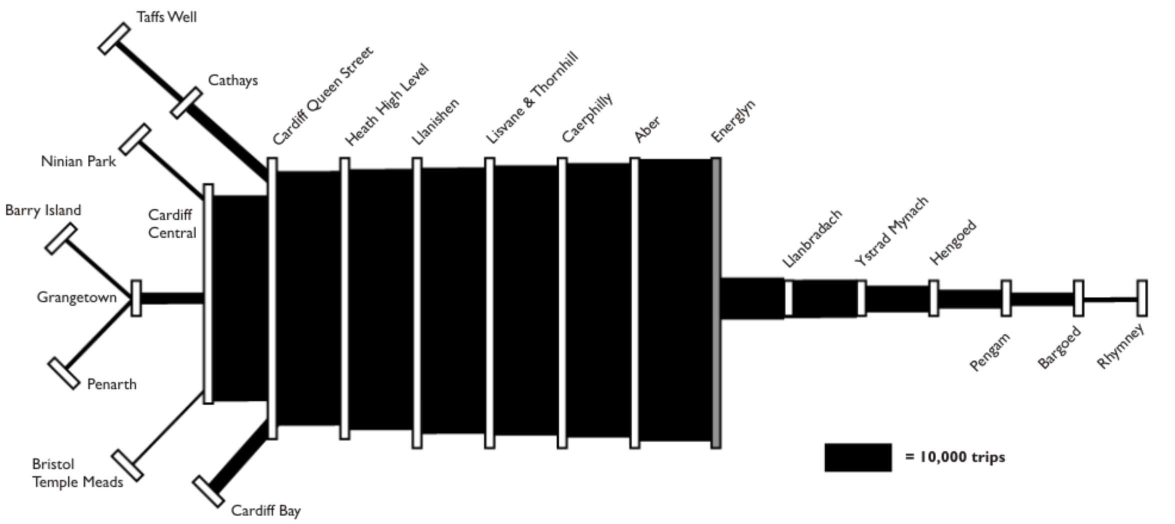

If Model 5.29 was constrained to produce a total number of trips comparable to that from the trip end models then all stations should have exhibited a similar pattern to St Mellons, with the sum of the flow level trips slightly lower than the total trip origin prediction because not all minor destinations were included in the flow level model. However, Table 7.12 shows that for most stations this was not the case, with flow level forecasts combining to give total forecasts either well above or well below the trip end totals. The largest differences, at St Fagans and Tremorfa, occur because the model produces very large flow level forecasts for travel to Cardiff Central, because they are very close in terms of rail travel time to this destination which is represented in the model by a sizeable dummy variable. The logarithmic form of the model means that the effects of this proximity are exaggerated, but it was not obvious how this apparent problem could be corrected. Cardiff Central is predicted to be the overwhelmingly dominant destination for all of the new stations. While this might be realistic for many of the stations, it seems unlikely that Cardiff Central would have such extreme primacy as a destination for the stations to the east of Newport. Cardiff Central dominates forecasts in this way because its dummy variable was much larger than that for any of the other destinations, probably because of the characteristics of the dataset on which the model was calibrated (based on the Rhymney line). It is therefore possible that Model 5.29 is not suitable for forecasting demand from these stations without recalibration.

The easiest way to solve the problem of flow level forecasts not summing to trip end forecasts would be to scale the flow level forecasts so that they sum to give the trip end forecasts, but as described in Section 5.2.4 this methodology gave inferior results when tested on the Rhymney line. Another option was to just use the Model 5.29 forecasts of demand to the two most important destinations (by flow size), and then gross up these predictions to give a total trip forecast based on the proportion of total demand at neighbouring stations made up by trips to their two most important destinations. The resulting forecasts are shown in Table 7.13, which compares the grossed up totals with the forecasts from the trip end models and also for comparison purposes shows the difference between the trip end forecasts and the total flow level forecasts from Table 7.12. This shows that while this method gives a closer fit with the trip end forecasts for some stations, for other stations the total of all flow level forecasts gives a better fit, and as the latter forecasts give more information for revenue estimation they should be preferred.

Table 7.13: Comparison of trips forecast by grossing up flow level forecasts to top two destinations and by trip end models

Station	Grossed Flow Totals	Trip end Totals	% Difference	% Difference All Flows
Brackla	44150.309	29234	51.02%	60.70%
Caerleon	28765.13	30932	-7.01%	7.24%
Coed y Pia	57319.363	70685	-18.91%	-22.16%
Energlyn	98832.058	79727	23.96%	32.28%
Glyncoch	29088.469	39338	-26.06%	-15.29%
Hirwaun	93304.272	87012	7.23%	-30.33%
Liswerry	120756.59	73244	64.87%	56.78%
Llantarnam	73691.631	61131	20.55%	39.47%
Llwydcoed	33165.66	28802	15.15%	-25.65%
St Athan	12193.391	34119	-64.26%	-70.41%
St Fagans	41772.697	16504	153.11%	120.56%
St Mellons	34331.278	37940	-9.51%	-2.40%
Tremorfa	327055.81	58281	461.17%	385.32%
Undy	49025.902	39958	22.69%	25.01%

An alternative solution to the problem of summed flow level forecasts differing widely from trip end forecasts, and one which might also deal with the problem of Cardiff Central's primacy being overestimated, is to use the IOTD Model 5.42 in conjunction with trip end Model 4.41 to forecast flow-level demand. It is necessary to assume that the destinations selected for modelling will include the top 95% of flows from each origin, and while it is not certain that this is the case, to enable comparison with the results from Model 5.29, and in the absence of a better methodology, the same set of destinations was used as before. Blainey (2009f) (see Appendix 1) again gives full details of the forecasts from this model, and the predicted flow sizes for Energlyn are shown in Figures 7.5 and 7.6.

$$P\left[V(j)\middle|V(J)\right] = \frac{1 - e^{[-\beta V(j) + \gamma W(j)]}}{1 - e^{[-\beta V(J) + \gamma W(J)]}} \tag{5.42}$$

Figure 7.5: Geographic representation of Model 5.42 predicted flow sizes from Energlyn

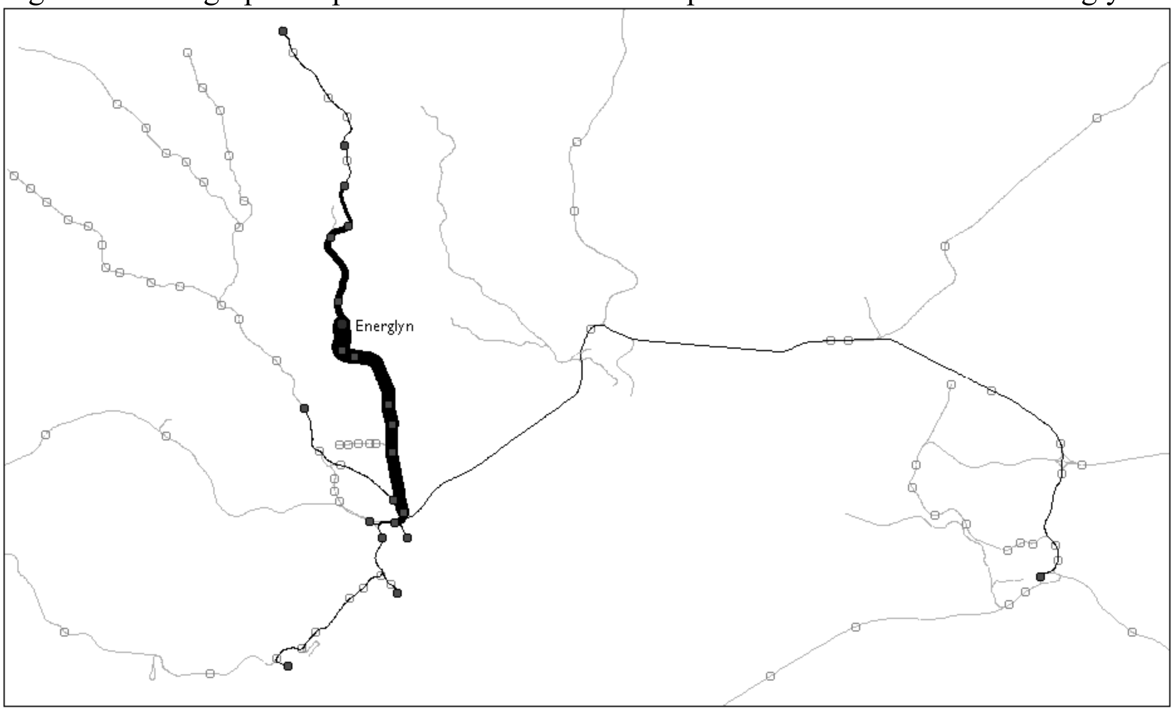

Figure 7.6: Schematic representation of Model 5.42 predicted flow sizes from Energlyn

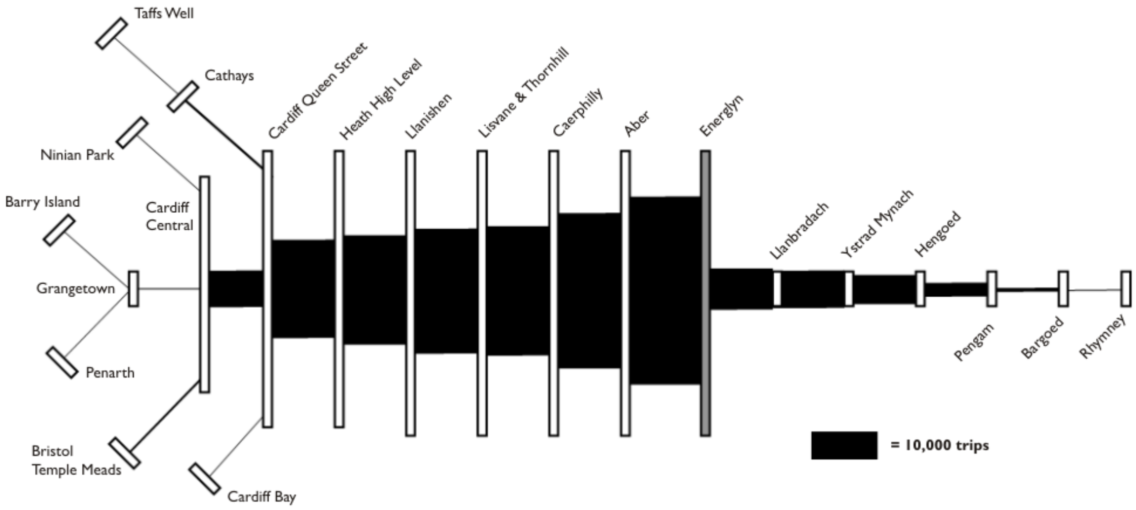

Intuitively the results from Model 5.42 seemed to be more believable, but as Table 7.14 shows there were some major differences between the ranking of destinations from Model 5.29 and that from Model 5.42. There was no way of establishing which set of model forecasts were more realistic, and therefore both were retained for use in the appraisal procedure.

Table 7.14: Mean difference in destination rank between Model 5.29 and Model 5.42

Station	Destinations	Mean rank difference
Brackla	17	4.59
Caerleon	36	8.00
Coed y Pia	19	6.11
Energlyn	21	6.48
Glyncoch	37	10.54
Hirwaun	22	5.36
Liswerry	14	2.71
Llantarnam	36	7
Llwydcoed	22	5
St Athan	17	5.53
St Fagans	17	5.06
St Mellons	36	9.44
Tremorfa	36	7.28
Undy	14	3.29

7.5 Demand build-up and abstraction

7.5.1 Demand build-up over time

The issues of demand build-up over time at new stations and of abstraction of demand by new stations from existing stations are closely related. Demand is likely to grow faster at new stations than at neighbouring stations for several years after they open as people gradually adjust their trip patterns to account for the existence of the new station. Such ‘ramp-up’ effects have been investigated in the past, most recently by Preston & Dargay (2005), who found that it could take up to five years for patronage to reach a long-run steady state. However, there has been an underlying and continuous growth in rail usage in recent years, and while this may be slowed by the current economic downturn, it is unlikely that demand at new stations will ever stabilise at an absolute equilibrium level. This phenomenon is not unique to rail, as it is unlikely that any transport system will reach a state of equilibrium (Dargay & Goodwin, 1995). This means that patronage at a new station is only likely to stabilise in relative terms compared to neighbouring stations, and the differences in rates of growth between new stations and neighbouring preexisting stations were therefore investigated to establish the time taken for such relative stabilisation in demand to occur. This process was complicated by the likelihood that new stations will abstract demand from the stations which are their immediate neighbours, meaning that demand at these stations is likely to grow more slowly or even decline in contrast to the general trend of underlying growth. It was therefore necessary to consider stations beyond the range of possible abstraction in addition to these immediate neighbours when investigating the rate of demand build-up after station opening.

Station usage data from the ORR was only available for years since 2002-3, and as it was desirable to have data from as long a time period as possible to assess demand build-up five stations opened in 2001, 2002 and 2003 were chosen as the initial case studies for this work. The total number of trips from each of these stations was obtained from ORR station usage spreadsheets for the years 2002-3, 2004-5, 2005-6, 2006-7 and 2007-8, together with similar data for local stations in the area around the new stations. Table 7.15 compares the growth rates for demand at the new stations with the mean growth rates for the preexisting stations over the same period.

Table 7.15: Comparison of demand growth rates at new stations and existing stations

Station	2003-05 Growth	2005-06 Growth	2006-07 Growth	2007-8 Growth
Beauly	24.74%	6.64%	26.34%	16.78%
Inverness area mean (6)[1]	23.94%	22.03%	18.24%	14.60%
Chandlers Ford	n/a	15.95%	10.15%	7.25%
South Hampshire area mean (11)	27.55%	5.24%	4.85%	12.57%
Howwood	49.68%	19.48%	17.77%	-3.93%
West Renfrewshire area mean (11)	8.96%	9.95%	3.98%	1.42%
Brunstane	35.15%	33.21%	1.53%	-10.03%
Newcraighall	72.56%	16.30%	10.73%	7.40%
Edinburgh area mean (11)	23.80%	7.39%	1.14%	12.88%

[1] Figures in brackets give number of stations included in area mean

Table 7.15 suggested that an equilibrium level of patronage relative to other stations was likely to be reached within the maximum period of six years investigated here, but that demand would be increasing relative to other stations for most of this period. However, the small size of the dataset used meant that it was not possible to place a great deal of confidence in this hypothesis. Similar analysis was therefore carried out for an additional 14 stations which opened between 1998 and 2000. Because the extent of the differences between the growth rates at the new stations and the area means was the crucial element in this analysis rather than the magnitude of growth, the results were summarised to produce Table 7.16, which gives the number of years since station opening and the difference in growth rates between the new stations and the area means.

Table 7.16: Difference in growth rates between new stations and local stations in surrounding area

Station		2004-05	2005-06	2006-07	2007-08
Beauly	Growth difference	0.80%	-15.39%	8.10%	2.18%
	Years since opening	2	3	4	5
Chandlers Ford	Growth difference	n/a	10.71%	5.31%	-5.31%
	Years since opening	1	2	3	4
Howwood	Growth difference	40.72%	9.53%	13.79%	-5.35%
	Years since opening	4	5	6	7
Brunstane	Growth difference	11.35%	25.82%	0.38%	-22.91%
	Years since opening	2	3	4	5
Newcraighall	Growth difference	48.76%	8.92%	9.59%	-5.48%
	Years since opening	2	3	4	5
Wavertree Technology Park	Growth difference	9.38%	6.77%	13.46%	8.42%
	Years since opening	4	5	6	7
Lea Green	Growth difference	23.37%	9.45%	0.79%	-1.65%
	Years since opening	4	5	6	7
Dunfermline Queen Margaret	Growth difference	22.97%	1.42%	-4.48%	-7.00%
	Years since opening	5	6	7	8
Brunswick	Growth difference	26.36%	-4.17%	2.56%	7.21%
	Years since opening	7	8	9	10
Conway Park	Growth difference	12.00%	-4.78%	5.93%	-6.15%
	Years since opening	6	7	8	9
Creswell	Growth difference	-24.31%	-2.23%	-2.37%	-3.28%
	Years since opening	6	7	8	9
Dalgety Bay	Growth difference	2.22%	-1.00%	-0.50%	0.36%
	Years since opening	7	8	9	10
Drumfrochar	Growth difference	-9.01%	6.92%	-7.11%	-9.24%
	Years since opening	6	7	8	9
Langwith Whaley-Thorns	Growth difference	-33.44%	-9.54%	-12.12%	-3.23%
	Years since opening	6	7	8	9
Shirebrook	Growth difference	-34.45%	3.23%	-3.70%	-0.15%
	Years since opening	6	7	8	9
Whitwell	Growth difference	-39.33%	-1.79%	-14.99%	-1.48%
	Years since opening	6	7	8	9
Braintree Freeport	Growth difference	19.64%	-10.84%	27.63%	27.45%
	Years since opening	5	6	7	8
Brighouse	Growth difference	22.77%	-0.81%	14.86%	-6.57%
	Years since opening	4	5	6	7
Horwich Parkway	Growth difference	10.08%	-2.02%	6.09%	5.78%
	Years since opening	5	6	7	8

The differences in growth rates were plotted against the number of years since the station opened, giving Figure 7.7. This indicated that there might be a slight negative correlation between the number of years since a station opened and the difference in growth rates at that station and stations in the surrounding area, but this pattern was far from clear and the fit of the 'best fit' line was very poor ($R_{adj}^2 = 0.079$). This probably results from the absence of data in most cases for the years immediately after station opening, which may have obscured any pattern which existed in reality. There is little that can be done about this problem until further station usage data becomes available from the ORR allowing stations opened more recently to be added to the dataset. However, it does appear that in general demand seems to have stabilised relative to other stations in the area six years after station opening, and at some stations it appears to stabilise within the first year.

Figure 7.7: Relationship between difference in new station demand growth rates and years since station opening

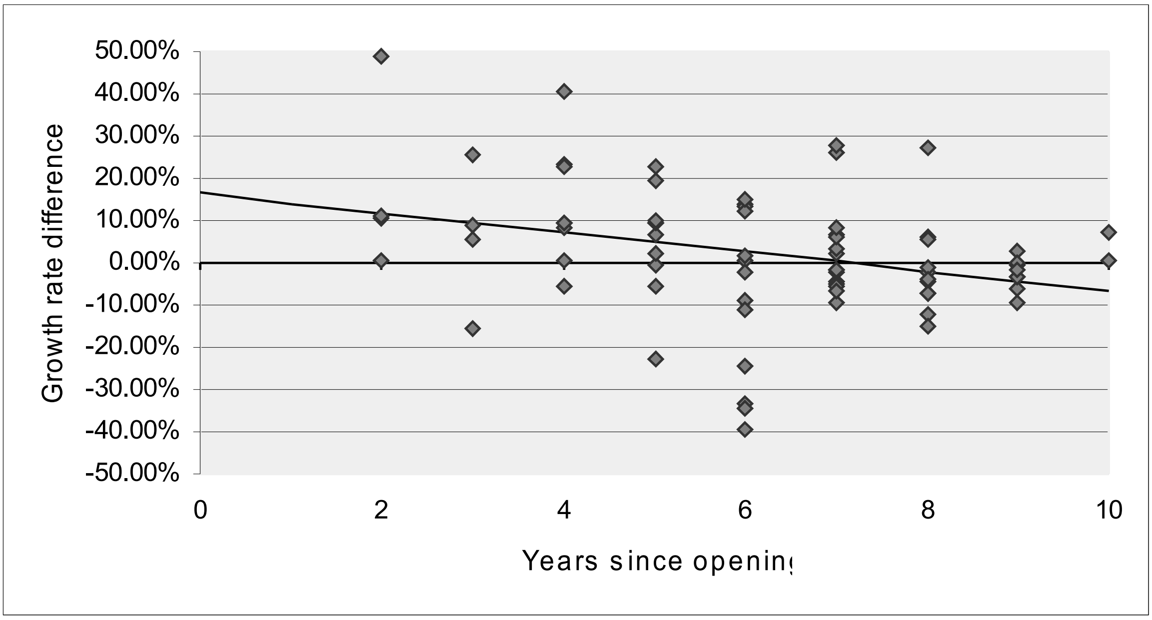

However, the issue of abstraction may also be affecting the results shown in Figure 7.7, as this might mean that demand growth at adjacent stations was negative in the years immediately following station opening as passengers switched to use the new station. The demand build-up analysis was therefore repeated with stations immediately adjacent to the new stations removed from the area means. Table 7.16 was revised to take account of this change, with additional more distant stations used in the area mean calculations to compensate for the removal of adjacent stations, giving the results shown in Table 7.17 and Figure 7.8.

Table 7.17: Difference in growth rates between new stations and local stations in surrounding area using revised area means

Station		2004-05	2005-06	2006-07	2007-08
Beauly	Growth difference	0.09%	-17.33%	6.53%	1.83%
	Years since opening	2	3	4	5
Chandlers Ford	Growth difference	n/a	10.88%	6.07%	-2.30%
	Years since opening	1	2	3	4
Howwood	Growth difference	40.14%	10.92%	14.36%	-4.85%
	Years since opening	4	5	6	7
Brunstane	Growth difference	11.12%	26.57%	0.40%	5.09%
	Years since opening	2	3	4	7
Newcraighall	Growth difference	48.53%	9.67%	9.60%	-4.48%
	Years since opening	2	3	4	7
Wavertree Technology Park	Growth difference	6.16%	11.79%	14.23%	-71.83%
	Years since opening	4	5	6	8
Lea Green	Growth difference	8.04%	1.14%	6.24%	2.69%
	Years since opening	4	5	6	10
Dunfermline Queen Margaret	Growth difference	23.96%	1.05%	-5.33%	-5.81%
	Years since opening	5	6	7	9
Brunswick	Growth difference	21.42%	0.39%	3.00%	-1.75%
	Years since opening	7	8	9	9
Conway Park	Growth difference	9.87%	-6.53%	5.81%	-0.59%
	Years since opening	6	7	8	10
Creswell	Growth difference	-14.30%	2.60%	1.93%	-9.02%
	Years since opening	6	7	8	9
Dalgety Bay	Growth difference	3.22%	-1.37%	-1.35%	-1.70%
	Years since opening	7	8	9	9
Drumfrochar	Growth difference	-7.93%	8.15%	-8.41%	1.38%
	Years since opening	6	7	8	9
Langwith Whaley-Thorns	Growth difference	-23.42%	-4.71%	-7.82%	0.04%
	Years since opening	6	7	8	9
Shirebrook	Growth difference	-24.43%	8.05%	0.60%	-18.84%
	Years since opening	6	7	8	5
Whitwell	Growth difference	-29.31%	3.04%	-10.69%	-1.41%
	Years since opening	6	7	8	5
Braintree Freeport	Growth difference	15.90%	-13.70%	11.58%	27.94%
	Years since opening	5	6	7	8
Brighouse	Growth difference	23.88%	-0.74%	14.45%	-7.44%
	Years since opening	4	5	6	7
Horwich Parkway	Growth difference	11.41%	-0.18%	10.11%	7.49%
	Years since opening	5	6	7	8

Figure 7.8: Relationship between difference in new station demand growth rates and years since station opening using revised area means

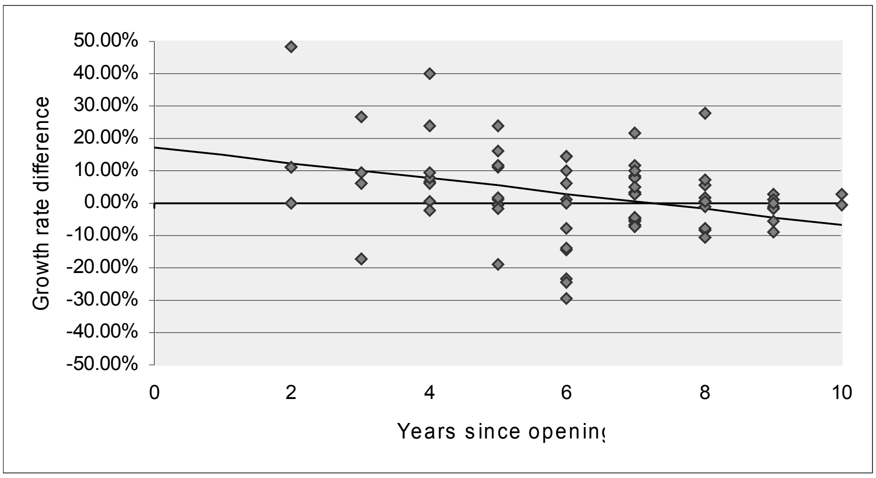

This modification to the area means did not affect the general pattern exhibited by the data, with the best fit line still having a poor fit ($R_{adj}^2 = 0.088$). A loglinear best fit line was tested but this did not give any improvement. A caveat must therefore be attached to the results from the demand models developed in this study, stating that the results they give are estimates of demand following any 'ramp-up' after the station opens.

7.5.2 Abstraction from existing stations

While the inclusion of adjacent stations in the area means made no apparent difference to the analysis of demand build-up over time in Section 7.5.1, the issue of abstraction was deserving of more detailed attention. The crucial period for establishing the extent of abstraction seemed likely to be the year immediately following the opening of the station, and therefore analysis concentrated on stations opened recently where usage data around the opening date is available. Data was required for the year before opening, the year of opening, and the year after opening, which given the data available from the ORR meant that only stations opened between 1 April 2003 and 31 March 2006 could be considered. The stations at Chandlers Ford, Chatelherault, Edinburgh Park, Gartcosh, Glasshoughton, Kelvindale, Larkhall, Llantwit Major, Merryton and Rhoose were therefore selected for further investigation.

Usage data was collated for the existing stations adjacent to the new stations, and the difference between usage in the year before opening of the new station and usage in the two following years was calculated for each of these stations to establish the level of abstraction which had occurred. Surprisingly, demand was found to have grown at all but one of the existing stations over this period, meaning that no abstraction could be assumed. This phenomenon presumably occurred because the underlying growth in demand for rail travel was greater than the reduction in demand at particular stations due to abstraction by the new stations. To account for this underlying demand, area mean demand growth levels were calculated based on non-adjacent stations in the same general area. The difference between this area mean growth and the growth recorded at the adjacent stations was then calculated to establish whether growth at these stations was lower than might have been expected. The results of this analysis are summarised in Table 7.18, divided into those from Year O (the year the new station opened) and Year O+1 (the year following that when the new station opened). For each year the mean relative growth rate is given, which is the mean of the growth rate at each station expected to experience abstraction minus the mean growth rate at non-adjacent stations in the same area. The maximum and minimum relative growth rates are also given for each new station to illustrate the range of the results.

Table 7.18: Relative growth rates for stations immediately adjacent to new stations

New Station	Opening date	Neighbouring stations	Year O			Year O+1		
			Mean relative growth	Max relative growth	Min relative growth	Mean relative growth	Max relative growth	Min relative growth
Chandlers Ford	18/05/03	4	No data	No data	No data	-11.15%	-2.19%	-23.68%
Chatelherault	01/12/05	4	10.32%	16.45%	4.18%	21.76%	42.18%	6.16%
Edinburgh Park	01/12/03	5	No data	No data	No data	12.29%	31.65%	-10.11%
Gartcosh	09/05/05	8	-3.77%	1.86%	-11.50%	1.24%	12.63%	-15.29%
Glasshoughton	21/02/05	6	13.61%	52.56%	-14.36%	29.91%	79.12%	-3.25%
Kelvindale	26/09/05	3	-3.86%	0.52%	-7.64%	-14.37%	-10.59%	-17.40%
Larkhall	01/12/05	4	10.32%	16.45%	4.18%	21.76%	42.18%	6.16%
Llantwit Major	12/06/05	4	4.15%	15.56%	-3.71%	-7.61%	7.31%	-14.49%
Merryton	01/12/05	4	10.32%	16.45%	4.18%	21.76%	42.18%	6.16%
Rhoose	12/06/05	4	4.15%	15.56%	-3.71%	-7.61%	7.31%	-14.49%

Table 7.18 shows that while abstraction of passengers from existing stations may have occurred around some new stations, in particular Chandlers Ford and Kelvindale, it was far from a universal phenomenon. The reasons for this variation are unclear, as while it is possible that differing distances between the new station and its neighbours are responsible for some of the variation, Chatelherault for example is extremely close to the existing stations of Airbles and Shieldmuir yet no demand appears to have been abstracted on aggregate from these stations. Mapping the variations in relative growth rates around the new stations might give some indication of the causes of the variation, and this was tested for two areas, giving Figures 7.9 and 7.10.

Figure 7.9: Relative growth rates around Glasshoughton

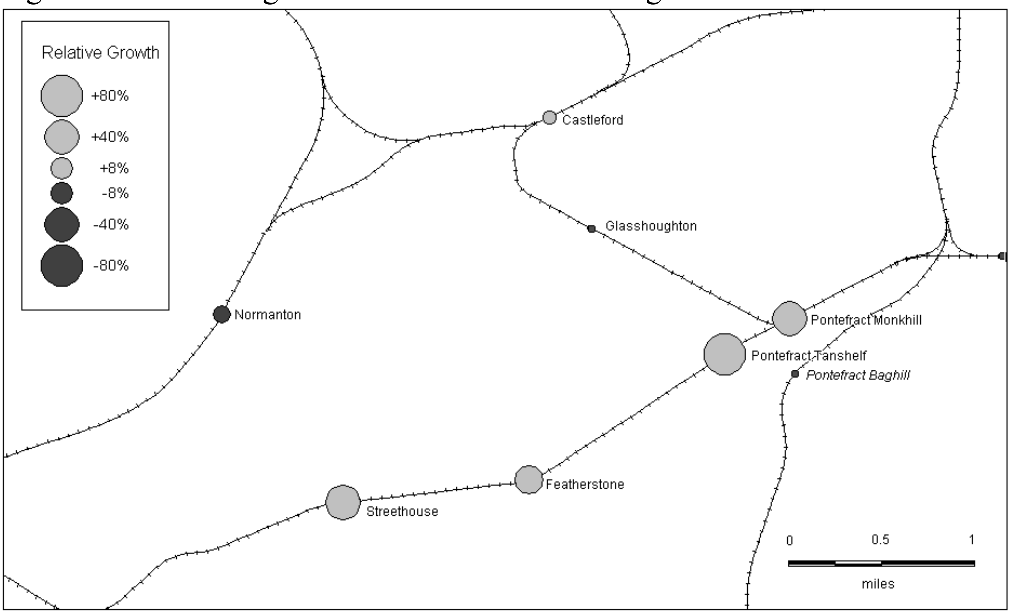

Figure 7.10: Relative growth rates around Llantwit Major and Rhoose

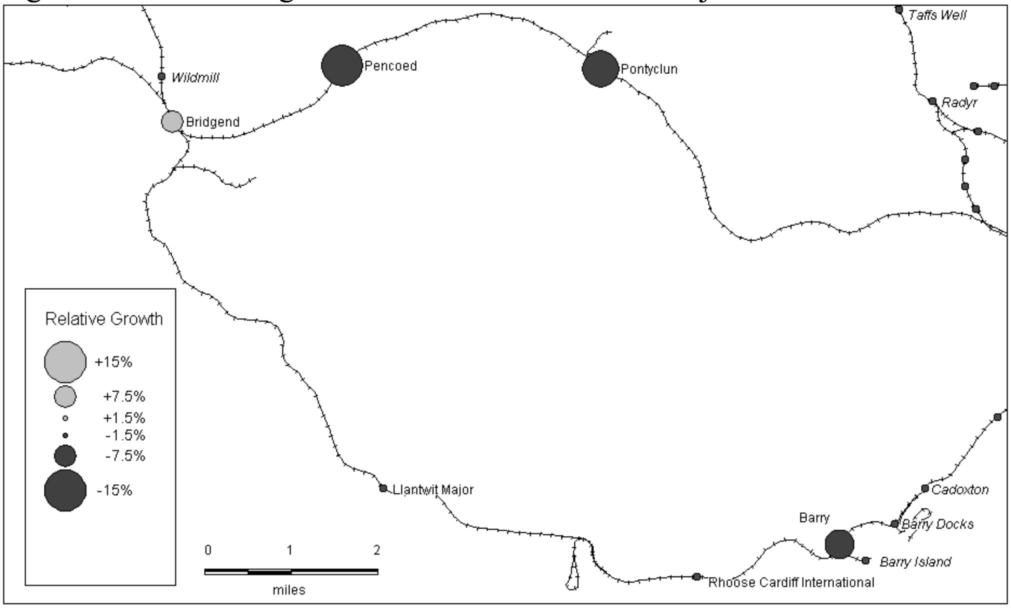

These maps show the relative growth rates, calculated by subtracting the mean growth in demand at stations in the surrounding area from the growth at the neighbouring station in question. This means that, for example, demand at Pontefract Monkhill grew by 39% more than the area mean growth rate. The maps did not though greatly aid understanding of the causes of variations in abstraction levels. The only station at which abstraction seemed likely to have occurred near Glasshoughton was Normanton, one of the most distant of the competing stations, while demand had grown faster than average at the other adjacent stations. The situation around Rhoose and Llantwit Major was slightly clearer, with apparent abstraction of travellers from Pencoed, Pontyclun and Barry, while demand grew faster than average at Bridgend. The latter station was a likely destination for travellers from the two new stations, whereas the other three adjacent stations may previously have been used as railheads for travel to Cardiff by travellers from the area around the new stations. However, the distances between Llantwit Major/Rhoose and their adjacent stations are much greater than those between Glasshoughton and its adjacent stations, which intuitively suggests that Glasshoughton should have been responsible for more abstraction. Glasshoughton may attract trips from nearby stations as it serves a major leisure centre, and this generation of new trips may have masked any abstraction of trip origins from neighbouring stations. There was no obvious way to investigate this further without more detailed usage data for the stations in question, and the wide variations in abstraction rates found here meant that abstraction from other stations was not taken into account in the appraisal procedure developed as part of this study.

7.6 Appraisal for new stations in South-East Wales

7.6.1 Benefits accrued

7.6.1.1 Fare revenue

The largest source of revenue from new local stations obviously comes from the fares paid by the passengers using the station. Some previous new station studies have estimated this revenue by simply multiplying the average single fare (for stations in the local area or for the TOC) by the number of trips forecast (Scott Wilson, 2006; Halcrow Group, 2006). However, when flow level forecasts are made a more sophisticated forecast of fare revenue is possible. Fares were estimated for each flow based on the fares at adjacent stations (as in Section 7.4.2) and then multiplied by the number of trips forecast for that flow. This process is complicated by the existence of a number of different fares as the demand models give no indication of the proportion of passengers using each ticket type. However, data were available from a previous project (Preston et al., 2008) which gave the proportion of passengers using each ticket type from six stations in South Wales in 2006-7, and this is shown in Table 7.19.

Table 7.19: Proportion of passengers using different ticket types from South Wales stations in 2006-7

	Ticket type			
Station	Standard full	Standard reduced	Standard season	Other
Barry Island	25.06%	46.57%	28.34%	0.02%
Cardiff Bay	33.83%	34.91%	31.20%	0.06%
Grangetown	46.42%	30.79%	22.78%	0.01%
Pantyffynnon	41.12%	58.88%	0.00%	0.00%
Pembrey & Burry Port	34.73%	42.08%	23.00%	0.19%
Trefforest	30.71%	44.94%	24.31%	0.03%
Mean	30.84%	42.21%	26.91%	0.04%

These data have limitations in that they do not distinguish between single and return fares. However, in the absence of any other data, it was assumed that the ticket type split for flows from the new stations would be the same as the 'mean' proportions from this dataset. Standard full fare tickets were assumed to be anytime day returns if such tickets were available, and anytime returns otherwise, with standard reduced tickets assumed to be off peak day returns if available and off peak returns otherwise. Season ticket fares per trip were estimated as being 57.05% of the respective standard full fare, as this was the average level found in the revenue dataset. The estimated fare revenue generated using the forecasts from each of the two modelling procedures is given in Table 7.20.

Table 7.20: Estimated fare revenue per annum from proposed stations in South-East Wales

Station	Model 5.29 Revenue	Model 5.42 Revenue
Brackla	£295,080.28	£71,543.54
Caerleon	£358,669.77	£73,124.21
Coed y Pia	£213,890.34	£200,277.09
Energlyn	£404,624.93	£225,036.58
Glyncoch	£267,341.20	£100,154.84
Hirwaun	£248,938.96	£274,175.68
Liswerry	£899,447.22	£186,025.66
Llantarnam	£843,653.87	£164,175.76
Llwydcoed	£92,632.98	£90,690.04
St Athan	£144,559.81	£103,103.14
St Fagans	£127,841.80	£32,579.07
St Mellons	£334,681.22	£64,666.24
Tremorfa	£584,710.78	£75,736.23
Undy	£523,823.28	£156,068.09

Table 7.20 shows that in some cases the differences in revenue forecast between the two models are huge. These differences result from the greater prominence given by Model 5.29 to more distant destinations where more revenue will be generated per trip, meaning that a relatively small difference in trips forecast will lead to a large difference in predicted revenue. On average 7.5% of trips predicted by Model 5.29 are to destinations outside South-East Wales (to which fares are generally higher), whereas the corresponding figure for Model 5.42 is just 2.7%. The implied mean fare per trip from both models is given in Table 7.21, along with the estimated mean fare for the comparator stations used in destination selection based on the LENNON data. This suggests that, overall, Model 5.29 gives a more accurate estimate of revenue in terms of mean fare than Model 5.42, although there is a lot of variation from station to station. This does not necessarily mean that Model 5.29 gives a more accurate estimate of total revenue than Model 5.42, as this will also depend on the accuracy of the flow-level demand predictions. It is important to note that these mean fares are for return trips, and should therefore be divided by two to give the mean fare per single trip if they are to be used in conjunction with data from trip end models. As discussed in Section 7.3, Table 7.6 suggests that both models underpredict revenue on the Rhymney line, and this implies that the higher cost forecasts of Model 5.29 are likely to be more accurate than the more conservative forecasts of Model 5.42.

$$\hat{T}_{ij} = \alpha\, En_i^{\beta} \prod_{j}^{n} D_j^{\gamma_j} D_{ij}^{\omega} Rs_{ij}^{\delta} Cs_{ij}^{\kappa} H_{ij}^{\eta} Rfkm_{ij}^{\lambda} \tag{5.29}$$

$$P[V(j)|V(J)] = \frac{1 - e^{[-\beta V(j) + \gamma W(j)]}}{1 - e^{[-\beta V(J) + \gamma W(J)]}} \tag{5.42}$$

Table 7.21: Estimated mean fare per passenger from model predictions and observed data

Station	Model 5.29 Mean Fare	Model 5.42 Mean Fare	Comparison Mean Fare
Brackla	£6.28	£2.58	£3.69
Caerleon	£10.81	£2.49	£7.94
Coed y Pia	£3.89	£2.98	£3.91
Energlyn	£3.84	£2.97	£4.02
Glyncoch	£8.02	£2.68	£4.24
Hirwaun	£4.11	£3.32	£2.82
Liswerry	£7.83	£2.67	£6.75
Llantarnam	£9.90	£2.83	£7.94
Llwydcoed	£4.33	£3.31	£2.82
St Athan	£14.32	£3.18	£3.92
St Fagans	£3.51	£2.08	£3.69
St Mellons	£9.04	£1.79	£7.94
Tremorfa	£2.07	£1.37	£7.94
Undy	£10.49	£4.11	£6.75
TOTAL	£5.49	£2.78	£5.05

7.6.1.2 Other revenue

In addition to fare revenue, new stations may generate other more minor revenue streams. If a station car park is provided and parking spaces are charged for then this will generate additional revenue. Some feasibility studies for local stations include this revenue, such as the Scott Wilson study of a possible station at Beeston Castle, which assumed a £1 charge and 75% occupancy rate. However, such revenue was not included in the appraisal carried out here, as parking is usually free at local stations, largely because enforcing parking charges at unstaffed stations would not be cost-effective.

There will often be potential to provide and charge for commercial advertising space at new stations, and this may offset some of the station maintenance costs. It was not possible to find any estimates of the amount of revenue generated in this way, and as the amounts involved are likely to be small they were not included in the appraisal process.

Revenue can also sometimes be generated at stations by renting out station buildings and retail units to private businesses. However, this usually occurs at larger stations or at stations where railway buildings are no longer required for their original purpose, and is therefore unlikely to be an option at new local stations.

7.6.1.3 User benefits

Integrating the area under the demand curve to calculate user benefits (see Section 3.4.6.1) does not give a finite result for Model 5.29, as with the double logarithmic form the demand curve never crosses the axis, meaning that the implied user benefit is infinite. In order to calculate user benefit it was therefore necessary to assume an alternative functional relationship, and the negative exponential form (7.2) was used here.

$$Q = \alpha e^{-\beta F} \tag{7.2}$$

Where:
Q is the predicted demand level
F is the fare
α is the intercept
$-\beta F$ is the fare elasticity

The calibration of Model 5.29 on the Rhymney line dataset gave a fare elasticity of -1.381, and values of F and Q were provided by the demand forecasts made for the proposed stations using Model 5.29. Calculation of the value of α for each flow, and then the equivalent values of β was therefore straightforward. The user benefit for the flow could then be estimated by dividing Q by the associated β value, but because this approximates the demand curve to a straight line it will slightly overestimate user benefit, as the demand curve is in fact convex to the origin. It is therefore more accurate to calculate user benefit by integrating equation (7.2) between two fare limits, as shown in equation (7.3).

$$UB = \int_{F=fQ}^{F=f1} \alpha e^{-\beta F} dF \tag{7.3}$$

Where:
UB is the user benefit
fQ is the fare which gives the demand level Q predicted by the model
$f1$ is the fare which gives a demand level of 1 passenger

The value of $f1$ can be found by rearranging (7.2), while fQ is simply the actual fare level used to make the demand forecast. Integrating (7.3) gives (7.4) which can then be evaluated using the values of $f1$ and fQ to give the user benefit.

$$UB = \left[\frac{\alpha}{-\beta} e^{-\beta F} \right]_{F=fQ}^{F=f1} \tag{7.4}$$

Estimating the user benefit from the demand forecast by Model 5.42 was more complicated, as this model does not include fares and therefore did not generate a fare elasticity. The fare elasticity from Model 5.29 was used as the best available alternative, but this is not an ideal solution to the problem. To give an indication of the sensitivity of the user benefit estimates to the level of the fare elasticity, user benefits were reestimated for both sets of forecasts using the PDFH recommended fare elasticity of -0.9. Table 7.22 summarises the user benefit estimates for each station using both sets of demand forecasts and both fare elasticities. This shows that using the PDFH fare elasticity gives much higher user benefits than using the elasticity from Model 5.29. There are though arguments for retaining the elasticity obtained here (not least for consistency), as a significant degree of bus competition would be expected to exist in South-East Wales, meaning that demand would be more elastic with respect to fare changes than might be the case elsewhere. The estimates of non-user benefit made using a fare elasticity of -1.381 were therefore taken forward for use in the appraisal.

Table 7.22: Estimated user benefit per annum from proposed stations in South-East Wales

Fare elasticity	-1.381		-0.9	
Station	Model 5.29	Model 5.42	Model 5.29	Model 5.42
Brackla	£213,264.82	£51,486.73	£327,243.01	£79,003.52
Caerleon	£258,213.24	£52,616.97	£396,213.87	£80,737.81
Coed y Pia	£154,751.97	£144,894.44	£237,458.31	£222,332.47
Energlyn	£292,835.76	£162,793.51	£449,340.21	£249,797.60
Glyncoch	£193,069.29	£72,149.91	£296,254.10	£110,710.02
Hirwaun	£180,101.15	£198,375.38	£276,355.21	£304,396.01
Liswerry	£650,712.23	£134,335.98	£998,481.76	£206,131.10
Llantarnam	£609,439.45	£118,462.75	£935,150.98	£181,774.51
Llwydcoed	£66,917.96	£65,511.06	£102,681.89	£100,523.08
St Athan	£104,291.90	£74,272.58	£160,030.12	£113,967.15
St Fagans	£92,180.49	£23,351.73	£141,445.85	£35,831.93
St Mellons	£240,769.11	£46,531.88	£369,446.82	£71,400.58
Tremorfa	£422,004.41	£54,529.48	£647,542.32	£83,672.45
Undy	£378,719.86	£112,664.85	£581,124.59	£172,877.96

7.6.1.4 Non-user benefits

The diversion rates in Table 3.5 were used to calculate the proportion of trips abstracted from each mode for all flows from the proposed stations. The urban diversion rates were used for flows shorter than 40 km and the interurban diversion rates for all other flows. The non-user benefits (or costs) of trips diverted from air, cycle or walk were not considered, as the number of passengers involved was likely to be minimal. Assessing the non-user benefits of trips diverted from bus services was not straightforward, as the values given in Section 3.4.6.2 are expressed in vehicle kilometres rather than passenger kilometres. Data on average bus loadings is not readily available, and a reduction in bus passenger numbers after a new railway station opens will not necessarily lead to a reduction in bus service frequency. For the purposes of this appraisal it was assumed that for every 50 passengers diverted from bus, the equivalent of one bus vehicle trip from the origin station to the destination station would be saved, but this assumption is entirely arbitrary. Furthermore, it does not take into account the disbenefit to remaining bus users of the reduced service frequency if bus trips were removed. Estimating the non-user benefits of trips diverted from car was more straightforward, as it was assumed that each passenger was equivalent to one car journey. Net cost per kilometre values half way between the high and low values given in Table 3.7 (12.18p per km for car journeys and 48.12p per km for bus journeys) were used to calculate the non-user benefits. The values of the non-user benefits for each of the proposed stations are shown in Table 7.23.

Table 7.23: Estimated non-user benefits from proposed stations in South-East Wales

Station	Model 5.29	Model 5.42
Brackla	£150,368.41	£43,755.69
Caerleon	£173,554.04	£35,343.91
Coed y Pia	£89,663.49	£73,502.25
Energlyn	£143,389.33	£73,637.83
Glyncoch	£111,231.32	£40,886.95
Hirwaun	£276,131.48	£264,416.36
Liswerry	£377,388.8	£104,216.26
Llantarnam	£405,370.39	£63,961.62
Llwydcoed	£95,136.05	£72,990.58
St Athan	£56,015.05	£51,813.60
St Fagans	£58,623.31	£17,711.77
St Mellons	£175,432.71	£37,004.63
Tremorfa	£385,513.45	£41,724.34
Undy	£221,249.70	£75,169.84

7.6.1.5 Other benefits

The construction of new stations may lead to other benefits, including journey time savings for passengers abstracted from other modes, land value uplift in the area around the station, and limited job creation as a result of the construction work. Agglomeration benefits might be obtained in central Cardiff, Swansea and Newport as their effective labour markets expanded as a result of the increased accessibility brought by the new stations. Some non-users of the new stations may place an option value or a non-use value on their existence. The former describes the value people place on having the option of using the station available in the future, even if they never take up this option, and the latter describes the value they place on the continued existence of a station that they do not ever expect to use themselves. Laird et al (2009) have shown that such values do exist with relation to rail schemes and that they are particularly important for schemes where user benefits are low, but that limited data availability makes their magnitude difficult to assess. While it did not prove possible to quantify any of these benefits in this project, these are issues which should be considered if this appraisal procedure was to be extended.

7.6.2 Costs incurred

7.6.2.1 Construction costs

The cost of constructing the new stations was estimated based on the mean cost per platform unit (£470,609.97) obtained from the analysis in Section 3.4.7.1. The number of platform units required at each of the new stations was based on the lengths of platforms at adjacent stations measured using Google Earth. This data is summarised in Table 7.24.

Table 7.24: Number and size of platforms at potential stations in South-East Wales

Station	Platforms	Units per platform
Brackla	2	4
Caerleon	2	5
Coed y Pia	2	5
Energlyn	2	5
Glyncoch	2	4
Hirwaun	1	2
Liswerry	2	4
Llantarnam	2	5
Llwydcoed	1	2
St Athan	2	4
St Fagans	2	4
St Mellons	2	5
Tremorfa	2	5
Undy	2	4

7.6.2.2 Station maintenance, operating and staffing costs

Because of the similarity of the three estimates of station maintenance and operating costs detailed in Section 3.4.7.2, it was assumed that maintenance and operating costs for all the proposed stations would be the mean of these estimates, giving a figure of £38,329 per annum. It was assumed that all stations would initially be unstaffed, meaning that no additional staffing costs would be incurred by their construction.

7.6.2.3 Fuel costs

The additional fuel costs associated with services calling at the new stations were simply calculated by multiplying the service frequency by the average figure per halt estimated in Section 3.4.7.5 (£2.44) and then multiplying the resulting figure by 365 to give the total costs per year. The only exceptions to this procedure were the stations at Hirwaun and Llwydcoed, which would require the extension of existing services. Fuel and other costs for these stations are considered in Section 7.6.2.5.

7.6.2.4 Costs to existing passengers

All flows which would be expected to use the services passing through each of the new stations were isolated from the South Wales LENNON dataset. The number of passengers on each flow was then multiplied by 1.64 minutes (the average additional journey time as a result of the new station opening) and by the mean value of travel time calculated during investigations into aggregate logit modal split models (Blainey, 2009b; see Appendix 1), converted into 2008 quarter three prices (9.9p per minute). Individual values of time were not calculated based on journey distance for each flow because over 18,500 flows were involved meaning that the time involved would have been prohibitive. The costs per year incurred for existing passengers by the opening of the proposed new stations are summarised in Table 7.25, showing that such costs can be substantial.

Table 7.25: Estimated costs to existing passengers of opening proposed stations in South-East Wales

Station	Cost to existing passengers
Brackla	£41,891.44
Caerleon	£66,020.74
Coed y Pia	£160,898.07
Energlyn	£168,078.91
Glyncoch	£102,455.20
Hirwaun	£0.00
Liswerry	£87,209.55
Llantarnam	£66,020.74
Llwydcoed	£0.00
St Athan	£47,803.21
St Fagans	£63,192.79
St Mellons	£100,394.05
Tremorfa	£100,394.05
Undy	£87,209.55

7.6.2.5 New Service Costs

The proposed stations at Hirwaun and Llwydcoed would require the extension of Aberdare line services, incurring additional costs which were calculated using the figures in Table 3.10. The costs for each station were calculated incrementally, with costs for Llwydcoed based on an extension of existing services from Aberdare and costs for Hirwaun based on a further extension from Llwydcoed. The vehicle operating costs include rolling stock and traincrew charges, which would not necessarily be incurred if the service could be extended using existing rolling stock, but the current six minute turnaround at Aberdare makes this unlikely. Congestion costs were not included, as no additional services would be provided on lines already served by passenger trains. It also seemed unlikely that there would be any Mohring effect benefits from the extension of the services as journey frequency and quality would remain unchanged on existing routes, and this element was also therefore excluded from the new service costs for these stations. It was assumed that no capital expenditure on infrastructure would be necessary other than station construction costs, as the line is already used by freight trains, but if such expenditure was required then the benefit-cost ratio (BCR) for these stations would obviously be reduced.

7.6.2.6 Costs to bus users

A significant proportion of demand at new local stations is likely to be abstracted from competing bus services, which as a result are likely to be reduced in frequency or even withdrawn. However, not all passengers using these bus services will be able or willing to transfer to rail. For example, rail travel may be too expensive, or they may live in areas served by the bus route but not by new or existing railway stations. There will therefore be a disbenefit to the remaining bus passengers as a result of the reduction in bus service quality, and this should ideally be accounted for in the appraisal. However, in practice it is impossible to quantify this without detailed modelling of the bus market, and these costs will therefore not be included in the appraisal procedure outlined here.

7.6.3 Cost-benefit analysis

The cost and benefit figures obtained using the above methods were used to carry out a cost-benefit analysis for the 14 proposed stations in South-East Wales. Both a financial appraisal (using equation (7.5)) and a social CBA (using equation (7.6)) of the schemes were carried out, to establish which sites have both a financial and a social case for construction, which have only a social case and which have no positive case for construction at all.

$$NPV_f = \sum_{i=0}^{N} \frac{R_i - VC_i - MC_i - K_i}{(1+r)^i} \tag{7.5}$$

$$NPV_s = \sum_{i=0}^{N} \frac{R_i + UB_i + NUB_i - VC_i - OC_i - K_i - UC_i}{(1+r)^i} \tag{7.6}$$

Where:

NPV_f is the financial net present value of the scheme

R_i is the fare revenue in year i

VC_i is the vehicle related costs in year i

OC_i is the station maintenance and operating costs in year i

K_i is the capital cost in year i

r is the interest rate

N is the project life

NPV_s is the social net present value of the scheme

UB_{ia} is the user transport benefits in year i

NUB_{ia} is the non-user benefits in year i

UC_i is the cost to existing users in year i

A 60 year appraisal period was used in line with the recommendations of WebTAG, based on the Treasury 'Green Book' (HM Treasury, 2003), as while most previous studies have used a 30 year appraisal period these were undertaken using guidelines which have now been superseded. Both costs and benefits were discounted to obtain their present value using equation (7.7) and the Green Book discount rates of 3.5% for the first thirty years of the project and 3.0% for the remainder of the appraisal period.

$$PV = \sum_{0}^{60} \frac{S}{(1+r)^n} \tag{7.7}$$

Where:

PV is the present value of the cost or benefit

S is the cost or benefit being discounted

r is the discount rate

n is the year in which the cost or benefit is incurred

Station construction costs were assumed to commence in year 0 of the project, with other costs and benefits commencing in year 1 and continuing to year 60. It is unlikely that station construction costs would be paid in total at the start of the project. A more realistic scenario is that money would be borrowed to pay for these capital costs, and they were therefore amortised using Equation (7.8) and then discounted before being used in the cost-benefit analysis. Amortisation was carried out using the Treasury test interest rate of 3.5% recommended in WebTAG. In reality the amortisation rate will depend on the available financing arrangements, which may therefore have an impact on station viability.

$$P = \frac{\mu V}{1 - (1 + \mu)^{-n}} \qquad (7.8)$$

Where:

P is the annual payment

V is the total sum being amortised

μ is the interest rate

n is the project life (in years)

While most new station studies assume that rail demand will grow over time, the current economic situation means that limited confidence can now be placed in such assumptions. As the analysis in Section 7.5.1 found no conclusive evidence of a gradual build-up of demand over time at new stations, it was assumed that station usage would stabilise within the first year of operation and remain constant over time, and this meant that costs to existing passengers would also remain constant over the project lifespan. Similar uncertainty applies to fuel costs, which might be expected to rise over the life of the project, but as forecasting this price rise is not straightforward fuel costs were also assumed to remain constant for the purposes of this appraisal. The total benefits, total costs, NPVs and BCRs are shown in Table 7.26, with full details of the benefits and costs for each of the stations given in Blainey (2009f) (see Appendix 1).

Table 7.26: BCR for proposed stations in South-East Wales (all values in £)

Model 5.29	Financial				Social			
Station	PVB	PVC	NPV	BCR	PVB	PVC	NPV	BCR
Brackla	7809939	6032199	1777740	1.29	17434281	7140947	10293334	2.44
Caerleon	9492973	7492886	2000087	1.27	20920630	9240268	11680362	2.26
Coed y Pia	5661072	8647897	-2986826	*0.65*	12130057	12906414	-776357	*0.94*
Energlyn	10709276	7846461	2862815	1.36	22254919	12295034	9959885	1.81
Glyncoch	7075764	7375784	-300019	*0.96*	15129740	10087483	5042257	1.50
Hirwaun	6588709	16509851	-9921142	*0.40*	18663895	16509851	2154043	1.13
Liswerry	23805820	6715777	17090043	3.54	51016743	9023967	41992775	5.65
Llantarnam	22329128	7492886	14836242	2.98	49188271	9240268	39948003	5.32
Llwydcoed	2451733	13798790	-11347057	*0.18*	6740843	13798790	-7057947	*0.49*
St Athan	3826089	6008628	-2182539	*0.64*	8068960	7273843	795116	1.11
St Fagans	3383610	6055771	-2672160	*0.56*	7374960	7728305	-353345	*0.95*
St Mellons	8858064	7492886	1365178	1.18	19873747	10150032	9723715	1.96
Tremorfa	15475638	7492886	7982752	2.07	36848349	10150032	26698317	3.63
Undy	13864119	6715777	7148341	2.06	29743614	9023967	20719647	3.30

Table continued on next page

Model 5.42	Financial				Social			
Station	PVB	PVC	NPV	BCR	PVB	PVC	NPV	BCR
Brackla	1893555	6032199	-4138644	*0.31*	4414352	6032199	-1617847	*0.73*
Caerleon	1935391	7492886	-5557495	*0.26*	4263466	7492886	-3229420	*0.57*
Coed y Pia	5300767	8647897	-3347130	*0.61*	11081109	8647897	2433212	1.28
Energlyn	5956081	7846461	-1890380	*0.76*	12213749	7846461	4367288	1.56
Glyncoch	2650815	7375784	-4724969	*0.36*	5642580	7375784	-1733204	*0.77*
Hirwaun	7256654	16509851	-9253198	*0.44*	19505440	16509851	2995589	1.18
Liswerry	4923572	6715777	-1792205	*0.73*	11237374	6715777	4521597	1.67
Llantarnam	4345267	7492886	-3147619	*0.58*	9173524	7492886	1680638	1.22
Llwydcoed	2400309	13798790	-11398481	*0.17*	6066054	13798790	-7732735	*0.44*
St Athan	2728848	6008628	-3279780	*0.45*	6065992	6008628	57365	1.01
St Fagans	862276	6055771	-5193495	*0.14*	1949110	6055771	-4106661	*0.32*
St Mellons	1711532	7492886	-5781354	*0.23*	3922507	7492886	-3570379	*0.52*
Tremorfa	2004524	7492886	-5488362	*0.27*	4552090	7492886	-2940796	*0.61*
Undy	4130680	6715777	-2585097	*0.62*	9102133	6715777	2386355	1.36

Table 7.26 shows that when using the demand forecasts from Model 5.29 eight stations have a financial BCR greater than 1 and a positive financial NPV, and a further three stations have a social BCR greater than 1 and a positive social NPV. In all cases the social NPV is higher than the financial NPV, as would be expected. The results suggest that the most promising sites for new stations in this area are at Liswerry and Llantarnam. If the forecasts from Model 5.42 are used, however, none of the stations give a positive financial NPV, although seven have a positive social NPV and a social BCR greater than 1. Liswerry again appears to be the most promising site, followed in this case by Undy. Only three of the stations (Coed y Pia, Llwydcoed and St Fagans) give a negative NPV in all cases, suggesting that there may be potential for station construction at all the other sites. Furthermore, it should be noted that the viability of the stations at Hirwaun and Llwydcoed would be greatly improved if they could be served using the rolling stock already used for Aberdare services, as the additional rolling stock-related costs are the dominant component of the total costs for these stations.

The large differences between the NPVs and BCRs produced by the two models are obviously due to the differences in the flow level forecasts they produce. Model 5.29 tends to give a higher PVB because it forecasts more trips to secondary, more distant, destinations, which in turn leads to a higher NPV and BCR. While the forecast flow sizes for these stations from Model 5.29 are not high, because the fares charged to these stations are much higher than those to the dominant local destinations they form a disproportionately large component of the total revenue generated. Model 5.42, because of its different form, tends to forecast only one or even no trips to these secondary destinations, and therefore the total revenue predicted by this model is much lower.

Observation of demand patterns at existing stations suggests that Model 5.29 is more realistic in this respect, although it may exaggerate the importance of some secondary destinations.

The DfT uses BCR values to assess the value for money of a scheme, with a BCR greater than 2 indicating high value for money, a BCR between 1.5 and 2 medium value for money, a BCR between 1 and 1.5 low value for money, and a BCR less than 1 poor value for money (DfT, 2007c). Table 7.27 shows how many of the proposed stations in South-East Wales fall into each category using both sets of demand forecasts.

Table 7.27: Value for money of proposed stations in South-East Wales

Value for money	Financial BCR		Social BCR	
	Model 5.29	Model 5.42	Model 5.29	Model 5.42
High	4	0	6	0
Medium	0	0	3	2
Low	4	0	2	5
Poor	6	14	3	7

All the cost-benefit analyses described above include the costs of station construction, but some new station studies do not include such capital costs when calculating the BCR. Table 7.28 gives the BCR and annual costs and revenues (in 2008 prices) when capital costs are excluded. It is obviously not valid to compute an NPV in this case as only a single year is being considered.

Table 7.28: Annual BCR for proposed stations in South-East Wales when capital costs excluded

Station	Model 5.29				Model 5.42			
	Benefits	Costs	Financial BCR	Social BCR	Benefits	Costs	Financial BCR	Social BCR
Brackla	658714	113173	4.14	5.82	166786	113173	1.00	1.47
Caerleon	790437	153333	4.11	5.16	161085	153333	*0.84*	1.05
Coed y Pia	458306	291849	1.63	1.57	418674	291849	1.53	1.43
Energlyn	840850	268750	4.02	3.13	461468	268750	2.24	1.72
Glyncoch	571642	224501	2.19	2.55	213192	224501	*0.82*	*0.95*
Hirwaun	705172	584628	*0.43*	1.21	736967	584628	*0.47*	1.26
Liswerry	1927548	184318	9.26	10.46	424578	184318	1.92	2.30
Llantarnam	1858464	153333	9.66	12.12	346600	153333	1.88	2.26
Llwydcoed	254687	482197	*0.19*	*0.53*	229192	482197	*0.19*	*0.48*
St Athan	304867	118194	2.05	2.58	229189	118194	1.46	1.94
St Fagans	278646	135365	1.77	2.06	73643	135365	*0.45*	*0.54*
St Mellons	750883	187706	3.83	4.00	148203	187706	*0.74*	*0.79*
Tremorfa	1392229	187706	6.70	7.42	171990	187706	*0.87*	*0.92*
Undy	1123793	184318	5.39	6.10	343903	184318	1.61	1.87

Table 7.28 shows that the Model 5.29 forecasts give a positive financial BCR for all

stations except Hirwaun and Llwydcoed (which are still affected by high rolling-stock costs), and a positive social BCR for all stations except Llwydcoed. The lower Model 5.42 forecasts still give a positive financial BCR for six of the stations, and a positive social BCR for a further three stations. This indicates that if construction could be funded externally the stations with a positive financial BCR would directly cover their operating costs.

7.6.4 Break-even demand levels

While the procedure summarised in Section 7.6.3 allows a detailed appraisal of new local station schemes to be carried out, it inevitably requires a significant quantity of data to be collected before it can be used. In some cases it may not be possible to justify such extensive data collection, and therefore a simplified procedure to estimate the breakeven demand level for new local stations was produced. This combines the amortised station construction costs for a single year with mean annual operating costs (from section 7.6.2), and then divides this total by the estimated mean fare, as shown by equation (7.9). This formula was used to produce breakeven demand levels for a range of station sizes and mean fare levels (assuming a 60 year project life), and these are summarised in Table 7.29.

$$T_{Bi} = \frac{K_i + VC_i + OC_i}{F_{\mu i}} \tag{7.9}$$

Where:

T_{Bi} is the breakeven number of trips at station i

$F_{\mu i}$ is the mean single fare at station i

Table 7.29: Estimated financial breakeven demand levels for new local stations (trips per year)

Mean fare / Platform units	£1	£1.40 (Model 5.42 mean)	£2.00	£2.50 (Observed mean)	£2.75 (Model 5.29 mean)	£3.50
2	130313	93081	65157	52125	47387	37232
3	149179	106557	74590	59672	54247	42623
4	168045	120032	84023	67218	61107	48013
6	205778	146984	102889	82311	74828	58794
8	243510	173936	121755	97404	88549	69574
12	318974	227839	159487	127590	115991	91136
16	394439	281742	197219	157775	143432	112697

The financial breakeven formula (7.9) can be extended to give the social breakeven formula (7.10), if it is assumed that the mean user and nonuser benefit will vary in direct proportion to the mean fare. For the 14 South-East Wales stations, the user benefit was found to be 72.2% of the fare, and the non-user benefit 54.4% of the fare. It is also necessary to assume a mean level of cost to existing users, and for the 12 South-East Wales stations served by existing services this was £90,964 per annum. These figures were used to estimate social breakeven demand levels, as shown in Table 7.30. It should be noted that the cost to existing users of a new station will vary widely, and will be heavily dependent on the position of a station on the route of the services which call there. For example, the opening of a station close to the outer terminus of a commuter route will have less effect on existing passengers than an additional station call just prior to the main trip attractor. This means that the estimates in Table 7.30 should be treated with caution. Nonetheless, the trip end demand forecasting spreadsheet described in Section 4.4 was extended to give an indication of whether or not station construction would be financially and socially justified.

$$T_{Bi} = \frac{K_i + VC_i + OC_i + EC_i}{2.266 F_{\mu i}} \qquad (7.10)$$

Where:

EC_i is the cost to existing passengers resulting from the opening of station i

Table 7.30: Estimated social breakeven demand levels for new local stations (trips per year)

Mean fare / Platform units	£1	£1.40 (Model 5.42 mean)	£2.00	£2.50 (Observed mean)	£2.75 (Model 5.29 mean)	£3.50
2	97651	69751	48826	39060	35509	27900
3	105977	75698	52988	42391	38537	30279
4	114302	81645	57151	45721	41565	32658
6	130954	93539	65477	52382	47620	37415
8	147605	105432	73803	59042	53675	42173
12	180908	129220	90454	72363	65785	51688
16	214211	153008	107106	85684	77895	61203

Consideration was given to extending this analysis by estimating breakeven population levels based on the breakeven demand levels and the parameters from the trip end models. However, this did not prove to be feasible because the number of variables included in the preferred trip end model and in particular the representation of population using a weighted term meant that computation of meaningful breakeven population figures was not possible.

7.7 Spreadsheet-based appraisal tool

The various stages of the appraisal procedure were brought together in an Excel spreadsheet which will produce a benefit-cost ratio for planned new stations once all relevant data has been input. This tool does have some limitations, the most notable of which is that it is restricted to the South-East Wales area, because of the limited spatial transferability of the flow level models as currently calibrated. However, it would be relatively straightforward to extend the spatial applicability of the tool by recalibrating these models. The spreadsheet does also require a large amount of data to be input by the user, but this is unavoidable given the complexity of the appraisal process. Despite these limitations, though, the tool should still prove useful in new station appraisal, and has the advantage that individual elements of the appraisal procedure can be updated without affecting other elements if improved techniques become available. An illustration of the tool is provided by Figure 7.11.

Figure 7.11 Spreadsheet-based appraisal tool

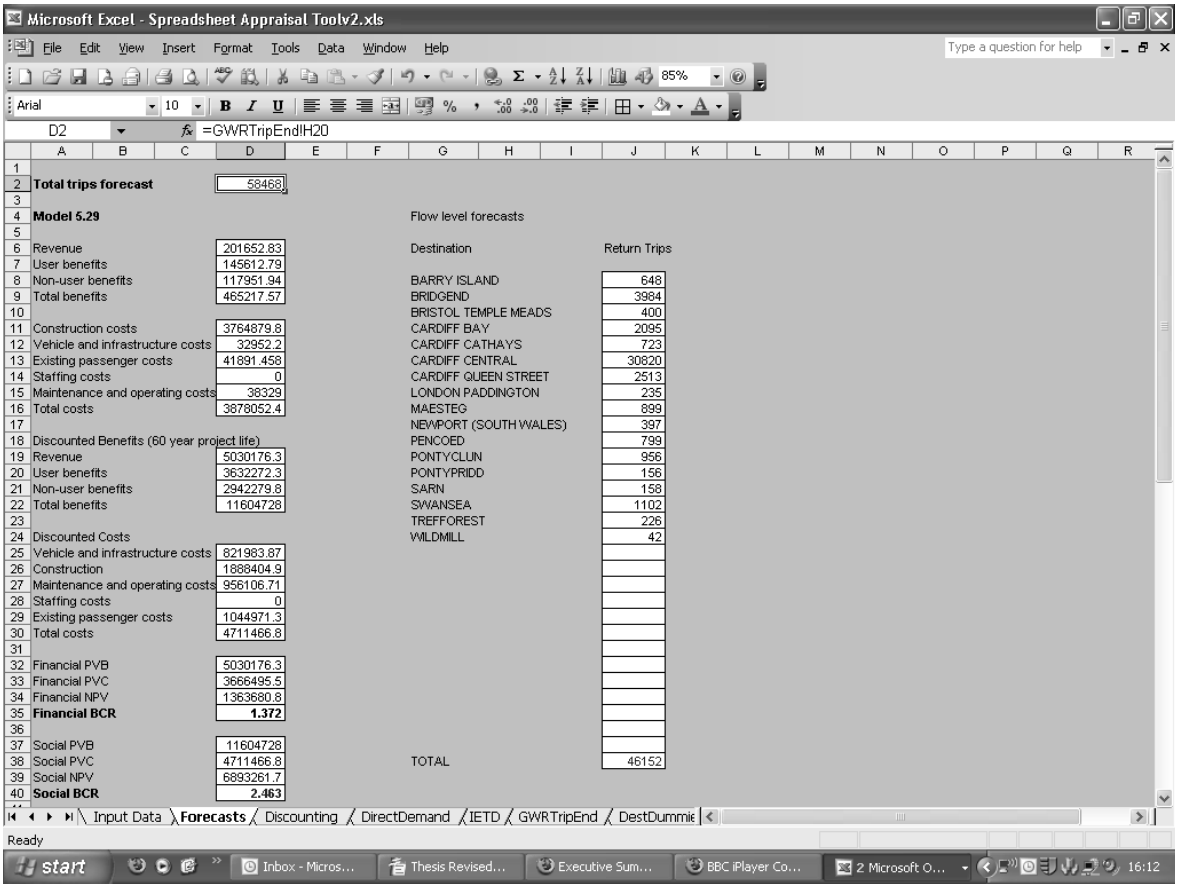

7.8 Conclusions

In this chapter the best demand models developed during this study have been compared, and brought together in a synthesised demand modelling methodology for new local railway station sites. This was then combined with the site search procedure developed in Chapter 6 and with techniques for estimating user and non-user benefits and costs (discussed in Chapter 3) to produce a procedure for appraising potential sites for new stations. To the author's knowledge this is the first time that such an integrated procedure has been developed.

Trip end models (developed in Chapter 4) were recalibrated for the South-East Wales area, but gave inferior results to the models calibrated over the whole of England and Wales, and GWR Model 4.41 was therefore retained as the preferred trip end model. After results from the various flow level models were compared a combination of Model 4.41 and Model 5.29 was selected as the preferred demand forecasting methodology. Model 4.41 was first used to forecast the total trip ends over a one year period at all 421 sites in England and Wales identified in the site search procedure. These sites were ranked by forecast demand, with 42 predicted to generate more than 200,000 trips per year (equivalent to the top 44% of existing stations), and a further 33 predicted to generate more than 150,000 trips per year (equivalent to the top 50% of existing stations). Model 5.29 was then used to forecast flow-level demand at 14 sites in South-East Wales, with destinations based on travel patterns at adjacent stations. However, the sum of these flow level forecasts was much greater than the predicted total trip ends and Cardiff Central appeared to be excessively dominant as a destination. In an attempt to deal with this problem Model 5.42 was used to forecast flow-level demand from the same set of stations, and this give results which intuitively seemed more believable. However, because there was no way of determining which set of forecasts was more realistic both were retained for use in the appraisal procedure.

Demand growth after new station opening was investigated using ticket sales data for a number of stations opened in the past decade. Background demand changes and possible abstraction from neighbouring stations were taken into account but it was not possible to establish any clear 'ramp-up' effects at the new stations. For the purposes of the appraisal procedure it was therefore assumed that demand would stabilise during the first year of operation. Demand abstraction from neighbouring stations was investigated by analysing demand changes at existing stations following the opening of new stations nearby. No clear pattern of abstraction could be found, with demand at some stations increasing after the opening of a new station close by. The assumption that a proportion of trips at new stations would be abstracted from neighbouring stations could not be justified, so in the appraisal procedure it was assumed that all trips at new stations were new to rail.

Estimations of the costs associated with new stations (using methods discussed in Section 3.4) were brought together with forecasts of revenue based on the flow level forecasts and estimated user and non-user benefits to calculate the financial and social net present value (NPV) and benefit-cost ratio (BCR) over a 60 year project life of the 14 proposed stations in South-East Wales. As expected, more sites had a positive social NPV than a positive financial NPV, with Liswerry and Llantarnam having the best case for construction based on the appraisal carried out here. Six sites were identified as giving high value for money by the DfT's criteria based on forecasts from Model 5.29. Operating BCRs (excluding construction costs) were also estimated, with only Llwydcoed having a social operating BCR below 1 when based on Model 5.29 forecasts.

As in some cases it may not be feasible to collect sufficient data to allow a full appraisal to be carried out, estimates were made of breakeven demand levels for new local stations based on a range of mean fares and station sizes. These were incorporated in the trip end demand forecasting spreadsheet, and the various elements of the full appraisal procedure were brought together in another spreadsheet to create a tool for assessing the case for constructing new local stations.

Chapter Eight: Conclusions

8.1 Introduction

This chapter outlines the conclusions that can be drawn from the work described in this thesis, and also suggests some potential extensions to the research. Section 8.2 summarises the theoretical findings from the research including the advantages and disadvantages of the preferred models and briefly explains the problems experienced with the less successful methods. Section 8.3 then gives the empirical findings from the application of the models and the appraisal procedure and from the survey of ultimate passenger origins and destinations. Section 8.4 explores some policy implications of the study findings, before recommendations for future work are made in Section 8.5. The final conclusions of the study are given in Section 8.6, where the results obtained are compared to the objectives set out for the work.

8.2 Theoretical findings

8.2.1 Trip end models

This study has shown that using GWR to calibrate trip end models at a national scale gives a clear improvement over conventional regression methods, with Model 4.41 the preferred model form. Significant spatial variation was found in a number of independent variables, which highlighted the advantages of GWR and also indicated that trip end models are unlikely to be transferable if calibrated locally on smaller datasets. While using the results from a GWR calibration to forecast demand is a somewhat complex process, the development of the spreadsheet-based demand forecasting tool (see Section 4.4) solves this problem, allowing the rapid production of demand forecasts for any site in England and Wales. The use of GWR to enhance trip end models represents a significant advance in local rail demand modelling. No national trip end model of any form had previously been developed, and the existence of such a model which also takes account of local variations in the effect of independent variables should make the assessment of the case for new local stations much more straightforward than was previously the case.

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \rho \ln Pk_i + \eta B_i + \kappa Te_i + \nu El_i \quad (4.41)$$

While the national GWR model gave very good results, the possibility of using locally calibrated trip end models such as those developed here for South Hampshire should not be ruled out. They allow locally important independent variables to be considered in the model form, and may therefore sometimes give better results than the national model. This was illustrated in recent work carried out for Transport for London on forecasting demand for two new stations in West London (Blainey & Preston, 2009a). To obtain a good model fit it proved necessary to incorporate several variables which had not been included in the national model, with the original model form giving a relatively poor model fit. This work also highlighted a potential limitation of the trip end models developed here, which is that they are unlikely to accurately forecast demand at stations with extraordinary catchment characteristics (for example those adjacent to major out-of-town shopping centres). However, no demand model is likely to be accurate in such circumstances, with the best solution either to 'benchmark' demand using similar comparator stations or to scale up trip end forecasts based on estimated usage of the facilities within the station catchment.

8.2.2 Flow level models

While developing reliable flow level models proved more difficult than developing trip end models, good results were eventually obtained. The preferred method is the direct demand Model 5.29, which incorporates generalised origin variables and is therefore suitable for forecasting demand from new stations. However, the transferability of this model is restricted, as the current calibration only contains a limited range of destination dummy variables, and recalibration would therefore be necessary if the model was to be applied outside the South-East Wales area. This type of model also tended to overpredict demand to secondary destinations, probably as a result of insufficient calibration flows being used to give realistic parameter values for some of the destination dummy variables. The sum of flow level forecasts from this model tended to be much higher than the corresponding forecasts from the trip end models, although attempts to constrain the flow level forecasts reduced model fit.

$$\hat{T}_{ij} = \alpha\, En_i^{\beta} \prod_{j}^{n} D_j^{\gamma_j} D_{ij}^{\omega} Rs_{ij}^{\delta} Cs_{ij}^{\kappa} H_{ij}^{\eta} Rfkm_{ij}^{\lambda} \qquad (5.29)$$

An alternative model type, the intervening opportunity trip distribution (IOTD) model, was given detailed consideration and was also found to give good results. This type of model had never previously been applied to rail demand modelling, and there were a number of points in its favour. Its spatial transferability is (at least in theory) not limited as it does not rely on destination dummy variables, and it gives much less prominence to secondary destinations. It also explicitly accounts for the presence of intervening and competing opportunities, which could not be adequately described by the direct demand models. However, the best model of this type (Model 5.42) was in general less accurate at making flow level predictions, and in many cases did not assign enough importance to secondary destinations. The forecasts from this model were almost always much lower than those from Model 5.29 as they were constrained to sum to the total trip end forecasts from Model 4.41.

$$P\left[V(j)\middle|V(J)\right] = \frac{1 - e^{[-\beta V(j) + \gamma W(j)]}}{1 - e^{[-\beta V(J) + \gamma W(J)]}} \qquad (5.42)$$

It did not prove possible to account for the effects of bus competition in the models, although in the case of the direct demand models this was probably due to deficiencies in the available data rather than to problems with the model form. No measure of intermodal competition could be incorporated in the IOTD models. This is a disadvantage of such models, as the model form means that the relative importance of this competition would need to be arbitrarily determined before calibration.

While Model 5.29 in general gives more accurate flow level predictions than Model 5.42, the issues highlighted above mean that it is not possible to give a conclusive recommendation as to which should be preferred. However, both of the models highlighted here give an improvement over previous flow level demand models for local rail services, and the use of either model should give a reasonable indication of the flow level demand which would be expected if a new station were to be opened.

8.2.3 Synthesised modelling and appraisal procedure

The preferred demand models were combined with the GIS-based site search procedure described in Chapter 6 to create a synthesised procedure for locating and forecasting demand at new station sites. An exhaustive procedure of this type has never before been developed, and this represents a major step forward in the analysis of local rail networks. This was combined with an appraisal procedure capable of estimating the financial and social BCR of new station construction at the sites identified, and the incorporation of this procedure in a spreadsheet-based tool provides a straightforward and consistent methodology for assessing the relative merits of station construction in different areas. Detailed investigations were carried out into the phenomena of demand build-up over time and abstraction by new stations from neighbouring existing stations. However, it did not prove possible to obtain conclusive evidence for the existence of either effect, and they were not therefore included in the appraisal procedure. This conflicts with the findings of some earlier studies, and both topics would therefore merit further study.

8.2.4 Unsuccessful methodologies

Not all of the methodologies tested were found to be suitable for incorporation in the final modelling procedure. Problems with data availability and quality and mediocre levels of model fit meant that aggregate logit modal split models were not found to be a reliable method for assessing demand at new local stations. While the use of hierarchical cluster analysis to partition the calibration dataset for national trip end models gave a marginal improvement in model predictions, because there was no obvious way to assign new stations to clusters the unpartitioned dataset was retained. The spatial expansion method was initially identified as a possible means of improving trip end models but, because it can only deal with unidirectional spatial variations and because the number of model parameters escalates rapidly with the number of independent variables, GWR was adopted as the preferred local analysis method.

8.3 Empirical findings

8.3.1 Findings from statistical modelling

The best trip end model (GWR Model 4.41) was able to account for over 82% of the variation in the data, which is an extremely good result given the size of the calibration dataset and the disparate nature of the stations included in it. The best conventional trip end model (4.44) also performed very well, explaining over 78% of the variation in the observed data. While the R^2 values are not directly comparable, an ANOVA test confirmed that the GWR model had a superior fit to the conventional model. As shown in the conclusions to Chapter 4, the fit of both the GWR and conventional model compares favourably to previous large-scale applications of trip end models.

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \rho \ln Pk_i + \eta B_i + \kappa Te_i + \nu El_i \quad (4.41)$$

$$\ln \hat{V}_i = \alpha + \beta \ln \sum_a P_a w_a + \delta \ln F_i + \lambda_S \ln T_S + \tau \ln J_{id4} + \rho \ln Pk_i + \upsilon_2 L_2 + \upsilon_3 L_3 + ... \\ \gamma Ma_i + \xi Bi_i + \omega Ca_i + \eta B_i + \kappa Te_i + \nu El_i \quad (4.44)$$

Assessing the relative performance of the flow level models was less straightforward. The direct demand model with the best fit (the dummy variable Model 5.27) explained over 85% of the variation in the data when calibrated on the Rhymney line, while the preferred direct demand model for forecasting purposes (Model 5.29) explained 83% of the variation. The best IOTD model (Model 5.42) explained nearly 88% of the variation in the same dataset, but because it forecasts probabilities rather than absolute flow sizes these results are not directly comparable. AD values are however equivalent, and these suggest that Model 5.27 is most accurate with forecasts on average within ±49% of actual usage, rising to ±57% for Model 5.29. This compares favourably to the AD value of ±57% achieved by Preston (1987) over a much smaller number of flows. Model 5.42 is though far less accurate when assessed using this measure, with forecasts only within ±131% of actual usage. However, this measure of fit gives more weight to prediction errors for smaller flows, and because IOTD models tend to underpredict such minor flows the inaccuracy of such models will be exaggerated. The relative forecasting performance of the two models is therefore likely to be much closer than this measure of fit would suggest.

$$\hat{T}_{ij} = \alpha En_i^{\beta} \prod_j^n D_j^{\gamma_j} D_{ij}^{\omega} Rs_{ij}^{\delta} Cs_{ij}^{\kappa} H_{ij}^{\eta} Rfkm_{ij}^{\lambda} \quad (5.29)$$

$$P[V(j)|V(J)] = \frac{1 - e^{[-\beta V(j) + \gamma W(j)]}}{1 - e^{[-\beta V(J) + \gamma W(J)]}} \quad (5.42)$$

8.3.2 Findings from OD survey and catchment definition investigations

The results from the survey of ultimate passenger origins and destinations show that the accuracy of theoretical catchment methods still leaves something to be desired. The best theoretical catchments only included 62-69% of observed trip ends. The survey indicated that if arbitrary catchment boundaries are to be defined then a 3 km boundary should be assumed for walk trips, with a 10 or 12.5 km boundary necessary for motorised trips. These boundaries are much larger than those assumed in the PDFH, and are also larger than the boundaries which gave the best model fit when trialled with trip end and direct demand models. This difference may occur because there is a significant amount of overlap between observed catchments, whereas the theoretical catchments used in this study were assumed to be non-overlapping. While the observed catchments were mapped using GIS spatial interpolation methods, it did not prove possible to generalise the features of these catchments.

A new catchment definition method was developed as part of the work on flow level models, with flow-specific catchments defined based on minimising overall journey time. While this method was not incorporated in the preferred modelling procedure, it has considerable potential and would merit further investigation.

However, the best balance between model accuracy and ease of implementation was found to be obtained by allocating all output areas to their nearest station by road access time, weighting their populations using function 4.6, and summing these weighted populations to give the station catchment population. Such catchments are recommended for both trip end and direct demand models.

$$w_a = (t + 1)^{-3.25} \tag{4.6}$$

8.3.3 Findings from forecasting and appraisal

Trip end forecasts were made at the 421 sites in England and Wales identified by the site search procedure, and these were ranked by demand. 42 sites were predicted to generate more than 200,000 trips per year and a further 33 predicted to generate more than 150,000 trips per year (equivalent to the top 50% of existing stations in Great Britain). Flow level forecasts were made for 14 sites within South-East Wales, although there was some variation between the predictions from the two flow level techniques.

Various cost estimations were made based on evidence from recently opened stations, including an estimated construction cost per platform unit of £470,609.97 for new stations. When these cost estimations were balanced against the revenue predicted to be generated by the new stations up to 8 sites were found to have a positive financial BCR and up to 11 sites a positive social BCR depending on the demand forecasting procedure used. It should be noted that none of the sites in South Wales were included in the 'most promising' group of 42 sites identified above, and therefore the BCRs for these sites could be even more positive. Breakeven demand levels were also estimated based on a range of mean fare levels and station sizes, indicating that anything between 37,000 and 395,000 trips per year could be required to make a new station viable, depending on its precise characteristics.

8.4 Policy implications

The procedures developed during this study should enable better evaluation of schemes for new stations and services, ensuring that those schemes which gain approval deliver the best value for money and meet the needs of the communities they serve. The importance of such appraisal is likely to increase in the future, as the challenges posed by climate change and oil depletion mean that rail is likely to grow in significance as a transport mode, making the existence of a reliable decision support system for new local railway stations increasingly crucial.

The results from the study indicate that there are a large number of sites where serious consideration should be given to the construction of a new railway station. The procedures included in this study make the appraisal of such sites relatively straightforward and, if complemented by a detailed assessment of any local conditions which may affect the likely performance of the station, should provide adequate support to enable the decision on construction to be taken.

There is a strong case for suggesting that the demand forecasting procedures contained in this thesis should replace those currently recommended for local stations in the Passenger Demand Forecasting Handbook. In terms of temporal and spatial transferability, of scope and of ease of use they offer a marked improvement over previous methodologies.

8.5 Recommendations for further research

The flow level models developed here have only been tested in one area of the country, and it would be sensible to investigate their spatial transferability by applying them in at least one other area. The direct demand models used to forecast demand in Chapter 7 appeared to overpredict demand to secondary destinations, but this problem could be solved by including some flows with very low usage levels in the calibration dataset, which might result in more realistic parameter values being obtained for the destination dummy variables. Such models could potentially be enhanced by using non-linear calibration methods to introduce an origin station choice element, allowing competition between stations to be better represented. Improvements in theoretical catchment definition methods should be sought, either through the use of nonlinear calibration to incorporate competition between stations or by other means, given that the Rhymney line survey indicated that around one third of trip ends are not accounted for by current methods. The IOTD models would also merit further investigation, as incorporating more variables in the measure of destination attractiveness might improve model results, as might using an alternative decay function.

GWR was successfully used in this study to enhance trip end models at a national scale, enabling spatial variations in the effect of independent variables on rail demand to be identified and incorporated in the model form. However, there may also be scope for GWR to be applied at a more local level to explore and explain variations in rail use across city regions, allowing more accurate demand forecasting, and this would merit investigation. The relatively low density of railway stations may preclude the identification of such variation based on station usage data as insufficient observations are available to give significant results. However, a possible alternative would be to use GWR to model rail travel to work using ward or output area data from the 2001 census, as this would provide a much larger calibration dataset within the same area.

Attempts were made in this study to quantify the level of abstraction of demand by new stations from existing stations, but the results from this analysis were inconclusive. Further work might allow this phenomenon to be better quantified, perhaps using journey to work data from the 1981, 1991 and 2001 censuses to establish the size of modal shifts in travel to work following the opening of new stations. Alternatively passenger surveys could be carried out at recently opened stations to quantify the extent to which trips from these stations were newly generated and the extent to which they were abstracted from other stations or from other modes. The relationship between station quality and abstraction levels should also be investigated, as recent work suggests that this can have a significant impact on station choice (Preston et al., 2008). If levels of abstraction from other stations could be better quantified this would increase confidence in the appraisal procedure, as would more detailed analysis of abstraction from other modes as the diversion rates identified by Balcombe (2003) may not be accurate for local rail services. A better understanding of abstraction would also allow a reliable assessment of the carbon savings delivered by new local railway stations through reductions in car use. Abstraction of trips from other stations could have the same affect by reducing access distances, allowing the use of non-motorised modes for station access.

There is potential to extend the appraisal procedure by considering the impact of station construction on land use, population and employment levels in the areas around new stations. These impacts were not included in the appraisal procedure outlined in this study, but their extent might be a significant determinant of the success or failure of station schemes. Probably the best way to investigate them would be to examine the short and long-run impact of previous new local station schemes on the communities they were designed to serve. Land use changes around a selection of stations opened in recent years could be studied using GIS and compared to changes around pre-existing stations and in areas where no station exists.

While the application of the demand forecasting and appraisal procedures in this thesis was restricted to new stations on existing lines, there is no reason why the same procedures should not be applied to assess the potential for new stations on new lines. While costs for such stations will obviously be much higher, all other forecasts should be equally valid.

8.6 Final conclusions

If the results detailed above are compared to the research objectives set out in Section 1.5, then the following conclusions can be drawn:

1) Simple aggregate trip rate and trip end models were recalibrated in Section 4.2 and direct demand models in Section 5.2. All of these calibrations made use of the most recent data available at the time.

2) A large number of additional explanatory variables have been tested in both trip end (see Section 4.2) and direct demand models (see Section 5.2), giving a major increase in model explanatory power while retaining the simplicity of the basic model form.

3) Extensive use of GIS has been made throughout the study. MapInfo was the primary system used although occasional use of ArcGIS and Google Earth was also required. GIS enabled the wide variety of datasets required for the study to be integrated into the modelling procedure in a way that would otherwise have been impossible. They also

proved valuable in the presentation and analysis of model results, enabling spatial patterns to be identified and described, and allowing further model variables to be identified.

4) GWR was applied to trip end models and was found to enhance model performance both through improvements in model fit and through the use of associated diagnostic measures to identify both data errors and potential model enhancements (see Section 4.3). The use of cluster analysis to enhance trip end models was also investigated, although this was not included in the preferred forecasting procedure (see Section 4.3). An untried procedure for flow level modelling, the IOTD model, was tested and was found to give extremely promising results (see Section 5.4).

5) A semi-automated GIS-based search procedure for potential new station sites has been developed (see Chapter 6).

6) A synthesised procedure for locating new station sites, forecasting demand and carrying out an appraisal has been developed (see Chapter 7). A spreadsheet tool has been developed which automatically carries out such an appraisal once the user inputs the required data.

7) The site search procedure has been applied to all existing railway lines in England and Wales, enabling 421 potential sites to be identified. Total demand forecasts have been made for all these sites, with the full demand modelling and appraisal procedure applied to 14 sites in South-East Wales.

The aims set out at the start of this research project have therefore been almost entirely fulfilled, and an up-to-date procedure now exists which can evaluate the impacts of constructing new local rail stations and which allows results to be easily communicated to stakeholders.

Appendix One: Summary of Associated Technical Notes

Blainey (2009b) **Aggregate Logit Modal Split Models: Technical Note**
This note describes the development of aggregate logit modal split models to forecast travel to work in the Southampton area. Nested and multinomial logit models are calibrated based on journey to work data from the 2001 census, and the problems of using this data for this purpose are discussed. In general results from the aggregate logit models were not as promising as those from the flow level aggregate models developed in Chapter 5 and this together with the associated data problems meant that they were not taken forward for use in the appraisal procedure.

Blainey SP (2009c) **Sensitivity Analysis of England and Wales Trip End Models: Technical Note**
This note gives full details of the tests carried out to establish the transferability of the trip end models over time and space and the sensitivity of the parameter estimates to the removal of individual observations.

Blainey SP (2009d) **Survey of Ultimate Trip Origins and Destinations on Rhymney Line: Technical Note**
This note describes the methodology and results from the Rhymney Line Travel Survey, giving full details of the analysis carried out on the data collected. It also contains maps of the observed catchment at each of the stations investigated.

Blainey SP (2009e) **Trip End Forecasts for 421 Potential Station Sites in England and Wales: Technical Note**
This note contains trip end forecasts made using the GWR calibration of Model 4.41 for all 421 potential sites for new stations identified in Chapter 6. It also gives the values of the associated independent variables for each of the sites.

Blainey SP (2009f) **Benefits and Costs of Construction of 14 Proposed Stations in South-East Wales: Technical Note**
This note contains full details of the flow level forecasts and benefit-cost ratios calculated for 14 sites in South-East Wales using the methodology outlined in Chapter 7.

Appendix Two: Category E and F stations in England and Wales excluded from trip end models

Group Stations:

Bedford St Johns
Dorchester West
Dorking Deepdene
Dorking West
Gainsborough Central
Gainsborough Lea Road
Liverpool James Street
Pontefract Baghill
Pontefract Monkhill
Pontefract Tanshelf
Thorne North
Thorne South
Wakefield Kirkgate
West Hampstead
West Hampstead Thameslink

Trains replaced by buses:

Barlaston
Norton Bridge
Stone
Wedgwood

No trains on weekdays:
Brigg
Kirton Lindsey
Pilning
Teesside Airport

Other:

Rice Lane
Walton (Merseyside)

These two stations are extremely close to each other, and were effectively treated as the same station when allocating output areas, which meant that Walton (Merseyside) appeared to have a catchment population of zero.

Further stations removed due to absence of car park data:

Barmouth
Garth (Powys)
Limehouse
Llandrindod
Llangadog
Llangammarch
Llanwrtyd
Pen-y-Bont

Further group stations removed as result of GWR diagnostic measures:

Farnborough North
Maidstone Barracks
Maidstone West

Appendix Three: Ultimate Origin-Destination Survey Form

This survey of passenger rail travel is being carried out by the University of Southampton, with the permission of Arriva Trains Wales. We would be grateful if you would tell us about the journey you were making when you received this questionnaire. It will only take a couple of minutes to complete. Please return it to the person who handed it to you, or leave it on your seat when you leave the train.

The information collected will only be used by the University of Southampton and Arriva Trains Wales for research into modelling rail use.

Many thanks for your help, which will assist in planning the locations of new railway stations.

1) What address have you just come from?

Property name…………..

Street name and number.

Town/city……………….

Postcode………………..

2) How did you travel from this place to the first National Rail (BR) station you used?
Please tick all methods used.

Walk……………………..○	Motorcycle…………….... ○	Underground………….....○
Bus/coach……………….. ○	Bicycle…………………... ○	Tram…………………......○
Car (driver)……………...○	Air/sea…………………... ○	
Car (passenger)………… ○	Taxi/minicab……………..○	

Other (please specify)……

3) How many minutes did it take you to travel to the first National Rail station you used?

minutes

4) Please write down every National Rail station which you are using on this journey, in the order in which you use them. Include all stations where you get on or off a train.

First station………………

Interchange station 1……...

Interchange station 2……...

Interchange station 3……...

Final station………………

5) How will you get from your final National Rail station to the address where you will finish your journey?
Please tick all methods used.

Walk……………………..○	Motorcycle…………….... ○	Underground……....…….○
Bus/coach……………….. ○	Bicycle……………….... … ○	Tram……………………. ○
Car (driver)……………...○	Air/sea……………….…... ○	
Car (passenger)…....…….○	Taxi/minicab……….....….. ○	

Other (please specify)……

6) How many minutes do you expect it to take to travel from your final National Rail station to the address where you will finish your journey?

minutes

PLEASE TURN OVER

7) What address will you travel to when you reach your destination station?

Property name…………..

Street name and number.

Town/city……………….

Postcode………………..

8) Why are you travelling to this place? Please tick ONE box only.

Going home…………….................. ……….. O
Shopping…………………....……...................O
Normal workplace……………….. ……….. O
Other workplace/meeting ………………… O
Personal business (eg doctor, bank)…...…..O
Visiting friends/relatives…………….......…. O
Sport or entertainment………………...….. O
Going to school/college………………....….O
Other (please specify)……

9) How often do you make this journey (identical to the one you are describing? Please tick ONE box only.

5 or more days a week……………..............O
2-4 days a week………....………................. O
Once a week…………...…………………...O
1-3 times a month………………………….O
Less than once a month……………………O
First time have made this journey………… O

10) What type of ticket are you using for this journey? Please tick ONE box only.

Anytime single…...……… O
Off Peak single………….. O
Anytime day return……...O
Anytime return…………. O
Off peak day return…….. O
Off peak return…………...O
Advance purchase……….O
Weekly season…………...O
Monthly season…………. O
Annual season……………O
Other (please specify)……

11) Did you use a railcard to buy your ticket?

Yes O (please answer question 11)

No O (you have finished the survey)

12) What type of railcard did you use? Please tick ONE box only.

16-25 Railcard…................................O
Senior Railcard…………...………..O
Family Railcard…………..................O
Disabled Persons Railcard.………..O
Other (please specify)……

Thankyou for taking the time to complete this questionnaire

List of References

AEA Technology Rail (2006) **Station Usage 2004/5**, AEA Technology, Derby

Aldenderfer MS & Blashfield RK, **Cluster Analysis**, Sage Quantitative Applications in the Social Sciences 44, Sage, Beverly Hills

Alker HS (1969) A Typology Of Ecological Fallacies, in Dogan M & Rokkan S (eds) **Quantitative Ecological Analysis**, MIT Press, Boston, p69-86

Armstrong J, Preston JM & Carlson M (2007) Automating the Import of Electronic Timetable Data to EMME/2-based Public Transport Models, **2nd International Seminar on Railway Operations Modelling and Analysis**, Hannover, Germany

Association of Train Operating Companies (ATOC) (2002) **Passenger Demand Forecasting Handbook**

Association of Train Operating Companies (ATOC) (2009) **Creating Communities: Expanding Access to the Rail Network**, ATOC, London

Atkinson PM, German SE, Sear DA & Clark MJ (2003) Exploring The Relations Between Riverbank Erosion And Geomorphological Controls Using Geographically Weighted Logistic Regression, **Geographical Analysis** 35(1):58-82

Atos Origin (2009) **Transport Direct**, available from URL: http://www.transportdirect.info/ , accessed 06/06/2009

Bagchi M & White PR (2005) The Potential of Public Transport Smart Card Data, **Transport Policy** 12:464-474

Balcombe R (Ed) (2003) **The Demand for Public Transport: A Practical Guide**, TRL Ltd, Crowthorne

Bates J, Polak J, Jones P & Cook A (2001) The Valuation Of Reliability For Personal Travel, **Transportation Research Part E** 37:191-229

Beckmann MJ & Golob TF (1971) A Utility Model for Travel Forecasting, **Transportation Science** 5(1):79-90

Beesley M & Foster C (1963) Estimating the Social Costs of Constructing an Underground Railway, **Journal of the Royal Statistical Society A** 126(1):46-58

Beimborn EA, Greenwald MJ & Jin X (2003) Accessibility, Connectivity and Captivity: Impacts on Transit Choice, **Transportation Research Record** 1835:1-9

Ben-Akiva M & Morikawa T (2002) Comparing Ridership Attraction Of Rail And Bus, **Transport Policy** 9:107-116

Bevan A (ed) (1998) **A-Z of Rail Reopenings**, Railway Development Society, Warwick

Bhat CR (1995) A Heteroscedastic Extreme-Value Model of Intercity Mode Choice,

Transportation Research B 29(6):471-483

Bhat CR, Eluru N & Copperman RB (2007) Flexible Model Structures for Discrete Choice Analysis, in Hensher DA & Button KJ (eds) **Handbook of Transport Modelling** (2nd ed), Elsevier, Oxford, ch. 5, pp.75-104

Bhat CR & Sardesai R (2006) The Impact Of Stop-Making And Travel Time Reliability On Commute Mode Choice, **Transportation Research Part B** 40:709-730

Blainey SP (2009a) Trip End Models of Local Rail Demand in England and Wales, **Journal of Transport Geography**, in press, doi:10.1016/j.jtrangeo.2008.11.002

Blainey SP (2009b) **Aggregate Logit Modal Split Models: Technical Note**, Transportation Research Group, University of Southampton

Blainey SP (2009c) **Sensitivity Analysis of England and Wales Trip End Models: Technical Note**, Transportation Research Group, University of Southampton

Blainey SP (2009d) **Survey of Ultimate Trip Origins and Destinations on Rhymney Line: Technical Note**, Transportation Research Group, University of Southampton

Blainey SP (2009e) **Trip End Forecasts for 421 Potential Station Sites in England and Wales: Technical Note**, Transportation Research Group, University of Southampton

Blainey SP (2009f) **Benefits and Costs of Construction of 14 Proposed Stations in South-East Wales: Technical Note**, Transportation Research Group, University of Southampton

Blainey SP & Preston JM (2009a) **Trip End Rail Demand Models for London Overground**, report to Transport for London by Transportation Research Group, University of Southampton, 25th June 2009

Blainey SP & Preston JM (2009b) Modelling Local Rail Demand in South Wales, **Journal of Transportation Planning and Technology**, in press

Boile MP, Spasovic LN & Bladikas AK (1995) Modelling Intermodal Auto-Rail Commuter Networks, **Transportation Research Record** 1516:38-47

Box D (1992) 'Costing Passenger Railways', in Harris NG & Godward EW, **Planning Passenger Railways**, Transport Publishing Company, Glossop

Bristow AL, Preston JM & Nash CA (1998) Investment Planning and Appraisal Issues in the Privatised Railway – the British Experience, **Transport Reviews** 18(4):353-362

Brown LA & Goetz AR (1987) Development-Related Contextual Effects And Individual Attributes In Third World Migration Processes: A Venezuelan Example, **Demography** 24(4):497-516

Brown LA & Jones JP III (1985) Spatial Variation In Migration Processes And Development: A Costa Rican Example Of Conventional Modelling Augmented By The Expansion Method, **Demography** 22(3):327-352

Brunsdon C, Aitken M, Fotheringham AS & Charlton ME (1999) A Comparison Of Random Coefficient Modelling And Geographically Weighted Regression For Spatially Non-Stationary Regression Problems, **Geographical & Environmental Modelling** 3:47-62

Buliung RN & Kanaroglou PS (2006) A GIS Toolkit For Exploring Geographies Of Household Activity/Travel Behaviour, **Journal Of Transport Geography** 14(1):35-51

Census Interaction Data Service (CIDS) (2007) **WICID – Information System**, available from URL: http://cids.census.ac.uk/cids/wicid/info.php?datatype=23 , accessed 30/05/2007

Clark SD (2007) Estimating Local Car Ownership Models, **Journal of Transport Geography** 15:184-197

Cole S (1998) **Applied Transport Economics: Policy, Management and Decision Making** (2nd ed), Kogan Page, London

Cox T, Houdmont J & Griffiths A (2006) Rail Passenger Crowding, Stress, Health & Safety In Britain, **Transportation Research Part A** 40:244-258

Cozens P, Neale R, Whitaker J & Hillier D (2003) Managing Crime And The Fear Of Crime At Railway Stations – A Case Study In South Wales (UK), **International Journal Of Transport Management** 1:121-132

Dargay JM & Goodwin PB (1995) Evaluation of Consumer Surplus with Dynamic Demand, **Journal of Transport Economics and Policy** 29(2):179-193

Department for Environment, Food and Rural Affairs (2007) **The Social Cost of Carbon and the Shadow Price of Carbon: What They Are, and How to Use Them in Economic Appraisal in the UK**, DEFRA, London

Department for Transport (2004) Transport Appraisal and the New Green Book, **Transport Analysis Unit 2.7.1**, available from URL: http://www.dft.gov.uk/webtag/webdocuments/2_Project_Manager/7_Transport_Appraisal_Green_Book/2.7.1.htm#1_8, accessed 18/08/2009

Department for Transport (2006a) **Transport Statistics Bulletin: Vehicle Speeds in Great Britain 2005**, Transport Statistics, London

Department for Transport (2006b) **Transport Statistics Great Britain:** Chapter 3: Energy and the Environment, available from URL:

http://www.dft.gov.uk/pgr/statistics/datatablespublications/energyenvironment/tsgbchapter3energyandtheenvi1863, accessed 06/06/2007

Department for Transport (2007a) **National Rail Travel Survey Provisional 2007 Report**, available from URL http://www.dft.gov.uk/pgr/statistics/datatablespublications/railways/NRTSProvisional2007 , accessed 25/04/2007

Department for Transport (2007b) Values of Time and Operating Costs, **Transport Analysis Guidance Unit 3.5.6**, available from URL: http://www.webtag.org.uk/webdocuments/3_Expert/5_Economy_Objective/3.5.6.htm#012, accessed 06/06/2007

Department for Transport (2007c) Guidance on Rail Appraisal, **Transport Analysis Guidance Unit 3.13.1**, available from URL: http://www.webtag.org.uk/webdocuments/3_Expert/13_Rail/3.13.1.htm#061, accessed 17/08/2009

Department for Transport (2008) **Accessible Train and Station Design for Disabled People: A Code of Practice**, Department for Transport and Transport Scotland, London/Glasgow

Deutsche Bahn AG (2009) **Planen & Buchen**, available from URL: http://reiseauskunft.bahn.de/bin/query.exe/e , accessed 06/06/2009

Dickins ISJ (1991) Park and Ride Facilities on Light Rail Transit Systems, **Transportation** 18(1):23-36

Dodgson J (2009) **Rates of Return on Public Spending on Transport**, RAC Foundation, London

Doi M & Allen WB (1986) A Time-Series Analysis of Monthly Ridership for an Urban Rail Rapid Transit Line, **Transportation** 13(3) 257-269

Du H & Mulley C (2006) Relationship Between Transport Accessibility and Land Value: Local Model Approach with Geographically Weighted Regression, **Transportation Research Record** 1977:197-205

Dueker KJ & Ton T (2000) Geographical Information Systems for Transport, in Hensher DA & Button KJ, **Handbook of Transport Modelling**, Pergamon Press, Oxford

Duncan C & Jones K (2000) Using Multilevel Models to Model Heterogeneity: Potential and Pitfalls, **Geographical Analysis** 32:279-305

Eldridge JD & Jones JP III (1991) Warped Space: A Geography of Distance Decay, **Professional Geographer** 43:500-511

Evans AW (1969) Intercity Travel and the London Midland Electrification, **Journal of**

Transport Economics and Policy 3:69-95
Evans AW (2006) **Cost-Benefit Analysis**, MSc Lecture, University College London
Everitt BS, Landau S & Leese M (2001) **Cluster Analysis**, 4th ed, Arnold, London
Experian (2006) **Mosaic United Kingdom: The Consumer Classification For The UK**, Nottingham, Experian Ltd
Faber Maunsell (2007) **Imperial Wharf Business Case Narrative**, 23/03/2007
Farhan B & Murray AT (2005) A GIS-Based Approach For Delineating Market Areas For Park And Ride Facilities, **Transactions in GIS** 9(2):91-108
Flyvbjerg B, Holm MKS & Buhl SL (2005) How (In)Accurate are Demand Forecasts in Public Works Projects?, **Journal of the American Planning Association** 71(2):131-146
Fotheringham AS (1981) Spatial Structure & Distance Decay Parameters, **Annals Of The Association Of American Geographers** 71:425-436
Fotheringham AS (1984) Spatial Flows and Spatial Patterns, **Environment & Planning A** 16:529-543
Fotheringham AS (1999) Guest Editorial: Local Modelling, **Geographical and Environmental Modelling** 3:5-7
Fotheringham AS & Pitts TC (1995) Directional Variation In Distance Decay, **Environment & Planning A** 27:715-729
Fotheringham AS, Brunsdon C & Charlton ME (2002) **Geographically Weighted Regression**, John Wiley & Sons, Chichester
Georgiou G (2008) **Station Usage 2006/07**, DeltaRail Ltd, London
Goldstein H (1987) **Multilevel Models in Education and Social Research**, Charles Griffin & Co, London
Gorr WL & Olligschlaeger AM (1994) Weighted Spatial Adaptive Filtering: Monte Carlo Studies and Application to Illicit Drug Market Modelling, **Geographical Analysis** 26:1:67-87
Grant-Muller SM, Mackie P, Nellthorp J & Pearman A (2001) Economic Appraisal of European Transport Projects: The State-of-the-Art Revisited, **Transport Reviews** 21(2):237-261
Greene WH (1993) **Econometric Analysis** (2nd ed), Macmillan, New York
Hadj-Chikh GJ & Thompson GL (1998) Reaching Jobs in the Suburbs: Tri-Rail in South Florida, **Transportation Research Record** 1618:14-21
Hagerstrand T (1965) A Monte Carlo Approach To Diffusion, **European Journal of Sociology** 6:43-67

Halcrow Group (2006) **RPP Case Study: South Hants Crossrail**, unpublished report

Harrigan KR (1985) An Application of Clustering for Strategic Group Analysis, **Strategic Management Journal** 6(1):55-73

Harrison PJ (1968) A Method of Cluster Analysis and Some Applications, **Applied Statistics** 17:226-236

Hauser RM (1970) Context and Consex: A Cautionary Tale, **The American Journal of Sociology** 75(4):645-664

Hawkins DM, Muller MW & ten Krooden JA (1982) ‘Cluster Analysis', in **Topics in Applied Multivariate Analysis** (Hawkins DM, ed), Cambridge University Press, Cambridge

Hearnshaw HM & Unwin D (eds) (1994) **Visualization in Geographical Information Systems**, John Wiley & Sons, Chichester

Heywood R (2007) Britain’s National Rail Network: Fit for Purpose in the 21st Century?, **Journal of Transport Geography** 15:198-216

Hicks JR (1942) The Rehabilitation of Consumers’ Surplus, **The Review of Economic Studies** 8(2):108-116

Hill M (1968) A Goals-Achievement Matrix for Evaluating Alternative Plans, **Journal of the American Institute of Planners** 34(1):19-29

Hirshleifer J (1958) On the Theory of Optimal Investment Decision, **The Journal of Political Economy** 66(4):329-352

HM Treasury (2003) **The Green Book: Appraisal and Evaluation in Central Government**, The Stationery Office, London

Holvad T (2004) **Cost-Benefit Analysis**, BA Lecture, School of Geography, University of Oxford

Horner MW & Grubesic TH (2001) A GIS-based Planning Approach to Locating Urban Rail Terminals, **Transportation** 28:55-77

Hsiao S, Lu J, Sterling J & Weatherford M (1997) Use of Geographic Information System for Analysis of Transit Pedestrian Access, **Transportation Research Record** 1604:50-59

Hurvich CM, Simonoff JS, & Tsai CL (1998) Smoothing Parameter Selection In Nonparametric Regression Using An Improved Akaike Information Criterion, **Journal Of The Royal Statistical Society: Series B (Statistical Methodology)** 60(2):271-293

Johnson D, Shires J, Nash CA & Tyler J (2006) Forecasting And Appraising The Impact Of A Regular Interval Timetable, **Transport Policy** 13:349-366

Jones JP III & Casetti E (1992) **Applications Of The Expansion Method**, Routledge, London

Jones JP III & Hanham RQ (1995) Contingency, Realism, and the Expansion Method, **Geographical Analysis** 27:185-207

Kaldor N (1955) **An Expenditure Tax**, Unwin University Books, London

Kanafani A (1983) **Transportation Demand Analysis**, McGraw Hill Book Co, New York

Kim S, Ulfersson GF & Hennessy JT (2007) Analysis of Light Rail Rider Travel Behaviour: Impacts of Individual, Built Environment and Crime Characteristics on Transit Access, **Transportation Research A** 41:511-522

Koppelman FS (2007) Closed Form Discrete Choice Models, in Hensher DA & Button KJ (eds) **Handbook of Transport Modelling** (2nd ed), Elsevier, Oxford, ch.13, pp.257-277

Koppelman FS & Wen CH (2000) The Paired Combinatorial Logit Model: Properties, Estimation and Application, **Transportation Research B**: 34(2):75-89

Krygsman S, Dijst M & Arentze T (2004) Multimodal Public Transport: An Analysis Of Travel Time Elements And The Interconnectivity Ratio, **Transport Policy** 11:265-275

Kuby M, Barranda A & Upchurch C (2004) Factors Influencing Light-Rail Station Boardings in the United States, **Transportation Research A** 38:223-247

Laird J, Geurs K & Nash CA (2009) Option and Non-use Values and Rail Project Appraisal, **Transport Policy** 16(4):173-182

Lane C, DiCarlantonio M & Usvyat L (2006) Sketch Models to Forecast Commuter and Light Rail Ridership: Update to TCRP Report 16, **Transportation Research Record** 1986:198-210

Lasch P, Haensch W, Naumann D & Diem M (2004) Imaging of Colorectal Adenocarcinoma Using FT-IR Microspectroscopy and Cluster Analysis, **Biochimica et Biophysica Acta** 1688(2):176-186

Layard R & Glaister S (eds) (1994) **Cost-Benefit Analysis** (2nd ed), Cambridge University Press, Cambridge

Levinson DM & Huang Y (1997) Windowed Transportation Planning Model, **Transportation Research Record** 1607:45-54

Lloyd C & Shuttleworth I (2005) Analysing Commuting Using Local Regression Techniques: Scale, Sensitivity and Geographical Patterning, **Environment and Planning A** 37(1):81-103

Lyons G, Jain J & Holley D (2007) The Use Of Travel Time By Rail Passengers in Great

Britain, **Transportation Research Part A** 41:107-120
Lythgoe WF (2004) **Enhancing Cross-Sectional Rail Passenger Demand Models,** PhD thesis, Institute for Transport Studies, University of Leeds
Lythgoe WF & Wardman M (2002) Demand for Rail Travel To and From Airports, **Transportation** 29:125-143
Lythgoe WF & Wardman M (2004) Modelling Passenger Demand for Parkway Rail Stations, **Transportation** 31:125-151
Mackie P & Preston JM (1998) Twenty-one Source of Error and Bias in Transport Project Appraisal, **Transport Policy** 5(1):1-7
Maddala GS (2001) **Introduction to Econometrics** (3rd ed), John Wiley & Sons, Chichester
Mandel B, Gaudry M & Rothengatter W (1997) A Disaggregate Box-Cox Logit Mode Choice Model of Intercity Passenger Travel in Germany and its Implications for High-Speed Rail Demand Forecasts, **Annals of Regional Science** 31:99-120
MapInfo (2009) **MapInfo Products Knowledge Base: Grid Interpolation**, available from URL: http://testdrive.mapinfo.com/TECHSUPP/MIPROD.NSF/0/ced04a3ded07815885256c4700597b25?OpenDocument , accessed 20/03/2009
Marshall N & Grady B (2006) Sketch Transit Modelling Based on 2000 Census Data, **Transportation Research Record** 1986:182-189
Maruyama T, Muromachi Y, Harata N & Ohta K (2001) The Combined Modal Split / Assignment Model in the Tokyo Metropolitan Area, **Journal of the Eastern Asia Society for Transportation Studies** 4(2):293-304
McCarthy PS (2001) **Transportation Economics, Theory and Practice: A Case Study Approach**, Blackwell, Oxford
McDonough CC (1973) The Demand for Commuter Rail Transport, **Journal of Transport Economics and Policy** 7:134-143
McFadden D (1978) Modelling the Choice of Residential Location, **Transportation Research Record** 672:72-77
McGeehan H (1984) Forecasting the Demand for Inter-Urban Railway Travel in the Republic of Ireland, **Journal of Transport Economics and Policy** 18:275-291
McMillen DP (1996) One Hundred Fifty Years of Land Values in Chicago: A Nonparametric Approach, **Journal Of Urban Economics** 40:100-124
Morisugi H (2000) Evaluation Methodologies of Transportation Projects in Japan, **Transport Policy** 7(1):35-40

Nakamura H (2000) The Economic Evaluation of Transport Infrastructure: Needs for International Comparisons, **Transport Policy** 7(1):3-6

Nash CA (1997) Cost-Benefit Analysis of Transport Projects, in Oum TH, Dodgson JS, Hensher DA, Morrison SA, Nash CA, Small KA & Waters WG, **Transport Economics: Selected Readings**, Harwood, Amsterdam, pp. 627-651

Nash CA & Preston JM (1991) Appraisal of Rail Investment Projects: Recent British Experience, **Transport Reviews** 11(4):295-309

Nellthorp J & Mackie P (2000) The UK Roads Review – a Hedonic Model of Decision Making, **Transport Policy** 7(2):127-138

Network Rail (2006) **South West Main Line Route Utilisation Strategy**, Network Rail, London

Network Rail (2007a) **Scotland Route Utilisation Strategy**, Network Rail, London

Network Rail (2007b) **North West Route Utilisation Strategy**, Network Rail, London

Network Rail (2007c) **Greater Anglia Route Utilisation Strategy**, Network Rail, London

Network Rail (2008a) **National Rail Timetable December 2008 – May 2009**, Network Rail, London

Network Rail (2008b) **East Coast Main Line Route Utilisation Strategy**, Network Rail, London

Nielsen OA (2003) Estimation of Random Coefficient Models on GPS Data, **European Transport Conference**, Strasbourg, 8-10 October 2003

Ochieng WO (2006) **Navigation Technologies II**, MSc Lecture, Imperial College, London

O'Fallon C, Sullivan C & Hensher DA (2004) Constraints Affecting Mode Choice By Morning Car Commuters, **Transport Policy** 11(1):17-29

Office of the Deputy Prime Minister (ODPM) (2004), **Planning Policy Guidance No. 11: Regional Planning: Supplementary Guidance on Regional Spatial Strategies,** ODPM, London

Office of the Rail Regulator (ORR) (2006) **The Stations Code**, ORR, London

Ord K (1975) Estimation Methods For Models Of Spatial Interaction, **Journal Of The American Statistical Association** 70:120-126

Ortuzar J (1983) Nested Logit Models for Mixed-Mode Travel in Urban Corridors, **Transportation Research A** 17:283-299

Ortuzar J de D & Willumsen LG (2001) **Modelling Transport** (3rd ed.), John Wiley & Sons Ltd, Chichester

Owen AD & Phillips GDA (1987) The Characteristics Of Railway Passenger Demand, **Journal Of Transport Economics And Policy** 21:231-253

Páez A (2006) Exploring Contextual Variations In Land Use And Transport Policy Using A Probit Model With Geographical Weights, **Journal of Transport Geography** 14:167-176

Paulley N, Balcombe R, Mackett R, Titheridge H, Preston JM, Wardman M, Shires J & White P (2006) The Demand For Public Transport: The Effects Of Fares, Quality Of Service, Income and Car Ownership, **Transport Policy** 13:295-306

Preston JM (1987) **The Evaluation of New Local Rail Stations in West Yorkshire**, School of Economic Studies (Institute for Transport Studies), University of Leeds, PhD thesis

Preston JM (1991a) Demand Forecasting For New Local Rail Stations And Services, **Journal Of Transport Economics And Policy** 25(2):183-202

Preston JM (1991b) **Passenger Demand Forecasting for New Rail Services: Manual Of Advice**, Institute for Transport Studies, University of Leeds, Working Paper 352

Preston JM (2001) **New Local Rail Stations And Services: Twenty Years On,** World Conference on Transport Research, Seoul

Preston JM & Dargay J (2005) **The Dynamics of Rail Demand**, Third Conference on Railroad Industry Structure, Competition and Investment, Stockholm School of Economics, 20-22 October 2005

Preston JM, Blainey SP, Wall G, Chintakayala PK & Wardman M (2008) **The Effect of Station Enhancements on Rail Demand**, European Transport Conference, Leiden, 6-8 October 2008

Price M (1999) A New Approach to the Appraisal of Road Projects, **Journal of Transport Economics and Policy**, 33(2):221-226

Quinet E (2000) Evaluation Methodologies of Transportation Projects in France, **Transport Policy** 7(1):27-34

Railtrack PLC (1996) **Great Britain Passenger Rail Timetable May – September 1996**, Railtrack, London

Railtrack PLC (2001) **National Rail Timetable May – September 2001**, Railtrack, London

Railtrack PLC (2002) **National Rail Timetable September 2002 - May 2003**, Railtrack, London

Rickard JM (1988) Factors Influencing Long Distance Rail Passenger Trip Rates in Great Britain, **Journal of Transport Economics and Policy** 22:209-233

Robinson WS (1950) Ecological Correlations And The Behaviour of Individuals, **American Sociological Review** 15(3):351-357

Rose G (1986) Transit Passenger Response: Short and Long Term Elasticities Using Time Series Analysis, **Transportation** 13(2):131-141

Rothengatter W (2000) Evaluation of Infrastructure Investments in Germany, **Transport Policy** 7(1):17-25

Rousseeuw PJ & van Zomeren BC (1990) Unmasking Multivariate Outliers and Leverage Points, **Journal of the American Statistical Association** 85(411):633-639

Ruiter ER (1967) Toward a Better Understanding of the Intervening Opportunities Model, **Transportation Research** 1(1):47-56

Sansom T, Nash C, Mackie P, Shires J & Watkiss P (2001) **Surface Transport Costs and Charges Great Britain 1998**, report commissioned by the Department of the Environment, Transport and the Regions, Institute for Transport Studies, University of Leeds

Sargious MA & Janarthanan N (1983) Forecasting Demand For The Park And Ride Mode And Determining The Optimal Location Of Stations, **Canadian Journal Of Civil Engineering** 10:695-702

Sayers TM, Jessop AT & Hills PJ (2003) Multi-criteria Evaluation of Transport Options – Flexible, Transparent and User-Friendly?, **Transport Policy** 10(2):95-105

Scott Wilson (2006) **Beeston Station Feasibility Study: Final Report**, Scott Wilson Railways Ltd, York

Self P (1970) "Nonsense on Stilts": Cost-Benefit Analysis and the Roskill Commission, **The Political Quarterly** 41(3):249-260

Shah K (2007) **Station Usage 2005/06**, DeltaRail Group Ltd, London

Sheskin IM & Stopher PR (1988) Spatial Variations in Attitudes Toward Expanded Public Transit Services, **Transportation** 15(3):211-232

Shilton D (1982) Modelling The Demand For High-Speed Train Services, **Journal Of The Operational Research Society** 33(8):713-722

Simma A & Axhausen KW (2001) Structures Of Commitment In Mode Use: A Comparison Of Switzerland, Germany And Great Britain, **Transport Policy** 8:279-288

Simpson EH (1951) The Interpretation of Interaction in Contingency Tables, **Journal of The Royal Statistical Society: Series B (Methodological)** 13(2):238-241

Sphaera Interactive Media (2002-6) **The Heart Of Wales Line**, available from URL: http://www.heart-of-wales.co.uk/howlframe.htm , accessed 20/11/2006

Spiekermann K & Wegener M (2000) Freedom From The Tyranny Of Zones: Towards New GIS-Based Spatial Models, in Fotheringham AS & Wegener M (eds) **Spatial**

Models & GIS: New Potential and New Models, Taylor & Francis, London, pp45-61

Stouffer SA (1940) Intervening Opportunities: A Theory Relating Mobility and Distance, **American Sociological Review** 5(6):845-867

Strategic Rail Authority (2003) **Appraisal Criteria**, SRA, London

Talvitie A (1973) A Direct Demand Model for Downtown Work Trips; **Transportation** 2(2):121-152

Tang KX & Waters NM (2005) The Internet, GIS and Public Participation in Transportation Planning, **Progress in Planning** 64:7-62

Timmermans H, Waerden P van d, Alves M, Polak J, Ellis S, Harvey AS, Kurose S & Zandee R (2003) Spatial Context And The Complexity of Daily Travel Patterns: An International Comparison, **Journal Of Transport Geography** 11(1):37-46

Tsamboulas D, Golias J & Vlahoyannis M (1992) Model Development for Metro Station Access Mode Choice, **Transportation** 19(3):231-244

Upchurch C, Kuby M, Zoldak M & Barranda A (2004) Using GIS To Generate Mutually Exclusive Service Areas Linking Travel On And Off A Network, **Journal of Transport Geography** 12(1):23-33

Van Vuuren D & Rietveld P (2002) The Off-Peak Demand For Train Kilometres And Train Tickets: A Microeconomic Analysis, **Journal Of Transport Economics And Policy** 36(1):49-72

Vickerman R (2000) Evaluation Methodologies for Transport Projects in the United Kingdom, **Transport Policy** 7(1):7-16

Voshva P (1997) Application of Cross-Nested Logit Model to Mode Choice in Tel Aviv, Israel, Metropolitan Area, **Transportation Research Record** 1607:6-15

Vuchic VR (1966) **Interstation Spacing For Line-Haul Passenger Transportation**, Institute of Transportation and Traffic Engineering, Berkeley, University of California, PhD thesis

Walters G & Cervero R (2003) **Forecasting Transit Demand in a Fast-Growing Corridor: The Direct-Ridership Model Approach**

Wardman M (1994) Forecasting the Impact of Service Quality Changes on the Demand for Inter-Urban Rail Travel, **Journal of Transport Economics and Policy** 28:287-306

Wardman M (1997) Inter-Urban Rail Demand Elasticities and Competition in Great Britain: Evidence From Direct Demand Models, **Transportation Research E** 33(1):15-28

Wardman M, Toner JP & Whelan GA (1997) Interactions Between Rail and Car in the

Inter-Urban Leisure Travel Market in Great Britain, **Journal of Transport Economics and Policy** 31:163-181

Wardman M & Whelan G (2001) Valuation Of Improved Railway Rolling Stock: A Review Of The Literature And New Evidence, **Transport Reviews** 21(4):415-447

Wardman M & Tyler J (2000) Rail Network Accessibility And The Demand For Inter-Urban Rail Travel, **Transport Reviews** 20(1):3-24

Wardman M, Shires J, Lythgoe WF & Tyler J (2004) Consumer Benefits And Demand Impacts Of Regular Train Timetables, **International Journal Of Transport Management**, 2(1):39-49

Wardman M, Lythgoe WF & Whelan G (2007) Rail Passenger Demand Forecasting: Cross-Sectional Models Revisited, in Dennis SM & Talley WK (eds) **Research in Transportation Economics,** Ch 5, Elsevier, Amsterdam

Waterson BJ (2009) **Cluster Analysis**, Transportation Research Group Seminar, University of Southampton, 17th March 2009

Wen CH & Koppelman FS (2001) The Generalised Nested Logit Model, **Transportation Research B**, 35(7):627-641

Westin RB (1974) Predictions From Binary Choice Models, **Journal of Econometrics** 2(1):1-16

Wheat P & Smith ASJ (2008) Assessing the Marginal Infrastructure Maintenance Wear and Tear Costs for Britain's Railway Network, **Journal of Transport Economics and Policy** 42(2):189-224

Whelan G & Johnson D (2004) Modelling The Impact Of Alternative Fare Structures On Train Overcrowding, **International Journal Of Transport Management** 2:51-58

Whelan G & Wardman M (1999a) **Disaggregate Models of Rail Travellers' Station and Access Mode Choices**, Working Paper 529, Institute for Transport Studies, University of Leeds

Whelan G & Wardman M (1999b) **Improved Rail Demand Models Using Geographical Information Systems**, Working Paper 527, Institute for Transport Studies, University of Leeds

Wootton HJ & Pick GW (1967) A Model for Trips Generated by Households, **Journal of Transport Economics and Policy** 1:137-153

Yang H & Bell MGH (1998) Models and Algorithms for Road Network Design: a Review and Some New Developments, **Transport Reviews** 18(3):257-278

Zhao F, Chow L-F, Li M-T, Ubaka J & Gan A (2003) Forecasting Transit Walk Accessibility: Regression Model Alternative to Buffer Method, **Transportation**

Research Record 1835:34-41

All map data ©Crown Copyright/database right 2009. An Ordnance Survey/EDINA supplied service.

This work is based on data provided through EDINA UKBORDERS with the support of the ESRC and JISC and uses boundary material which is copyright of the Crown.

This work is based on data provided with the support of the ESRC.

LENNON data supplied by Arriva Trains Wales

GMATS data provided by Greater Manchester PTE

Census output is Crown copyright and is reproduced with the permission of the Controller of HMSO and the Queen's Printer for Scotland. Source: 2001 Census; Key Statistics; Special Workplace Statistics (Level 2)