Fare Policy Analysis for Public Transport: A Discrete-Continuous Modeling Approach Using Panel Data

by

Hazem Marwan Zureiqat

B.A. in Economics
Macalester College, 2006

Submitted to the Department of Civil and Environmental Engineering
in Partial Fulfillment of the Requirements for the Degree of

Master of Science in Transportation
at the
Massachusetts Institute of Technology

June 2008

Signature of Author ……………………………………………………………………
Department of Civil and Environmental Engineering
May 16, 2008

Certified by …………………………………………………………….....................
Nigel H. M. Wilson
Professor of Civil and Environmental Engineering
Director, Master of Science in Transportation Program
Thesis Supervisor

Accepted by …………………………………………………………………………….
Daniele Veneziano
Chairman, Departmental Committee for Graduate Students

Fare Policy Analysis for Public Transport:
A Discrete-Continuous Modeling Approach
Using Panel Data

by

Hazem Marwan Zureiqat

Abstract

In many large metropolitan areas, public transport is very heavily used, and ridership is approaching system capacity in the peak periods. This has caused a shift in attention by agency decision-makers to strategies that can more effectively manage the demand for public transport, rather than simply increase overall demand. In other words, a need has arisen to understand not only *why* people use public transport as opposed to other modes but also *how* they use public transport, in terms of their ticket, mode, and time-of-day choices. To that end, fares become an increasingly important policy tool that can trigger certain behavioral changes among riders.

This thesis develops a methodology to model, at the disaggregate level, the response of public transport users to fare changes. A discrete-continuous framework is proposed in which ticket choice is modeled at the higher (discrete) level and frequencies of public transport use, based on mode and time-of-day, are modeled at the lower (continuous) level. This framework takes advantage of the availability of smartcard data over time, allowing individual-specific behavioral changes with various fare policies to be captured.

This methodology is applied to London's public transport system using Oyster smartcard data collected between November 2005 and February 2008. The results indicate a strong inertia effect in terms of ticket choice among public transport users in London. An individual's prior ticket choice is found to be a very important factor in determining their future ticket choice. This is also evident when we simulate the effects of two policy changes on ticket choices. We find that the impact of changing the prices of period tickets may take several months or more to fully materialize.

In terms of the frequency of public transport use, the results indicate estimated short and long-run fare elasticities of -0.40 and -0.64, respectively, for travel on the London Underground and equivalent estimates of -0.08 and -0.13 for travel on bus. The estimated Underground fare elasticities are comparable to those in the literature. The bus fare elasticities, on the other hand, are relatively smaller, in absolute value, than prior estimates. This difference reflects the small variations in bus fares in the dataset on which the model was estimated and the low fare sensitivity for users under such variations.

Furthermore, we apply the model, in conjunction with related assumptions and findings from previous research, to evaluate an AM peak pricing scheme on the London Underground, in which travelers are charged £2.00 between 8:30am and 9:15am, rather than the current fare of £1.50. This application estimates that such a policy could potentially decrease AM "peak-of-the-peak" demand on the Underground by about 9%, with the reduction in ridership shifting either to a different mode or to a different time period.

Thesis Supervisor: Nigel H. M. Wilson
Title: Professor of Civil and Environmental Engineering

Acknowledgements

First, I would like to thank my research advisors, Professor Nigel Wilson and John Attanucci, for their guidance over the past year. Our weekly meetings not only challenged and enriched my research experience but also encouraged me to go beyond the technical aspects of this work and to present it as a practical tool that can truly serve its intended purpose. Thank you both for providing me with the opportunity to work with you on this topic.

I would also like to express my deepest gratitude to Professor Moshe Ben-Akiva, who was always willing to meet and discuss my work. His advice was always thorough and rigorous, and his thoughtful insight provided the basis for the methodology and modeling framework in this thesis.

I am extremely grateful to Fred Salvucci, with whom I did research during my first year at MIT. It was an honor for me to work with and get to know such an inspirational person and a renowned public servant. My interactions with Fred over the past two years truly made me passionate about the field of transportation, in general, and about public transport, in particular.

I would also like to thank Mikel Murga, George Kocur, Chris Zegras, and members of the Transit Research Group at MIT for their help and support, as well as their insightful feedback over the past year. Special thanks to my officemates, David Block-Schachter and Martin Milkovits, and to Maya Abou Zeid, Justin Antos, Michael Frumin, Zhan Guo, Tejus Kothari, Carlos Mojica, Catherine Seaborn, Sean Sweat, David Uniman, and Jinhua Zhao. And many thanks to Ginny Siggia for all the administrative support she has provided throughout the past two years.

This research would not have been possible without the generous support of Transport for London. I would like to especially thank Tony Richardson, who took the time to go over this thesis and send me his feedback, and Lauren Sager Weinstein, Shashi Verma, Malcolm Fairhurst, Geoff Hobbs, Pauline Matkins, Mike Collop, Tej Hunjan, Dale Campbell, and Colin Shepherd for all their help both during my summer internship at TfL and as I was working on this thesis.

To all my friends at MIT and elsewhere, thank you for making the last two years in Boston very enjoyable.

Finally, I would like to thank my parents, Marwan and Abla, my brother, Tarek, and my sister, Dana, and her family—Ammar, Jad, and Eddie—for always being there. Without my family's love and continuous support and encouragement, I would not be where I am today.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

With the introduction of new ticketing technologies and the increasing complexity of fare structures in public transport systems worldwide, fares have become an ever more important tool in managing the demand for public transport. This thesis presents a fare policy-sensitive framework to model that demand and applies the framework to the public transport system in London using a dataset of smartcard users over time.

This chapter presents the research motivation, objectives, and approach, as well as a brief introduction to London's public transport system.

1.1 Motivation

Most studies and models that have examined public transport demand are concerned with the overall level of that demand in the context of a multi-modal transport system. This involves looking into the factors that affect mode choice (public transport versus other modes), for example, or how overall demand for public transport responds to changes in factors such as fares or service levels.

In many large metropolitan areas, public transport is very heavily used, and ridership is approaching system capacity in the peak periods. This has caused a shift in attention to strategies that can more effectively manage the demand for public transport, rather than simply increase the overall level of that demand. In other words, a need has arisen to understand not only *why* people use public transport as opposed to other modes but also *how* they use public transport. To that end, fares become an increasingly important policy tool that can trigger certain behavioral changes. Such changes include, but are not restricted to:

1. The choice among the various public transport modes: If rail and bus, for example, are two viable options for a certain trip, adjusting relative fares between these two modes can influence an individual's public transport mode choice.

2. The choice of ticketing medium and ticket type: Fare differentials can be used to influence people's choice among smartcards and paper tickets, for example, and among the different ticket types available on a public transport system. These ticket types may include a per journey fare, a weekly period ticket, or a monthly period ticket.

3. Time-of-travel choice: Introducing fares that vary by time of day can shift people's departure times for certain types of trips.

4. Path choice: Although this has not been widely discussed and may not be applicable in the foreseeable future, charging path-based fares may be an effective way to spread public transport demand across different routes in a public transport system. Currently, fares are either flat or are charged based on the origin and the destination of a journey; however, given the available technologies and by installing additional equipment, it is conceivable to charge people different fares for the same origin and destination based on the paths they choose.

This thesis will study the use of fares as a policy instrument that affects three of the choices listed above: mode choice, ticket choice (i.e. the choice of a ticket type), and time-of-day choice. The choice of a ticketing medium will be addressed in the analysis but will not be modeled explicitly. Path choice will not be considered in this thesis.

To conclude, the motivation behind this research lies in the need for effective demand management strategies for public transport systems—strategies using fares as a policy instrument to create behavioral changes that would lead to more effective use of existing networks. In order to develop such strategies, a modeling framework is needed to assess people's responsiveness to changes in fare policy. (For a comprehensive, non-technical overview of public transport fare policy, see Fleishman *et al.*, 1996.)

1.2 Research Objectives

The objective of this thesis is to develop a robust, policy-sensitive tool to assess the impacts of changes public transport fares on mode, ticket, and time-of-day choices. This will be done through the development of a modeling framework that can capture people's responsiveness to changes in the fare structure with respect to these three choice dimensions. We will specifically address the following questions:

1. Why do people choose certain ticket types over others? How important are the monetary prices of the various ticket types in making that choice?
2. How do holders of the various ticket types use public transport? What modes do they use? At what times of day do they choose to make their trips? And how important are fares in determining these decisions?

Answering these questions will provide a better understanding of what fare changes needed to be made in order to make more effective use of the existing public transport network and infrastructure.

1.3 Research Approach

The methodology developed in this thesis is based on using longitudinal smartcard (or automated fare collection systems--AFC) data to measure people's responsiveness to fare changes. Having such disaggregate data allows us to 'follow' individuals over time and observe changes in their mode, ticket, and time-of-day choices.

The premise behind using such an approach is that individuals' past behavior plays an important role in determining their current and future behavior. To that end, changes in fare policy should not be expected to have an immediate effect on all users of the system, as past behavior (in terms of ticket, mode, and time-of-day choices) can represent a form of inertia that causes a delayed effect or, in aggregate terms, a gradual shift in behavior.

The methodology we propose uses statistical models that have been applied in the past in transportation as well as other disciplines. Applying our methodology requires preprocessing the raw smartcard data by constructing a panel structure in

which individuals (represented by cards) are observed over multiple pre-defined time intervals. The data preprocessing, described in detail in Section 5.1, requires some aggregation from the journey level to the individual level. This aggregation allows for conducting the analysis without the need to consider choices and attributes at the origin-destination level, which are often not necessary in a fare policy context, while maintaining a disaggregate dataset.

Since the methodology will be applied to the public transport system in London in this thesis, it may be useful to present a brief overview of public transport services in London. This is the topic of the next section.

1.4 Overview of London's Public Transport System

Public transport in London consists of an extensive network of bus and rail lines, the majority of which are planned and regulated and, in some cases, operated by Transport for London (TfL), a local government body created in 2000 as part of the Greater London Authority.

The public transport network in London includes the following modes:

1. Buses—consisting of a fleet of 6,800 scheduled buses[1].
2. The Underground (commonly known as 'the Tube')—consisting of eleven heavy rail lines that operate mostly underground.
3. National Rail—consisting of rail lines that run beyond the boundaries of the Greater London area and are operated by various train operating companies.
   – In November 11, 2007, TfL became the contracting authority for a portion of National Rail services that primarily serves East and North London. These services constitute what is known as the 'London Overground', a brand which has now become part of the TfL family.
4. The Docklands Light Rail (DLR)—a light rail line that serves East London.
5. Tramlink—consisting of three tram routes that operate in South London.

According to the 2007 London Travel Report, around 27.6 million journey segments were made in London on an average weekday in 2006 (a journey segment

---

[1] Source: http://www.tfl.gov.uk/corporate/modesoftransport/1548.aspx, Retrieved February 15, 2008.

is defined as a one-way movement in which a single mode is used). As shown in Figure 1-1, a significant proportion (around 37%) of these journey segments was made on public transport modes. The figure shows that bus is by far the most widely used public transport mode in London, followed by the Underground, National Rail, and then DLR. (Note that Tramlink journeys, which constitute a very small proportion of total journey segments, are included in the 'bus' category in Figure 1-1.) Figure 1-2, on the other hand, illustrates the significance of public transport in central London. The figure shows the mode split as a percentage of the 1.1 million people entering central London on an average weekday in 2006. It indicates that close to 90% of those people entering central London in the morning peak do so on a public transport mode.

It should also be noted that the mode share of public transport in London has been steadily increasing over time. This is shown in Figure 1-3, which indicates that public transport mode share has risen from about 29% in 1993 to 37% in 2006. This increase is partly the result of service improvements throughout the system but is also due to various policies that were meant both to encourage the use of public transport and to discourage the use of the car. (Figure 1-3 shows a corresponding significant decrease in the mode share for car between 1993 and 2006.) One such policy involves a congestion charge which was imposed in 2003 (and revised in 2007) on all cars entering central London between 7am and 6pm on weekdays.
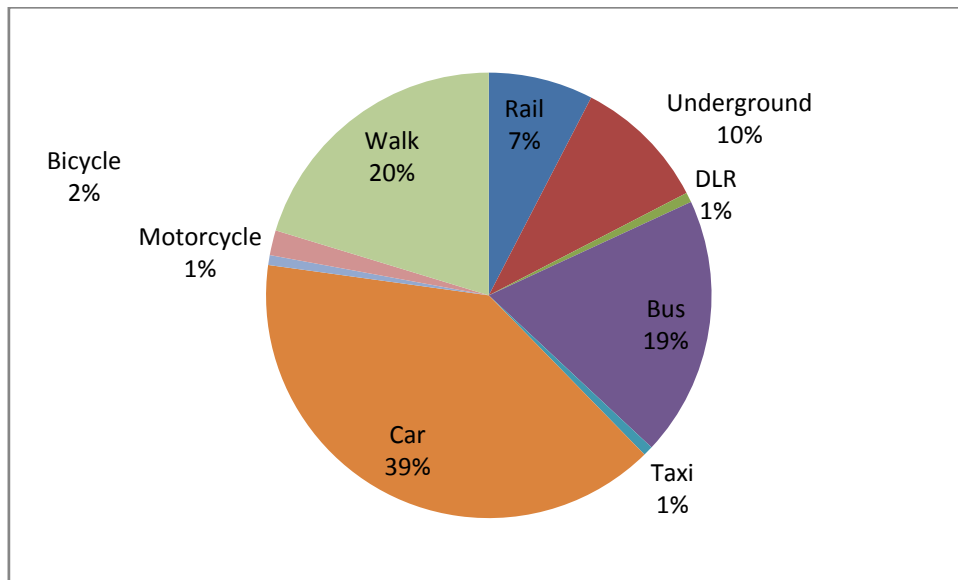
Figure 1-1: Mode split in London on an average weekday in 2006
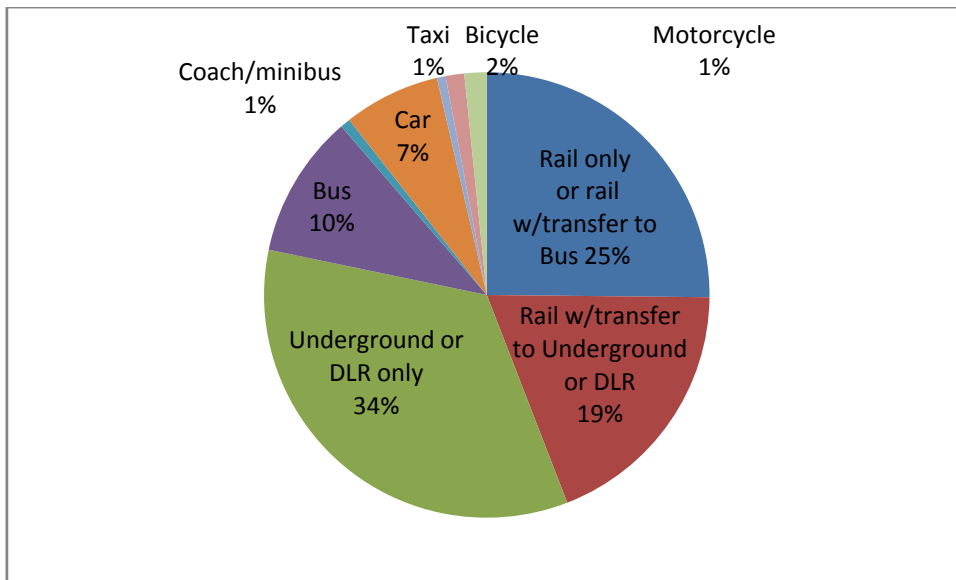
Figure 1-2: Mode split for journeys to central London in the morning peak (7am-10am) on an average weekday in 2006
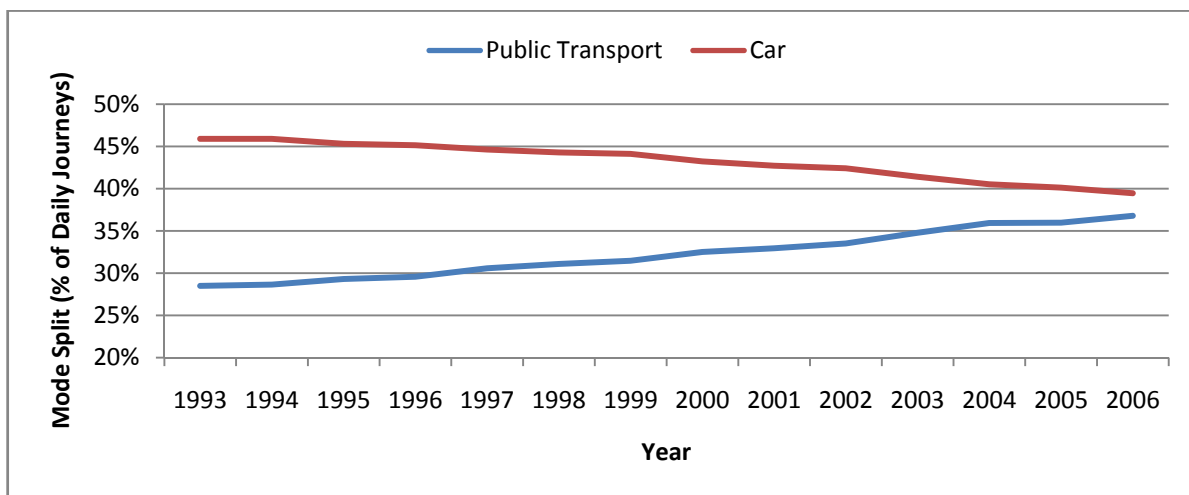
Figure 1-3: Public transport and car mode shares in London over time

Another policy meant to encourage the use of public transport was the launching of the Oyster smartcard as a new ticketing medium in London in December 2003. The launch of the Oyster card, which is now accepted on the Underground, buses, the DLR, Tramlink, and in several National Rail stations, was accompanied by subsequent fare policy changes which were meant to encourage its adoption. These changes will be discussed in more detail in Section 3.2. It should be noted at the outset, however, that all the data on which the models developed in

this thesis are estimated come from transactions made using Oyster cards at Underground station gates and on board buses.

1.5 Thesis Organization

This thesis is divided into six chapters. Chapter 2 reviews previous research on public transport fare policy and models. We discuss papers from the literature that are most relevant to the research objectives and present a brief review of the methodologies used in these papers.

Chapter 3 discusses fare policy at Transport for London. It includes a detailed description of the current fare structure in London, presents patterns and trends relevant to this analysis, and explains how TfL currently assesses the impacts of fare changes on demand and revenues.

In Chapter 4, we present our proposed methodology. The chapter starts with a conceptual framework, after which we present the proposed model structure, and the theoretical background that pertains to that structure.

Chapter 5 includes the application of the proposed methodology to London's public transport system. We describe the preprocessing of the Oyster data and present the empirical model specification and estimation results. The chapter concludes with some policy applications and an overview of how the estimation results could be incorporated into models currently used at TfL.

Chapter 6 concludes this thesis with a summary of the findings, as well as limitations and suggestions for future research.

Chapter 2

Literature Review

This chapter reviews previous research on public transport fare policy analysis. Assessing the impacts of fare changes on public transport demand has been studied in a vast number of research papers and reports over the years. The selection of papers reviewed here is by no means comprehensive but is representative of the various methodologies used in this area of study.

We should note that the research presented here pertains mostly to public sector pricing and fare policy, which is most relevant to public transport in North American and western European metropolitan areas. In many other cities, public transport is run by private operators to which concepts of profit-maximizing private sector pricing are more applicable.

In public sector pricing, the objective is to maximize social welfare, which is the difference between social benefits and social costs. These benefits and costs include those of the public transport agency, the users of the public transport system, and others, including the residents of the metropolitan area. Given the variety of stakeholders involved, public sector fare policy often involves tradeoffs between simplicity, financial performance, and equity (Colin Buchanan, 2006).

In Section 2.1, we present the different contexts in which public transport fare policy has been studied and then discuss the contexts that are most relevant to this thesis. In Section 2.2, we use the first section as a basis to review the methodologies applied in some of the prior work.

2.1 Evolution of Public Transport Fare Policy Analysis

2.1.1 Overview of Fare Policy Issues

Public transport fare policy has evolved greatly over the past few decades. Setting fares in the past was often just a matter of defining one flat fare level for each mode of public transport. Period passes were introduced in many public transport systems as a way to guarantee a fixed revenue stream but also to provide an incentive for users to switch to public transport given the convenience and the potential savings that can be incurred by using these passes. Fare policy has also addressed issues of equity by offering different fares to the elderly and, in some cases, to students or to other segments of the community (Fleishman *et al.*, 1996).

In recent years, fare policy has become a tool in managing public transport demand. This has been mainly due to the changing needs of public transport agencies, which have become more focused on reducing costs and increasing revenues, but it is also due to the emergence of smartcards and new ticketing technologies that have allowed these agencies to apply more complex fare structures and to adopt demand management strategies that can result in more productive use of their systems (Fleishman *et al.*, 1998). Some of the uses of fare policy as a demand management tool were listed in Section 1.1 and will be addressed later in this thesis.

2.1.2 Types of Models Used in Fare Policy Analysis

This evolution of fare policy has resulted in the use of various types of models to assess the impacts of fare changes on public transport demand. For systems with a simple fare structure in which there is a flat cash fare and two or three period passes (e.g. weekly pass, monthly pass, and annual pass), a certain set of models would provide satisfactory results. However, for more complex fare structures, one might need to use a different set of models that reflect these structures and relax some of the assumptions made in the simpler models. The various types of fare policy models are well-researched and documented in the literature. In this section, we will briefly review some of these models.

We start with the simplest of these models or, more precisely, models that predict public transport demand under a simple fare structure. Such models include the four-step transportation planning model, which is a cross-sectional, trip-based model that consists of trip generation and attraction, trip distribution, mode split, and trip assignment. The 'mode split' step in these models generally includes the different public transport and non-public transport modes available in the metropolitan area being studied. Inputs to this step usually include travel time and travel cost variables, as well as public transport service attributes (such as frequency, speed, and capacity) and attributes of the road network. The output is the mode split of the trips produced in the first two steps. In such models, there is usually one measure of fare for each public transport mode, and the interest is typically in modeling the total demand for public transport in the context of the overall transportation system (Horowitz, 1993, and Bhat, 1995). Since these models are cross-sectional, they do not provide a full picture of how future fare changes will affect public transport demand. They are, however, still widely used in metropolitan areas worldwide.

As period passes became more ubiquitous, a new set of models emerged to try to explain the factors affecting people's choices among the various ticket types. Such models were not meant to be a substitute for the four-step model, as they did not directly predict public transport use but only considered choices among different "products" and how important factors such as income or potential savings from period passes are in determining that choice. These models were also applied on cross-sectional datasets, usually obtained from surveys (Hensher, 1998).

Time-series models have also been developed in the literature to assess the impacts of fare changes over time on public transport demand. Such models were typically linear regression equations that varied in complexity and were estimated on datasets spanning several years (Mitrani *et al*, 2002). They usually produced a set of demand elasticities (i.e. percentage changes in public transport demand resulting from a one percent increase in fares) and, in some cases, cross elasticities (i.e. the percentage change in the demand for one public transport mode resulting from a

one percent increase in the fare level of another public transport mode or in the cost of travel on some other non-public transport mode).

The elasticities produced by these models are often used as inputs to another set of models used at public transport agencies. These latter models are typically composed of a number of spreadsheets in which all the different ticket types or 'fare classes' are accounted for. The spreadsheets include a 'base case' scenario with the current fare and demand levels for each ticket type. One can then enter the elasticities of demand for the various ticket types and, in some cases, the cross elasticities (e.g. the demand for bus resulting from a one percent change in the train fare). The model then produces estimated demand and revenue levels for proposed fare policies.

The advantage of these "spreadsheet models" is that they can account for complex fare structures in which there are many ticket types and fare levels. The disadvantage is that the elasticities used as inputs to these models are exogenous and need to be estimated separately. This becomes harder if the spreadsheet includes more ticket types, as that would mean that more elasticities are needed (one for each ticket type and, ideally, cross elasticities between each pair of ticket types). Furthermore, these models are mostly aggregate and do not account for heterogeneities among different individuals and demand segments.

To summarize, we have outlined the following types of models in which public transport fare policy has been analyzed:

1. Cross-sectional mode split models
2. Ticket choice models
3. Aggregate time-series models
4. "Spreadsheet models"

In this thesis, we are interested primarily in modeling the effects of fare changes on *how* people use public transport services. To that end, we will focus on the last three types of models listed above. Ticket choice models developed in the prior literature provide useful insight into the methodologies used to capture such choices, as well as the factors that are presumed to influence them. Time-series models, although they are mostly aggregate, unlike the model proposed in this

thesis, introduce a temporal aspect to the analysis, which is crucial if one wants to have accurate forecasts of demand under future fare policies. Finally, "spreadsheet models" used by public transport agencies give a better understanding of how these agencies account for their complex fare structures in forecasting demand.

Section 2.2 presents a review of the methods used in developing and estimating ticket choice and aggregate time series models in the literature. Models used in London will be discussed in detail in the next chapter.

2.2 Review of Methodologies

2.2.1 Prior Literature on Fare Modeling

In this subsection, we review the methodologies used in previous studies to develop and estimate ticket choice models and aggregate time-series models for public transport.

There are very few papers in the literature on public transport ticket choice. These papers vary in the methodologies they use. Some use survey data and present descriptive statistics on the profile of pass users as opposed to cash fare users, for example, while others develop more complex discrete choice models to measure the effect of different factors on ticket choice.

An example of the first methodology can be found in Dittmar (1983). The paper provides a profile of monthly pass users in the San Francisco Bay Area. It gives an overview of the design of a survey in which respondents were asked about their public transport use patterns, the reasons for which they bought a monthly pass, their interest in purchasing an intersystem pass (which did not exist at the time), and other questions related to their preferences for marketing and distribution of fare media. The survey results indicate that convenience and savings were important reasons for people to choose monthly passes. Respondents who used local service also cited the option of using the pass for discretionary trips as another important factor.

Hensher and King (1998), Hensher (1998), and Taplin, Hensher, and Smith (1998) use a more rigorous methodology to analyze the choice among different

public transport ticket types. These papers develop and estimate discrete choice models using revealed and stated preference data. The model structure they use is the heteroskedastic extreme value logit (HEVL), which differs from classical multinomial logit models by allowing for varying cross elasticities among different ticket types. Hensher (1998), for example, used a combination of revealed and stated preference data to estimate a ticket choice model for the Sydney (Australia) metropolitan area. The choice set in the model consisted of three options for train users: single, weekly, and TravelPass (which allows travel on trains, buses, and ferries) and four options for bus users: single, TravelTen (which includes ten one-way trips), bus TravelPass (which allows travel on buses and ferries only), and TravelPass. Note that since the choice set in this paper was defined by mode and ticket type, mode switching, in additional to ticket switching, was accounted for in the results. The dataset included respondents' public transport use patterns at the time the survey was conducted (these are the 'revealed preference' data) as well as their potential use patterns under two hypothetical scenarios, a 'low fare' scenario and a 'high fare' scenario—'low' and 'high' being relative to the fare levels at the time of the survey (these are the 'stated preference' data). The explanatory variables in the model include travel time, fares, car ownership, and 'train captivity' and 'bus captivity' dummy variables for individuals who cannot switch modes. The result of the analysis in this paper is a matrix of elasticities and cross elasticities for the different ticket types.

The methodology used in the above three papers provides a good framework for modeling the factors that affect ticket switching as well as public transport mode switching. The disaggregate models in these papers allow for capturing individual-specific heterogeneities. Also, the use of combined revealed and stated preference data and the inclusion of 'captivity' dummies capture people's inertia when it comes to ticket and mode choice. The methodology, however, assumes that public transport use is an exogenous variable that does not depend on ticket and mode choice. In reality, people may have different use patterns depending on their ticket choices.

Another paper worth mentioning, although it was written many years before the ones just discussed, modeled the choice between cash fares and monthly passes in Sacramento, California using a simple binary logit model which was estimated on a cross-sectional dataset (Page, 1981). The explanatory variables included savings or losses incurred by using a pass, the ratio of the number of work trips to all trips for each person, the number of cars per worker in a household, income, and gender. The results showed that the initial cash outlay to purchase a monthly pass is a deterrent to people with limited incomes and that low income people are more sensitive to the savings made by using a pass. It should be noted that the model assumed different effects for savings compared to losses (i.e. positive and negative values of the difference between the cash fare cost of monthly use and the price of a monthly pass). The results were consistent with behavioral economic theory and indicated that a dollar lost caused more disutility than a dollar saved.

Ticket choice has also been studied in the literature in the context of peak pricing. Given the increasing levels of peak-period crowding on many public transport systems, the idea of charging a premium for traveling during that period has become more appealing. One recent study examined the possibility of spreading the peak by implementing peak pricing on National Rail services in the United Kingdom (UK Department for Transport, Transport for London, and Network Rail, 2007). The study involved focus groups and a travel survey to assess people's flexibility in terms of the timing of the trip and the level of crowding during the peak. The responses obtained from these focus groups and surveys were used to estimate a discrete choice model, in which individuals choose between thirty-minute intervals during which they travel. The utility of each of these intervals included the following explanatory variables: travel time, fare, crowding level, and time displacement (i.e. the amount of time displacement relative to the preferred or current time of travel).

The results in the report, obtained from the focus groups, surveys, and by simulating travel on a number of corridors, have significant policy implications. They indicate that 45% of passengers have no flexibility in their arrival times in central London. A significant proportion of those who do are high-income

individuals, who are generally senior managers and professionals. Furthermore, the results show that significant percentage changes in fares and in seating capacity are required in order to produce behavioral change. For example, the penalty cost for traveling 60 minutes later was estimated at £12. The report also found there to be significant interactions between fares, crowding levels, and time displacement. In other words, a fare surcharge may not be sufficient to induce changes in the times of travel if crowding levels remain high.

The abovementioned report provides rich insight into the policy implications and the potential feasibility of implementing peak pricing schemes on public transport systems. The models developed in the report, however, cannot be readily applied to other public transport systems since they rely on very detailed survey data, as well as data on crowding levels on public transport vehicles.

The models discussed so far are based on cross-sectional data and, hence, cannot be used as reliable forecasting tools, as they do not provide temporal fare elasticities. Next, we will discuss time-series models that have been developed and estimated mostly at the aggregate level.

There is a vast literature in which aggregate analyses are conducted to assess the impacts of changes in fares or other variables over time on public transport demand, as well as to measure the change in the demand for one public transport mode given a change in fares of another. The methodologies applied in this literature are derived from basic economic theory and have been applied not only on public transport but also on a myriad of products and services. In most of these papers, a direct demand function is specified (the function can have different forms but is usually a linear, log-linear, or log-log function) and then estimated using an econometric method on a dataset that typically spans several years. The basic premise behind these models is that by controlling for different variables, one can capture the effect on demand of changes in prices (or fares in the case of public transport). The obvious disadvantage of these models arises from the fact that they are aggregate, which means that individual-specific differences are not captured.

Litman (2004) and Oum *et al.* (1992) provide fairly comprehensive surveys of the various studies done to estimate public transport demand elasticities and cross

elasticities (Oum also includes estimates of demand elasticities for other transport modes.) In many previous studies in which more detailed data are available, elasticities are estimated separately for peak travel and off-peak travel or for the short-run and the long-run. For example, Pham and Linsalata (1991) as referenced by Oum *et al.* (1992) estimated a peak-hour fare elasticity of -0.18 for bus travel in large cities with more than one million people. This implies that a 10% increase in peak-hour fares in those cities will result in a 1.8% decrease in peak-hour bus travel. The off-peak elasticity for those same cities was estimated at -0.39. This makes sense since off-peak travel is generally more sensitive to fare levels. The peak and off-peak elasticities for smaller cities with less than one million people were estimated at -0.27 and -0.46, respectively. The estimates in this study are based on short-run (less than two years) data from 52 US transit systems during the late 1980s.

Mitrani *et al.* (2002) specified a log-linear demand function for public transport services in London (bus and Underground). The model was estimated on data from 1970 to 2000. Given the lack of ridership data, the paper used an index that is a function of revenue as the dependent variable. In other words, revenues were used as a proxy to infer ridership. The estimated demand elasticities were -0.64 for bus and -0.41 for the Underground. The paper also includes estimates of demand elasticities with respect to income and levels of vehicle mileage operated.

The above methodologies are reviewed more extensively in TRL (2004). The report provides a comprehensive survey of public transport fare elasticities estimated in prior literature, including elasticities by income group, age group, trip purpose, distance traveled, and other criteria. The report also discusses the differences between short- and long-run fare elasticities and the methods and assumptions used in estimating each of them. It outlines two ways by which elasticities with different time horizons are estimated:

1. The use of cross-sectional data while making assumptions about the time-scale of the response to fare changes
2. The use of time-series data and dynamic econometric models in which a gradual response over time is represented explicitly

The choice of a methodology and the assumptions made may very well affect the estimated elasticity values, as indicated in the abovementioned report. In our methodology, we develop dynamic econometric models making use of disaggregate panel data (in which individuals are observed over time).

Table 2-1 shows some estimates of short- and long-run public transport fare elasticities estimated in prior research and referenced by TRL (2004).

Mode	Fare elasticities	
	Short-run	Long-run
Bus	-0.2 to -0.3	-0.7 to -0.9
Urban rail	-0.3	-0.65

Table 2-1: Short- and long-run fare elasticities for bus and urban rail in the UK

2.2.2 CTA Fare Model

This subsection presents the fare model used at the Chicago Transit Authority (CTA). We discuss this model in some detail since it is the most comparable to the methodology proposed in this thesis, in terms of accounting for ticket choice and the fact that public transport use is dependent on that choice. The review in this subsection is based on an earlier review by Hong (2006).

Figure 2-1 shows a flowchart that summarizes the CTA fare model. Before the model can be used, a survey must be conducted to collect data on demographics and to categorize CTA riders into "market groups" based on frequency of travel and the type of fare used (Full Fare or Senior/Disabled Reduced Fare). The average number of trips by time period and mode is then estimated, after which a macro is executed to estimate the initial (base) share of each fare type (or "fare media shares" as denoted in the flowchart).

In Step 2, the new fare structure is entered and a discrete choice nested logit model is estimated to determine the shares of the different fare types in each market group. The upper level of the nested structure includes three ticketing medium alternatives: pay-per-use, cash, and passes, while the lower level includes the fare types within each ticketing medium (Transit Card and Chicago Card under the pay-

Step 1: Establish a Baseline Scenario
- Enter base fare structure
- Run macro to generate initial fare media share
- Calculate monthly cost of travel for riders in each market group using base fares

Step 2: Estimate New Fare Media Share
- Enter proposed fare structure
- Calculate new monthly cost of travel for riders in each market group using proposed fares
- Nested logit model calculates new fare media share

Step 3: Compare the Base and New Fare Media Shares for Each Market Group
- If there is a loss in fare media share, the monthly cost of travel using proposed fares used in Step 4 is as calculated in Step 2
- If there is an increase in fare media share, the monthly cost of travel using proposed fares used in Step 4 is calculated based on a weighted average of the previous monthly costs across all market groups

Step 4: Estimate the Elasticity Impact and Ridership Change for Sample Ridership
- Mid-point elasticity is used to calculate the percentage change in ridership resulting from the fare change
- Calculate the predicted number of trips for each fare medium and market group using the elasticity impact factor, base number of trips, and new fare media share calculated in Step 2

Predict Total Revenue
- Calculate the revenue for sample ridership by multiplying new monthly cost for all trips by elasticity impact and new fare media share calculated in Step 2
- Expand this revenue to obtain annual estimates for entire CTA ridership

Predict Special Fare Media Ridership
- Changes in the Link-Up Pass are based on the change in the cost of that pass and changes in the number of Metra riders
- Changes in the use of Student Fares are based on changes in the student fares and changes in the cost of the high school permit
- Changes in the use of the U-Pass are based on changes in the number of students eligible and changes in the average cost per day of the pass
- Changes in the use of the Visitor Passes are based on changes in the cost of these passes.
- Changes in the use of Child Fares are based on changes in this fare and changes in the total ridership of the CTA.

Figure 2-1: Flowchart of CTA fare model

per-use alternative and monthly, weekly, and visitor passes under the 'passes' alternative).

After the fare type shares are determined, fare elasticities are estimated in Step 3. Elasticities are calculated using different methods based on whether there is a gain or loss in each fare medium and market group compared to the base case.

In Step 4, the elasticities calculated in the previous step are applied to estimate changes in frequency of use based on the changes in fare types. This step represents one of the key similarities between our proposed methodology and the CTA fare model: changes in ridership, or frequency of use, are modeled as being dependent on changes in fare type shares.

Step 5 converts trips to unlinked boardings (or journey segments) based on previously measured transfer patterns of the different types of trips (e.g. peak bus trips generally include two boardings). After Step 5, changes in ridership on special fare media are estimated.

The CTA fare model is a good example of how public transport fare policy analysis should be done in an agency with a fairly complex fare structure. The decision on ticket type or fare type should not be disassociated from the frequency of public transport use. This is the basis of our proposed methodology which will be presented in detail in Chapter 4.

The following points should also be noted about the CTA fare model:

- The macro used to generate baseline data does not explicitly require any baseline smartcard data, as it uses the survey data that are input to the model. This can be especially useful to predict the adoption rate of smartcards before introducing them as a fare medium, although this would require making some initial assumptions, sometimes arbitrarily, on the expected penetration rate. (Smartcard data can be used later as they become available to produce more accurate baseline shares for the different fare media.)

- The nested logit model structure allows for incorporating various fare policies and fare structures and to better account for ticket switching, including switching among different fare media. A change in fares for any ticket type can be easily accounted for in the empirical specification of the model, and the

addition of a new ticket type can also be addressed by adding that ticket to the nest of its corresponding fare media. One should keep in mind, however, that major changes to the fare structure require making additional assumptions regarding the initial penetration rate for new ticket types.

The CTA fare model differs from our methodology in that it requires survey data to establish a baseline scenario (as opposed to our model, which relies on smartcard data). Despite the advantage discussed above regarding the ability to predict smartcard shares before such a fare medium is introduced, relying on surveys can be costly, especially if the public transport agency wants to update the baseline shares on a regular basis.

Furthermore, the CTA model is cross-sectional and, thus, makes some assumptions on people's responsiveness to fare changes when used for forecasting purposes. In other words, predicting the effects of a change in the fare structure requires adjusting cross-sectional attributes to produce fare media shares and ridership estimates under 'before' and 'after' scenarios, without taking into account any dynamics that can influence the choice of fare media. Our model, on the other hand, uses a panel data structure that accounts for such dynamics and for changes in individuals' behavior over time. Such a data structure, as we explain later, can even be useful in predicting the behavior among new adopters of smartcards without having any information on their past use patterns under the cash fare. This is because smartcard data from past time periods can be used to develop proxy measures that capture the effects of fare policies causing switches among fare media.

To conclude, this chapter has presented a review of the various contexts in which fare policy models have been developed and reviewed prior work that is most relevant to the motivation for this thesis. Next, we turn to fare policy in London.

Chapter 3

Fare Policy in London

Public transport fare policy has become integral in guiding short and long-run transport strategies in London. The Mayor of London's Transport Strategy states that fare policy should be used "to make public transport more attractive and affordable, with more consistency between modes, greater simplicity and convenience for passengers, shorter queues and quicker journeys". The document, most recently revised in July 2006, also includes short-run proposals to freeze bus fares and to cap Underground fares in real terms.

The introduction of the Oyster smartcard in December 2003 has made fare policy an increasingly powerful tool in managing public transport demand. This is also evident in the Mayor's Transport Strategy, which states that "Transport for London (TfL) will develop targeted fare options using Smartcards to offer benefits to passengers, increase use and reduce delays".

In this chapter, we first introduce the fare structure on London's public transport system. We then present some spatial and temporal patterns and trends based on this fare structure and conclude with a discussion on how the effects of fare changes on demand are currently modeled at TfL. The material in this chapter raises the questions and hypotheses that underly our proposed methodology, which is presented in Chapter 4.

3.1 Introduction to London's Fare Structure

3.1.1 Overview

Public transport fare structures are generally classified into the following categories (Colin Buchanan, 2006):

1. Graduated fares, in which fares vary by origin-destination pair,

2. Zonal fares, in which fares vary by zones or regions of travel, and

3. Flat fares, in which a flat fare is charged for all trips.

The fare structure in London is a combination of the second and third categories listed above. Bus, DLR, and Tramlink trips are charged on a flat fare basis, while Underground and Overground trips are charged based on the zones of travel. Examples of public transport systems with different types of fare structures include those in Boston and New York, both of which have flat fares, and Washington, DC, which has graduated fares for rail and flat fares for bus. Clearly, graduated and zonal fares require a mechanism for validating both entries and exits into stations or onto public transport vehicles. This is not the case for flat fare structures, in which the cost of travel is not a function of the destination.

Figure 3-1 shows the London Underground and Overground networks with the underlying zonal map. Zones 1 thru 6 fall within the boundaries of the Greater London area, while Zones 7 thru 9 fall beyond these boundaries[2]. As mentioned above, fares on the London Underground and Overground vary by zones of travel. Specifically, a premium is charged for travel that includes Zone 1. This means that travel between two Underground stations, for example, is charged a premium if the path between those two stations passes through Zone 1 (even if neither station is in Zone 1)[3]. Furthermore, Underground and Overground fares increase as the number of zones between origins and destinations increases.

In addition to the modal and zonal variations discussed above, fares in London also vary by ticketing media and, in some cases, by time of day. This is shown in Table 3-1, which presents the various criteria (or attributes) by which fares vary on London's public transport network.

---

[2] On January 2, 2008, Zones A, B, C, and D were replaced by Zones 7, 8, and 9 (shown on the map in Figure 3-1). Zone 7 replaced both Zones A and B, Zone 8 replaced Zone C, and Zone 9 replaced Zone D.

[3] To our knowledge, TfL currently uses an ad-hoc approach to determining the most likely chosen path between each origin-destination pair on the London Underground in order to set the fare for that pair.
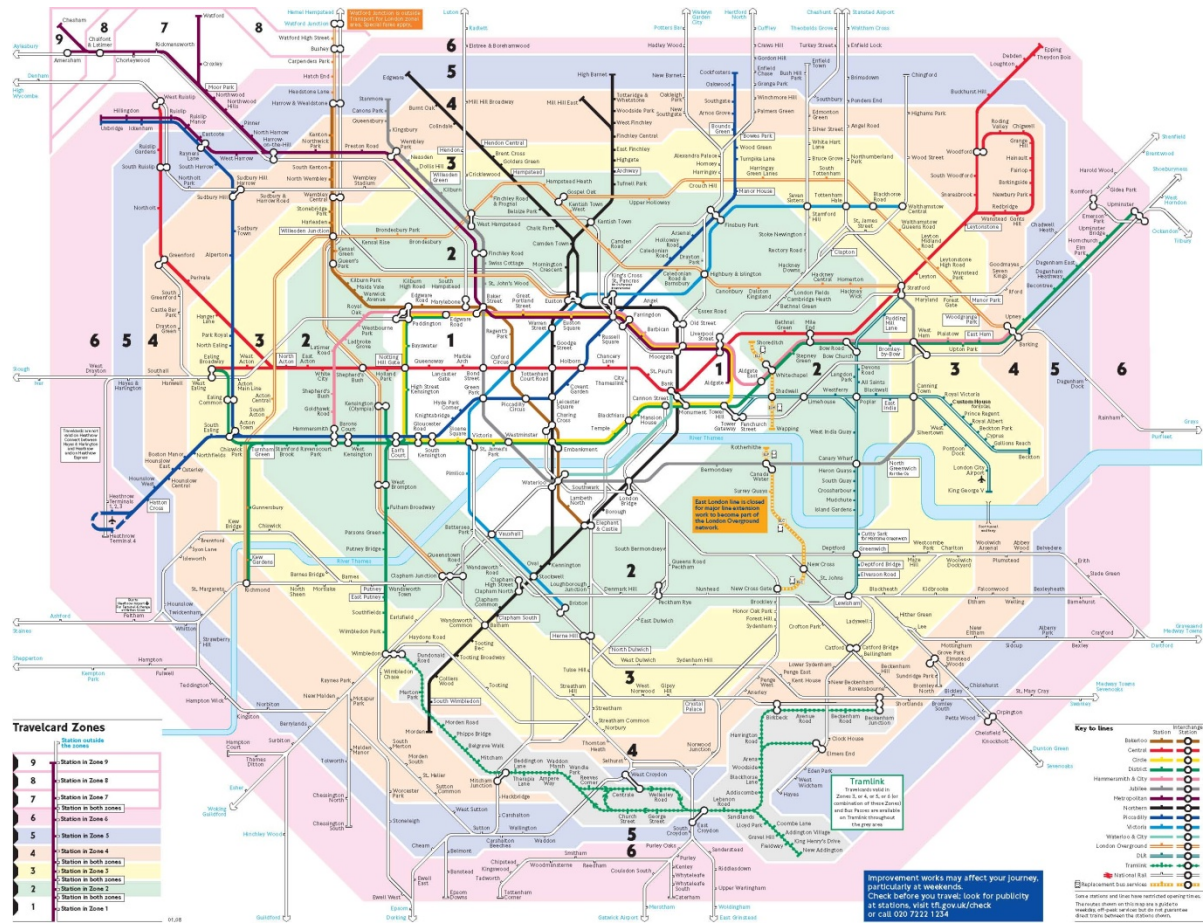
Figure 3-1: Zonal map of the London Underground and Overground networks *(TfL Website)*

Criterion	Fare Differential(s)
Ticketing medium	Cash (magnetic stripe) fares > Oyster fares.
Zones (Underground and Overground only)	Premium charged on travel that includes Zone 1. Fares increase as number of zones of travel increases.
Time-of-Day	For some ticket types: Peak fares > Off-peak fares.
Mode	Underground and Overground fares > bus and Tramlink fares.

Table 3-1: Fare differentials in London

3.1.2 Ticket Types[4]

As is the case with many public transport operators worldwide, TfL offers its customers a wide range of ticket types. These ticket types vary as described in Table

---

[4] The information in this section was obtained through correspondence with Tony Richardson at TfL as well as from the latest edition of the booklet "Your Guide to Fares and Tickets" issued by TfL in January 2008.

3-1, and some are also issued with different time validities. The number of ticket types offered in London is very large, as there are tickets for many of the possible ticketing medium, zonal validity, time-of-day, mode, and time validity combinations. An exhaustive list of all ticket types in London is not required for the purposes of this thesis. Instead, we provide the following summary in which the ticket types are grouped into three categories:

1. **Single (per journey) Fare Tickets:** available on both Oyster cards (on which these tickets are known as "Pay-As-You-Go") and magnetic stripe cards. Fares among these tickets vary as follows:

   - Oyster fares are significantly lower than cash fares.
   - Underground, Overground, and DLR fares are greater than bus and Tramlink fares.
   - On the Underground and Overground, different fares are charged depending on the zones of travel. A premium is charged for travel through Zone 1, and the fare increases as the number of zones of travel between the origin and the destination increases. Bus and Tramlink journeys are charged a flat fare.
   - Fares charged for Underground and Overground journeys made on Mondays thru Fridays between 7:00am and 7:00pm are greater than those charged for travel during all other times (including public holidays).

   In February 2005, TfL introduced daily price capping on all Oyster Pay-As-You-Go (PAYG) journeys. Under this pricing scheme, PAYG users are not charged more than a certain amount during a 24-hour period (from 4:30am on a given day to 4:30am on the next day). This amount (or price cap) is £0.50 below the price of a Day Travelcard or bus Pass and varies by time of day (peak vs. off-peak) and zones of travel. Price capping was introduced to encourage the adoption of Oyster by ensuring that users will always get the best value out of their Oyster cards, at least on a daily basis.

2. **Travelcards:** allow unlimited travel on all public transport modes, with a few exceptions[5]. Travelcards can be issued with different time validities. Day and

---

[5] Travel on Tramlink is permitted only with Travelcards that cover a zone beyond Zone 2.

Three-Day Travelcards are available only on magnetic stripe cards, while Weekly, Monthly, and Annual Travelcards are available only on Oyster cards. Furthermore, within each of these time validities, there are multiple types of Travelcards, as follows:

- For Day Travelcards: One can choose to buy a Zone 1-2, 1-3, 1-4, 1-5, 1-6, 1-9, 2-6, or 2-9 Travelcard. For each of these zonal validities, one can purchase a peak Travelcard, which allows travel at all times of the day, or an off-peak Travelcard, which allows travel at all times except from 4:30am to 9:30am on working days.
- For Three-Day Travelcards: One can choose to buy a Zone 1-2 or Zone 1-6 Travelcard.
- For Weekly, Monthly, and Annual Travelcards: One can choose to buy a Zone 1-2, 1-3, 1-4, 1-5, 1-6, 1-7, 1-8, or 1-9 Travelcard or one that allows travel between any two zones or among any three, four, five, six, or eight zones as long as they do not include Zone 1.

Zonal restrictions apply only to Underground, DLR, and Overground travel, in addition to the Tramlink restrictions noted previously.

3. **Bus Passes:** allow unlimited travel on bus and Tramlink. Bus Passes can be issued with different time validities, similar to those of Travelcards (although there is no three-day bus Pass). There are no variants of Bus Passes based on the zones of travel.

It should be noted that the prices of monthly and annual period tickets (be they Travelcards of Bus Passes) are always set at 3.84 and 40 times the price of the corresponding weekly period ticket, respectively.

TfL also provides a large number of concessionary tickets. Children, the unemployed, and other segments of society are offered a discounted selection of ticket types, some with zonal and time validities and modal restrictions similar to the ones discussed above. So-called "Freedom Passes" are also issued for the elderly and disabled and allow them free travel on London's public transport network.

In the remainder of this thesis, we will focus mostly on travel on the London Underground (LU) and London Buses, both because these modes are by far the

most widely used on London's public transport network (see Figure 1-1) and because fare policy decisions pertaining to other modes are generally tied to those of the Underground and buses. DLR fares are generally set within the same fare structure as the London Underground, while Tramlink fares are set to be equal to bus fares. The London Overground, although it has a separate fare structure, has only been part of TfL since November 2007, so including it as part of a longitudinal or panel data analysis is not feasible.

3.1.3 Chronology of Fare Changes[6]

Before we discuss the spatial and temporal trends in London's public transport system, it is important to highlight fare changes and major fare policy decisions that have taken place over the past few years and that have shaped the current fares environment in London discussed in the previous subsection.

First, we present a chronology of fare policy decisions through which the Oyster card was rolled out into the public transport system, and then we highlight some specific fare changes that have taken place.

1. Fare policies related to the introduction of Oyster:
   - September 2002: Oyster card first introduced on TfL staff passes.
   - May-June 2003: Monthly and Annual Travelcards and Bus Passes became available on Oyster.
   - September 2003: Monthly and Annual Travelcards and Bus Passes were mandated to Oyster.
   - October 2003: Weekly Travelcards and Bus Passes became available on Oyster.
   - January 2004: Oyster PAYG was launched on the Underground and DLR.
   - February 2004: Oyster Freedom Passes were introduced.
   - May 2004: Oyster PAYG was launched on buses and Tramlink.
   - February 27, 2005: Oyster daily price capping was introduced. As mentioned above, this fare policy guaranteed best value pricing on a daily basis.

---

[6] The information in this section was obtained through correspondence with Tony Richardson at TfL.

- September 2005: Weekly Travelcards were mandated to Oyster.
- November 19, 2006: The Underground entry charge on Oyster PAYG was increased to the maximum cash fare (£4.00). This meant that users who do not tap out at the end of their Underground journey get charged the maximum cash fare, while those who do are refunded on exit the difference between the maximum cash fare and the actual fare based on the zones of travel.
- January 2007: Weekly Bus Passes were mandated to Oyster.

2. Specific changes to fares and costs of period tickets (from January 2005 onwards):
- Prices of all Travelcards and Bus Passes were increased in January 2005, January 2006, and January 2007 to reflect overall inflation levels. Prices of Travelcards were also increased in January 2008.
- Oyster PAYG fares and cash fares were increased in January 2005.
- PAYG Underground fares were reduced in January 2006 by £0.10-£0.20 and have not changed since.
- In November 2006, Underground cash fares were set at £4.00 for travel that includes Zone 1 and £3.00 for all other travel and have not changed since.
- On September 30, 2007, Oyster bus fares were reduced from £1.00 to £0.90, and the price of the Weekly Bus Pass was reduced from £14.00 to £13.00. (This also affected the prices of Monthly and Annual Bus Passes given the relationship between those and the price of a Weekly Bus Pass mentioned earlier.)

Some of the above fare changes were not implemented consistently across all the different zonal validities available on some ticket types. However, they generally reflect the policies outlined in the Mayor's Transport Strategy, as well as a clear policy objective at TfL to encourage the adoption of the Oyster card. This is very clear in the example shown in Table 3-2, which compares the Oyster single fare and the equivalent cash fare over time for Underground travel within Zone 1.

	Zone 1 Underground Oyster single fare	Equivalent cash fare
January 2005	£1.70	£2.00
January 2006	£1.50	£3.00
January 2007	£1.50	£4.00

Table 3-2: Oyster and cash fares over time for Zone 1 Underground travel

3.2 Patterns and Trends

This section describes some patterns and trends from London based on the fare structure described above. The figures in this section highlight variations in public transport demand that form the basis for the proposed methodology presented in the next chapter. Data for these figures are based on a random (unbiased) sample of Oyster cards[7].

Figures 3-2 and 3-3 show the number of journey segments made using Oyster cards on the Underground and bus, respectively, by ticket type. The figures shows a generally increasing trend, with more people adopting Oyster cards, either by switching from magnetic stripe tickets or by switching to public transport from other modes. (Note that the dips in the figure correspond to holidays, such as Christmas and Easter.)

In Figure 3-2, we see that Weekly Travelcards had been the most popular ticket type on the Underground before they were overtaken by PAYG in early 2006. This increasing popularity in PAYG can be attributed to policies such as daily price capping, as well as the increasing differentials between cash and Oyster PAYG fares.

Figure 3-3 shows a similar pattern. PAYG has also become the most frequently used Oyster ticket type on buses, followed by Weekly Travelcards and Bus Passes. The use of PAYG on buses has also been increasing at a much faster rate compared to other ticket types, on which the number of journey segments has remained fairly constant since the beginning of 2007. Also of note in Figure 3-3 is the sudden increase in the use of (Oyster) Weekly Bus Passes in early 2007. This is mostly because Weekly and Monthly Bus Passes, as noted earlier, were mandated

---

[7] Details on how this sample was constructed are presented in Chapter 5.

to Oyster in January 2007, forcing people who had previously used the paper equivalents of those tickets to switch to Oyster. Finally, in both Figure 3-2 and Figure 3-3, the number of journey segments on annual and, to a certain extent, monthly period tickets has been fairly constant over the recent past.
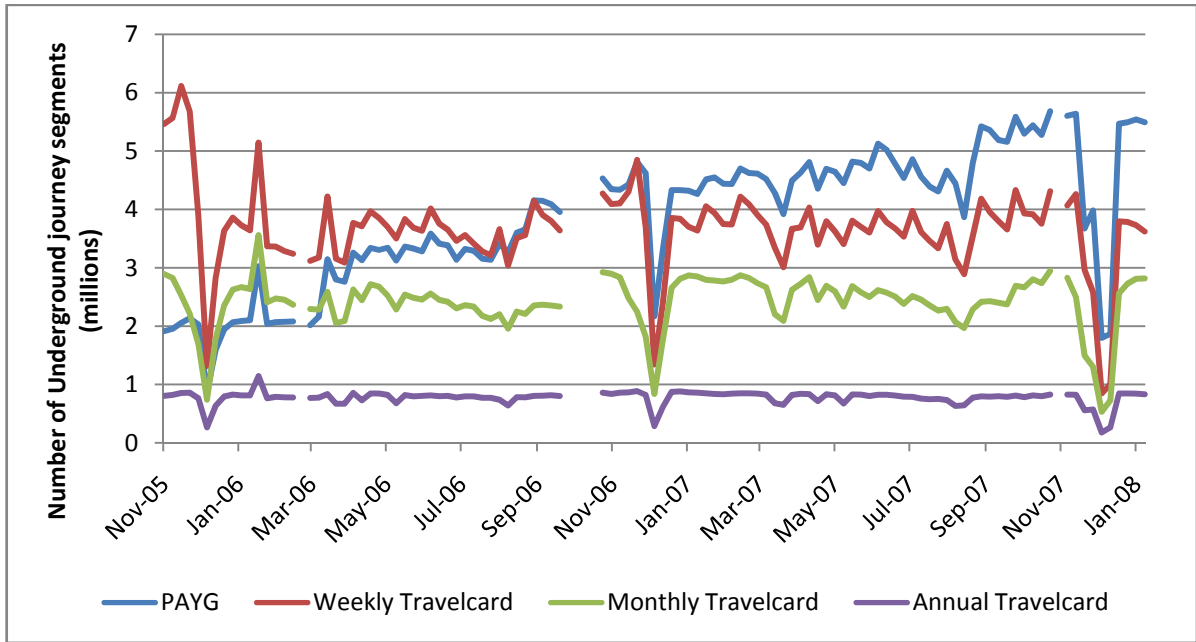
Figure 3-2: Underground Oyster journey segments over time by ticket type
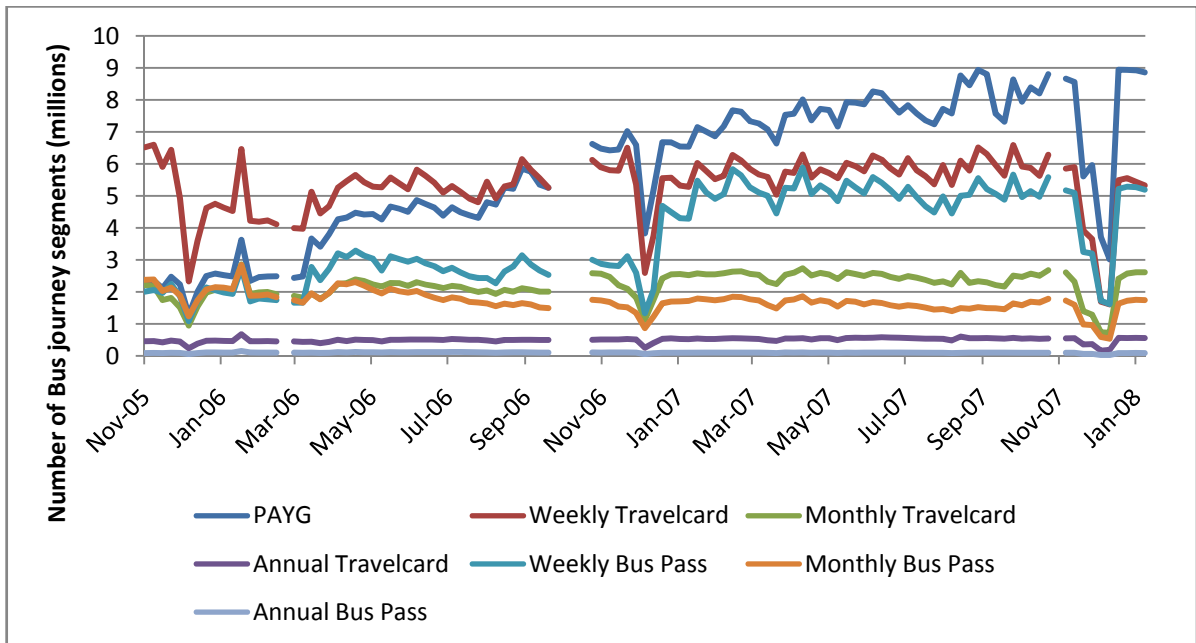
Figure 3-3: Bus Oyster journey segments over time by ticket type

41

These trends are also highlighted in Figures 3-4 through 3-6, which present the relative shares of Oyster ticket types on Underground and bus in November 2005, November 2006, and November 2007, respectively. The figures show a clear increase in the share of PAYG in this two-year period. This is due to the adoption of PAYG by new Oyster users during this time period (e.g. occasional system riders who had previously paid in cash or purchased single or return tickets for each trip) but is also due to switching to PAYG by customers who had previously used other ticket types. PAYG has become an increasingly attractive alternative for both existing and new Oyster users, given the high differential with cash fares, as well as other factors, such as daily price capping and automatic top-up for registered Oyster cards, which was introduced on the Underground in September 2005 and on buses in June 2006.

As mentioned in the introduction of this thesis, fare policy can be used not only to influence ticket choice but also to address time-of-day variations. Figures 3-7 and 3-8 show variations in times of travel on the Underground and bus, respectively, by ticket type. (Note that the vertical axes in these two figures are the percentage of daily journeys on each mode.) What is striking about the patterns in Figure 3-7 are the clearly sharp peaks in Underground weekday travel—the AM peak being between 8am and 9am and the PM peak between 5pm and 6pm. Such patterns suggest an opportunity for fare policy to play a role in spreading these peaks and possibly allowing more users into the system. On the other hand, bus travel, as shown in Figure 3-8, exhibits time-of-day variations that are much more spread out.

The peaking patterns in the figures also indicate that a majority of both Underground and bus travel in the peak is on period tickets (be they Weekly, Monthly, or Annual Travelcards or Weekly or Monthly Bus Passes in the case of bus travel), whereas off-peak travel is more equally divided between PAYG and period tickets. This difference implies that any fare policy that aims to spread peak demand should not only target PAYG fares but also take into account that the majority of peak travel is on period tickets.
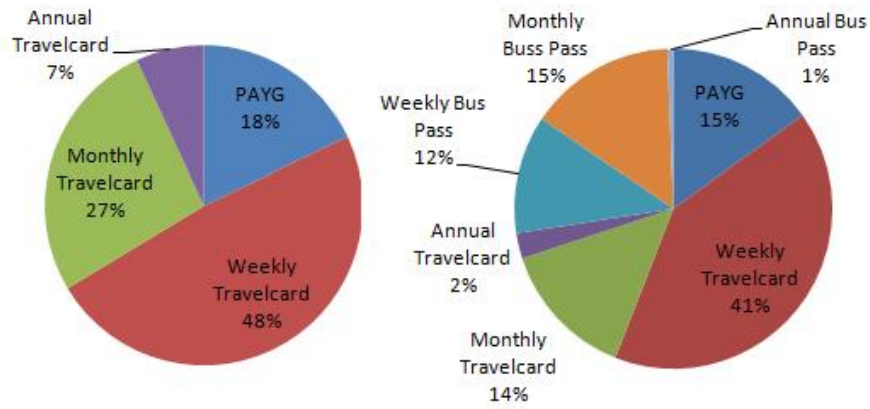
A. Underground **B. Bus**

Figure 3-4: Journey segments by ticket type: November 2005
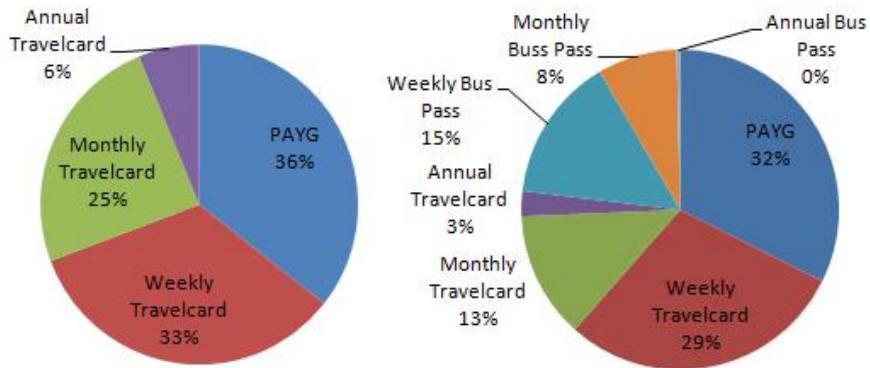
A. Underground **B. Bus**

Figure 3-5: Journey segments by ticket type: November 2006
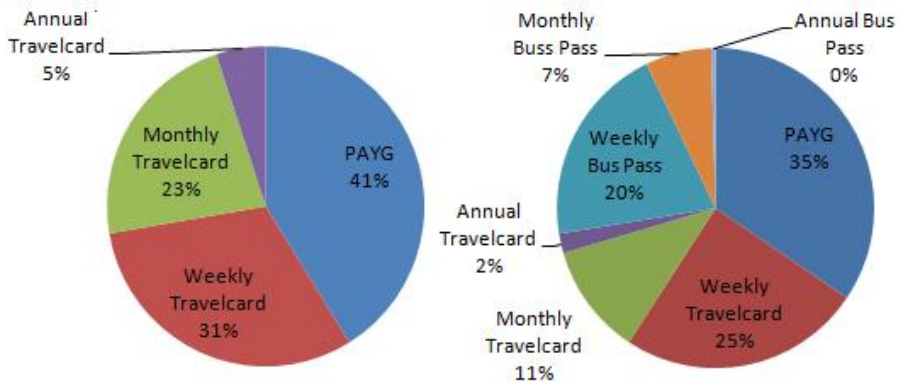
A. Underground **B. Bus**

Figure 3-6: Journey segments by ticket type: November 2007

43
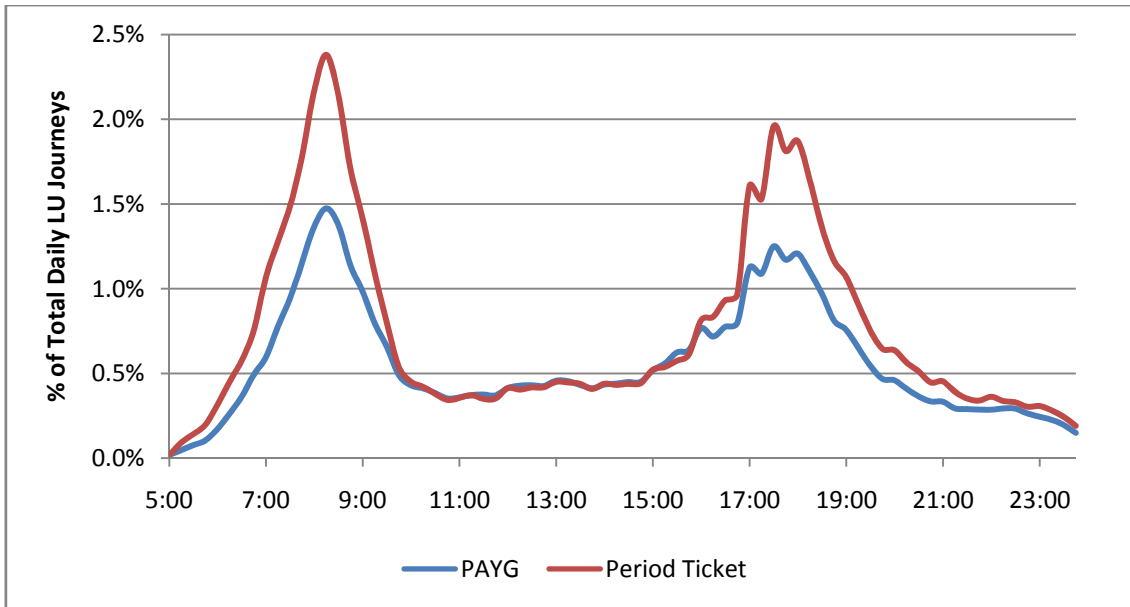
Figure 3-7: Time-of-day variations in Underground weekday travel by ticket type by 15-minute interval (January 2008)
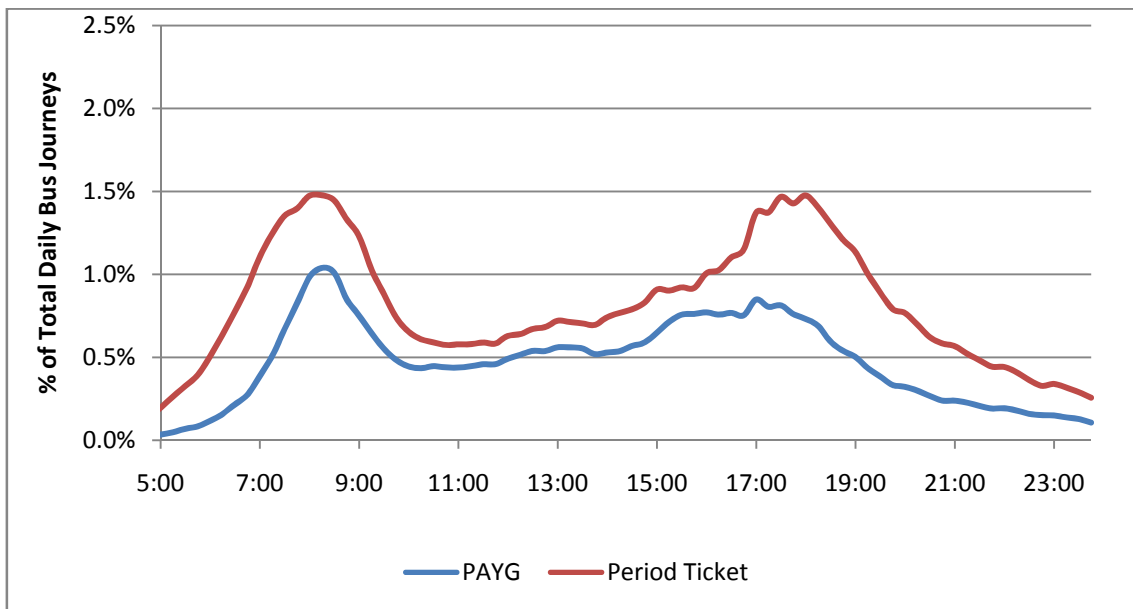
Figure 3-8: Time-of-day variations in bus weekday travel by ticket type by 15-minute interval (January 2008)

The figures in this section highlight the following patterns in public transport use in London:

- The continuing adoption of Oyster as the primary ticketing medium
- The increasing use of Oyster PAYG

44

- The sharp peaking patterns on the London Underground and the differences in ticket types between peak and off-peak users
- The differences between Underground and bus use, especially by time of day

These patterns shed light on some of the variations in public transport demand that would not be obvious if one looked only at overall public transport ridership over time, for example. These variations also raise some questions and hypotheses that we will use to formulate our proposed model structure in the next chapter. Before doing so, however, it is useful to review the current fare modeling practice at TfL.

3.3 Current TfL Fare Models

The previous two sections highlighted the complexity of the fare structure for London's public transport system. This complexity requires models that can accurately capture the effects on public transport demand of changes in the multiple dimensions of fare.

In Chapter 2, we presented the different types of fare models and the contexts in which each is used. We discussed the use of "spreadsheet models" at many public transport agencies and how such models can account for complex fare structures with multiple modes and ticket types. Inputs to these models typically include sets of elasticities and cross elasticities, base fare levels, base demand or revenues, and new fare levels. These inputs are used to produce a set of demand or revenue estimates predicted under the new fare levels.

In this section, we present a detailed review of TfL's fare modeling process, which is a type of spreadsheet model. A similar review was done in a previous report (Colin Buchanan, 2006), from which we obtained the figures in this section.

TfL's spreadsheet models are stored in two Microsoft Excel files: the Underground fare model and the bus fare model. The inputs and outputs in the two models are similar, but the Underground fare model is more complex, given the zonal variations in fares. Given this, we will focus this review on the Underground model.
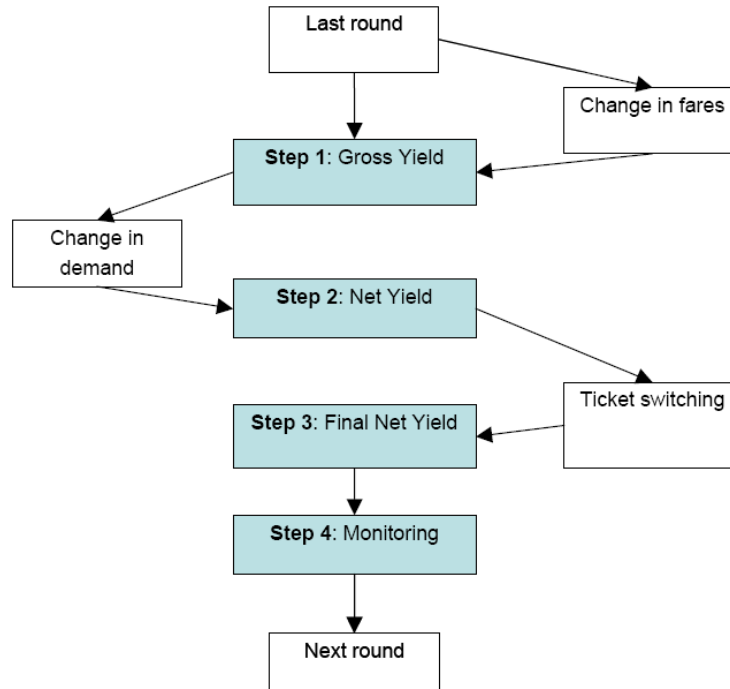
Figure 3-9: Outline of fare models currently used at TfL

Figure 3-9 presents an outline of TfL's spreadsheet models. According to the figure, there are four steps by which revenue and demand estimates are produced and reviewed:

1. **Gross Yield calculation:** This step uses base fares, base demand levels, and new fares to calculate the gross yield, which is the change in revenues assuming no change in demand and no ticket switching. This is highlighted in Figure 3-10, which shows that revenues under the new fare structure (labeled 'B' on the figure) are calculated by multiplying new fares by existing demand.
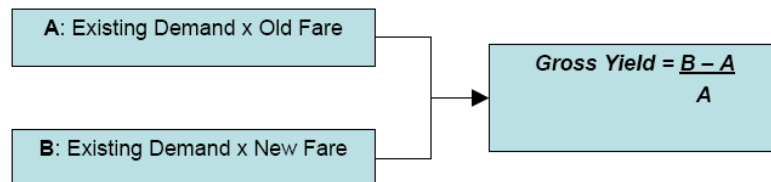
Figure 3-10: Gross Yield calculation in the current TfL model

2. **Net Yield calculation:** This step updates the gross yield estimate produced in the first step by accounting for changes in demand within each ticket type but not across ticket types (i.e. it does not account for ticket switching). In this step, own price elasticities are required for each type. The own price elasticity of a

46

ticket type is the percentage change in demand of that ticket type given a one percent increase in its own price. The elasticities are entered manually into the spreadsheet. Figure 3-11 summarizes the process by which the net yield is calculated. As shown in the figure, new fare levels are multiplied by new demand (as opposed to 'existing demand' in the first step). The new demand, in turn, is estimated by multiplying the price elasticity for each ticket type by the percentage change in fares for that ticket type.
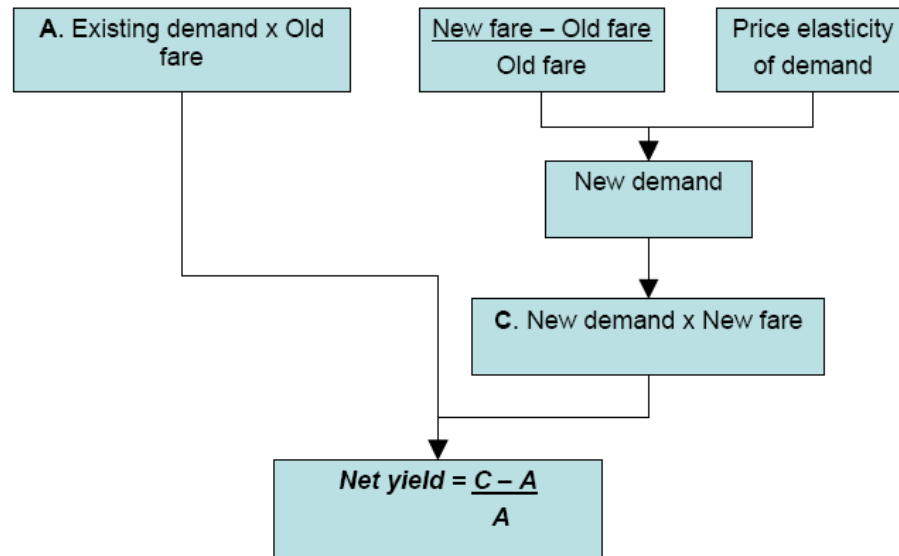
Figure 3-11: Net Yield calculation in the current TfL model

3. **Final Net Yield calculation:** In this step, ticket switching is accounted for using a set of cross elasticities. The cross elasticity of demand for ticket type X with respect to the price of ticket type Y is the percentage change in demand of ticket type X resulting from a one percent increase in the price of ticket type Y. Figure 3-12 gives an example of how cross elasticities can be used to compute the change in demand for Oyster single fares given a change in cash single fares. The numbers in the figure are purely illustrative. In TfL's fare models, switching to and from Oyster PAYG and to and from Travelcards is considered. Switching between different types of Travelcards and between different modes is not explicitly accounted for in the model.
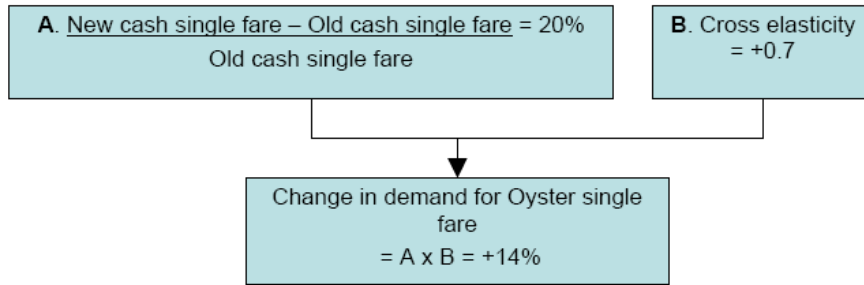
Figure 3-12: Example of the use of cross elasticities to determine switching from cash single fares to Oyster single fares

4. **Monitoring:** The estimates produced by the model are monitored on a continuing basis in order to assess the performance of the model and to determine whether certain elasticity estimates need to be recalculated.

The spreadsheet model described above includes Oyster and non-Oyster tickets, all the possible zonal validities provided by TfL, and the adult and child fare for each ticket type. Oyster PAYG price capping is also accounted for. Although price caps are currently set at £0.50 below the price of a Day Travelcard (which varies by zones and by time-of-day), the model offers the flexibility of using different price caps.

Furthermore, it is important to note that the measure of base demand that is input to the model is the total monetary value of ticket receipts for each ticket type. Using pre-determined revenue apportionments, these receipts are allocated among the Underground, buses, and the train operating companies (which operate National Rail services). The values of the ticket receipts are also converted to journeys using a set of values for 'trips per ticket' for each mode. However, elasticities and cross elasticities are always applied on ticket receipts (i.e. they represent the percentage change in ticket receipts, not journeys, resulting from a one percent increase in fares).

To summarize, TfL's spreadsheet models include the following parameters, which are all pre-determined and can be modified manually by the user:

- Own price elasticities for each ticket type
- Cross elasticities that measure potential (aggregate) switching to and from PAYG and to and from Travelcards

- Revenue apportionments among the Underground, London Buses, and train operating companies
- Underground, bus, and National Rail trips per ticket per day for each ticket type
- Price caps for Oyster PAYG

Given base demand levels, base fare levels, and new fare levels, the above parameters are applied to produce a set of predicted demand levels.

The models described above clearly account for the complex fare structure of London's public transport system by including the various ticket types. However, the performance of the model depends heavily on the values of the elasticities and cross elasticities that it uses. Some of the values currently used are based on rules of thumb, while others were estimated in past studies, which were mostly aggregate and only considered elasticity values by mode (e.g. Bus own elasticity and Underground own elasticity) without accounting for the differences among ticket types.

Furthermore, the current TfL fare models do not explicitly account for mode switching. As mentioned earlier, the models are stored in two separate Excel spreadsheets, one for the Underground and the other for bus. The effect of a change in fares on one mode on demand for the other is, therefore, not modeled.

Also, as is the case with most spreadsheet models, the TfL models are highly customized for the current fare structure. This makes it harder to predict changes in demand due to the introduction of new ticket types, for example, or the introduction of fare policies that are not consistent with the current structure. For example, the current model does not allow for modeling switching between period tickets (e.g. from monthly period tickets to weekly period tickets), since it assumes fixed monthly-to-weekly and annual-to-weekly price ratios that cannot be easily modified.

Another weakness of the current spreadsheet models is that they are cross-sectional and, thus, do not account for certain dynamics when used for forecasting purposes. As is the case with the CTA fare model discussed in Chapter 2, the TfL model uses cross-sectional data to produce 'before' and 'after' scenarios, whereas

the methodology we propose in the next chapter uses a panel data structure that accounts for dynamics resulting from past behavior.

Finally, the methodology proposed and applied in this thesis aims to provide more accurate and robust measures of passengers' responsiveness to fare changes. In the short run, some of these measures may be input and used as part of the spreadsheet model described above. The integration of our results with TfL's models will be discussed in Chapter 5.

Chapter 4

Proposed Methodology

The objective of this thesis, presented in Chapter 1, is to develop a robust, policy-sensitive tool that can accurately predict the effect of fare changes on public transport demand. Chapter 2 reviewed some of the models that have been used in the past for that purpose and discussed their advantages and disadvantages in the context of the changing needs of public transport agencies and the increasing complexity of public transport fare structures. The previous chapter used London as an example to illustrate this complexity. We gave some evidence on how public transport demand varies by the different ticket types, modes, and times of travel and briefly discussed how fare policy has played a role in shaping public transport demand in London and how it can continue to do so in the future.

In this chapter, we build on the material covered so far to present our proposed methodology for modeling the effects of fare changes on public transport demand. We start by discussing the conceptual framework on which our model is based. Then, we present the proposed model structure, followed by the theoretical background and finally a brief summary of the chapter. The methodology we develop here is applied to London's public transport system in the next chapter.

4.1 Conceptual Framework

Model development typically begins by specifying in conceptual terms the variables (observed or unobserved patterns or trends) that are to be modeled and the factors that are thought to influence these variables. In this thesis, we want to develop a model that accurately captures the effects of fare changes on public transport demand. In other words, we want to model public transport demand using fares as one of the factors that influence that demand.

In order to have an accurate measure of the effect of fares on public transport demand, one needs to control for other factors. Public transport demand is influenced not only by fares but also by other factors that are specific to individuals or to the transport system. By controlling for such factors, one can 'isolate' the effect of fares on an observed change in public transport demand and, thus, provide a more robust policy tool and a more effective means for public transport demand management.

Furthermore, it is important to define at the outset what we mean by "public transport demand". In Chapter 2, we reviewed models that used aggregate ridership over time, for example, to capture this demand. However, such measures, as discussed earlier, may no longer be adequate given the complexity of public transport fare structures. The patterns and trends in Chapter 3 emphasized that complexity by highlighting the demand variations in the following three dimensions, which we use, in turn, to define public transport demand in our framework:

1. Mode choice
2. Ticket choice
3. Time-of-day choice

Demand in London and many other public transport systems varies by the above three dimensions because of many factors, including fare variations. A fare differential between two modes affects mode choice. Likewise, a higher peak fare affects time-of-day choice.

Also, note that we use the word "choice" to define the above three dimensions since the analysis in this thesis is done at the disaggregate (i.e. individual) level. One can alternatively use aggregate shares of modes, ticket types, and time segments to characterize public transport demand. We should also note that the choices listed above are interdependent, as clearly shown in Chapter 3. Time-of-day choice, for example, affects mode and ticket choice. Similarly, ticket choice affects both mode and time-of-day choice.

Given the above discussion, we will use the following conceptual model for our analysis:

Public transport demand, characterized by individuals' mode, ticket, and time-of-day choices, is a function of the following:

1. Fares: Overall fare levels and fare differentials across tickets, modes, and times of day affect public transport demand. Exploring and quantifying this relationship is the objective of this thesis.

2. Service attributes on public transport modes: Travel times, station access times, crowding levels, and many other public transport service attributes influence demand in any of its three dimensions. Peak-hour crowding, for example, may influence time-of-day or mode choice and, in some cases, ticket choice.

3. Service attributes on non-public transport modes: Although we are studying variations within public transport demand, it is important to keep the overall transport system in mind. In the context of the choices listed above, one could think of "no ticket" and "non-public transport mode" as being alternatives in the ticket and mode choice sets, respectively. In other words, people always have the option of not using public transport at all.

4. Individuals' socioeconomic characteristics: Public transport demand is affected by individual-specific socioeconomic characteristics. Factors such as income may influence all three choice dimensions. A recent study, cited in Chapter 2, found that a significant proportion of UK National Rail users who are flexible in their arrival times in central London are high-income senior managers and professionals (UK Department for Transport, Transport for London, and Network Rail, 2007). This example illustrates how income can influence time-of-day choice on public transport.

5. Individuals' travel patterns and preferences: Public transport demand is also influenced by other individual-specific factors. These include the tolerance of crowding, mode preferences, and the frequency and purpose of travel.

6. Institutional factors: Strict work schedules, the provision of public transport subsidies by employers, and other institutional factors may have a clear effect on public transport demand in all its dimensions.

It is not possible to account for all the factors listed above in the empirical specification of our model, due mostly to data constraints. Rather, the conceptual

model represents an "ideal dataset" scenario and serves only to guide the development and specification of our model.

4.2 Model Structure

As described in the research approach in Chapter 1, the model we develop in this thesis is based on disaggregate, trip-level smartcard data. The availability of these data allows the use of discrete choice modeling methods in the analysis. A question that arises from this is the extent to which we can use such methods and the level of aggregation that is necessary in order to make the analysis feasible and allow the resulting model to serve its intended purpose.

When developing a discrete choice model, one needs to have information on some, or all, of the alternatives available in the choice set, as well as their attributes. Consider a trip that a person makes by bus from point A to point B, using a weekly period ticket, and departing point A at 8:00am. Developing discrete choice models for the mode, ticket, and time-of-day choices for that trip requires answers to the following questions, in addition to the attributes of the trip that was made:

- What other modes could have been used? If a rail alternative is available, could that have been used for the trip (i.e. are points A and B accessible by rail)? What are the service attributes (e.g. travel times) on the other modes? How does the mode choice for this trip affect the ticket and time-of-day choices?

- What other ticket types could have been used (e.g. a per journey fare or a monthly period ticket)? What would the cost of this trip be under the different alternatives? How does the ticket choice for this trip affect the mode and time-of-day choices?

- What are the other times of day at which this trip could have been made? Did the person consider 3:00am, for example, when making their time-of-day choice? What are the service attributes at other times? How does the time-of-day choice for this trip affect the mode and ticket choices?

Answers to these questions cannot be directly inferred from smartcard data alone. Some require either the development and administration of a survey or conducting an analysis at the origin-destination level (to determine, for example,

the available modes and the travel times on each mode), which is more useful in an operations, rather than a fare policy, context. On the other hand, the questions in the second bullet point above pertaining to ticket choice may easily be addressed. Public transport ticket types and their costs are easy to obtain. The effect of the ticket choice on the mode and time-of-day choices can be determined to a reasonable level of detail (e.g. a bus period ticket may only be used on buses and an off-peak fare may only be used during off-peak periods).

To resolve the issues related to modeling mode and time-of-day choices, we can aggregate the data to the individual (or smartcard) level. By doing so, we will have to model the frequency of trips made by a given individual during a given time period on each mode and at each time of day, rather than the discrete choice of mode and time of day for each trip. This aggregation eliminates the need for trip-specific attributes for each mode and time-of-day alternative but retains the ability to capture the tradeoffs between these different alternatives. For example, a fare differential that would make the 'off-peak' alternative more favorable than the 'peak' alternative in a time-of-day discrete choice model would also make the frequency of travel in the off-peak period higher than in the peak period in a model which considers trips aggregated to the individual level.

To maintain consistency across all three choice dimensions, we will also model ticket choice at the individual, rather than the trip, level. Note, however, that even with such an aggregation, we will still be able to model ticket choice within a discrete choice framework (rather than modeling the frequency of travel by a given individual during a given time period with each ticket type). This is because we can assume that, during a short enough time period, an individual makes one ticket choice (rather than one for each trip, as is the case with mode and time-of-day choices). An individual who purchases a weekly period ticket at the beginning of a week, for example, will most probably not use a different ticket type during that week for trips that can be made with the weekly ticket he/she already holds.

Given the above discussion, we now want to develop a framework to model the following:

1. A discrete ticket choice

2. The frequency of use on each available public transport mode and time-of-day segment (Time-of-day segments can be defined based on a number of criteria as discussed later in this section.)

As mentioned earlier, mode, ticket, and time-of-day choices are interdependent (i.e. one choice affects the other two). To account for this:

- Frequency of use will be modeled in mode-time-of-day segment combinations (e.g. Bus peak use and bus off-peak use). Doing so accounts for the relationship between mode and time-of-day choices.
- Frequency of use will be modeled based on ticket choice. Doing so accounts for the relationship between ticket choice and the two other choices.

C_j is a ticket type.
Q_{jis} is the frequency of usage on mode i and time segment s given ticket choice C_j.
J is the number of ticket types available.
M_j is the number of modes in which travel is permitted given ticket choice C_j
S_j is the number of defined time periods in which travel is permitted given ticket choice C_j.
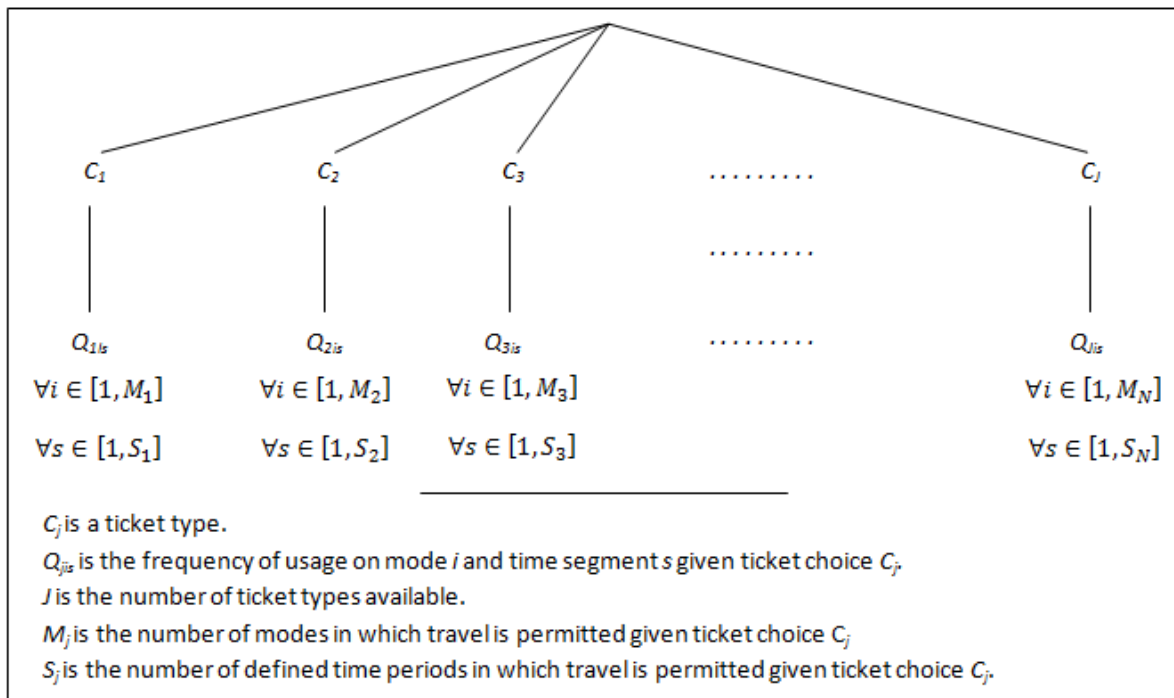
Figure 4-1: Proposed model structure

Figure 4-1 summarizes the proposed model structure. The top level in the figure captures the choice among the different ticket types. Under each ticket choice, there is a set of frequencies of use defined by a mode and time-of-day segment. Ideally, time-of-day segments would be defined based on the fare structure at hand since we are developing a policy-sensitive fare model. So, if there are two fare levels, peak and off-peak, for a given ticket type, then we would define

two time-of-day segments ('peak' and 'off-peak') for that ticket type under the assumption that the frequency of use would be different in each segment given the difference in fares. Note that the model structure has the flexibility of defining different modes and time-of-day segments for different ticket types. This allows accommodation of more complex fare structures.

4.3 Theoretical Background

The proposed methodology presented in the previous section involves modeling a discrete choice and one or more continuous quantities that are based on that discrete choice. A class of models, known as discrete-continuous models, has been developed specifically for that purpose. Discrete-continuous models are well documented in the literature and have been applied in many contexts, including car ownership and use (Train, 1993). The basic theory underlying discrete-continuous models is presented in Subsection 4.3.1.

In Subsection 4.3.2, we introduce a dynamic aspect to the analysis. So far, we have talked about ticket choice and frequency of use "by a given individual during a given time period". Given a suitable dataset, however, our methodology can be extended to allow for modeling individuals' behavior over time. In other words, we can use individuals' ticket choices and frequencies of use through multiple time periods in order to better understand the effects of fare changes on public transport demand and to produce a more powerful forecasting tool.

4.3.1 Discrete-Continuous Models

In many cases, individuals are faced with two interdependent choices: one that is made among a discrete set of alternatives and another that is made among a continuous set of alternatives. In the proposed methodology presented above, a discrete choice is made among ticket types and a continuous choice is made about the number of public transport journeys.

The interdependence among the two choices can be illustrated with basic microeconomic theory. Figure 4-2 shows an individual's indifference curves in a

world with two goods: (1) public transport journeys and (2) all other goods and services. The figure includes two budget constraints: one that is linear (in red) and another that is non-linear (in blue). The linear budget constraint applies if the individual chooses the per journey fare option. The individual in that case would make public transport journeys (i.e. move along the red line) up to a point that maximizes their utility, which is the point of tangency between the red line and the indifference curves. Note that the slope of the red line represents the cost of each journey in terms of 'all other goods and services'.

The non-linear budget constraint, on the other hand, applies if the individual chooses the period ticket option. In that case, the individual would first have to pay the fixed cost of the ticket (represented by the vertical portion of the blue budget constraint), after which they can make any number of journeys without additional cost. The horizontal portion of the budget constraint indicates that journeys under a period ticket are free. Naturally, it is not rational for an individual to purchase a period ticket without making any journeys, as this moves them down the vertical axis and reduces their utility. Also, note that after purchasing a period ticket, the utility of the individual keeps increasing as they make more public transport journeys. This implies that the individual should make an infinite number of public transport journeys. Of course, this does not occur in reality because of time constraints (not shown in the figure) and because the demand for transport is a derived demand which depends on a more complex set of decisions that pertain to other activities and choices.

In the simple example provided in Figure 4-2, we can find the optimal ticket choice for a given individual for any number of journeys. This is done by mapping every point on the horizontal axis to the budget constraints that provides the higher utility (i.e. the budget constraint that is higher on the graph above each point on the horizontal axis). At point 'x', the individual is indifferent between a per journey fare and a period ticket. At any point to the left of 'x', the individual would be better off with a per journey fare, and at any point to the right of 'x', he/she would be better off with a period ticket.
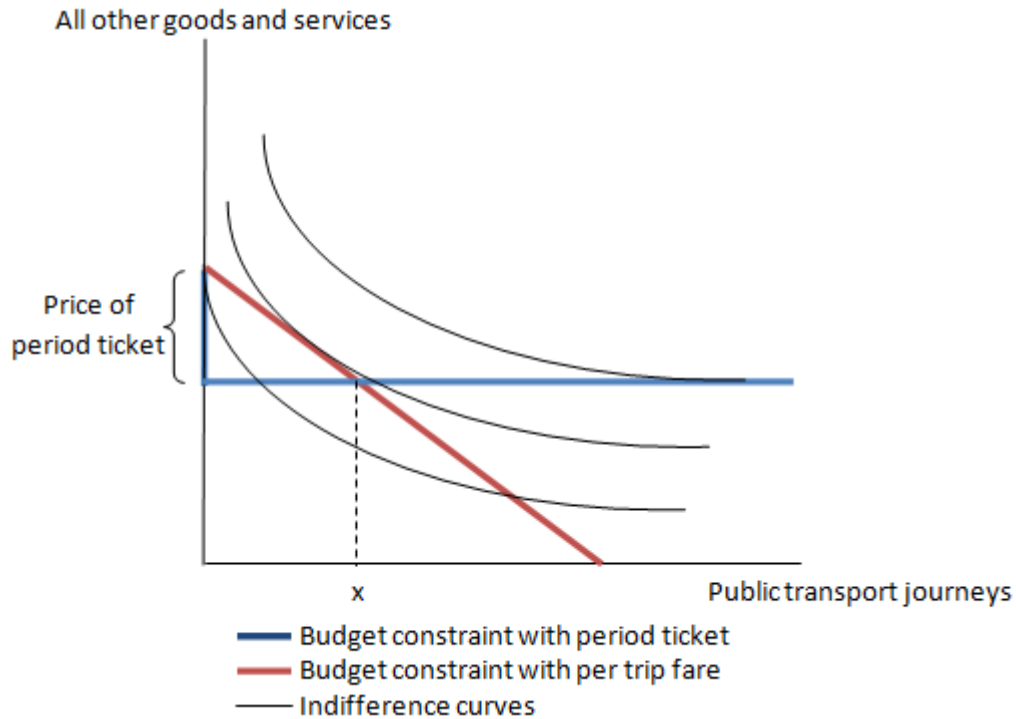
Figure 4-2: Microeconomic framework for ticket choice and public transport use

Figure 4-3 provides an example in which an individual is choosing between a period ticket and a per journey fare with best value pricing. In that case, an individual who chooses the per journey fare is guaranteed not to pay more than the price of a period ticket throughout the span of the time period covered by that ticket. (An example from London would be the choice between PAYG, which is capped on a daily basis, and a Day Travelcard.) In that case, the individual is indifferent between the two ticket types at point 'x' and any point to the right of it, but he/she is still better off with the per journey fare at any point to the left of 'x'.

These two figures highlight the interdependence between the discrete and continuous choices. Ticket choice is based on public transport use (an individual who expects to make more journeys would purchase a period ticket), and public transport use is based on ticket choice (an individual who chooses a period ticket will probably make a larger number of journeys than one who chooses a per journey fare because of the horizontal portions of the budget constraints shown above). The figures also show that both fixed and variable costs under each set of alternatives affect both ticket choice and public transport use. For example, a change in the per journey fare would change the slope of the red line in Figure 4-2.

This would naturally affect the number of journeys made, but it might also affect ticket choice, since the location of point 'x' on the horizontal axis would shift as the slope of the red line changed. In other words, a change in the per journey fare would change the minimum number of journeys that would make purchasing a period ticket worthwhile.
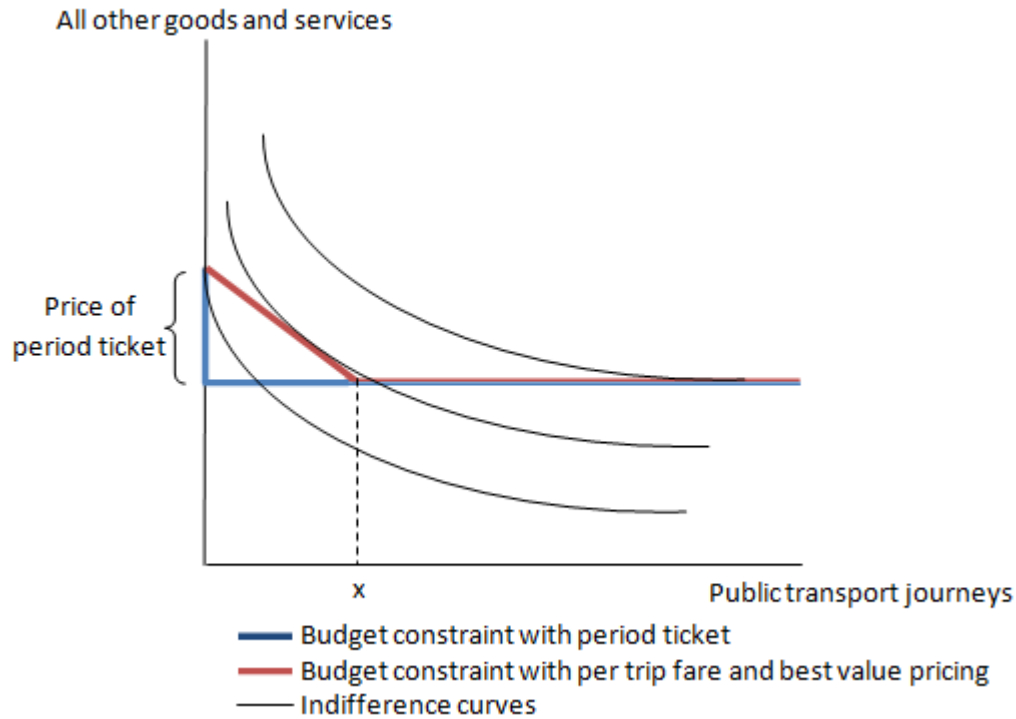
Figure 4-3: Microeconomic framework for ticket choice and public transport use with best value pricing

The framework presented in the two figures is simplistic. First, it ignores any temporal dynamics in ticket choice and public transport use. Second, it assumes that these two decisions are based solely on monetary costs. We relax these assumptions in Subsection 4.3.2. Before doing so, however, it may be useful to provide the analytical context for discrete-continuous models based on utility maximization theory. The notation and derivations below are based mostly on Train (1993).

Traditional, or so-called 'direct', utility functions measure an individual's utility given the quantity of goods and services consumed. $U(x_1,x_2)$, for example, is the utility of an individual who consumes x_1 of the first good and x_2 of the second

good in a bundle of two goods. The utility-maximizing bundle, call it (x_1^*, x_2^*), varies with price levels for x_1 and x_2 and with income levels for the individual.

If we were to map only the utility-maximizing bundles under different price and income levels (rather than the entire direct utility function with different quantities of x_1 and x_2), we would obtain the indirect utility function, call it Y. So, we have:

- Direct utility function $U(x_1, x_2)$: is the utility given quantities of x_1 and x_2
- Indirect utility function $Y(p_1, p_2, y)$: is the utility of the utility-maximizing bundle at price levels p_1 and p_2 and income level y

Using indirect utility functions is preferred in demand analyses, because one could easily derive the consumer's demand curve given such utility functions. This is done using Roy's identity which states that the demand for good i is defined as the negative of the partial derivative of the indirect utility function with respect to the price of i divided by the partial derivative of the indirect utility function with respect to the individual's income (Hausman, 1985). So, using the above notation, we have:

$$x_i^* = -\frac{\partial Y / \partial p_i}{\partial Y / \partial y}$$

for $i = 1$ or 2. See Train (1993) for a proof of Roy's identity.

Now, we extend the above discussion to account for the availability of two or more (discrete) alternatives to choose from before selecting the utility-maximizing (continuous) quantity of the good. Let the choice set available to an individual be J, the price of alternative i in J be p_i, the observed and unobserved characteristics of each alternative i in J be z_i and w_i, respectively, and the socioeconomic characteristics of the individual be s.

To put this in the context of ticket choice and public transport use, let us assume J includes the following three alternatives:

1. Per journey fare
2. Weekly period ticket
3. Monthly period ticket

Given the above set, we have the following indirect utility functions:

$$Y_1 = 0 \text{ (normalized to zero)}$$

$$Y_2 = f_2(p_2, y, z_2, s, w_2)$$
$$Y_3 = f_3(p_3, y, z_3, s, w_3)$$

Each of the above functions is called a *conditional* indirect utility function, since it represents the utility given that a certain alternative is chosen. In turn, alternative i in J is chosen if and only if $Y_i > Y_j$ for all j in J (where $i \neq j$). And since the unobserved characteristics w_i are not deterministic, we can say that the probability of choosing alternative i is:

$$P_i = \text{Prob}(Y_i > Y_j), \text{ for all } j \text{ in } J \text{ where } i \neq j$$

To summarize, modeling the interdependent decisions of ticket choice and public transport use requires first specifying either the indirect utility function of each alternative in the ticket type choice set or the demand function for public transport use and then deriving the other function using Roy's identity.

The interdependence between the two decisions is accounted for in the derivation of the demand functions, which is done from *conditional* indirect utility functions (i.e. utility functions that apply *given that a certain alternative is chosen*). On the other hand, accounting for the joint decision when specifying and (sequentially) estimating an empirical model requires taking further corrective actions, as explained below.

After estimating a discrete choice model for ticket choice, in which the alternatives are similar to the ones specified in the set J above, one cannot simply specify a linear regression model for public transport use under each alternative and estimate it using OLS. This is because each of these models will be estimated only on individuals who chose the respective alternative. As a result, estimating a simple OLS model without taking any corrective measures would result in what is known as "selectivity bias" or "self-selection bias" in the estimated coefficients. For example, in a regression equation for public transport use given the 'per journey fare' ticket type, the resulting fare coefficient (i.e. the estimated effect of fare changes on public transport use) will not reflect the true sensitivity to fare changes of the entire population. This is because such a regression model would not take into account any possible ticket switching to or from the 'per journey fare' ticket type that would also affect use (the dependent variable). In other words, a fare

change for ticket type *j* does not only affect the frequencies of use for those who chose that ticket type, as it could also cause other users to switch to ticket type *j*, which, in turn, would affect the observed frequencies of use for that ticket type.

This selectivity bias can be addressed using different methods which are highlighted in Train (1993) and Mannering and Hensher (1987). We will use one of those methods in our empirical example in Chapter 5.

4.3.2 Model Development with Panel Data

So far, we have presented our methodology in a static framework in which we model individuals' ticket choices and public transport use in a given time period based on fares and other factors. Introducing dynamics allows for capturing time-specific factors (e.g. seasonality effects) that may affect the two decisions and, more importantly, for including individual-specific factors and inertia in the model. Inertia implies that the decisions made by an individual in prior time periods affect their decisions in the current time period.

When developing and estimating a dynamic econometric model, it is important to distinguish between inertia and unobserved individual-specific factors. The occurrence of car accidents is sometimes used to illustrate this difference. If person X was in a car accident in both periods 1 and 2, then we cannot say that the occurrence of the accident in period 2 was a result of the accident in period 1 (i.e. a result of inertia). Rather, the second accident was a result of unobserved individual-specific factors, such as reckless driving. In public transport ticket choice and use, both inertia and unobserved individual-specific factors play a role. An individual may purchase a monthly ticket because they do so every month (inertia) and also because they have a high income that allows them to do so (an unobserved individual-specific effect—assuming income is unobserved and, hence, not explicitly controlled for in the model).

Modeling inertia and unobserved individual-specific effects requires having a panel dataset in which a number of individuals are observed over multiple time periods. Econometric models account for inertia and unobserved individual-specific effects as follows (Heckman, Statistical Models for Discrete Panel Data, 1981):

- Inertia is known as "state dependence" in discrete choice models and is addressed by including alternative-specific dummy variables (or functions of these variables) indicating whether or not each alternative was chosen in the previous period. In regressions (i.e. continuous models), inertia is addressed by including lagged dependent variables. So, if y_t was the dependent variable, we could add y_{t-1} as an independent variable to indicate that the observed value of the dependent variable in period t-1 influences the observed value of that variable in period t.

- Unobserved individual-specific effects (also known as serial correlation or unobserved heterogeneities) are addressed either by assuming "fixed effects" (i.e. a constant individual-specific effect for each individual in the sample) or "random effects" in which each individual-specific effect is modeled as a distribution, rather than as a constant. In the latter case, the error term for each observation (in both the utility and regression specifications) is split into a term that does not vary by individual and one that varies with each observation.

One can account for either inertia or unobserved individual-specific effects or both when developing models using panel data. If neither effect is accounted for, the model is said to be static and would treat each observation separately.

4.4 Summary of Methodology

We conclude this chapter by summarizing the basic formulation and data requirements for the methodology outlined above.

Our proposed methodology consists of developing the following set of models:

1. A disaggregate discrete choice model for public transport ticket choice: In this model, the systematic utility of each ticket type alternative is specified. The result is a set of choice probabilities (one for each alternative) for each individual-time period observation in the sample.

2. Aggregate regression models for public transport use: A set of models is developed for each ticket choice alternative. The dependent variable in each

model is the frequency of public transport use for a given mode and time of day conditional on the ticket choice.

Estimating the above set of models requires a panel dataset, in which individuals are observed over time. This will allow controlling for inertia and/or unobserved individual-specific effects that could determine both ticket choice and the frequency of public transport use.

To conclude, this chapter has presented our proposed methodology for modeling the effects of fare changes on public transport demand. We started by presenting a conceptual framework. This was followed by the proposed model structure and an overview of the theory underlying that structure. In the next chapter, we apply our methodology to London's public transport system and present a policy analysis of the results.

Chapter 5

Model Application and Policy Analysis

In this chapter, we apply the methodology developed in Chapter 4 to London's public transport system using disaggregate Oyster card data. We specify an empirical model based on the structure presented earlier, present the estimation results, and discuss the policy implications of those results. In Section 5.5, we present four applications of the model illustrating how it can be used to evaluate different fare policy decisions in London. The chapter concludes by highlighting ways in which the results of the analysis could be incorporated into models currently used at Transport for London.

Before delving into the empirical specification of the model, however, we will discuss, in some detail, the preprocessing of the Oyster data. The dataset in its raw format cannot directly be used to apply our model, since it does not represent a panel dataset in the traditional sense. In a panel dataset, there are discrete time periods in which individuals are observed. In the raw Oyster dataset, however, we only observe a series of taps-in (and taps-out for Underground journeys) by Oyster cards over a time continuum. The preprocessing of the data is the topic of the first section of this chapter.

5.1 Data Preprocessing

In order to apply the methodology proposed in Chapter 4, we will need a panel dataset with a structure similar to that shown in Table 5-1. Such a data structure

includes observations for individuals (or cards[8]) over a set of discrete time periods. In Table 5-1, which is purely illustrative, we observe card n_1 in three time periods (t_1, t_2, and t_3) and card n_2 in two time periods (t_1 and t_2). This example represents an unbalanced panel dataset, in which not all individuals are observed in every single time period (as opposed to a balanced panel dataset, in which each individual in the sample is observed in every time period). The dataset we will construct in this section will also be an unbalanced panel dataset, since Oyster cards appear and disappear from the sample in different time periods.

Card	Time period	Ticket choice	Frequency of use (by mode and time of day)				Explanatory variables		
n_1	t_1	. . .	. . .	. . .	. . .	. . .	. . .	. . .
n_1	t_2	. . .	. . .	. . .	. . .	. . .	. . .	. . .
n_1	t_3	. . .	. . .	. . .	. . .	. . .	. . .	. . .
. . .	. . .	. . .	. . .	. . .	. . .	. . .	. . .	. . .
n_2	t_1	. . .	. . .	. . .	. . .	. . .	. . .	. . .
n_2	t_2	. . .	. . .	. . .	. . .	. . .	. . .	. . .
. . .	. . .	. . .	. . .	. . .	. . .	. . .	. . .	. . .

Table 5-1: Panel data structure required to apply the proposed model

In the raw Oyster dataset used in this analysis, each observation represents a journey segment made on London's public transport system. Some of the fields (or data items) in each observation include the following:

- The date on which the journey segment was made
- An encrypted ID for the Oyster card which was used to make the journey segment
- The mode on which the journey segment was made
- The boarding and alighting stations for Underground journey segments and the bus route for bus journey segments.
- The start time of the journey segment (and the end time for Underground journey segments)
- The ticket type used for the journey segment

[8] We will use "card" and "individual" interchangeably, assuming that each Oyster card in the sample represents one individual. This assumption ignores the fact that some cards may be shared by friends or members of the same family.

In order to transform the raw data into a panel with a structure similar to that in Table 5-1, we need to address two issues: (1) the definition of the discrete time periods of the panel and (2) the definition of "ticket choice" as it pertains to these time periods. These issues have been resolved as follows:

1. A week starting on Sunday and ending on Saturday has been defined as a time period in the panel. Given this definition, we are assuming that every individual in the sample makes one distinct ticket choice every week. There are certainly flaws in such an assumption, since individuals can change the ticket type they use halfway through a week or use multiple ticket types in a given week; however, as Table 5-2 shows, the assumption is generally valid. In other words, we do not lose much by aggregating the data to the week level. The table shows the distribution of cards observed in the week between November 30 and December 6, 2007 based on the percentage of journey segments made with the most frequently used ticket type during that week. The fact that 91% of those cards used only one ticket type during the week studied supports our definition of a week as the time period in the panel.

Percentage of journeys made using most frequently used ticket type	Percentage of cards
100%	91%
75-100%	6%
60-75%	2%
50-60%	1%
<50%	0%

Table 5-2: Distribution of Oyster cards based on the percentage of journeys made with the most frequently used ticket type in the week of 11/30-12/6/2007

2. Given our above definition of a time period, the definition of "ticket choice" becomes straightforward. The ticket type "chosen" by an individual in a given week is the ticket type by which most of that individual's journey segments were made during that week. Given the information in Table 5-2, the ticket choice for a given individual in a given week actually represents the only ticket type used in that week for 91 percent of individuals and the ticket type used for more than half of the journeys for the remaining 9 percent of individuals.

In this analysis, we will consider only the following ticket types:

1. Pay-As-You-Go (PAYG)
2. Weekly Travelcard
3. Monthly Travelcard
4. Annual Travelcard
5. Weekly Bus Pass
6. Monthly Bus Pass

Although there are many more ticket types available on Oyster cards, we only consider the above ticket types because they are by far the most widely used, as shown in Chapter 3.

This section discussed the preprocessing of the raw dataset to make it usable in the context of our proposed methodology. In order to estimate the discrete-continuous model on the Oyster dataset, some further preprocessing is needed. This mostly relates to constructing the explanatory variables and will, thus, be discussed in the next section, which presents the empirical specification of the model.

5.2 Model Specification

Using the procedure outlined in the previous section, we can now use Oyster data to specify and estimate a discrete-continuous model for ticket choice and public transport use in London. This section presents the empirical specifications of both the discrete and continuous sub-models.

5.2.1 Discrete Choice Sub-Model

As mentioned earlier, the six ticket choices that will be considered in this analysis are PAYG, Weekly Travelcard, Monthly Travelcard, Annual Travelcard, Weekly Bus Pass, and Monthly Bus Pass. The ticket choice sub-model of the discrete-continuous model will be specified and estimated as a multinomial logit model. For a detailed overview of multinomial logit models and other discrete choice models, see Ben-Akiva & Lerman (1985).

Since we are working with smartcard data, the information we have on the different users of the system is limited. We also have little information on service

attributes on public transport and non-public transport modes. As a result, our model does not include measures for all the variables listed in the conceptual framework in Chapter 4. The only factors we use to explain the ticket choice that an individual makes in a given week are (1) the individual's previous choice (this captures the inertia effect discussed earlier) and (2) the expected cost of travel under the different ticket alternatives. The expected cost of travel under each available ticket type in a given week is estimated using that individual's previous travel patterns. More specifically, an individual's expected travel patterns in week t are assumed to be similar to their observed travel patterns in week t-1. One shortcoming in using this approach is that it does not account for other factors that may affect expected cost. For example, if week t-1 did not include any holidays, whereas week t included one holiday, the measure of expected cost for week t would not be very accurate, since frequency of use (and, as a result, cost) in week t would probably be significantly lower than in week t-1. (This may be especially problematic in capturing cases where people temporarily switch between tickets, for example from Monthly Travelcards to Weekly Travelcards, because of upcoming holidays.) Addressing this can be done either by simply accounting for holidays when estimating expected cost (by scaling down the previous week's costs depending on the number of holidays in the current week) or by using more advanced techniques to model expected cost (see Section 6.2).

Furthermore, using the choice and the cost of travel in the previous week as factors that explain ticket choice raises some questions regarding the first observation of each Oyster card in the sample. In other words, having no "previous choice" or a measure for expected cost, how do we predict the very first ticket choice that an Oyster user makes? To address this, we use information on what other "first-time users" have done in the past. More precisely, we use an aggregate measure of ticket shares among previous first-time users to better predict what current first-time users will choose as their ticket type. This can be thought of as an aggregate extension to the disaggregate inertia effect. For an individual whom we have observed in the past, the current ticket choice is a function of that individual's previous choice; for a "new" individual whom we have not observed in the past,

the current ticket choice is a function of previous choices made by other individuals. Such an approach helps predict the first choice of each card and also captures some policies not explicitly accounted for that change the relative attractiveness of the various ticket types over time.

Given the above discussion, we have the following specifications for the systematic utilities of each ticket type[9]:

$$V_{PAYG,n,t} = \beta_{PAYGFirstObs}(PAYGMovAvg_{t-1} \times FirstObs_{n,t}) + \beta_{PAYG}PAYG_{n,t-1}$$
$$+ \beta_{PAYGCost}Cost_{PAYG,n,t}$$
$$V_{WTC,n,t} = \alpha_{WTC} + \beta_{WTC}WTC_{n,t-1} + \beta_{WCost}Cost_{WTC,n,t}$$
$$V_{MTC,n,t} = \alpha_{MTC} + \beta_{WTC}MTC_{n,t-1} + \beta_{MCost}Cost_{MTC,n,t}$$
$$V_{ATC,n,t} = \alpha_{ATC} + \beta_{WTC}ATC_{n,t-1} + \beta_{ACost}Cost_{ATC,n,t}$$
$$V_{WBP,n,t} = \alpha_{WBP} + \beta_{WBP}WBP_{n,t-1} + \beta_{WCost}Cost_{WBP,n,t}$$
$$V_{MBP,n,t} = \alpha_{MBP} + \beta_{MBP}MBP_{n,t-1} + \beta_{MCost}Cost_{MBP,n,t}$$

where

PAYG corresponds to the Pay-As-You-Go alternative,

WTC corresponds to the Weekly Travelcard alternative,

MTC corresponds to the Monthly Travelcard alternative,

ATC corresponds to the Annual Travelcard alternative,

WBP corresponds to the Weekly Bus Pass alternative,

MBP corresponds to the Monthly Bus Pass alternative,

$V_{i,n,t}$ is the systematic utility experienced by individual *n* in week *t* if he/she chooses alternative *i*,

$i_{n,t}$ is a dummy variable equal to 1 if alternative *i* was chosen by individual *n* in week *t* and 0 otherwise,

PAYGMovAvg_t is a weighted moving average of observed shares of PAYG among all previous first-time users from week 1 up to week *t*,

FirstObs_{n,t} is a dummy variable equal to 1 if week *t* is the first week in which individual *n* is observed and 0 otherwise,

---

[9] The utility of each ticket type is the sum of the systematic utility and an error term. In multinomial logit models, the error terms are assumed to have an extreme value distribution (Ben-Akiva & Lerman, Discrete Choice Analysis, 1985).

$Cost_{i,n,t}$ is the expected weekly cost of travel using alternative *i* for person *n* in week *t*.

The variable *PAYGMovAvg*, which, when interacted with the dummy *FirstObs*, addresses the issues raised regarding the first observation of each card, is estimated as follows:

$$PAYGMovAvg_t = \frac{tPAYGShare_t + (t-1)PAYGShare_{t-1} + \cdots + 2PAYGShare_2 + PAYGShare_1}{t + (t-1) + \cdots + 2 + 1}$$

where

$PAYGShare_t$ is the aggregate share of PAYG among individuals who are observed for the first time in week *t*.

From the above expression, we see that *PAYGMovAvg*_t is a weighted moving average in which the most recent week has the highest weight and in which weights decrease arithmetically for preceding weeks. Note that the model includes only the share of PAYG among previous first-time users. This is because fare policies in London and behavioral changes over time generally tend to affect the share of PAYG relative to all other ticket types and not, for example, the share of Weekly Travelcards relative to Monthly or Annual Travelcards. The interaction term $PAYGMovAvg_{t-1} \times FirstObs_{n,t}$ can be thought of as an "adjustment factor" that modifies the attractiveness of the various ticket types over time for first-time users. Without this term, the attractiveness of each ticket type to first-time users is restricted by the alternative-specific constant to be the same throughout the entire panel. In other words, excluding this term would make the probability of a new user in 2005 choosing PAYG equal to the probability of a new user in 2007 choosing PAYG. This is not realistic given the changing policies in London that have aimed to make PAYG a more attractive alternative (see Chapter 3).

$PAYGShare_t$ is meant to be an aggregate measure of the ticket choices of previous first-time Oyster users. A simpler measure that could have been used is the share of PAYG among first-time Oyster users in the previous week (*t*-1). However, in our exploratory data analysis, we found some week-to-week fluctuations in 'first-time ticket shares'. The overall trend, however, indicated an increasing share of PAYG among new adopters of the Oyster card. In order to capture this trend, we, therefore, chose to use a weighted moving average measure,

which, despite being a function of the shares of PAYG among first-time users in all previously observed weeks in the sample, gives the highest weights to the most recent weeks.

Next, we discuss the cost variables used in the above specifications. Calculating the cost variables for each observation (i.e. for each individual in a given week) is not straightforward. PAYG cost is not simply the number of journeys made during the week multiplied by a fare level and the cost of travel under a Weekly or Monthly Travelcard, for example, is not the same for all individuals in a given week. This is because Underground fares in London differ by zones of travel, as discussed in Chapter 3. PAYG Underground fares are different for travel that includes Zone 1 and varies by the number of zones through which an individual passes during a journey segment. Travelcard prices vary in a similar way. Given these variations, it s important to have a fairly accurate estimate of the costs of travel under the different ticket type for each observation in the sample.

Our ticket choice model is not meant to capture the choices among Weekly, Monthly, or Annual Travelcards with different zonal validities (e.g. the choice between a Zone 1-2 Weekly Travelcard and a Zone 1-4 Weekly Travelcard). This is because we are not modeling destination choice but, rather, studying the choice among different fare options *given* an individual's destination choice. The approach we use to calculate the cost variables should, therefore, be based on finding the most likely zones of travel for each observation in order to produce a realistic choice set that an individual faces in a given week.

Given this, we use the following approach to calculate the cost variables:

1. Every observation is classified into an LU fare category. Fare categories are defined by (1) whether or not travel to Zone 1 is included and (2) the difference between the innermost and outermost zones through which Underground travel is permitted. This corresponds to the criteria by which fares vary in London. The LU fare category to which an observation, corresponding to individual n in week t, is assigned depends on the patterns of use for individual n in the previous week in which they are observed. So, if most Underground journey segments made by individual n in week t-1, for example, were from Zones 1 to

2, then the observation for individual n in week t will be assigned to the LU fare category of "Zone 1-2". In other words, the expected LU fare category for an individual in a given week is the most frequently used fare category for that individual in the previous week in which he/she is observed. Individuals who either do not make any Underground journeys in their previous observed week or who do not tap out after most of their Underground journeys in that week are assigned to a special LU fare category in which the PAYG Underground fare is the average of that fare for all other categories and in which the price of a Weekly, Monthly, and Annual Travelcard is the average of the prices of those Travelcards for all other fare categories.

2. The cost of the Weekly, Monthly, and Annual Travelcard is assigned to each observation based on its LU fare category (which, in turn, is based on use in the previous week). Note, however, that if an individual is first observed with a Monthly Travelcard in week t, then their Monthly Travelcard cost would be zero for weeks $t+1$, $t+2$, and $t+3$ (because they would already have a valid Monthly Travelcard in those weeks, so their cost of travel under a Monthly Travelcard alternative would be zero). Time validities are addressed in a similar fashion for Annual Travelcards and Monthly Bus Passes.

3. The costs of the Weekly and Monthly Bus Passes are the same for all individuals in a given week, since there are no zonal variations of these period tickets.

4. PAYG cost for each observation is calculated by multiplying the Underground fare corresponding to the observation's fare category by the number of Underground journeys made in the previous week and the bus fare (which is the same across all fare categories) by the number of bus journeys in the previous week. To account for daily price capping, the number of journeys used in the PAYG cost calculations is modified, such that if an individual made an average of more than three journeys per day in a given week, only three of those journeys are counted.

Note that the assignment of a fare category and the calculation of the costs of the different alternatives use patterns of public transport use in the previous observed week. This is because we are developing a model that attempts to explain

people's behavior. When an individual is assumed to make a ticket choice at the beginning of week t, their expected use is something that is not observed (because it would not yet have occurred). Our model uses past use patterns as a measure of expected use and uses those patterns to calculated expected costs. In summary, we are assuming that use in period t is similar to use in period t-1 in terms of the fare category (i.e. zones of travel), frequency of travel on each mode, and the distribution of travel (in terms of journeys per day). The first observation of each card has an expected cost of zero for all alternatives (since we would not have information on the past use of the individual using that card). The dummy variables corresponding to the previous choice would also be equal to zero for the first observation of each card.

Finally, after having discussed the definitions of the explanatory variables in the ticket choice sub-model, we present our *a priori* expectations on the signs of the coefficients:

- The dummies indicating previous choice are all expected to have positive coefficients. Choosing ticket type i in week t would make an individual more likely to choose that ticket type in week t+1. (This is the inertia effect or the effect of 'state dependence' discussed in Chapter 4.) In fact, we would expect the dummy variables in the utility functions of monthly period tickets (i.e. the Monthly Travelcard and Monthly Bus Pass) to be larger in magnitude than those in the utility functions of weekly period tickets (and, similarly, we expect the dummy variable in the Annual Travelcard utility function to be larger than the dummy variables in the other utility functions). This is because the inertia factor is stronger on a weekly basis for period tickets with longer time validities. For example, an individual with an Annual Travelcard is more likely to be observed choosing that ticket type for multiple consecutive weeks compared to a Monthly Travelcard, simply because an Annual Travelcard is valid for 52 weeks, whereas a Monthly Travelcard is only valid for about 4 weeks.
- The interaction term ($PAYGMovAvg_{t-1} \times FirstObs_{n,t}$) is expected to have a positive coefficient. As the adoption of PAYG increases among first-time users, future first-time users are more likely to choose PAYG.

- The coefficients of the cost variables are expected to be negative. As the cost of travel under a given ticket type increases, the probability of choosing that ticket type decreases (holding all other factors constant). Notice that period tickets with the same time validities have the same cost coefficient. This is based on the assumption that a £1 in the cost of a weekly period ticket, for example, causes the same disutility across all types of weekly period tickets, since such tickets share the same time validity and payment schedules (in terms of upfront payments).

Before presenting the specifications of the continuous sub-models, it may be useful to note that the ticket choice sub-model developed above represents what is known as a "holding model", as opposed to a "transaction model". Our definition of ticket choice is based on use patterns and not on some transaction that is observed in which an individual purchases a certain ticket type. In a transaction model, observations in which an individual holds a valid Annual Travelcard that was purchased earlier, for example, are not included (because they are not associated with any transaction). In our model, on the other hand, we include these observations with the Annual Travelcard cost equal to zero (as discussed above). In other words, our model is explaining the decision of "holding" different ticket types on a weekly basis as opposed to the decision to purchase a ticket. The difference between the two types of models is subtle but is, nevertheless, worth noting.

5.2.2 Continuous Sub-Models

The second part of our proposed methodology includes a set of continuous regression models in which the dependent variable is the frequency of public transport use based on the ticket choice. Multiple types of 'frequency of use', defined by mode and time of day, can be modeled under each ticket type. Ideally, the modes and times of day by which the frequencies of use are defined would be based on the fare structure at hand. More precisely, each type of frequency of use defined under a ticket choice should have one value for the per journey fare.

Ticket type	Weekly frequencies of use to be modeled
PAYG	Uncapped Underground journeys
	Uncapped bus journeys
	Capped Underground journeys
	Capped bus journeys
Weekly Travelcard	Underground journeys
	Bus journeys
Monthly Travelcard	Underground journeys
	Bus journeys
Annual Travelcard	Underground journeys
	Bus journeys
Weekly Bus Pass	Bus journeys
Monthly Bus Pass	Bus journeys

Table 5-3: Frequencies of use to be modeled for London

Table 5-3 shows the frequencies of use we defined for each ticket type in London. For each Travelcard, we model the number of Underground journeys and the number of bus journeys, and for each Bus Pass, we model the number of bus journeys. (Again, the unit of time in the panel we have constructed is a week, so we are modeling the frequency of use by an individual in a given week.) For PAYG, we have defined four types of frequency of use: uncapped Underground journeys, uncapped bus journeys, capped Underground journeys, and capped bus journeys. Note that these types are defined such that there is one value for per journey fare for each. This is clear for uncapped bus journeys, for which the per journey fare is the bus fare, and capped Underground and bus journeys, for which the per journey fare is zero. Uncapped Underground journeys, on the other hand, are assumed to have the same per journey fare for each individual in a given week. This fare is defined as the PAYG per journey Underground fare in the fare category to which that individual is assigned in that week (i.e. the fare category most frequently used by that individual in that week). Furthermore, since there are actually two PAYG per journey fares for each fare category (one that is effective from 7am to 7pm on weekdays and another that is effective at all other times, as discussed in Chapter 3), we also assume that the per journey fare of uncapped Underground journeys is that of the most frequently used time period (between the two time periods defined above). Note that we do not separately model the journeys made in these different time periods (as suggested by our model structure), because they do not clearly

represent a 'peak' and an 'off-peak' period in the traditional sense. As a result, one cannot make straightforward *a priori* hypotheses on the differences in fare elasticities within these two periods and on any tradeoffs or trip time switching that is made between them. In fact, when we estimated two separate models, one for Underground peak use and another for Underground off-peak use (based on the peak and off-peak fares defined in the current fare structure), we obtained an off-peak fare coefficient that was statistically insignificant. We also obtained counterintuitive results when we added the fare of each time period in the use equation of the other period to account for any time-of-day tradeoffs.

Also, note that we do not include the Underground fare as a factor that explains bus travel and vice versa. This is because the dataset on which we estimate the model includes frequencies of use measured in journey segments and not in complete journeys or trips that include interchanges. As a result, the Underground and bus can be either substitutes or complements, so it is not clear how fares on each mode would affect demand, as defined by the dataset, on the other.

Since we are modeling frequency of use based on ticket choice, we need to correct for selectivity bias in the regression equations, as discussed earlier. To do so, we will use the following selectivity bias correction term defined by Hay, Dubin, and McFadden and referenced by Hensher and Milthorpe (1987):

$$SBC_{n,t} = \frac{|J_{n,t}| - 1}{|J_{n,t}|} \ln\left(\hat{P}_{i,n,t}\right) + \sum_{\substack{j=1 \\ j \neq i}}^{|J_{n,t}|} \left(\frac{\ln\left(\hat{P}_{j,n,t}\right)}{|J_{n,t}|} \frac{\hat{P}_{j,n,t}}{1 - \hat{P}_{j,n,t}}\right)$$

where

$SBC_{n,t}$ is the selectivity bias correction term for the observation corresponding to individual *n* and week *t*,

i is the chosen ticket alternative,

$\hat{P}_{n,i}$ is the probability of individual *n* choosing ticket alternative *i* in week *t* estimated by the ticket choice sub-model,

$|J_{n,t}|$ is the size of the choice set for individual *n* in week *t*.

The above term is defined for each observation (i.e. individual-week) in the sample and is a function of the probabilities of the chosen alternative and the non-

chosen alternatives, as well as the size of the ticket type choice set. The term is added to each 'frequency of use' regression equation to correct for selectivity bias. The estimated coefficient for this term is expected to have a negative sign[10].

Given the above discussion and based on data availability, we specify the following empirical regression models for frequencies of use for each of the ticket types defined in the ticket choice sub-model:

For the PAYG alternative:

$$LU_{Uncap,n,t} = \alpha_1 + \beta_1 LU_{Uncap,n,t-1} + \beta_2 Holidays_t + \gamma_1 LUFare_{n,t} + \tau_1 SBC_{n,t} + \varepsilon_{1,n,t}$$

$$Bus_{Uncap,n,t} = \alpha_2 + \beta_7 Bus_{Uncap,n,t-1} + \beta_8 Holidays_t + \gamma_2 BusFare_{n,t} + \tau_2 SBC_{n,t} + \varepsilon_{2,n,t}$$

$$LU_{Cap,n,t} = \alpha_3 + \beta_9 LU_{Cap,n,t-1} + \beta_{10} Holidays_t + \tau_3 SBC_{n,t} + \varepsilon_{3,n,t}$$

$$Bus_{Cap,n,t} = \alpha_4 + \beta_{11} Bus_{Cap,n,t-1} + \beta_{12} Holidays_t + \tau_4 SBC_{n,t} + \varepsilon_{4,n,t}$$

where

$LU_{Uncap,n,t}$ is the number of uncapped Underground journey segments made by individual n in week t,

$Bus_{Uncap,n,t}$ is the number of uncapped bus journey segments made by individual n in week t,

$LU_{Cap,n,t}$ is the number of capped Underground journey segments made by individual n in week t,

$Bus_{Cap,n,t}$ is the number of capped bus journey segments made by individual n in week t,

$Holidays_t$ is the number of holidays in week t (0, 1, or 2),

$\varepsilon_{i,n,t}$ is the unexplained error term for the observation corresponding to individual n and week t in regression equation i.

---

[10] The selectivity bias correction term defined above is negatively related to the error term in the utility specifications of the ticket choice sub-model. For two observations in which all the explanatory variables in the utility functions are equal, the one with the larger error terms (i.e. the more inherent inclination to choose one alternative over the others) is expected to have a higher frequency of use. So, a higher *SBC* implies a lower frequency of use (holding all else constant).

For each of the Weekly, Monthly, and Annual Travelcard alternatives:

$$LU_{n,t} = \alpha_{j,5} + \beta_{j,13}LU_{n,t-1} + \beta_{j,14}Holidays_t + \tau_{j,5}SBC_{n,t} + \varepsilon_{j,n,t}$$
$$Bus_{n,t} = \alpha_{j,6} + \beta_{j,15}Bus_{n,t-1} + \beta_{j,16}Holidays_t + \tau_{j,6}SBC_{n,t} + \varepsilon_{j,n,t}$$

where

$LU_{n,t}$ is the number of Underground journey segments made by individual n in week t,

$Bus_{n,t}$ is the number of bus journey segments made by individual n in week t,

j corresponds to the type of Travelcard (Weekly, Monthly, or Annual).

For each of the Weekly and Monthly Bus Pass alternatives:

$$Bus_{n,t} = \alpha_{k,7} + \beta_{k,17}Bus_{n,t-1} + \beta_{k,18}Holidays_t + \tau_{k,7}SBC_{n,t} + \varepsilon_{k,n,t}$$

where

k corresponds to the type of Bus Pass (Weekly or Monthly).

Following are some *a priori* hypotheses we can make about the signs of the coefficients in the above specifications:

- The coefficients for past use (β_1, β_7, β_9, β_{11}, $\beta_{Weekly,13}$, $\beta_{Monthly,13}$, $\beta_{Annual,13}$, $\beta_{Weekly,15}$, $\beta_{Monthly,15}$, $\beta_{Annual,15}$, $\beta_{Weekly,17}$, and $\beta_{Monthly,17}$) are all expected to be positive. An increase in the observed number of journeys for an individual in week t-1 results in an increase in the expected number of journeys for that individual in week t.

- The coefficients for the number of holidays in a given week (β_2, β_8, β_{10}, β_{12}, $\beta_{Weekly,14}$, $\beta_{Monthly,14}$, $\beta_{Annual,14}$, $\beta_{Weekly,16}$, $\beta_{Monthly,16}$, $\beta_{Annual,16}$, $\beta_{Weekly,18}$, and $\beta_{Monthly,18}$) are all expected to be negative. An individual will tend to make fewer journeys in a week with holidays compared to one with no holidays.

- The fare coefficients (γ_1 and γ_2) are expected to be negative. An increase in the per journey fare for PAYG users will result in a lower frequency of use.

- As mentioned above, the coefficients for the selectivity bias correction terms (τ_1, τ_2, τ_3, τ_4, $\tau_{Weekly,5}$, $\tau_{Monthly,5}$, $\tau_{Annual,5}$, $\tau_{Weekly,6}$, $\tau_{Monthly,6}$, $\tau_{Annual,6}$, $\tau_{Weekly,7}$, and $\tau_{Monthly,8}$) are expected to have negative signs.

Finally, we should note that for the first observation of each card, we follow a similar approach to that used in the ticket choice sub-model. The 'past use' in mode i for an individual who is first observed in week t with ticket choice j is the

weighted moving average up to week t-1 of the aggregate frequency of use on mode i for individuals who had chosen alternative j.

The next section describes the dataset on which the discrete-continuous model was estimated.

5.3 Dataset

The raw dataset used in this analysis includes journey segments for 5% of the Oyster cards in London. At the end of every TfL period[11], the dataset is updated with the journey segments of 5% of new cards that joined the Oyster system during that period.

The sample used in this analysis spanned the time period between Sunday, November 20, 2005, and Saturday, February 2, 2008. Using the procedure described in Section 5.1, a panel was constructed with 115 weeks (with every week starting on a Sunday and ending on a Saturday). Oyster cards which were not sold at full-price or included some form of a special ticket (e.g. Freedom Pass or Staff Pass) were excluded from the dataset. The remaining sample included 642,985 cards (appearing in either some or all of the 115 weeks) and was too large to work with using the available processing power. To address this, we randomly sampled about 5% of those cards (29,239 cards) and used this smaller sample to estimate the model. We later used another 5% random sample to validate our model.

Data for the four weeks starting on October 15, 2006 (week 48) and ending on November 11, 2006 (week 51) were missing from the raw dataset. To make our panel dataset complete, the following simple procedure was followed to impute the missing values of the dependent variables (ticket choice and frequencies of use) and the explanatory variables:

1. Cards that appeared in the sample in both weeks 47 and 52 (i.e. before the after the four 'missing' weeks) were assumed to have appeared in all four missing weeks. The values of the dependent and explanatory variables for those cards in weeks 48 and 49 were assumed to be equal to what they were observed to be in

---

[11] A TfL period is four weeks long. The first period (P1) of every calendar year starts on April 1, and the last period (P13) for that year ends on March 31 of the next calendar year.

week 47, and the values of those variables for this group of cards in weeks 50 and 51 were assumed to be equal to what they were observed to be in week 52.

2. Cards that appeared in the sample in week 47 but did not appear in week 52 or in any subsequent week were assumed not to have appeared in the sample in any of the four missing weeks.

With the above imputation, we have a full panel dataset with 115 weeks and just over 29,000 Oyster cards. The cards were observed in as few as one week and as many as 115 weeks. Figure 5-1 shows the distribution of the sampled cards based on the number of weeks in which they were observed. According to the figure, about 50% of the cards in the sample are observed in 35 weeks or less. In other words, the median turnover time for an Oyster card is just under nine months. The relatively short turnover time is expected, since many people, especially those who often use PAYG or weekly period tickets, tend to either lose or discard their Oyster cards and, thus, use multiple cards over time. This is not expected to happen as often for people who use period tickets with a time validity of one month or longer, since they are required by TfL to register their Oyster cards (Transport for London, 2008).

In the next section, we present the estimation results for the discrete-continuous model.
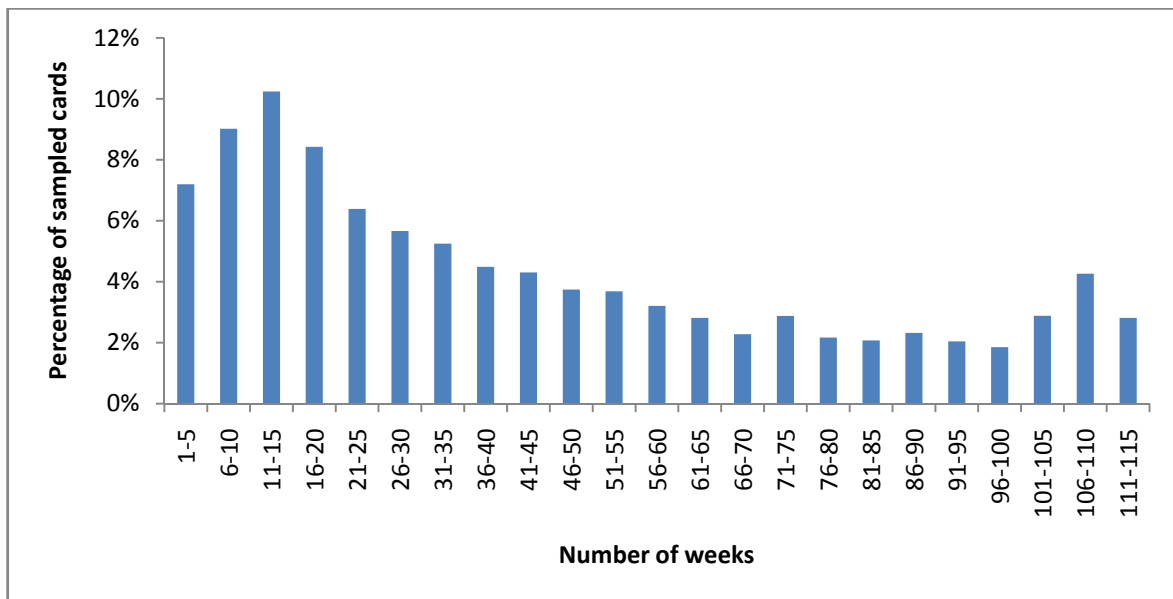
Figure 5-1: Distribution of sampled cards based on the number of weeks in which they were active

5.4 Estimation Results

In this section, we present the results we obtained estimating the models specified in Section 5.2 on the dataset described in the previous section.

5.4.1 Discrete Choice Sub-Model

The ticket choice sub-model specified earlier was estimated as a multinomial logit model using Biogeme[12]. The model specification includes dummies for previous choice (as shown earlier) but does not account for serial correlation. The ticket choice sub-model was, therefore, estimated assuming pure state dependence and no unobserved heterogeneities (see Subsection 4.3.2).

All six alternatives (PAYG, Weekly Travelcard, Monthly Travelcard, Annual Travelcard, Weekly Bus Pass, and Monthly Bus Pass) were made available to the first observation of each Oyster card in the sample. For every other observation, Bus Passes were included only if 70% or more of public transport journeys in the previous observed week were made on bus. The size of the choice set for each observation was, therefore, either four or six.

Table 5-4 shows the estimation results. The signs of the estimated coefficients are all in accord with our *a priori* expectations. What stands out in these results are the high levels of significance for the coefficients of the inertia variables. This has significant policy implications for TfL, as is illustrated in the next section. Any changes to the relative costs of the various ticket alternatives will result in a slow, belated response because of the strong inertia effect.

The increasing magnitude, in absolute terms, of the cost coefficients as the time validity increases indicates the disutility associated with large upfront payments. A £1 increase in the weekly cost of travel with a Monthly Travelcard, for example, corresponds to an increase in the upfront payment of about £4 when the

---

[12] Biogeme is a software package developed by Professor Michel Bierlaire at the Ecole Polytechnique Fédérale de Lausanne in Switzerland. Biogeme version 1.6 was used to estimate the ticket choice sub-model.

ticket is purchased. This disutility does not have direct policy implications but is a result of the way our holding model and the cost variables were defined.

Variable	Estimated Coefficient (t-stat)
Weekly Travelcard constant (Weekly Travelcard)	0.731 (25.34)
Monthly Travelcard constant (Monthly Travelcard)	-0.183 (-6.47)
Annual Travelcard constant (Annual Travelcard)	-1.67 (-31.42)
Weekly Bus Pass constant (Weekly Bus Pass)	-0.507 (-19.30)
Monthly Bus Pass constant (Monthly Bus Pass)	-1.66 (-45.33)
For first-time users: Share of PAYG among previous first-time users (PAYG)	2.75 (58.70)
Previous observed choice was PAYG (PAYG)	2.97 (154.25)
Previous observed choice was Weekly Travelcard (Weekly Travelcard)	2.69 (140.85)
Previous observed choice was Monthly Travelcard (Monthly Travelcard)	4.15 (153.59)
Previous observed choice was Annual Travelcard (Annual Travelcard)	7.35 (78.78)
Previous observed choice was Weekly Bus Pass (Weekly Bus Pass)	3.50 (136.24)
Previous observed choice was Monthly Bus Pass (Monthly Bus Pass)	5.17 (105.96)
PAYG cost (PAYG)	-0.051 (-64.47)
Weekly Period Ticket cost (Weekly Travelcard and Weekly Bus Pass)	-0.091 (-79.21)
Monthly Period Ticket cost (Monthly Travelcard and Monthly Bus Pass)	-0.132 (-109.75)
Annual Period Ticket cost (Annual Travelcard)	-0.250 (-43.51)
Number of observations	530,020
Final log likelihood	-130,619.36
Adjusted rho-squared	0.837

Table 5-4: Estimation results for the ticket choice sub-model

Furthermore, the alternative-specific constants estimated above (all relative to the PAYG constant, which was fixed at zero) show an inherent preference for Weekly Travelcards over PAYG. In other words, the model shows that in the case where the weekly cost of travel is the same across all alternatives and when there is no inertia effect, the Weekly Travelcard is a more attractive alternative than PAYG. This might be because holding a Weekly Travelcard provide individuals with unlimited journeys without having to worry about paying the per journey fare under PAYG or about the longer-term commitment associated with a Monthly or Annual Travelcard.

Finally, the variable measuring the share of PAYG among previous first-time users (having a value of zero for any observation that is not the first for a given card) has a significant positive coefficient, as expected. This indicates that as PAYG becomes more popular (which, in fact, has been the case because of marketing

efforts and policies such as daily price capping), new Oyster users become more likely to choose PAYG as their first ticket.

In order to validate the performance of the ticket choice sub-model, the estimation results were used to predict choice probabilities for a validation sample (i.e. a sample other than the one used for the estimation). The observed (denoted 'Obs.') and predicted (denoted 'Pred.') aggregate shares are shown in Table 5-5[13]. The shares are shown both for the entire panel and for the first observations of each card in the panel. (Note that these shares are not cross-sectional and, thus, have no intuitive meaning, since they measure aggregate ticket shares over multiple time periods. They are included here only for the purposes of validating the model.)

Subset		PAYG	Weekly Travelcard	Monthly Travelcard	Annual Travelcard	Weekly Bus Pass	Monthly Bus Pass
Full panel	Obs.	57.5%	17.40%	11.30%	4.50%	6.10%	2.70%
	Pred.	57.7%	17.70%	11.50%	4.50%	5.90%	2.70%
First observations	Obs.	58.2%	24.0%	6.1%	1.8%	7.7%	2.0%
	Pred.	57.7%	22.6%	9.1%	2.0%	6.5%	2.1%

Table 5-5: Comparison of observed and predicted shares in the validation sample

The numbers in Table 5-5 are promising and indicate that the ticket choice sub-model performs very well. The absolute difference between the observed and predicted shares does not exceed 3%. The table also shows that the model performs well in predicting the aggregate shares of ticket types among 'first observations', for which the choice probabilities are calculated based solely on the alternative-specific constants and past aggregate shares of PAYG. Having a model that produces fairly accurate predictions for new Oyster users is important, as it provides a useful forecasting tool that can accommodate not only changes in the relative costs among the various ticket types but also the increasing attractiveness or awareness of some ticket types due to unobserved factors. (One should keep in mind, however, that the validation sample used above was obtained from the same time period as the sample on which the model was estimated. Ideally, the model should be validated

---

[13] Predicted aggregate shares for a ticket type are calculated by taking the average of the predicted individual choice probabilities for that ticket type.

using a sample that spans a different time period in order to determine whether there are any time-specific factors that affect the performance of the model.)

We will use the above estimation results to conduct some simulated policy analyses in Subsection 5.5. Before doing so, however, we present the estimation results for the continuous frequency of use sub-models.

5.4.2 Continuous Sub-Models

Including the selectivity bias correction term in the regression equations of the continuous sub-models allows us to estimate these models using common econometric models to obtain unbiased and consistent estimators for the coefficients. The regression models specified in Section 5.2 were estimated using Stata and assuming random effects. The results for the PAYG sub-models and period ticket sub-models are shown in Tables 5-6 and 5-7, respectively.

Variable	Estimated coefficient (t-stat)			
	PAYG			
	Uncapped		Capped	
	LU	Bus	LU	Bus
Constant	2.96 (125.30)	2.90 (40.75)	0.156 (68.47)	0.181 (101.55)
Past use	0.369 (219.66)	0.385 (228.72)	0.123 (67.34)	0.362 (210.66)
Holidays	-0.574 (-56.37)	-0.575 (-38.48)	-0.019 (-7.89)	-0.013 (-3.31)
SBC	-0.134 (-18.06)	-0.276 (-25.15)	-0.005 (-2.91)	-0.031 (-12.10)
Fare	-0.601 (-58.65)	-0.329 (-4.48)	-	-
Overall R^2	0.480	0.49	0.07	0.13
Number of Observations	308,788			

Table 5-6: Estimation results for the frequency of use sub-models for PAYG

Variable	Estimated coefficient (t-stat)							
	Weekly Travelcard		Monthly Travelcard		Annual Travelcard		Weekly Bus Pass	Monthly Bus Pass
	LU	Bus	LU	Bus	LU	Bus	Bus	Bus
Constant	5.06	5.46	5.04	3.69	4.22	0.946	14.38	13.91
	(112.61)	(94.06)	(119.66)	(71.85)	(75.51)	(33.15)	(82.41)	(64.87)
Past use	0.349	0.467	0.410	0.516	0.447	0.821	0.309	0.342
	(112.92)	(160.73)	(117.92)	(154.58)	(81.14)	(224.10)	(51.83)	(47.91)
Holidays	-1.10	-1.09	-1.48	-0.850	-1.52	-0.633	-2.47	-2.71
	(-34.77)	(-21.33)	(-44.69)	(-21.24)	(-31.77)	(-11.82)	(-18.62)	(-20.61)
SBC	-0.029	0.191	-0.019	-0.002	-0.15	-0.137	0.176	-0.065
	(-1.93)	(7.90)	(-1.30)	(-0.14)	(-5.19)	(-4.19)	(3.62)	(-1.23)
Overall R^2	0.450	0.587	0.471	0.691	0.483	0.673	0.303	0.460
Number of observations	90,862		63,607		24,541		26,518	15,704

Table 5-7: Estimation results for the frequency of use sub-models for period ticket alternatives

Again, the results are mostly in accord with our *a priori* expectations discussed earlier. The coefficients for past use are positive; the coefficients for the number of holidays are negative, and the coefficients for most of the selectivity bias correction terms are negative. The fare coefficients in the uncapped LU journeys and uncapped bus journeys regression equations are also both negative. These latter coefficients can be used to calculate fare elasticities (i.e. the percentage change in the number of uncapped weekly LU journeys and uncapped weekly bus journeys given a one percent increase in the respective fares). In order to calculate such elasticities, we use the point-slope formula, which is typically specified as follows:

$$Elasticity_{Y\ w.r.t\ X} = \hat{\beta}\frac{\overline{X}}{\overline{Y}}$$

where

$Elasticity_{Y\ w.r.t\ X}$ is the elasticity of Y with respect to X, or the percentage change in Y given a one percent increase in X,

$\hat{\beta}$ is the estimated coefficient for the variable X in a linear regression equation in which Y is the dependent variable and X is one of the explanatory variables.

By substituting the estimated fare coefficient, the average fare, and the average number of uncapped journeys, we can obtain fare elasticities for both LU and bus. Since the regression equations include past use, these elasticities represent short-run fare elasticities. In order to obtain long-run elasticities, or measures of

people's responsiveness to fare changes in the long run, we assume that past use is equal to current use. So, the original specifications of the PAYG uncapped sub-models from Subsection 5.2.2 can be rewritten as follows:

$$LU_{Uncap,n,t} = \alpha_1 + \beta_1 LU_{Uncap,n,t} + \beta_2 Holidays_t + \gamma_1 LUFare_{n,t} + \tau_1 SBC_{n,t} + \varepsilon_{1,n,t}$$

$$Bus_{Uncap,n,t} = \alpha_2 + \beta_7 Bus_{Uncap,n,t} + \beta_8 Holidays_t + \gamma_2 BusFare_{n,t} + \tau_2 SBC_{n,t} + \varepsilon_{2,n,t}$$

(Notice that the subscript in the 'past use' variables is now t, not $t-1$.)

By rearranging the terms in the above equations, we obtain:

$$(1 - \beta_1)LU_{Uncap,n,t} = \alpha_1 + \beta_2 Holidays_t + \gamma_1 LUFare_{n,t} + \tau_1 SBC_{n,t} + \varepsilon_{1,n,t}$$

$$(1 - \beta_7)Bus_{Uncap,n,t} = \alpha_2 + \beta_8 Holidays_t + \gamma_2 BusFare_{n,t} + \tau_2 SBC_{n,t} + \varepsilon_{2,n,t}$$

Dividing both sides of the first and second equations by $(1 - \beta_1)$ and $(1 - \beta_7)$, respectively, we now have long-run fare coefficients of $\frac{\gamma_1}{(1-\beta_1)}$ and $\frac{\gamma_2}{(1-\beta_7)}$. By substituting the estimated values of these coefficients in the point-slope formula, we obtain the long-run fare elasticities. Table 5-8 shows the short and long-run elasticities for uncapped PAYG journeys on both LU and bus.

Mode	Short-run elasticity	Long-run elasticity
Underground	-0.40	-0.64
Bus	-0.08	-0.13

Table 5-8: Short and long-run fare elasticities for uncapped PAYG journeys

The long-run elasticities in the table are comparable to but generally lower than those currently used at TfL: -0.48 for LU and -0.25 for bus[14]. The elasticity values are also comparable to those in the literature. Mitrani *et al.* (2002), referenced in the literature review, estimated a fare elasticity on the London Underground of -0.41, which is very close to our short-run elasticity estimate shown in Table 5-8. The authors also calculated a fare elasticity for bus travel in London of -0.64, which is larger in absolute terms than our estimate. One possible reason for this difference, as well as for the difference between our estimated bus fare elasticity and the one currently used at TfL, is that the dataset on which we conducted our analysis had little variation in bus fares (bus fares were either £1.00 or £0.90 in the dataset). The

---

[14] These elasticity estimates were obtained from TfL's spreadsheet model. Since 'short run' in the context of our model represents a period in the order of a few weeks, we compare TfL's elasticities to our long-run estimates.

relatively small (in absolute value) bus fare elasticities, therefore, reflect the low fare sensitivity for bus users under such small variations.

Finally, we should note that the estimation results for the period ticket sub-models shown in Table 5-7 have no direct fare policy implications. They do, however, provide some insight into the differences in travel patterns between the different period ticket holders.

5.5 Policy Applications

In this section, we apply the estimation results of the discrete-continuous model to some fare policy scenarios for London.

5.5.1 Increasing the Prices of Monthly Travelcards and Bus Passes

The prices of Monthly Travelcards and Bus Passes are currently set at 3.84 times the price of their weekly equivalents. Here, we forecast the response of public transport users in London to a fare policy in which that multiple is increased to 4.5. The reasoning behind such a policy change would be the fact that the convenience—perceived by some—of holding a Monthly Travelcard compared to other tickets can override the effect of a relatively small increase in the monetary cost of that ticket for some users.

In order to produce the forecasts, we simulate the ticket choice sub-model over 66 weeks, starting with the last week in our sample (1/27 to 2/2/2008) and ending with the week of 4/26/2009 to 5/2/2009. The simulation is repeated under two scenarios: (1) no changes in fares and (2) an increase from 3.84 to 4.5 in the monthly-to-weekly price ratio. While the ticket choice sub-model is simulated, the 'previous choice' dummies are updated based on the simulated choices for the previous week. The validities of period tickets held by individuals are also updated. This is important in accounting for delayed responses to fare changes due to holding period tickets that have not yet expired.

Figure 5-2 shows the results of the simulation. The vertical axis indicates the difference between the simulated aggregate share of a ticket type under the policy

change and the simulated aggregate share of the same ticket type under the base policy. The trends in the figure show a forecast shift from Monthly Travelcards and Bus Passes to their weekly equivalents and, to an even greater extent, to PAYG. The simulations indicated that the aggregate share of PAYG in the last simulated week would be 68.1% without the policy change and 70.8% with the policy change. For Monthly Travelcards, these shares would be 7.1% and 3.3%, respectively. Note that we are simulating changes in ticket choice (i.e. the most frequently used ticket type) and not in frequency of use. The higher share of PAYG under the fare policy change will lead to a reduction in the total number of journeys, since PAYG users on average make a smaller number of trips compared to period ticket holders, as shown by the results of the continuous sub-models estimated above. This is also evident in the sample on which the model was estimated, where individuals who switched from Monthly Travelcards to PAYG made, on average, 5.3 fewer journeys per week using PAYG compared to the number of weekly journeys they used to make with a Monthly Travelcard.
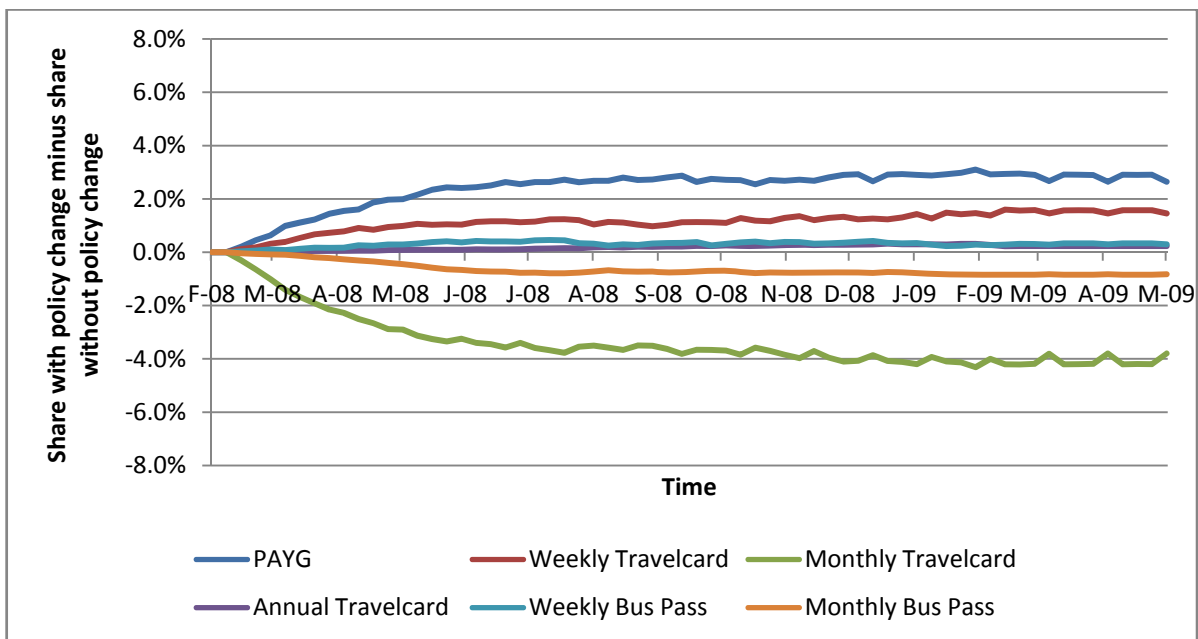

Figure 5-2: Simulation results for increasing the monthly-to-weekly price ratio

Figure 5-2 emphasizes the importance of simulation in making use of the estimation results presented in this thesis. The models we have developed are

largely based on inertia effects and past behavior. It is, therefore, expected that any policy change will have a delayed or gradual impact on public transport demand. This is clearly evident in the figure, which shows gradual changes in ticket shares that eventually stabilize five to six months after the new policy is implemented.

Aggregate ticket shares for Monthly Travelcards and Bus Passes with and without the policy change are summarized in Table 5-9, which shows that the reductions in the shares of both Monthly Travelcards and Bus Passes exceed the percentage increase in their respective prices. These differences indicate that demand for monthly period tickets in London is elastic in the long run. This is not surprising, as public transport ticket types represent very close substitutes. An individual can make the same public transport journey with almost any ticket type, so a change in the relative prices of these ticket types is expected to trigger a shift in demand towards the cheaper alternative.

Ticket type	Shares in last simulated week		Expected reduction in share
	Without price increase	With price increase	
Monthly Travelcard	7.1%	3.3%	54%
Monthly Bus Pass	2.5%	1.7%	32%
Price increase = (4.5-3.84) / 3.84 = 17%			
Expected reductions in shares are both greater than 17%			
Demand for Monthly Travelcards and Bus Passes is elastic			

Table 5-9: Changes in ticket shares resulting from increasing the monthly-to-weekly price ratio

Table 5-10 shows the revenue implications resulting from the policy change being studied and indicates a predicted net loss of £1.55 million in monthly Oyster fare revenues. This is due to a decrease in revenues from monthly period tickets (due to the increase in their price and to their elastic demand) that is larger than the corresponding increase in revenues from other period tickets.

Finally, the following should be noted about Table 5-10:

- The revenue estimates under 'before fare change' represent those realized by TfL in the last observed month in the sample (period 11 of 2007)[15].
- Revenues from period tickets were adjusted based on the ticket switching patterns shown in Figure 5-2. For PAYG, we also accounted for expected changes in use patterns due to switching from Monthly Travelcards or Bus

---

[15] Base case revenue estimates were obtained through correspondence with Pauline Matkins at TfL.

Passes to PAYG. To do so, we used observed changes in use from the sample of individuals who actually made the switch from monthly period tickets to PAYG. Changes in use due to ticket switching are expected and form the basis of our modeling methodology, in which ticket choice and public transport use are two interdependent decisions.

Revenue Source	Monthly Oyster revenues (millions)		
	Before Fare Change	After Fare Change	Net
PAYG	£66.78	£71.09	£4.31
Weekly Travelcards	£31.23	£34.55	£3.32
Monthly Travelcards	£21.08	£11.48	-£9.60
Annual Travelcards	£6.37	£6.88	£0.50
Weekly Bus Passes	£11.61	£12.21	£0.60
Monthly Bus Passes	£3.28	£2.59	-£0.68
Total	£140.35	£138.80	-£1.55

Table 5-10: Revenue implications of increasing the monthly-to-weekly price ratio

5.5.2 Increasing the Prices of Annual Travelcards

Increasing the prices of Annual Travelcards may be another way by which TfL could raise revenues, especially since users of Annual Travelcards are more likely to be high-income individuals who are less sensitive to price changes.

In order to test the effect of such a policy change, we simulate the ticket choice sub-model under a scenario in which the price of an Annual Travelcard is increased from 40 times the price of a Weekly Travelcard to 46 times that price. The simulation results are shown in Figure 5-3 and indicate a slow decrease in the share of Annual Travelcards over time compared to what it would have been with no policy change. The predicted shares of Annual Travelcards in May 2009 with and without the proposed policy change are 2% and 2.9%, respectively.

What stands out in Figure 5-3 is that most users who are predicted to switch away from Annual Travelcards are expected to use PAYG or Weekly Travelcards as an alternative. This is not intuitive, as one would have expected most of the increase to be in the share of Monthly Travelcards, since it may have represented the 'next best alternative' to those who had previously used an Annual Travelcard. The only justification we can provide for the switching patterns shown in Figure 5-3 are that (1) the model, given current ticket shares, produces more favorable predictions for

PAYG and (2) the dataset on which the ticket choice sub-model was estimated included no instances in which the annual-to-weekly price ratio was different from 40, possibly resulting in inaccurate predictions in terms of switching patterns across period tickets (but not in terms of reductions in the share of Annual Travelcards).
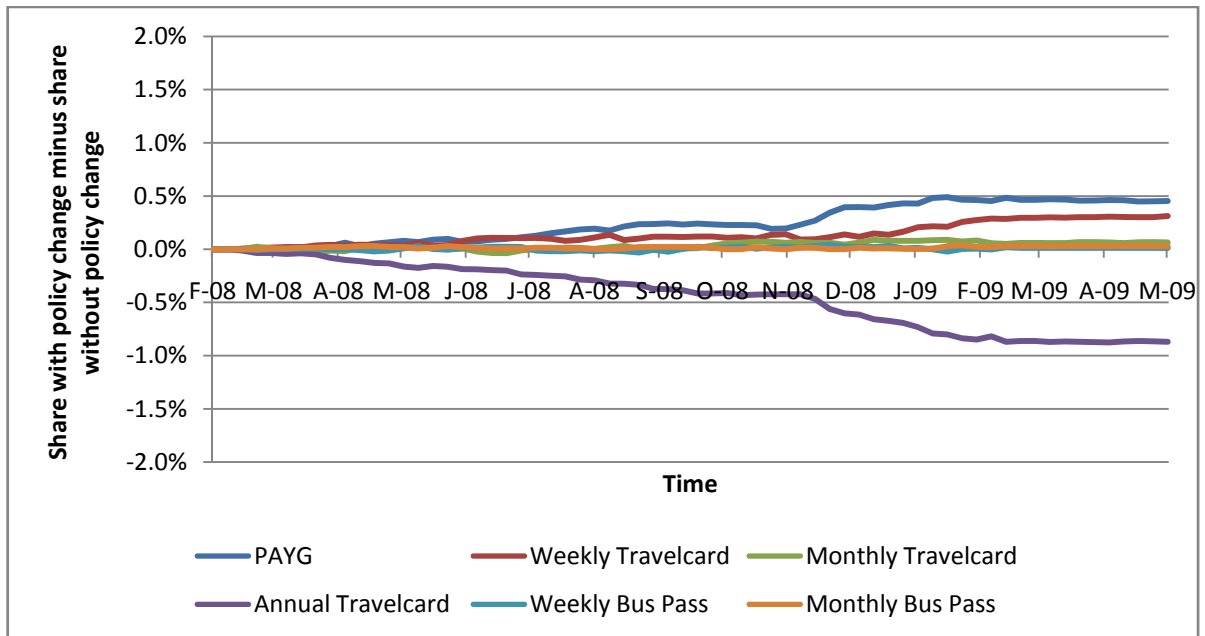
Figure 5-3: Simulation results for increasing the annual-to-weekly price ratio

Table 5-11 presents the predicted aggregate shares of Annual Travelcards in May 2009 with and without the price increase. The numbers, as is the case with the previous application, indicate a highly elastic demand for Annual Travelcards. Increasing the prices of these period tickets by 15% will result in a larger reduction in their aggregate share. We should note, however, that the observed price elasticity in this case (i.e. the ratio of the percentage change in shares to the percentage change in price) is about -2, which is smaller, in absolute terms, than the observed price elasticity of -3.18 for Monthly Travelcards indicated by the numbers in Table 5-9. This implies that, although the demand for both types of tickets is elastic, users of Annual Travelcards are less likely to switch to a different type than users of Monthly Travelcards.

93

Ticket type	Shares in last simulated week		Expected reduction in share
	Without price increase	With price increase	
Annual Travelcard	2.9%	2.0%	31%
Price increase = (46-40) / 40 = 15%			
Expected reduction in share is greater than 15%			
Demand for Annual Travelcards is elastic			

Table 5-11: Changes in ticket shares resulting from increasing the annual-to-weekly price ratio

Table 5-12 presents the revenue implications of increasing the annual-to-weekly price ratio from 40 to 46. The table shows a net loss of about £530,000 in monthly Oyster fare revenues. As in the previous subsection, the revenue estimates shown in the table account for changes in use for those who switch from Annual Travelcards to PAYG.

Revenue Source	Monthly Oyster revenues (millions)		
	Before Fare Change	After Fare Change	Net
PAYG	£66.78	£67.48	£0.69
Weekly Travelcards	£31.23	£31.93	£0.70
Monthly Travelcards	£21.08	£21.27	£0.19
Annual Travelcards	£6.37	£4.21	-£2.17
Weekly Bus Passes	£11.61	£11.63	£0.02
Monthly Bus Passes	£3.28	£3.31	£0.04
Total	£140.35	£139.83	-£0.53

Table 5-12: Revenue implications of increasing the annual-to-weekly price ratio

5.5.3 Increasing the Prices of All Travelcards

The third application we consider involves increasing the prices of all Travelcards by 10% (without changing the monthly-to-weekly or annual-to-weekly ratios). This application pertains to switching between PAYG and Travelcards, as opposed to the ones above, which modeled switching across all ticket types. We believe the results in this subsection may be more robust than the ones above, since the dataset on which the discrete-continuous model was estimated does include instances where prices of Travelcards have increased. (In the previous two applications, the predicted switching patterns may not be as accurate since, as mentioned earlier, the monthly-to-weekly and annual-to-weekly price ratios were constant throughout the observed time period.)

Again, we simulate the expected ticket shares with an increase in the prices of Weekly, Monthly, and Annual Travelcards by 10%. The results are shown in Figure 5-4 and Table 5-13. The figure shows a large increase in the predicted share of PAYG between February 2008 and February 2009. What is interesting about this change in the share of PAYG is that it appears to be the largest in magnitude in the first three to four months of the policy change, after which the rate of change decreases. This may be because most of the 'switchers' in the first few months initially used Weekly Travelcards, whereas in later time periods, the less price-sensitive Monthly and Annual Travelcard users switch to PAYG (again, time validities also play a role in this slow decrease). This is also evident in the changes in shares of the different Travelcards. The decrease in Weekly Travelcard shares seems to stabilize a few months after the price change, whereas the changes in Monthly and Annual Travelcard shares stabilize about one year after the price increase.

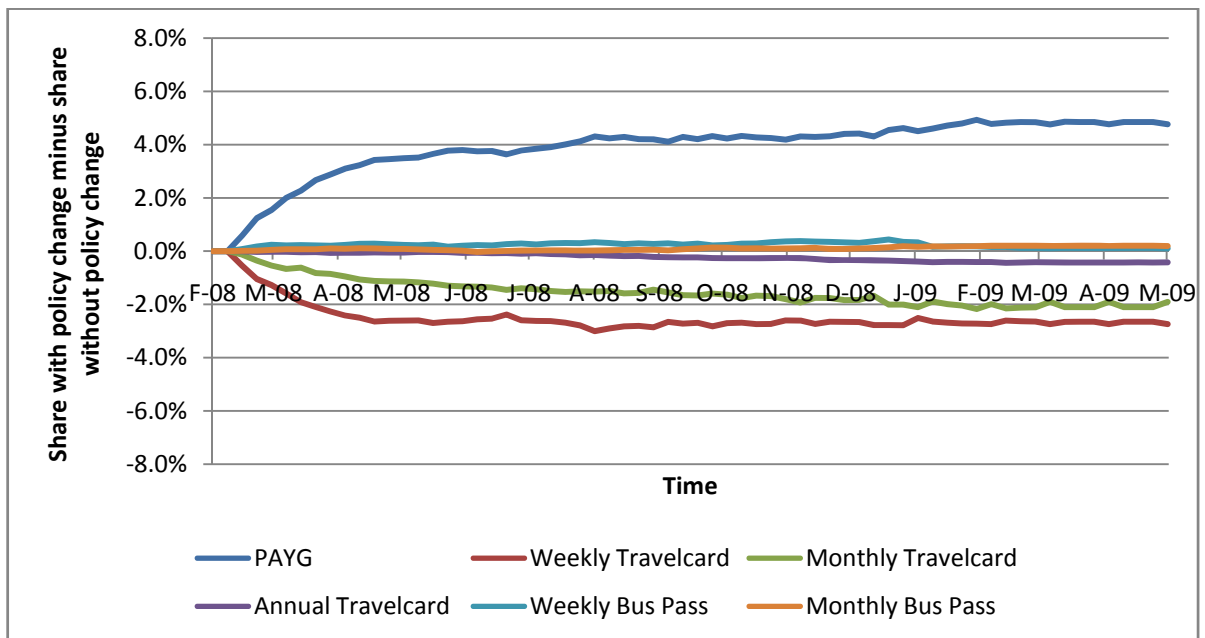**Figure 5-4:** Simulation results for increasing the prices of Travelcards

In Table 5-13, we see that the smallest percentage reduction in shares is in Annual Travelcards, followed by Weekly Travelcards and Monthly Travelcards. As mentioned earlier, Annual Travelcard users are expected to be less sensitive to price changes compared to others. This may explain the relatively low reduction in usage

95

among those users. On the other hand, the difference between the reductions in Weekly and Monthly Travelcards may be due to the high level of convenience of Weekly Travelcards perceived by users (this is also evident in the discrete choice sub-model results, shown in Table 5-4, where Weekly Travelcards had the highest alternative-specific constant). We should note, however, that despite these relative differences, the changes in shares for all Travelcards exceed 10%, which indicates an elastic demand for these Travelcards when PAYG fares remain constant.

Ticket type	Shares in last simulated week		Expected reduction in share
	Without price increase	With price increase	
Weekly Travelcard	13.7%	10.9%	20%
Monthly Travelcard	7.1%	5.2%	27%
Annual Travelcard	2.9%	2.5%	14%
Price increase = 10% Expected reductions in shares are all greater than 10% Demand for Travelcards is elastic			

Table 5-13: Changes in ticket shares resulting from increasing the prices of Travelcards

Table 5-14 presents the revenue implications from increasing the prices of all Travelcards by 10%. The table again shows a net reduction in revenues of about £5 million, indicating that the loss in revenues from Travelcards is not offset by the increase in revenues from PAYG (and Bus Passes). Despite the perceived convenience of using Travelcards, it seems that the increasing popularity of PAYG and policies such as daily price capping have caused this relatively large predicted switch from Travelcards to PAYG, which, in turn, is expected to lead to a net loss in monthly Oyster revenues. (Note that, as in the previous applications, the revenue estimates in Table 5-14 have been adjusted to account for changing use patterns among people who switch from the different Travelcards to PAYG.)

Revenue Source	Monthly Oyster revenues (millions)		
	Before Fare Change	After Fare Change	Net
PAYG	£66.78	£74.11	£7.32
Weekly Travelcards	£31.23	£24.97	-£6.26
Monthly Travelcards	£21.08	£15.43	-£5.65
Annual Travelcards	£6.37	£5.45	-£0.92
Weekly Bus Passes	£11.61	£11.82	£0.21
Monthly Bus Passes	£3.28	£3.53	£0.25
Total	£140.35	£135.31	-£5.05

Table 5-14: Revenue implications of increasing the prices of Travelcards

The three applications presented so far result in an expected loss in monthly Oyster revenues. This has significant policy implications for TfL, as it indicates that changing the relative prices of tickets (as opposed to increasing overall price levels) may not be a suitable strategy to raise revenues. This is because, as discussed earlier, the different ticket types are seen by public transport users in London as close substitutes.

Finally, the following should be noted about the above three applications:

1. We expect that many users who were predicted to switch to PAYG were initially making a financially irrational choice by using a period ticket, even before the simulated fare changes took effect. Given the complexity of the fare structure in London, many period ticket (as well as PAYG) users were found to be making an irrational choice when comparing what they had paid for their tickets to what they would have paid under PAYG (see Frumin, 2008). The simulated fare changes, which all involved increases in the prices of some period tickets, exacerbated the already irrational behavior of some users (i.e. they increased the difference between the expected cost of travel under a period ticket and the expected cost of travel under PAYG) causing them to switch to PAYG.

2. The simulations did not account for new Oyster cards entering the system or current Oyster cards dropping out of the system. We still believe, however, that the results in terms of predicted aggregate shares of the various ticket types are fairly accurate, especially since the relative differences among these Oyster ticket types have generally been stable over the past few months (after the disproportionate increase in the share of PAYG due to marketing efforts and the introduction of policies such as daily price capping).

3. The sample on which the simulations were performed included both the sample on which the discrete-continuous model was estimated and the validation sample (used in Table 5-5). The reason behind pooling the two samples was to obtain a large dataset on which to perform the simulations, given that the only observations included in those simulations were those from the last week in the panel (1/27 to 2/2/2008).

5.5.4 AM Peak Pricing on the London Underground

The fourth and final policy application we consider in this thesis is the implementation of AM peak pricing on the London Underground. The patterns in Chapter 3 showed sharp morning and afternoon peaks on the Underground. In other words, the highest levels of demand on a typical weekday are observed during relatively short time intervals. Such patterns suggest an opportunity for fare policy to play a role in spreading peak demand. By charging a premium for travel in the peak periods, one would expect some of the demand to shift to other modes (be they public transport or non-public transport modes) and to time periods preceding and following the one during which a premium is charged. Peak spreading can be an effective means for allowing more users into the public transport system, especially given the constraints on the infrastructure (in terms of the ability to run additional trains).

In order to implement an effective peak pricing strategy on the Underground, the following should be taken into consideration:

1. With the current fare structure, charging a premium on peak travel for PAYG customers without modifying the pricing structure for period tickets would probably not result in the desired change in ridership levels. This is because (1) most peak travel on both Underground and bus is made on period tickets, as shown in Chapter 3, and (2) implementing such a policy may induce some PAYG users to switch to period tickets in order to avoid either paying a premium or changing their travel schedules.

2. Queues may form at entry or exit gates (depending on how the peak premium is charged) where people would wait for the peak period to end before entering or exiting the station to avoid paying the premium. Addressing this issue requires careful consideration of the timing and structure of the peak pricing scheme, which may involve, for example, a gradual increase in fares before the peak period begins and a gradual decrease after the peak period ends, rather than a sudden change in fares at certain points in time.

3. Peak pricing should be implemented on a short time period during which demand is at its highest level. Applying a peak pricing strategy on a wide time

interval may induce many public transport users to switch to different modes rather than change their times of travel.

The example we provide here does not address the first two points above. We present expected ridership changes under a peak pricing strategy that is applied only on PAYG, without considering ticket switching or any changes to the fare structure. We assume a one-time, rather than a gradual, change in fares at the start and end of the peak period and assume that fares are charged based on exit (i.e. tap-out) times. The third point is addressed by applying peak pricing in the example below on a relatively short time period in which demand is at its highest levels. Furthermore, our example is applied only to Underground travel in the AM peak.

The methodology developed in this thesis theoretically allows for modeling tradeoffs among modes and times of day. This is done by adding the fare on a given mode and time of day as an explanatory variable in the regression sub-models measuring demand on other modes and times-of-day. This was not done in this thesis because (1) as mentioned earlier, given that we measure frequency of use by journey segments and do not account for interchanges, tradeoffs between modes are not clear, and (2) fares in the current fare structure do not really have a role in any time-of-day tradeoffs experienced by travelers. Using the model specified in this thesis to assess the impacts of a peak pricing strategy, therefore, requires making certain assumptions about peak period travel behavior on public transport in London. To that end, we do the following:

1. The PAYG sub-model for uncapped LU trips is re-estimated with the dependent variable being the number of (weekly) uncapped journeys (based on exit times) on the Underground made on weekdays between 8:30am and 9:15am, the time period with the highest demand in the AM peak. (Note that we also tried an alternative approach, in which *three* PAYG sub-models are estimated, one for each 15-minute time interval in the period between 8:30am and 9:15am. We opted for the single regression approach, because the current fare structure for the Underground does not include any time differentials in the peak period, so estimating fare elasticities at the 15-minute level does not produce any useful

results with regards to assessing people's sensitivity to fares within relatively short intervals.)

2. Using the results, we estimate an AM peak-period fare elasticity. This elasticity measures the effect of a one percent increase in fares between 8:30am and 9:15am on the frequency of use on the Underground during that time period.

3. A cross elasticity of bus use with respect to Underground fare obtained from the literature and is adjusted to represent the cross elasticity in the AM peak period.

4. Observed PAYG demand levels on the Underground are adjusted as follows:

    a. Ridership levels in each of the three 15-minute intervals in the time period during which a premium is charged (8:30-8:45am, 8:45-9:00am, and 9:00-9:15am) are reduced based on the estimated fare elasticity.

    b. The reduction in ridership in the second 15-minute interval (8:45-9:00am), which represents the peak-of-the-peak, is assumed to have shifted to a different mode (be it a public transport or a non-public transport mode). The assumption made in this step is that peak-of-the-peak users are not willing to shift their times of travel.

    c. Based on the AM peak-period cross elasticity of bus use with respect to Underground fares, we find the percentage (X) of the reductions in ridership in the first and third 15-minute peak intervals (8:30-8:45am and 9:00-9:15am) that shifted to bus. The remaining percentage ($100-X$) of the reductions is assumed to have shifted to a different time period. Passengers who initially traveled between 8:30am and 8:45am are assumed to have shifted to the 8:15-8:30am time interval, while those who initially traveled between 9:00am and 9:15am are assumed to have shifted to the 9:15-9:30am time interval. The assumptions made in this step are (1) the AM peak-period cross elasticity is assumed to be smaller, in absolute value, than the frequency of use elasticity estimated by the regression sub-model, (2) 'mode shifters' in the first and third 15-minute interval remain on public transport but switch to bus, and (3) 'time shifters' are only willing to shift one 15-minute interval from their original time of travel.

Table 5-15 presents the estimation results for the AM peak period regression sub-model (non-peak-period travelers are excluded from this analysis). Using the point-slope formula, we find a long-run AM peak-period fare elasticity of -0.28[16], which, as expected, has a smaller absolute value compared to the long-run fare elasticity of -0.64 estimated for all PAYG users and shown in Table 5-8. More precisely, the peak-period elasticity is about 44% the value of the overall elasticity. This relationship between the two elasticities is used to find a peak-period cross elasticity of bus demand with respect to Underground fares. Mitrani *et al.* (2002) estimated an overall cross elasticity of 0.13, which indicates that a 10% increase in Underground fares would result in a 1.3% increase in bus ridership. Multiplying this elasticity by 44%, we obtain a peak-period cross elasticity of 0.06.

Variable	Estimated coefficient (t-stat)
Constant	0.606 (20.20)
Past use	0.590 (153.95)
Holidays	-0.248 (-21.38)
SBC	1.294 (5.72)
Fare	-0.058 (-5.03)
Overall R^2	0.582
Number of Observations	49,712

Table 5-15: Estimation results for the frequency of use sub-model for uncapped Underground PAYG use between 8:30am and 9:15am

We should note that our estimated peak fare elasticity (-0.28) is similar to the fare elasticities obtained in a recent report (referenced in Chapter 2) looking into possible peak pricing schemes on National Rail (UK Department for Transport, Transport for London, and Network Rail, 2007). The report used survey data to estimate a time-of-day choice model and simulated travel, based on the results of that model, on a number of corridors. Several peak pricing scenarios were tested, and the observed fare elasticities on the various corridors ranged between -0.35 and -0.15.

We now test a peak pricing strategy in which a fare of £2.00 is charged for Underground travel made between 8:30am and 9:15am. This represents a £0.50 (or

---

[16] For the purpose of comparison, we also estimated a long-run off-peak fare elasticity equal to -0.78.

33%) increase from current fare levels. The results are shown in Figures 5-4 and 5-5. The first figure shows overall ridership levels before and after the peak pricing strategy is implemented, while the second figure shows the breakdown of the assumed switching patterns. The minus signs in the second figure indicate a shift away from the corresponding time intervals (or from the Underground), and a plus sign indicates a shift to the corresponding time interval.
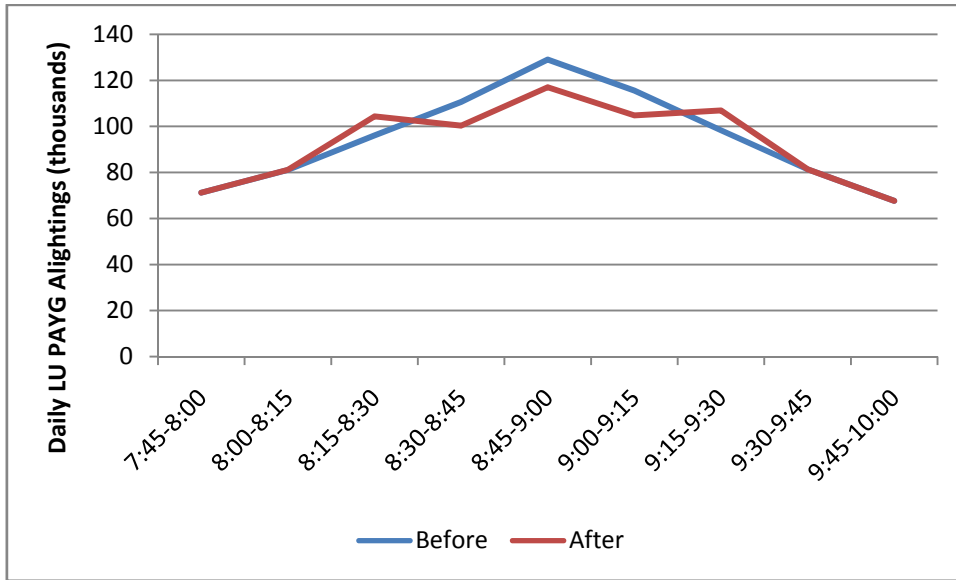

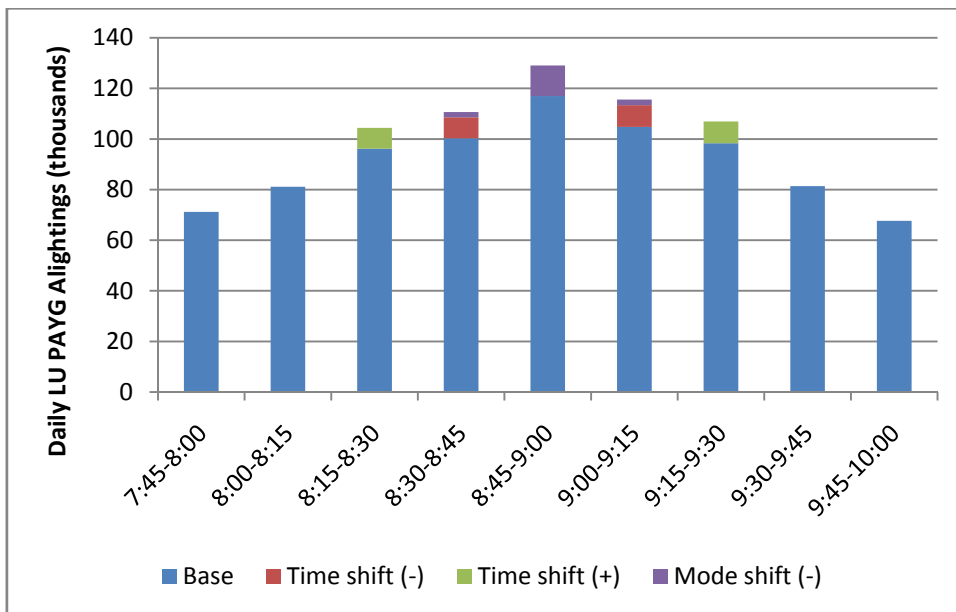**Figure 5-5:** Underground PAYG ridership in the AM peak before and after peak pricing

Figure 5-6: Underground PAYG ridership in the AM peak with peak pricing

The results above show that the suggested peak pricing strategy would reduce ridership between 8:30am and 9:15am by about 9%. (Note that this does not take into account any users who would start using the Underground during that time because of the reduction in crowding levels, for example.) Of the passengers who no longer travel in this time interval, 51% shift to a different time interval, and 49% shift to a different mode. The majority of those who switch to a different mode are presumed to have originally travelled in the peak-of-the-peak, represented in this example by the 8:45-9:00am time interval.

This section presented some policy applications for the discrete-continuous model developed earlier in this thesis and estimated on data from London. The final section of this chapter discusses ways in which the estimation results could be incorporated into models currently used at TfL.

5.6 Integrating the Results into Current Fare Models

This section explains how the methodology developed and applied in this thesis can be integrated into models currently used at TfL.

In the short run, the outputs of the model estimated above could be used as inputs to TfL's spreadsheet fare models described in Section 3.3. The spreadsheet models take ticket type and demand (i.e. usage) elasticities as inputs. Some of these elasticities can be readily estimated from the results of our model and input into the spreadsheet model. Table 5-16 shows the PAYG fare elasticities estimated by our model for each fare category currently included in TfL's spreadsheet models.

In addition to providing more robust elasticity estimates, our model also has the advantage of being estimated at the individual level. This allows for producing separate elasticity estimates, for both ticket choice and use, for any group of users (e.g. peak users, frequent users, weekend users, or users of the various fare categories, as shown in the above table). These estimates, in turn, can be used in the spreadsheet model to produce more accurate demand and revenue forecasts.

Fare category used in TfL Model	Estimated PAYG fare elasticity
LU: Zone 1 only	-0.37
LU: Zone 1-2	-0.40
LU: Zone 1-3	-0.48
LU: Zone 1-4	-0.43
LU: Zone 1-5	-0.60
LU: Zone 1-6	-0.65
LU: 1 zone (excluding Zone 1)	-0.23
LU: 2 zones (excluding Zone 1)	-0.20
LU: 3 zones (excluding Zone 1)	-0.32
LU: 4 zones (excluding Zone 1)	-0.35
LU: 5 zones (excluding Zone 1)	-0.38
Bus	-0.08

Table 5-16: PAYG fare elasticities for fare categories included in TfL's spreadsheet models

We should note, however, that not all outputs of our model are comparable to the inputs to TfL's spreadsheet models. In terms of ticket switching, for example, TfL's current models capture switching from cash fares to Oyster PAYG and other Oyster period tickets—something which is not explicitly accounted for in our model. The spreadsheet models, on the other hand, do not allow for switching among the different period tickets (e.g. switching from Weekly Travelcards to Monthly Travelcards) due to changes in their relative costs.

In the longer run, our methodology—including the ability to model ticket switching—can be further integrated into TfL's fare modeling capabilities by continuously making use of smartcard data as they become available. By developing automated computer procedures to preprocess the data, construct the panel structure, and calculate the expected costs of travel under each ticket type, TfL could continuously keep track of individuals' ticket choices, as well as their public transport use patterns. The discrete-continuous models specified above could be re-estimated on a regular basis in order to update the various elasticity estimates and validate the model's performance in terms of the differences between predicted and observed ticket shares and frequencies of use. This could all be done as part of a larger effort to make more effective use of most recent smartcard data. Developing a clear system in which the data are analyzed regularly has benefits in many areas beyond fare policy.

In summary, our methodology can be integrated in the short run into the current spreadsheet model used at TfL by (1) providing a more robust, policy-sensitive tool for calculating some of the exogenous inputs to the spreadsheet model and (2) enhancing the performance of that model by making it more disaggregate thus accommodating different subgroups or types of users. In the longer run, however, realizing the full potential of our methodology requires integrating it within a framework that makes use of smartcard data on a regular basis.

Chapter 6

Summary and Conclusions

This thesis proposed a methodology for modeling the impacts of fare changes on public transport demand. The methodology is based on using smartcard data over time, where individuals' past behavior can be used to infer their future behavior and to obtain a more accurate estimate of the true impact of fare changes on their ticket choices and public transport use patterns. We developed a discrete-continuous framework in which ticket choice is modeled at the higher (discrete) level and frequencies of use, based on mode and time-of-day, are modeled at the lower (continuous) level. We then specified an empirical model for London's public transport system and estimated that model using smartcard data that were collected between November 2005 and February 2008. The estimation results were mostly in accord with our expectations and provided insight into some fare policy options currently being considered by TfL.

In this chapter, we summarize our findings, first in terms of general methodological conclusions and then in terms of specific conclusions for London. This is followed by a discussion of the limitations of this research and some directions for future work.

6.1 Summary of Findings

6.1.1 Methodological Conclusions

The following methodological conclusions can be extracted from the theoretical framework and model estimation results contained in this thesis. These conclusions are general in nature and can be used as guidelines in developing models that are more complex than those presented in this thesis.

- Effective public transport demand management strategies require models that capture the increasing complexity of fare structures in public transport systems. Measuring public transport cost in one 'average fare' variable may no longer be adequate. Accounting for complex fare structures allows for having better predictions of people's choices and, hence, a more powerful forecasting tool.

- An important part of capturing the complexity of fare structures is the ability to model the choice among the various ticket types available on a public transport system. Although this has been done in the past, it did not account for the interdependence between ticket choice and public transport use, which is key in determining overall ridership levels. Changing ticket choices cause changes in public transport use patterns and vice versa.

- The availability of smartcard data over a relatively long period of time allows for developing models that use individuals' past behavior as an explanatory factor in determining their future behavior. Doing so accounts for 'inertia' effects that can explain the delayed or gradual impacts that certain fare changes have on public transport demand. Longitudinal smartcard data can also be used to capture unobserved factors, such as the changing awareness (or perceived attractiveness) of the various ticket types.

6.1.2 Conclusions for London

Given that we have applied our methodology using data from London, it may be useful to present some conclusions that pertain to the city's public transport system. The following conclusions are in addition to the ones listed in the previous subsection, which are also applicable to London:

- At many public transport agencies, analyzing available smartcard data is becoming useful in various planning contexts, including fare policy. In London, the high adoption rate of the Oyster card, as well as the availability of Oyster data over time, provides a unique opportunity to use this dataset in order to evaluate fare policy options and to understand how people use public transport.

- In evaluating any new fare policy, TfL must take into account the time horizon during which the effects of that fare policy are expected to take place. This time

horizon may vary due to (1) the different inertia effects among ticket types and (2) the different time validities for the tickets affected by the policy (i.e. the time before a ticket expires, allowing its holder to switch to a different ticket type). The policy applications shown in the previous chapter indicate that the impact of a change in the price of monthly period tickets will take about five to six months to fully materialize. On the other hand, an increase in the price of Annual Travelcards triggers behavioral changes that stabilize about one year after the price increase.

- PAYG has become increasingly popular in London, especially among first-time Oyster users. This is partly due to policies such as daily price capping and to marketing efforts. Observing smartcard data over time can account for this pattern. In addition to the changing cost of travel under PAYG due to fare changes and the introduction of different policies, observing people's behavior, especially among new Oyster users, can give some insight into factors that cannot be easily measured and that are contributing to the increasing popularity of PAYG. (In our analysis, we included a measure of prior aggregate shares of PAYG among first-time Oyster users in the ticket choice sub-model to account for the changing 'awareness' or 'attractiveness' of that ticket type.)

- The sharp peaks on the London Underground present an opportunity to implement a peak pricing scheme that would spread the demand during those time periods. The framework developed in this thesis allows for using smartcard data to gauge users' sensitivity to fare changes in the peak. This would help in determining the most effective peak pricing scheme to implement. One should keep in mind, however, that since peak pricing, and what it entails in terms of changes in the pricing structure for period tickets, represents a radical change in the fare structure in London. Using the available Oyster data, although it presents some useful insight, may not provide the full picture when evaluating such a policy. To that end, surveys and focus groups, such as those conducted for the purpose of studying peak pricing on National Rail (UK Department for Transport, Transport for London, and Network Rail, 2007), can be helpful in complementing the analysis.

6.2 Limitations and Directions for Future Research

To conclude this thesis, we discuss the limitations of our model and present, based on these limitations, some directions for future research in the area of public transport fare policy analysis.

Some of the limitations we present here pertain to the modeling framework we developed, while others are specific to the application to London presented in the previous chapter. Most of the limitations in the latter group have already been discussed in the context of the analysis. They include the following:

1. Given the current fare structure, the model does not fully capture time-of-day tradeoffs for public transport users in London. In order to test policies that involve such tradeoffs, certain assumptions have to be made, such as the ones used in the peak pricing analysis in Section 5.5. In the future, if such policies are implemented and new data become available, some of these assumptions could be relaxed.

2. The model, applied to London, does not capture tradeoffs between Underground and bus. As mentioned earlier, this is because our analysis was done at the journey segment level, without accounting for trips with interchanges, in which bus, for example, represents a complement, rather than a substitute, to the Underground. Capturing modal tradeoffs requires preprocessing the smartcard data to determine what groups of journey segments made by an individual represent a single journey (see Seaborn, 2008). After doing so, the model could be used to produce cross elasticities between modes.

3. The model uses a panel data structure in which one ticket type is assumed to have been used by an individual during a given week. As mentioned earlier, we have found that about 90% of Oyster cards in London were found to use exactly one ticket type per week. It is conceivable, however, that as PAYG gains popularity in the near future, more people may start using multiple ticket types on their cards (e.g. a Zone 1-2 Weekly Travelcard for journeys between and within Zones 1 and 2 and PAYG for all other journeys). If such a pattern develops, our definition of ticket choice in the model becomes less valid.

4. In the ticket choice sub-model, 'expected cost' is estimated using the observed use in the previous week. This is a very simplistic approach and may result in inaccurate estimates, especially in a week that includes holidays (and that is preceded by a week with no holidays). Addressing this can be done either by simply accounting for holidays when estimating expected cost (by scaling down the previous week's costs depending on the number of holidays in the current week) or by using a more advanced approach that develops latent variable models to estimate expected cost. (Expected cost for individual n in week t is an unobserved, or 'latent', variable.) Such models may include structural equations in which expected cost is a function of use in previous weeks, holidays, seasonality effects, and other variables. (For an overview of latent variable models and how they can be integrated with discrete choice models, see Bollen, 1989, and Ben-Akiva *et al.*, 2002.)

5. The ticket choice sub-model was applied to London assuming pure state dependence and not accounting for serial correlation. However, this does not mean that in the long run, the minimum cost ticket type will be chosen. This is because the model includes measures of the relative attractiveness of each alternative (represented by the alternative-specific constants). To account for both state dependence and serial correlation (i.e. unobserved heterogeneities), one would need to use more advanced statistical methods (see Heckman, 1981 and Keane, 1997).

The above limitations present some opportunities for future research that retain the proposed methodology but provide some adjustments that better capture London's public transport system and fare structure. This may involve revising the empirical specifications for London to better incorporate time-of-day tradeoffs when peak pricing is introduced, restructuring the model such that full journeys with interchanges are accounted for, redefining 'ticket choice' based on a new panel data structure, improving the measure of 'expected cost' in the ticket choice sub-model, and using more advanced statistical methods to account for both state dependence and unobserved heterogeneities.

We should also note that capturing individual-specific variations need not be done only through accounting for unobserved heterogeneities, as some of these variations may be captured with actual data. For example, many Oyster cards in London are currently registered. Using the postcode associated with each Oyster card can provide estimates for some socioeconomic variables, such as income (assuming the postcode within which the card is registered is for the home address). Using these data present another direction for further research.

We now turn to more general limitations that pertain to the methodology itself. One such limitation relates to the generation and suppression of smartcards. The current methodology does not account for this, as illustrated in the simulations in Section 5.5 in which we only used cards that appeared in the last observed week in the panel without accounting for any cards disappearing from the sample or for new cards appearing in the sample in the near future. One way in which this could be accounted for in future models is to develop a framework in which future smartcard generation and suppression are simulated over time based on past patterns.

Using such an approach, however, may still be insufficient. As shown earlier, Oyster cards in London have a relatively short lifespan, since many users either lose their cards or get a new card in order to switch to a different ticket (as some may be unaware of the ability to use multiple ticket types on the same card). This means that the observed 'generation and suppression' of smartcards is due not only to new users entering the system or existing users switching to non-public transport modes but also to existing users obtaining multiple cards over time. The methodology we developed, in addition to assuming that each card represents one individual (as noted earlier), also implicitly assumes that each individual is represented by only one card. Methodologically, if one wants to control for individual-specific effects, the fact that the same individual may use multiple cards needs to be somehow captured not only when simulating forecasts but also when specifying and estimating the model.

Furthermore, there are other methodological limitations of this research that relate to the continuous portion of the discrete-continuous model. Estimating

multiple regression equations for frequency of public transport use by mode and time of day introduces the following problems, which were summarized in a paper that modeled the demand for local telephone service options (Train, McFadden, & Ben-Akiva, 1987):

1. As the number of 'frequency of use' equations increases, so does the number of parameters, since each equation theoretically includes its own 'fare' variable, as well as those of all the other equations.

2. The dependent variables in the 'frequency of use' equations are truncated at zero (i.e. the take a value greater than or equal to zero). This may introduce what is known as 'truncation bias' in the estimated parameters (Amemiya, 1974). Correcting for this bias is complex, especially with a larger number of equations.

3. In our methodology, the 'frequency of use' equations are defined such that there is one measure of the per journey fare for each. This definition may be problematic in some cases, where it is not clear what that fare is. For example, in our application of the model to London, the fare variable in the PAYG uncapped Underground journeys regression did not represent the actual fare for each of these journeys; rather, it represented the per journey fare for the LU fare category on which the individual most frequently traveled during a given week. Policies such as price capping also do not fit well with our definition for the frequency of use equations. In our application, we used an ad-hoc approach to define what journeys would have been capped and then modeled capped journeys in two separate regression equations (capped Underground journeys and capped bus journeys). In reality, however, such journeys should be modeled as part of the total number of Underground and bus journeys, since they do not really represent a different 'class' of use that individuals are choosing.

Developing a fully discrete model, in which mode, ticket, and time of day are modeled as discrete choices, would address the above limitations. However, as mentioned in the beginning of this thesis, such an approach would require knowing the non-chosen alternatives for each journey and the attributes of those alternatives. A fully discrete model, therefore, would require significant data preprocessing and some analysis at the origin-destination level.

In summary, the proposed model structure, in its current form, can be improved to better capture factors specific to London that were not addressed in this thesis. More generally, however, there are certain limitations to the model structure itself that may call for the use of more advanced models for ticket, mode, and time-of-day choices. The tradeoff between such models and the one proposed in this thesis is between simplicity, in terms of the model structure and data requirements and preprocessing, and the level of disaggregation. Our model requires a dataset that is aggregated to the individual-week level, while a fully discrete set of models would capture the impacts of fare changes at the journey level.

Given the limitations discussed above, we conclude this thesis by highlighting some key points on how our methodology could be applied at TfL, as well as other public transport agencies:

- The methodology developed in this thesis could be readily used to model public transport ticket choice and frequency of use. In order to do so, the available longitudinal smartcard data must be preprocessed to produce the panel data structure presented earlier. Given the available ticket types, as well as knowledge of the modes and times of day on which travel is permitted for each ticket type, the model can produce estimates for the effects of inertia and cost on ticket choice, as well as the effects of past use, holidays, and the per journey fare (where applicable) on frequency of use based on that ticket choice.

- When specifying the frequency of use sub-models, each regression equation should be defined such that there is one constant fare value for its corresponding journeys. So, if rail and bus fares are different for a given ticket type, then rail use and bus use should be modeled in separate equations.

- By simulating the ticket choice sub-model over time, one could get accurate measures for the effect of a fare change on ticket shares and the time horizon during which that effect would fully materialize. The frequency of use sub-models could also be simulated over time when evaluating changes in per journey fares.

- Given that the ticket choice sub-model is discrete, the model produces satisfactory results that capture switching among different ticket types. On the other hand, the frequency of use sub-models, being continuous, can capture tradeoffs between modes and times of day only when certain assumptions and conditions are met:
  – To capture tradeoffs between modes: The raw data should be preprocessed to account for interchanges. Doing so allows for producing measures for tradeoffs among the various modes. (At the journey segment level, these tradeoffs are not clear, since some modes can be either complements or substitutes to others.)
  – To capture tradeoffs between times of day: The fare structure should include tickets or fare classes whose costs vary by time of day and where such variations present a viable time-of-day choice that is made by users of the system. The fare structure in London, for example, includes an Oyster PAYG fare for travel between 7am to 7pm on weekdays and a different Oyster PAYG fare for travel at all other times. Although these fares do vary by time, they are both applied to relatively long time periods, such that individuals are not really making a time-of-day choice based on the difference between these two fare levels.

(Note that in the AM peak pricing application presented earlier, we made some assumption on mode and time switching to address the above issues.)

Finally, as discussed in Section 5.6, the results of our application of the model to London could also be used and integrated into TfL's current fare models. The estimated frequency of use elasticities, for example, could be incorporated into TfL's spreadsheet models, while keeping in mind that our estimated bus fare elasticity is significantly smaller than previously estimated values, given the small variations in bus fares in the dataset. Measures that relate to ticket switching, on the other hand, may require a longer-term integration strategy to incorporate into TfL's fare modeling capabilities, given the new definition we propose for 'ticket choice' and the use of inertia as a factor that explains that choice.

Bibliography

Amemiya, T. (1974). Multivariate Regression and Simultaneous Equation Models when the Dependent Variables are Truncated Normal. *Econometrica , 42* (6).

Ben-Akiva, M., & Lerman, S. (1985). *Discrete Choice Analysis.* Cambridge, MA: MIT Press.

Ben-Akiva, M., Walker, J., Bernardino, A. T., Gopinath, D. A., Morikawa, T., & Polydoropoulou, A. (2002). Integration of Choice and Latent Variable Models. In H. S. Mahmassani, *In Perpetual Motion: Travel Behavior Research Opportunities and Application Challenges.* Elsevier Science.

Bhat, C. R. (1995). A Heteroscedastic Extreme Value Model of Intercity Travel Mode Choice. *Transportation Research B , 29* (6).

Bollen, K. A. (1989). *Structural Equations with Latent Variables.* New York, NY: John Wiley and Sons.

Colin Buchanan. (2006). *TfL Fares Study.* London, UK: London Assembly Budget Committee.

Dittmar, H. (1983). Profile of Monthly Pass Users in the San Francisco Bay Area. *Transportation Research Record , 877.*

Fleishman, D., Schweiger, C., Lott, D., & Pierlott, G. (1998). *TCRP Report 32: Multipurpose Transit Payment Media.* Washington, DC: National Academy Press.

Fleishman, D., Shaw, N., Joshi, A., Freeze, R., & Oram, R. (1996). *TCRP Report 10: Fare Policies, Structures, and Technologies.* Washington, DC: National Academy Press.

Frumin, M. (2008). *The Choice between Period Tickets and Pay As You Go on London's Public Transport System: Implications for Efficiency, Revenue, and Best Value.* MIT. Submitted to Transport for London.

Hausman, J. (1985). The Econometrics of Nonlinear Budget Sets. *Econometrica , 53* (6).

Heckman, J. J. (1981). Heterogeneity and State Dependence. In S. Rosen, *Studies in Labor Markets.* Chicago, IL: University of Chicago Press.

Heckman, J. J. (1981). Statistical Models for Discrete Panel Data. In C. F. Manski, & D. L. McFadden, *Structural Analysis of Discrete Data and Econometric Applications.* Cambridge, MA: MIT Press.

Hensher, D. A. (1998). Establishing a Fare Elasticity Regime for Urban Passenger Transport. *Journal of Transport Economics and Policy , 32* (2).

Hensher, D. A., & King, J. (1998). Establishing Fare Elasticity Regimes for Urban Passenger Transport: Time-Based Fares for Concession and Non-Concession Markets Segmented by Trip Length. *Journal of Transportation and Statistics* .

Hensher, D. A., & Mannering, F. (1987). Discrete/Continuous Econometric Models and their Application to Transport Analysis. *Transport Reviews , 7* (3).

Hensher, D. A., & Milthorpe, F. (1987). Selectivity Correction in Discrete-Continuous Choice Analysis. *Regional Science and Urban Economics , 17* (1).

Hong, Y. (2006). *Transition to Smart Card Technology: How Transit Operators Can Encourage the Take-Up of Smart Card Technology.* Thesis, Master of Science in Transportation Program, MIT.

Horowitz, J. L. (1993). Semiparametric Estimation of a Work-Trip Mode Choice Model. *Journal of Econometrics , 58*.

Keane, M. P. (1997). Modeling Heterogeneity and State Dependence in Consumer Choice Behavior. *Journal of Business & Economic Statistics , 15* (3).

Litman, T. (2004). Transit Price Elasticities and Cross-Elasticities. *Journal of Public Transportation , 7* (2).

Mayor of London. (2006). *Mayor's Transport Strategy.* London, UK.

Mitrani, A., Kincaid, I., Edwards, D., & Hobbs, G. (2002). *London Underground and Bus Demand Analysis: 1970-2000.*

Oum, T. H., Waters, S., & Yong, J.-S. (1992). Concepts of Price Elasticities of Transport Demand and Recent Empirical Estimates. *Journal of Transport Economics* .

Paaswell, R. E., & Stuart, D. G. (1993). Evolution of Fare Policy: A Product of Modern Transit Management. *Transportation Research Record , 1395*.

Page, E. (1981). Factors that Influence Choice Among Transit Payment Methods: A Study of Pass Use in Sacramento. *Transportation Research Record , 799*.

Pham, L., & Linsalata, J. (1991). *Effects of Fare Changes on Bus Ridership.* Washington, DC: American Public Transit Association.

Seaborn, C. (2008). *Application of Smart Card Fare Payment Data to Bus Network Planning in London, UK.* Thesis, Master of Science in Transportation Program, MIT.

Taplin, J. H., Hensher, D. A., & Smith, B. (1999). Preserving the Symmetry of Estimated Commuter Travel Elasticities. *Transportation Research Part B , 33.*

Tegner, G., Loncar-Lucassi, V., Nilsson, C., & Holmberg, I. (1998). The Demand for Public Transport Trips in Stockholm County: A Two-Stage Aggregate, Non-Linear Time-Series Model. *Economics and Institutions of Transport.* Borlange, Sweden.

Train, K. E., McFadden, D. L., & Ben-Akiva, M. (1987). The Demand for Local Telephone Service: A Fully Discrete Model of Residential Calling Patterns and Service Choices. *RAND Journal of Economics , 18* (1).

Train, K. (1993). *Qualitative Choice Analysis: Theory, Econometrics, and an Application to Automobile Demand.* Cambridge, MA: MIT Press.

Transport for London. (2007). *London Travel Report 2007.*

Transport for London. (2008). *Your Guide to Fares and Tickets.* London, UK.

TRL. (2004). *The Demand for Public Transport: A Practical Guide (TRL Report 593).* London, UK: TRL.

UK Department for Transport, Transport for London, and Network Rail. (2007). *Demand Management Techniques - Peak Spreading: Final Report.*