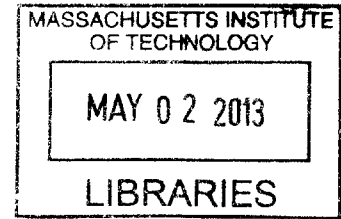Assessing Journey Time Impacts of Disruptions
on London's Piccadilly Line

by

Yonah Freemark
B.A. in Architecture
Yale University
New Haven, CT (2008)

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

February 2013

Author: _____
Department of Civil and Environmental Engineering
January 30, 2013

Certified by: _____
Nigel H.M. Wilson
Professor of Civil and Environmental Engineering
Thesis Supervisor

Certified by: _____
Harilaos N. Koutsopoulos
Visiting Professor of Civil and Environmental Engineering
Thesis Supervisor

Accepted by: _____
Heidi M. Nepf
Chair, Departmental Committee for Graduate Students

Assessing Journey Time Impacts of Disruptions on London's Piccadilly Line

by

Yonah Freemark

Submitted to the Department of Civil and Environmental Engineering
on January 30, 2013 in partial fulfillment of the requirements for the degree of
Master of Science in Transportation

Abstract

Public transport users depend on a reliable level of service on a daily basis. But system disruptions, caused by infrastructure problems, passenger events, and crew duty constraints, can result in reduced reliability for users. Understanding the impacts of those disruptions on customers is vital to evaluating the performance of the system and appropriately communicating delays to passengers.

The goal of this thesis is to investigate the impact of certain disruptions on passenger journey times using several new metrics. The thesis has three primary components: First, a description and categorization of incidents that occur on a urban rail transport line over the course of 29 days with some degree of disruption; second, the development of a new measure of impacts on passengers resulting from those incidents using automated fare collection (AFC) data; and third, an exploration of the potential use of AFC data in real-time applications to monitor service.

The proposed approach is applied to the Piccadilly Line, one of the London Underground's major rail lines. The line suffers from instances of significant disruption caused by aging technology and infrastructure, but it will not be upgraded for more than a decade. Therefore, insights from existing automated data sources, such as AFC, could play an important role in improving service without capital-intensive improvements.

The passenger impact analysis method developed in the thesis relies on dividing the line into sections and aggregating all AFC transactions on all origin-destination (OD) pairs within each section. The resulting disruption impact index summarizes the effects of a disruption on the average passenger for each section of the line. In addition, the accumulation of passengers on a line is introduced as an indicator of delays relating to a disruption. These metrics are each compared with information provided by train-tracking information systems.

The methods developed in the thesis were compared with actual passenger notifications on several study days. The results indicate that, despite the methods' limitations, there is potential for using AFC data, along with operational data, to provide more accurate and timely information to the users of the line. The application also leads to recommendations for how the method described for disruptions analysis could be extended to other types of analysis.

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

Thesis supervisor: Harilaos N. Koutsopoulos
Title: Visiting Professor of Civil and Environmental Engineering

Acknowledgements

Table of Contents

List of Figures

List of Tables

1

Introduction

This thesis describes research conducted on evaluating passenger response to disruptions and cus-tomer-facing communications on the London Underground Piccadilly Line. The aim of the research is to show how understanding incidents from a variety of perspectives offers potential for better under-standing disruption impacts and improving information provided to customers in the short term and is a first step toward establishing a framework for increasing reliability of transit service in the long term. In particular, it focuses on using data sources currently available—both automatically and manually generated—in order to show what insights are, and potentially can be gained from each type of data. It demonstrates several measures that can be used by transit agencies to better understand disruptions occurring on their lines and recommends a path forward for further research on the subject.

1.1 MOTIVATION

Disruptions frustrate riders, leading some to choose other modes of travel. They increase operating costs for transit agencies, which have to plan for unusual situations. Improving the quality of service on an existing transit line has the potential to bring significant benefits both to passengers—who desire more reliable service and more accurate real-time information about conditions—and to transit system managers, who want to minimize the length and cost of disruptions. Improvements may come in the

form of a significant renovation of the line, but capital projects are very expensive, take years to implement, and make assumptions about post-expenditure improvements that may not be realized. On the other hand, the effective use of automated data already collected by most large transit agencies presents a far less expensive and easier-to-implement model for understanding disruptions and better informing customers about them.

By their very nature, train service incidents are unpredictable and difficult to respond to. They cause passenger delays and complicate crew assignments, but they are hard to plan for. Yet analyzing them can be quite beneficial. By reviewing the types of incidents that occur, a transit system's vulnerabilities can be identified. If certain problems occur repeatedly, a change in operations strategy may be necessary. If certain responses to incidents are more successful than others, these can be identified as best practices. Systems currently used to keep track of trains and to charge passengers the correct fare also can be used to assess reliability of service during disruptions. If transit agencies become more expert in using the data they are already generating, they have the opportunity to communicate more effectively with both service managers and passengers, enabling both to make better decisions as a result. The question is how assessments of automated data can be made, and what information about disruptions can be derived.

Incidents pose major challenges for rail transit systems. Part of the explanation is that they simply represent the inevitable, occasional failures of the transportation system. But it is also true that there are sources of information already available that could aid in the response to problems on the line, but currently are not readily accessible to service controllers. Within the transit agency itself, differences in approaches between departments can limit the degree to which data is shared, thus true data integration rarely exists. Additionally, transit systems often place emphasis on average service quality, rather than the impacts of specific incidents, limiting the lessons that can be learned from case-by-case analysis. In some cases, riders are sometimes provided incorrect information about service status. This research identifies what kinds of incidents occur and examines how automated data collection offers an opportunity for improved response and passenger information.

This thesis suggests that proper analysis of automated farecard data can offer significant insight into the effects of disruptions on passenger service, for examples, by providing information about changes in journey times experienced by riders. Using aggregate metrics of disruption impacts on passengers, the overall effect on passengers of a disruption can be measured and compared with estimates derived from analysis of train-tracking data. This can broaden our understanding of which passengers were impacted by delays, in a retrospective analysis of data.

In addition, by offering insight into what kinds of conclusions can be made by evaluating real time data about disruptions, this research points to potential improvements in passenger communications. Rather than simply point to successes and failures in disruption response in retrospect, as is currently done to review the transit network's difficulties, it would be advantageous to develop techniques to understand disruptions as they occur. The advent of widely available real-time communication tools, from Twitter to electronic signboards at stations to email alerts, only raises the importance of ensuring the accuracy of whatever information is provided to customers in the event of a disruption.

Previous research has described the decision environment in which service managers determine how to alter operations in unexpected circumstances (Carrel 2009). Automated data collection offers a potential avenue towards significantly improving the information managers rely upon to make such decisions, and advanced techniques can be used to measure reliability of service (Frumin 2010; Hickey 2011; Uniman 2009). Efforts have been made to combine passenger-generated and train-generated data sources for the purpose of "assigning" passengers to trains (Paul 2010). However, these efforts have not examined how different sources of data provide alternative perspectives on the environment in the transit system before, during, and following a disruption.

This research is a step towards increasing the value of the automated data already available within the transit system. By offering a description of the service environment before, during, and after incidents, it offers insight into the types of information that might be easily derived from automated data but which would be of value to transit users and service managers. It integrates available data sources with the purpose of understanding impacts on passenger travel during disruptions.

1.2 OBJECTIVES

This research has three principal objectives, designed specifically to address concerns on the London Underground Piccadilly Line but in many cases these are broadly applicable to many metro rail systems.

1. **To categorize incidents on the Piccadilly Line**. While there are existing tools used to make such classifications, they are flawed from two perspectives. One, they are done retrospectively, not based on what service controllers know when they are making decisions. Controllers may have incorrect information about an incident that nonetheless is what they must use to respond to problems on the line. Two, they estimate the severity of an incident based on estimates of passenger travel, despite the fact that information about actual travel patterns using smartcard technology is available. An improved incident categorization system should not only reflect train movements, as does the current system, but also incorporate information about actual passenger movements where possible. This objective recognizes that actions controllers take in response to incidents affect outcomes.

2. **To integrate data sources available for analysis of disruptions on the system**. Though the Piccadilly Line has a plethora of data currently available, full integration of varying types of data has not yet occurred. Using archived information about performance on specific days, this analysis offers a reconstruction of operations and show how different sources of data provide different sorts of information, with varying degrees of accuracy. This analysis focuses on the use of automated fare collection (AFC) data to analyze the impact of individual disruptions on the journey times of riders and the accumulation of passengers in the system. In addition, analyzing the AFC data specifically on disrupted days may offer insight into customer route patterns, adding to the survey data currently used. If this data is currently not available to all interested staff of the transit system, the research indicates how the information should be passed on and who should have access to it and use it.

3. **To use data integration to provide customers better information about the current service status on the line.** The research shows how real-time information from a variety of data sources can be used to help guide what information is communicated to customers, and how different sources of data offer varying insight about conditions. The analysis demonstrates how variability in journey time and accumulation of passengers may be used as metrics for evaluating the onset and severity of incidents on the line. As such, the research concludes that using AFC data in a real-time context offers significant promise for monitoring the system state.

1.3 RESEARCH APPROACH

Though the conclusions of this research are broadly applicable to many transit systems due to the similarities in data available across many systems, this research focuses on the London Underground and specifically the Piccadilly Line. This case study provides a particularly informative example as it is an aging transit line that will not undergo significant upgrades for a decade or more. The central question is whether taking advantage of automatic data more effectively can produce improvements to service in the short term without the benefit of major line upgrades.

The research was conducted in three modes: observation of existing conditions, review of data, and analysis. Eleven weeks of on-the-ground presence in London provided the opportunity to conduct interviews, attend meetings, and watch a wide variety of personnel in the transit system conducting their daily activity. This qualitative review provided the chance to see how decisions are currently made when disruptions occur and question controllers on the best methods currently used. In addition, this review offered a view of how the transit agency interacts with customers during disruptions.

Once certain observations were made on the conditions on the line, a review of data available on the service was conducted. The data include that supplied by the smartcards with which passengers tap in and out to pay fares on the transit system, and the location of trains in the system at all times of the day. In addition, non-automated sources of data were compiled to provide a view of how operational response currently occurs. These latter data include reports written by service controllers about the decisions they made in response to incidents and processed information on the causes and effects

of disruptions compiled post-event by management.

The bulk of the research is the analysis that attempts to integrate the sources of data described above, keeping in mind the operational constraints in the system, as established in the observation period. The analysis was constrained to a specific period (three months in spring 2012). Within that period, a sample of days with "good" performance and a selection of days with minor and major incidents were selected. For each of these days, the automated and non-automated data were compiled to examine retrospectively how circumstances occurred. First, the incidents that occurred were categorized based on several criteria, and those categorizations were compared with the incident analysis methods currently used by the Underground (Objective 1).

Then, using a new approach, the data were analyzed to determine the changes in performance of groups of passengers traveling between various origin-destination pairs. The comparison of the two automated data sources with the non-automated sources provided insight into the different conclusions that might be made by reviewing data independently and indicated how two sources of data could be integrated to offer a more informed view of what is occurring in a disrupted situation. In addition, a method was developed to estimate the accumulation of passengers in the system to demonstrate that disrupted conditions result in more passengers on the line than under normal circumstances. The disrupted days information was used to evaluate existing assumptions about patterns of travel between specific origin-destination pairs (Objective 2).

Finally, the analysis method was used to describe how AFC data might be used in a real-time environment. Journey time variation and passenger accumulation metrics were evaluated as potential sources of information about service disruptions. These were then compared with train-tracking data, in the interest of demonstrating the comparative advantages of the different sorts of data (Objective 3).

A similar approach could be undertaken by other transit systems to understand how they currently respond to disruptions and what kinds of improvements they could make by using their data more effectively.

1.4 THESIS ORGANIZATION

- **Chapter 2** comprises a review of the literature and demonstrates what gaps the research aims at filling. In addition, this chapter offers an introduction to the issues of service control and management in transit systems in general.

- **Chapter 3** introduces the London Underground and the Piccadilly Line, describing the agency's objectives, decision factors, and constraints. This section provides a detailed description of the data sources used in the research and the manner in which information flows between departments of Transport for London.

- **Chapter 4** describes the case study period used for the majority of the research. It then categorizes incidents during that period, in terms of location, duration, type, impact, and frequency, among other concerns. Finally, it offers a comparison between the findings here and the disruptions information provided by an existing automated disruption evaluation tool.

- **Chapter 5** introduces the approach used in the research to integrate large amounts of automatically generated Oyster data. It defines a metric and baskets of origin-destination pairs to measure variability in journey times in order to better assess overall trends in passenger movement. It also shows how passengers accumulate over the course of a disrupted period. The chapter examines a set of disrupted days to show how these data compare with both manual and train-tracking data and illustrates how retrospective comparisons can be made. In addition, it documents how Oyster data from disrupted days provide the potential to enhance understanding of customer path choice, currently based only on passenger surveys.

- **Chapter 6** shows how the approaches in Chapter 5 can be used on a real-time basis to better inform customers of the disruptions occurring in the system. It provides an analysis of the comparative benefits of different sources of data available in real-time. It also offers potential guidelines for providing customer information, laying out when and how passengers could be notified of problems occurring on the line based on inputs from farecard transaction data.

- **Chapter 7** summarizes the findings of the research and notes the limitations of the methods described. It provides a set of recommendations to Transport for London and London Underground based on the research results. It also suggests potential avenues for future research.

2 Literature review

Disruptions are an important matter of concern to urban rail system managers around the globe because they affect reliability. A number of studies have assessed the occurrence of disruptions and provided analytical methods to examine disruption impacts. This chapter examines the literature related to incidents and their impacts and describes how the literature informs the research conducted for the rest of the thesis. In Section 2.1, studies related to customer response to disruptions on rail transit systems are reviewed. The studies show that while customers respond negatively to delays, they benefit from comprehensive passenger notifications and information. In Section 2.2, research on the use of automatic fare collection (AFC) smartcard data is reviewed to illustrate what information can be derived from AFC data and what techniques are used to do so.

Section 2.3 focuses on studies conducted on the reliability of rail systems. In particular, the review highlights metrics that have been developed to identify the overall reliability of transit lines using AFC and train tracking data. In addition, it provides a review of how such reliability statistics could be used for passenger information. Section 2.4 continues this discussion with a brief look at studies on controller approaches to disruption response. Though controller response is not a focus of this thesis, this section provides context for the work of service controllers in responding to incidents with changes in train and crew assignment. Finally, Section 2.5 suggests how this thesis extends the literature.

2.1 CUSTOMER RESPONSE TO DISRUPTIONS

Public transport, like most transportation systems, suffers from occasional incidents that impair the ability of customers to complete their journeys in the conditions and time indicated by the timetable. The resulting inconvenience may result in reduced ridership, which is contrary to the broad public policy goal of encouraging greater use of public transport. This section reviews several studies that discuss the impact of disruptions on the travelling public, and addresses the question of customer perceptions of quality specifically for Transport for London's passengers.

Andreassen (1995) examined the performance of public transport systems and demonstrates that repeated delays result in a decrease in demand. Bernal et al (2009) undertook a similar study on customers on Chicago's Blue Line rail service to document the impact of delays on riders. Using a comparison between ridership and "slow zones" traveled by trains in the system, the study demonstrates that increasing delays are strongly correlated with reduced ridership. The authors argue that this shows that as customers experienced increasing delays, they decreased their use of the rail system. That said, the slowdowns documented by Bernal are somewhat different from those described by Andreassen, since the former analyzes "planned" delays (regular customers knew about the "slow zones"), whereas the latter considers unplanned delays. Nonetheless, in both situations, disruptions clearly have negative effects on passenger use of the system.

Fonseca et al (2010) argue that for passengers, the concept of "quality" is correlated to customer satisfaction in terms of the services provided by public transport. They claim that the perceived quality of public transport services is defined by relatively objective measures such as reliability, security, speed, comfort, and punctuality. This indicates that the best way to encourage use of a transit service is to improve performance on these metrics.

Transport for London (TfL), which is the focus of this thesis, undertook a major study to determine how its riders assessed the reliability of its services, and what they perceived the primary components of reliability to be (TfL, 2011a). The study quotes customers who shared their views about the definition of reliability, such as:

"A reliable someone/something is dependable and trustworthy."

"Reliable is... something that's dependable, trustworthy, honest, reassuring, and something you can count on."

The study notes that customers value the predictability and dependability of services they receive. They feel more confident when they are provided additional information about their journeys, but they are not satisfied by line or system performance averages. Rather, customers say that they want personalized metrics that reflect specific information about the journeys they are likely to take. In the case of disrupted situations, this means that they want information about how a disruption will affect the specific services they are likely to use, not just any line on the system.

These studies imply that passengers are significantly negatively impacted by disruptions and that they are less likely to use transit systems if service quality declines. This is particularly true when there are either short- or long-term disruptions affecting operations. In addition, there are customer-friendly ways in which to improve the perception of a "quality" service. Customers are interested in speedier, more reliable transit; if the transit system can show that it has improved on these metrics, it is likely to improve rider satisfaction.

The ability of a system to respond to disruptions in a manner that is beneficial to customers, however, depends on the management of the line, and the degree to which incidents affect overall service. Valdés-Diaz et al (2005) conduct a review of disruptions to compare the causes and effects of incidents on the train service. They determine relationships between time between disruptions; disruption durations; and delay/headway ratio. They also document that disruptions affecting one section of the system (even just one line) may propogate throughout the entire system.

Valdés-Diaz et al's analysis approach is largely controller-dominated, with little consideration of the passenger impacts of disruptions on a rapid transit line. Jiang et al (2012), however, offer more insight into how a disruption might affect passenger flow by developing a simulation model to estimate the impact of train delays on passenger movements on a rail system. Their study shows that some short train delays may actually *not* result in any passenger delay and may improve conditions for cer-

tain passengers. This is because almost all passengers on an urban rail system take the first train that arrives, not the scheduled train, so a service with trains and crew out of place may not affect passengers (though it will of course affect crew management). That said, longer train delays may result in significant passenger delays and knock-on effects over the course of the day. In addition, disruptions affecting riders on one line are likely to spread to other parts of the system.

The conclusions produced by Jiang et al provide an interesting view into the customer impact of disruptions, rather than simply the effects on train service, which are those that are typically measured in the literature. The Jiang et al study, in particular, provides an important demonstration of how train service delays may not directly reflect passenger service delays. This has significant implications for studying the impact of disruptions.

2.2 USES OF SMARTCARD DATA FOR DATA ANALYSIS

Since the 1990s, public transport networks around the world have widely adopted automatic fare collection (AFC) technologies to improve their passenger revenue systems. AFC smartcards, such as the Oyster card in London, allow customers to store money or travel passes on a single card, which is then tapped on fare gates at the entrance and (sometimes) exits of transit stations (or on transit vehicles, in the case of buses). AFC technology facilitates transit systems charging fares based on travel distance or time of entry, as is the case in London. Smartcards have the added benefit of offering transit analysts more data than is possible with predecessor fare technologies, such as magnetic stripe paper cards, particularly when customers are required to "tap" out of the system (as is the case on the London Underground). This section provides an overview of several studies that describe how AFC data may be used to analyze performance on a transit system and document the movements of riders.

Bryan and Blythe (2007) note that analysis of smartcard data offers the potential to better understand passenger demand. Rather than rely on vehicle load estimates, AFC data show the exact number of people who have boarded at individual stations or onto individual vehicles. Smartcards also allow passenger journey histories and behavior patterns to be accumulated, since individual riders can be tracked using their unique smartcard numbers. Given the value of understanding how passengers

change their journeys over time and in response to incidents, this feature of AFC data is particularly useful—assuming that riders retain their cards.

Trépanier et al (2009) show that smartcard transactions allow continuous data collection over a long period of time, with transactions recording spatial and temporal details. Van der Hurk et al (2012) show that examination of AFC data can document changes in passenger demand over time. This demonstrates the value of using smartcards for panel studies that consider the actions taken by individual passengers in response to changes to transit service offered.

Wonjae (2010) evaluated smartcard data on the Seoul transport system for their potential to support transportation planning applications. Wonjae demonstrates that AFC data provide a far more comprehensive look at actual customer journeys than available through manual surveys. This expands the array of potential research into virtually any time period, at any date, and significantly reduces the cost of undertaking research, since data is recorded automatically. His research demonstrates how AFC information might be used for bus applications to identify customer travel times. This may offer an opportunity to improve planning, because if transfer points are identified through customer data, specific problem locations (for instance, those marked by long transfer times) can be targeted for improvement. Nassier et al (2011) have completed similar work that combines AFC data with transit schedule information to provide broad estimates of passenger origins and destinations.

Paul (2010) analyzes how AFC data can be compared to train arrivals and departures information on the London Underground. She considers whether it is possible to assign passengers moving through the system to individual trains. By cross-checking AFC records for passengers entering and leaving stations with train movement information, she makes several assumptions about how passengers choose to board trains (such as usually boarding the first train to arrive). She shows that connecting this information would be feasible with perfect data integrity and she develops a method by which this information might be used to estimate train loads and the number of passengers left behind at stations because of overloaded trains. But TfL's existing data, especially its train-tracking system, is insufficiently accurate and incomplete. These deficiencies limit the ability to determine the trains pas-

sengers take.

As the aforementioned studies demonstrate, the use of AFC data for an analysis of passenger flows on a public transport network offers considerable potential for further understanding the impact of disruptions on system customers. Much of the rest of this chapter describes studies that also use AFC data to analyze reliability, as does the thesis in general.

2.3 UNDERSTANDING RELIABILITY IN URBAN RAIL SYSTEMS

As discussed in Section 2.1, customers strongly value a reliable public transport system. Reliability, in turn, is often affected by disruptions on a particular transit line. Several studies, including a number that reference the London Underground specifically, review the value customers place on reliability and the ways in which transport systems might seek to improve customer views of reliability. These are described at the beginning of this section. Later, several metrics that attempt to measure the reliability of transit service using AFC and train-tracking data, are reviewed.

Bates et al (2001) conduct an extensive review of passenger views of rail travel. Their study shows that riders place a high value on punctuality of services. In particular, longer wait times (particularly compared to normal for regular passengers) increase stress. This provides empirical support for the notion that reliable transport services with useful information about the status of operations are likely to be preferred by passengers using the system. Thus there is a business case for improving reliability or at least improving the communications to customers about disruptions.

Nonetheless, customers are decision makers who frame their decisions within the framework of a wide array of options. As such, to assert that reliability "matters" to customers, one must put that value within the context of other potential travel options, which have different levels of reliability. Bates et al use the Expected Utility Theory, introduced by Von Neumann and Morgenstern (1944), to describe the relative utilities of different choices within the set of options provided to a decision-maker. Von Neumann and Morgenstern assert that the utility of an action is the sum of all utilities of potential outcomes, multiplied by their respective probabilities.

$$EU_a = \Sigma_O (U_{oa} \cdot P_{oa})$$ (2-1)

Where,
EU_a is the expected utility of course of action a;
U_{oa} is the utility of outcome o of action a; and
P_{oa} is the probability of outcome o of action a.

This would imply that, in the case of a transit system, customers simply include the lack of reliability of a line, for example, in their personal utility calculations and adjust their decisions according to the reliability of different options. Bates et al, however, argue that this model does not account for the emotional element of a decision-maker's thinking process and implies that people are indifferent to risk. Bates et al argue that this is not reflective of the way typical people make choices, since most people do in fact have quite visceral responses to the idea of risk.

Indeed, as noted by Kahneman and Tversky (1979), people encountering situations with outcomes that might involve losses are likely to increase their willingness to accept risk. In situations with outcomes that involve gains, on the other hand, people are likely to increase their proclivities toward risk aversion. The authors' *Prospect Theory* suggest that there are five strategies for people to minimize risk:

1. Seeking to reduce the uncertainty by accessing additional information;
2. Seeking to reduce the uncertainty by advance planning;
3. Seeking to reduce the consequences of the uncertainty;
4. Accepting the uncertainty and seeking to make the best decision in light of it; and
5. Seeking to capitalize on the uncertainty.

For transit systems, there are clear ways to aid people in undertaking these strategies. Better schedule and real-time information allow customers to make decisions in advance and en route. Improved communications provide the potential of reducing uncertainty and limiting the consequences of a disruption by informing customers about alternatives given the circumstances.

Chan (2007) shows how AFC data from London can be used to measure the compactness of the journey time distribution for individual origin-destination (OD) pairs. She compares the data available from Oyster smartcard data and the information provided by train arrivals and departures. She demonstrates that AFC data provides a clearer look at actual journey times than estimations made using

the train travel information, and is able to use that insight to produce an improved OD flow matrix for Underground system use.

Uniman (2009) expands on Chan's method, developing a measure of reliability for a rail transit system by examining operations on the London Underground. He documents how the reliability of journeys on a specific OD pair diverges from the schedule in order to clarify the degree to which services were being delivered as planned. To do so, he compares Oyster data across the Underground's most-traveled 800 OD pairs and compares the 95[th] percentile of journey times to the median, as shown in Equation 2-2. Uniman defines the "reliability buffer time" (RBT) here as the "amount of extra time that passengers must budget above the typical journey time in order to arrive on time at their destination with a specified level of certainty."

$$RBT_{OD} = (TT_{95\%} - TT_{50\%})_{OD} \hspace{3cm} \text{(2-2)}$$

Where,
$TT_{50\%}$ and $TT_{95\%}$ are the median and 95th percentile journey times, at OD pair level, for a time interval of 15 minutes to a full day, for a period of 1 to 20 days.

In order to measure the reliability of a line as a whole, Uniman proposes weighing the reliability of each OD pari within the line by the flow along that OD pair (for all OD pairs on the line), as shown in Equation 2-3.

$$RBT_{line} = \frac{\Sigma_{OD \in line} f_{OD} \cdot RBT_{OD}}{\Sigma_{OD \in line} f_{OD}} \hspace{3cm} \text{(2-3)}$$

Where,
f_{OD} is the passenger flow on the respective OD pairs; and
RBT_{OD} is the reliability buffer time for that respective OD, as established in Equation 2-2.

As an example of results from this analysis, Uniman finds that passengers on the Victoria Line of the London Underground had a median total travel time of 16.71 minutes with an RBT of 8.55 minutes, producing a total travel time of 25.26 minutes for a particular study period.

Schil (2012) extends Uniman's method and demonstrates how RBT information could be communicated to customers using, as an example, web-based travel planning software already available from many public transport systems, including TfL. Schil demonstrates that, in order to be utilized most effectively, AFC data should be aggregated spatially and temporally. When providing riders "maximum" potential journey times, Schil notes that these should reflect the time of day travelers are intending to make their journeys.

Hickey (2011) takes a somewhat different approach to Uniman and Schil, using train-tracking data to determine an improved estimate of platform waiting time for passengers at London Underground stations. Though TfL currently uses headway data to estimate the amount of time passengers wait at stations to board trains, Hickey argues that deficiencies in the available train tracking information make using headways problematic. As an alternative, he suggests using the count of trains arriving and departing from stations to make an assessment of how long passengers waited for trains.

This review of the literature on reliability indicates that since reliability is a heavily valued aspect of a typical customer's experience on public transport, system managers should place a high priority on improving the reliability of their systems. In addition, as highlighted by Kahneman and Tversky, there are customer-friendly ways to improve information provision to reduce the stress resulting from unreliable service. Finally, Uniman, Schil, and Hickey provide useful measures for assessing the reliability of transport systems using available AFC and train-tracking data.

2.4 RESPONDING TO INCIDENTS

The focus of this thesis is on understanding the customer impacts of disruptions, but the impact of incidents on riders must be seen in the context of the decisions made by transport system personnel on how to respond to those incidents. This section describes how service managers use train-tracking

information to alter the services provided to restore services to normal.

Pender et al (2012) compare disruption response management across 48 international rail transit agencies, with a particular focus on Australian agencies. They compare disruption responses on different systems by undertaking interviews with staff on response to service outages and changes required to the train service to respond to problems. Jesperson-Growth et al (2007) perform a similar analysis on the Dutch Railways and Copenhagen suburban commuter rail network. Their study notes that there are several specific ways to handle the impacts of disruptions, including timetable adjustment, rolling stock re-scheduling, and crew re-scheduling.

Carrel (2009) analyzes the manner in which line service controllers manage disruptions on the London Underground system, with a particular focus on the Central Line, one of the major Underground lines. He performs a qualitative study of train control decision making by visiting control centers and taking into account the decision environment of those determining what changes to operations should be made in response to specific disruptions on the line. He then uses signaling data to reconstruct train operators' interventions and proposed a measure for assessing the impact of those interventions on train operations.

2.5 EXTENDING THE LITERATURE: IMPROVING UNDERSTANDING OF CUSTOMER RESPONSE TO DISRUPTIONS

The literature review offered in this chapter provides an overview of the studies that have been performed on rail transport network reliability and disruption response, in addition to the use of smartcard technologies to better analyze available information. The literature indicates that there are significant advantages in using AFC data, as it provides a large amount of information that describes customer travel patterns at all points of the day on systems outfitted with the technology, such as the London Underground. Several of the studies in Section 2.3 note that AFC data can be used to assess the reliability of public transport lines. In the context of customer desire for better service reliability, as shown in Section 2.1, expanding our understanding of reliability may be particularly useful.

The literature provides a look at the manner in which rail service controllers choose to respond to

disruptions on the line, but it is less comprehensive in describing the actual effects of disruptions on passengers. Though Uniman and Schil, among others, show how reliability of a rail transport line can be measured using smartcard data, they do not detail how the specific effects of a disruption can be evaluated using that data. Because disruptions can reduce reliability significantly, this is a gap in our current understanding of the way public transport works and how customers can be better served by it.

This thesis extends the literature by using smartcard data to evaluate the impact of disruptions on customers. Rather than focusing on line reliability as a whole, the thesis identifies specific disruptions and then identifies what information is offered by smartcard data about impacts on customers. Though the thesis does not identify specific ways in which controllers can use this additional information to improve response to incidents, the literature indicates that current controller response focuses on information provided by train tracking data, not AFC data. Thus there is a gap in controller knowledge that could be expanded with more analysis of AFC data. If controllers want to adapt their response strategies to minimize customer impacts, this could be a useful direction for study.

3

Transport for London, the Underground and the Piccadilly Line

This chapter introduces the transport system in London in order to provide context for the services and data sources analyzed in the remainder of this thesis. Section 3.1 describes the transport system, emphasizing its extent and usage, focusing on the London Underground in particular. Section 3.2 describes the Piccadilly Line, one of the main Underground services and the focus of this research. Section 3.3 introduces the main data sources available for research on the transport system in London, which will be described in more detail in the following chapters in this thesis. Finally, Section 3.4 provides an overview of planning and analytical methods currently used in London to evaluate the performance of the transport network.

3.1 TRANSPORT FOR LONDON AND THE UNDERGROUND

Transportation in Greater London, an 8.2 million-person metropolitan area, is overseen by Transport for London (TfL), a public agency created in 2000 and overseen by the Mayor of London (TfL, 2013). TfL manages the region's arterial road network and sets budget and fare policies for the operation of most public transport services. It oversees local buses, the London Underground (LU), the London Overground rail line, the Tramlink light rail line, the Docklands Light Railway (DLR), the Barclays Bike Share network, and the Emirates Air Line aerial tramway, but does not manage operations on most of

the national rail lines serving the region, which are the responsibility of the national Department for Transport. TfL contracts out the operation of most of its services, with the major exception of the Underground, which is operated by London Underground Ltd. (LU), a subsidiary of TfL. LU is the world's oldest urban rail transit system, opening its first subway service in 1863. It is now one of the planet's largest urban transport systems, serving 270 stations, 11 lines, 402 line-kilometers, and about 3.7 million daily journeys. TfL has an annual budget of almost £10 billion and roughly 30,000 employees (TfL, 2012a).

Figure 3-1 illustrates TfL's rail network, which consists primarily of Underground lines. The Overground network is shown as dotted orange lines; Tramlink, which serves the southern part of the region, is shown as a dotted green line; and the DLR is shown as a dotted aqua line. The National Rail network, not managed by TfL, is shown as a series of gray lines. The Piccadilly Line, which is the focus of this research, is shown in dark blue. Altogether, the network is large, complex, and highly interconnected.

Under the leadership of the Mayor of London, TfL has substantially expanded and improved the region's public transportation network. Since 2000, several Underground lines have been upgraded, the Overground network has been created, and DLR has been substantially extended. Bus services have been expanded with increased operating hours and more frequent services. At the same time, the regional government, working with local councils, has expanded the number of available bike lanes, improved the walking environment, and introduced a congestion charge in central London.

These initiatives have contributed to significant changes in the commute patterns of residents of Greater London. As shown in Figure 3-2A, and documented in the Travel in London report, the public transport mode share of all journeys in the region has increased substantially over the past two decades (TfL, 2012c). The overall public transport mode share, including rail, bus, and Underground journeys, increased from 24% in 1993 to 35% in 2011, while that of private motorized vehicles (motor cycles, automobiles, and taxis) declined from 51% to 40% over the same period. The share for non-motorized modes, including walking and cycling, remained steady at around 25%. Figure 3-2B shows

| Piccadilly Line | London Overground | Tramlink | Docklands Light Railway |
| Other London Underground Lines | National Rail Network | Emirates Air Line | |
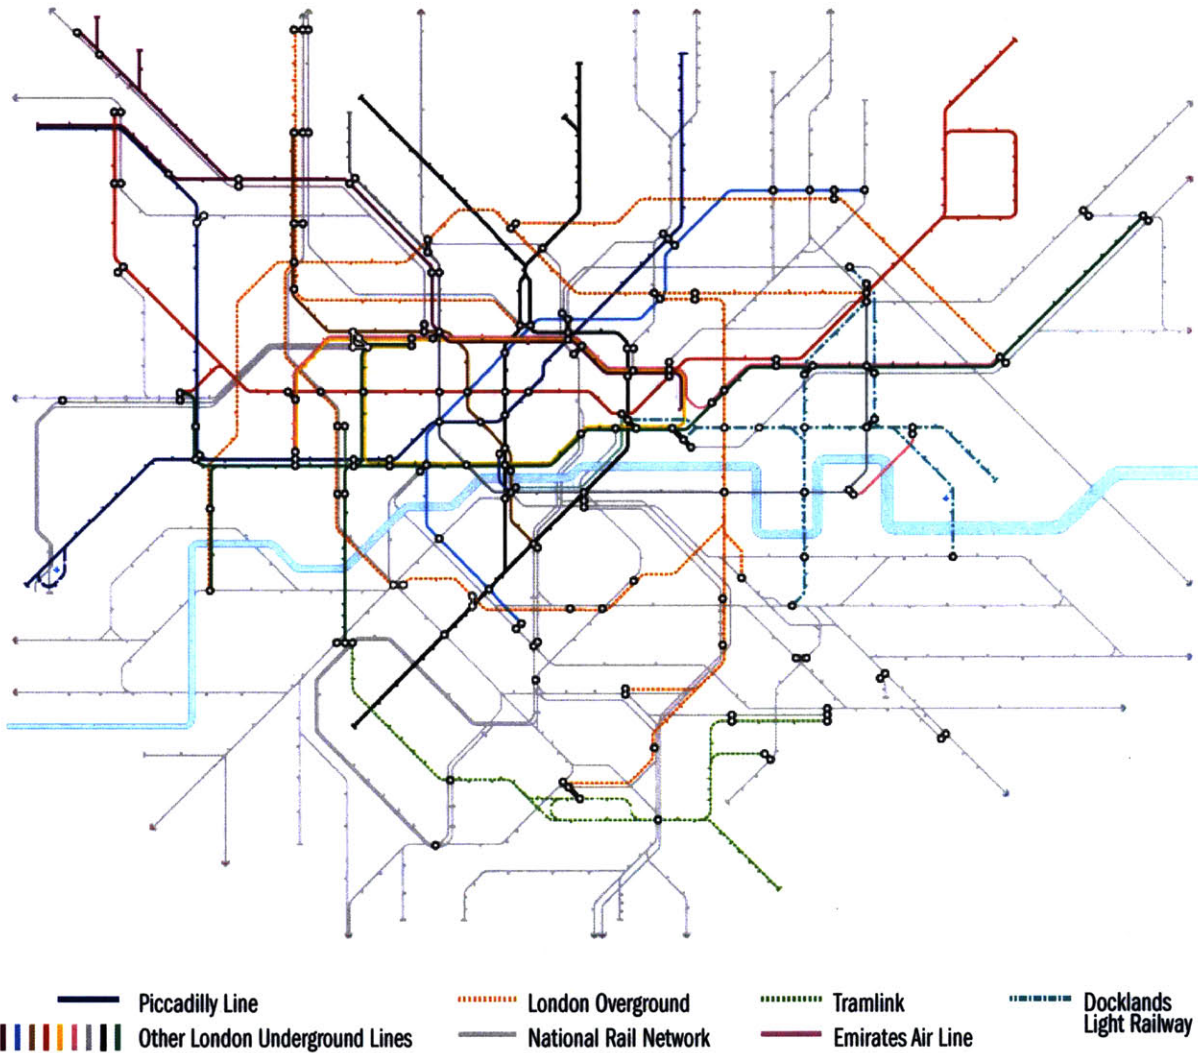
Figure 3-1: The TfL rail network. (Adapted from TfL, 2012b)

that while the number of trips made by public transport modes has expanded by more than 50% since 1993 and the number of walking trips has increased by about 20%, overall use of private vehicles has declined slightly. In this period, about 1 million more people have been added to the region's population. By 2031, the population is expected to increase by 1.25 million (Mayor of London, 2010).
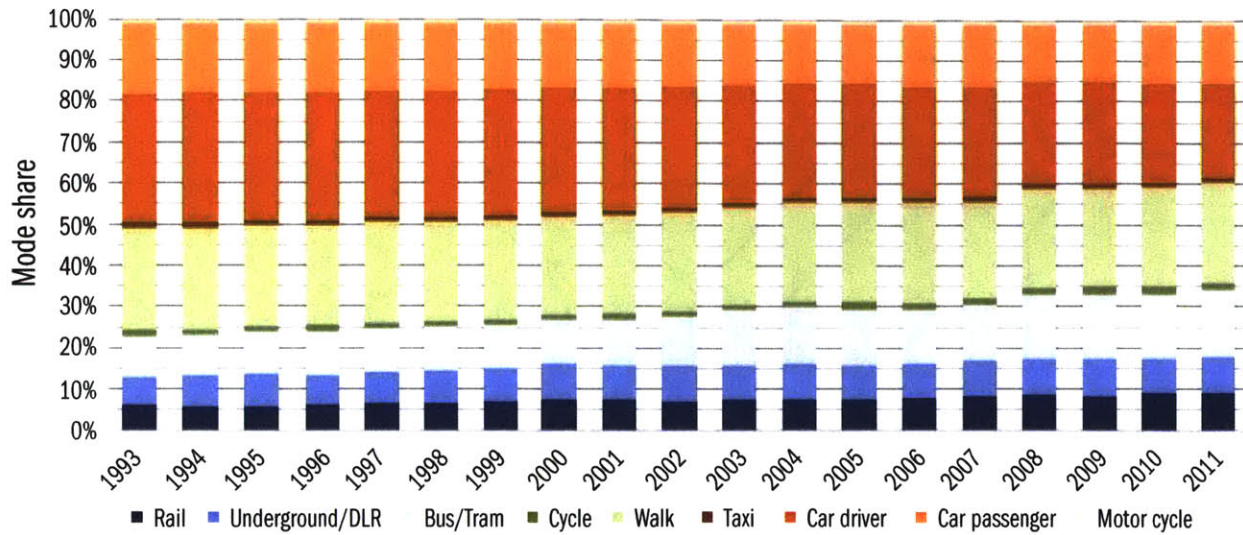
A. Mode share of all trips in London.
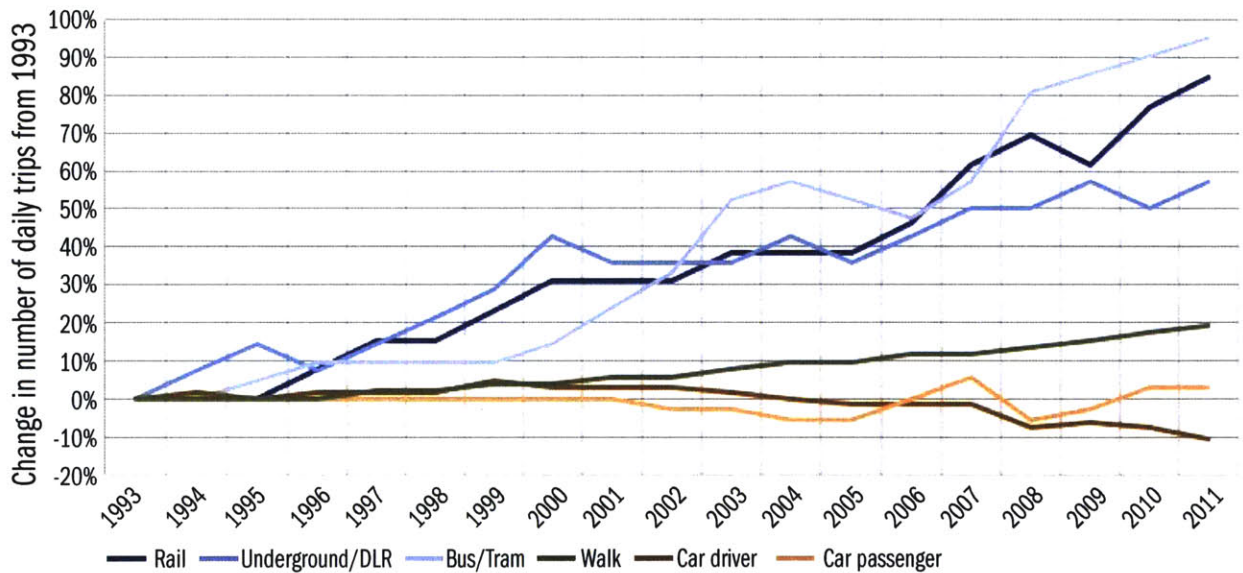
B. Change since 1993 in number of trips by Londoners by mode.

Figure 3-2: Travel modes of Greater London residents. (Adapted from TfL, 2012c)

The increased use of the public transport network can also be quantified in terms of passenger entries into the system. In 2011, the TfL system provided more than 3.5 billion rides, of which 63% were provided by bus and 31% by the Underground (see Figure 3-3). 2011 was a record year for the Underground in particular, with more than 1.1 billion rides. Current estimates are that annual ridership will increase to 1.3 billion by 2015 (TfL, 2012d). London's continued population growth and increasing use of the public transport network demonstrates the vital importance of providing efficient and comprehensive transportation services throughout the region.
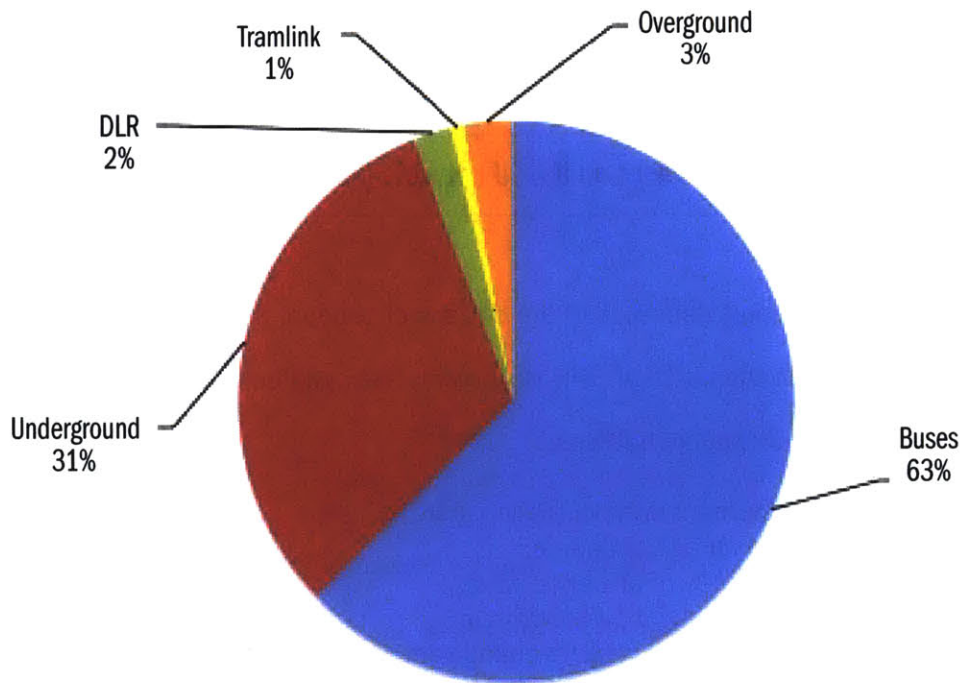
Figure 3-3: Distribution of ridership on TFL-managed transport modes. (Adapted from TfL, 2012e)

The distribution of ridership on the TfL network, however, is not even, either spatially or temporally. Ridership of the public transport network is highest in London's central area (TfL, 2012c). Weekday ridership is considerably higher across the TfL system at peak hours (roughly 7:00-9:00 and 16:00-18:00), than during the early morning, midday, or evening periods, as shown in Figure 3-4 (Hickey, 2011). The transit system must handle a heavy influx of passengers and very crowded conditions at these times.
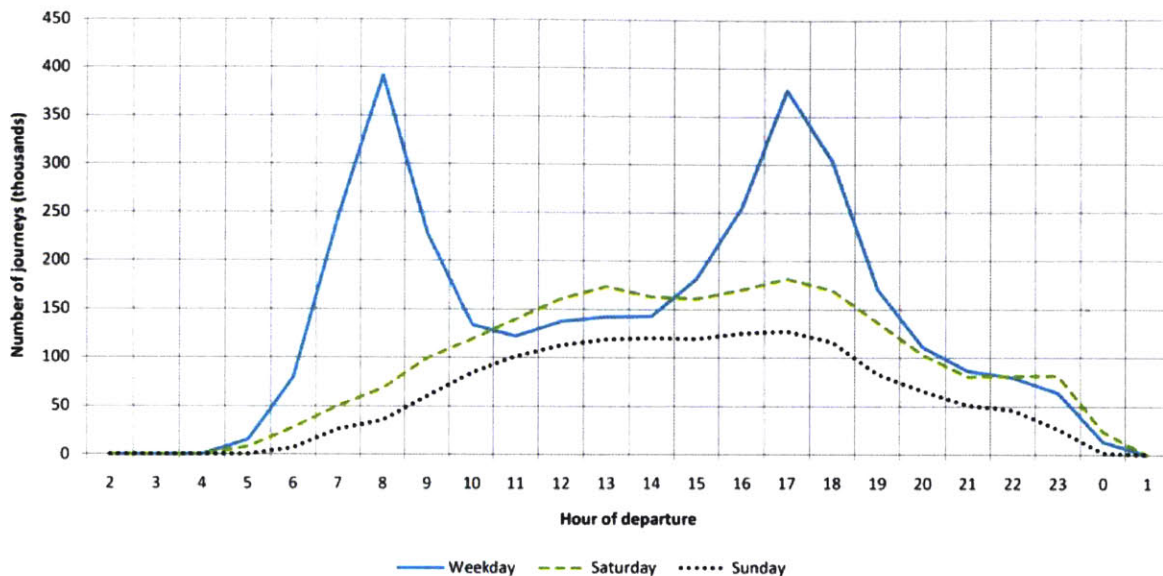
Figure 3-4: Temporal distribution of journeys on the LU network. (Source: Hickey, 2011, Figure 3-2)

In order to respond to the transport demands of the citizens of London, the Mayor has established a Transport Strategy (Mayor of London, 2010). This document, last updated in 2010, highlights six primary goals the transport network should address:

- Support economic development and population growth;
- Enhance the quality of life for all Londoners;
- Improve the safety and security of all Londoners;
- Improve transport opportunities for all Londoners;
- Reduce transport's contribution to climate change and improve its resilience; and
- Support the delivery of the London 2012 Olympic and Paralympic Games and its legacy.

These goals remain relevant. TfL's management, which is appointed by the Mayor, organizes transport policy to reflect the intentions of that office. Its 2012 business plan identifies three major goals that correspond to the Mayor's overarching policy objectives:

1. Driving London's employment and population growth;
2. Putting customers at the heart of the business; and
3. Making life in London better for all.

The focus of this thesis is on examining the effects of disruptions on LU customers. From this perspective, TfL's second goal is particularly relevant. Indeed, the Business Plan argues that "the projects laid out in this plan demonstrate a commitment to support customers when things go wrong. They are designed to keep customers up to date during planned or unplanned disruption, and ensure that they can make informed journey decisions" (TfL, 2012d, 29). This thesis responds to this objective and seeks to contribute to TfL's efforts to improve customer service.

3.2 THE PICCADILLY LINE

Though many of the methods discussed in the thesis are applicable throughout the London Underground, and, in the longer term, to other rail transit lines with AFC data, the focus throughout this research is the Piccadilly Line. This rapid transit service is 71 kilometers in length, the second longest in the system (after the Central Line), and serves 52 stations. As shown in Figure 3-5, a spatial representation of the system, the line runs from Cockfosters in North London, through the center of London, and then to two western branches—one service to Heathrow Airport and the other to Uxbridge. Along the way, the line serves many important stations, including King's Cross-St. Pancras, Piccadilly Circus, and South Kensington. The central section of the line is in a deep tube that first opened in 1869, but most of the branch sections run on the surface or on embankments.

As Figure 3-5 shows, LU provides a range of service at rush hour to different parts of the line. The central section, from Acton Town to Arnos Grove, is scheduled for up to 24 trains per hour. Other sections of the line see less service. Heathrow Airport stations, for example, see between 6 and 12 trains per hour during the rush period. At Uxbridge, there are only 3 trains scheduled per hour. One explanation for the reduced service to stations between Rayners Lane and Uxbridge is that this two-track section of the line is shared with the Metropolitan Line, which offers more direct—and more frequent—service to central London. The four-track section of the line between Ealing Common and Earl's Court is shared with the District Line, which operates as a "local" service to the "express" Piccadilly Line (the two services are on separate, but adjacent tracks), which serve fewer stations. Service is generally operated from 5:30 to 0:30, with 894 scheduled one-way trips on the line. As illustrated in Figure 3-5,
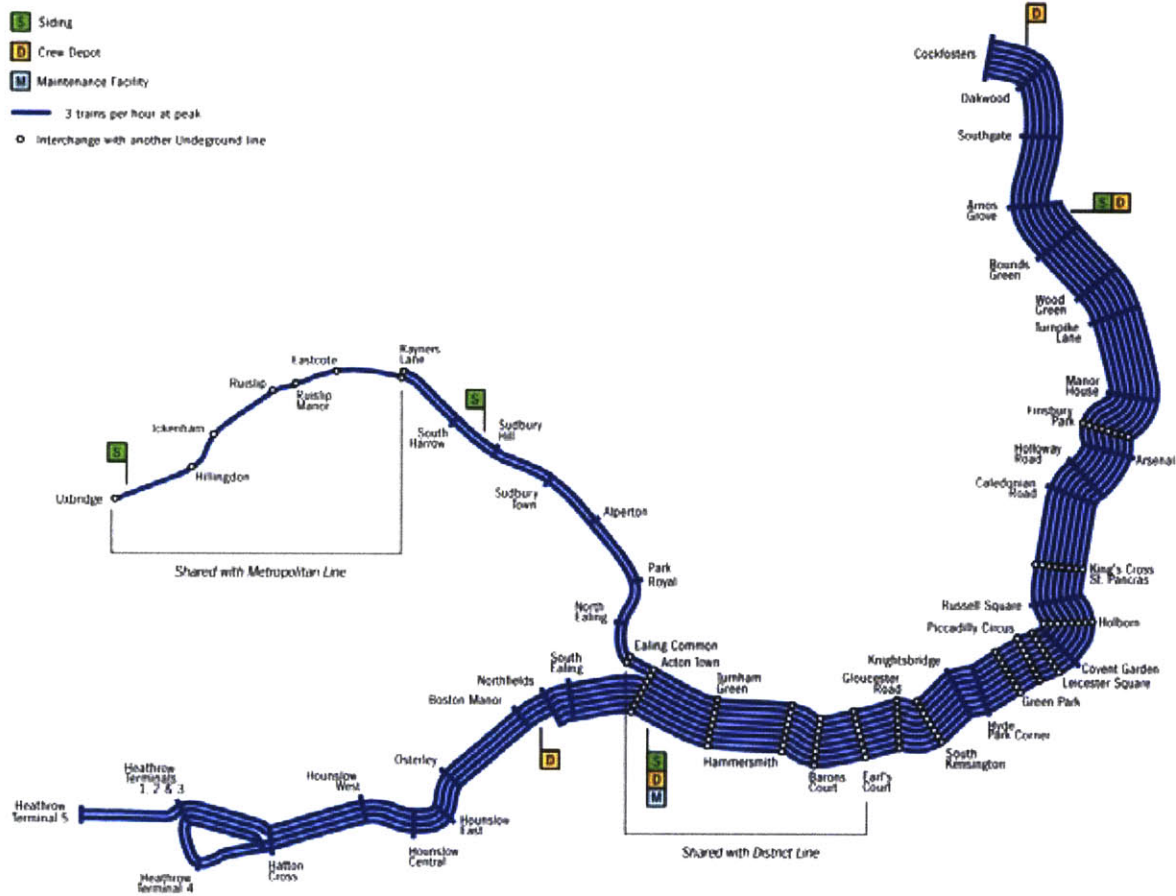
Figure 3-5: The Piccadilly Line. (Source: the author, based on various TfL sources)

many of the trips provided are not operated terminus-to-terminus and some, in fact, terminate quite a bit before the end of the line.

There are four crew depots on the line, located near Northfields, Acton Town, Arnos Grove, and Cockfosters. Although each of these depots allow for trains to be reversed and have facilities for staff, only the Arnos Grove and Acton Town depots are currently used. There are sidings for trains to be reversed at Uxbridge, South Harrow, Acton Town, and Arnos Grove. The main maintenance base for train upkeep is located at Acton Town. Finally, the line operations control facility is located at Earl's Court.

The Piccadilly Line, which runs through some of London's most popular residential districts as well as many of its major destinations, serves more than 210 million annual journeys, making it the fourth-busiest line on the LU network. Congestion, as on the TfL system in general, is heavily concentrated at

peak times and in the central area. In the morning peak, for example, eastbound trains between Acton Town and Knightsbridge operate at or above nominal capacity; westbound trains from Finsbury Park to Holborn operate at more than 4 standing passengers per square meter, This means that passengers are often crowded into trains or left waiting on platforms (Ravichandran, 2012; London First, 2012).

This thesis focuses on the Piccadilly Line in part because the service is one of TfL's most popular. Its ridership is likely to increase over the next decade as London's population continues to increase and Heathrow Airport and King's Cross-St. Pancras are expected to become even more popular than they are today. But just as important to the selection of the Piccadilly Line as the subject of this in-depth analysis is the fact that the line is years away from a major upgrade. Though the service has been maintained well for decades, it is handicapped by having some of the oldest trains in the system (they date from 1973), an antiquated signaling technology, and decaying capital infrastructure resulting from the line's age.

The TfL business plan suggests that capital improvements are many years off, though the Mayor's Transport Strategy notes that there is an effort to provide the Piccadilly Line "new trains, more capacity and quicker journeys" (Mayor of London, 2010, 10). Upgrades for the Piccadilly Line are expected in the "2020-30s" period. Moreover, any alternative investments that will lead to reduced stress on the line are unlikely in the medium term. The Victoria Line, which parallels the Piccadilly Line for much of its route, is often even more crowded. And the proposed "Crossrail 2"—a new service mentioned in the Mayor's Transport Strategy that would run between Chelsea and Hackney (and could parallel the Piccadilly Line from Wood Green and King's Cross to Piccadilly Circus)—has yet to be funded, let alone begun construction (London First, 2012).

These conditions—rising use of the London Underground and a wait of a decade or more before a major line upgrade—make the Piccadilly Line an especially important case study. This is particularly true in fulfilling TfL's objective to "put customers at the heart of the business" (TfL, 2012d, 28). If there are better ways to evaluate the effects of disruptions and to communicate that information to riders, the service provided by TfL will be improved, even if physical upgrades remain far off. The remainder of

the thesis explores methods by which such an evaluation might be undertaken.

3.3 DATA SOURCES

In order to conduct such an evaluation, it is necessary to understand and utilize the data sources that are available to TfL. Thanks to major advances in data collection and access, public transport agencies such as TfL are able to monitor the movement of customers and vehicles much more effectively, and sometimes even in real time. The datasets available at TfL, and used in this thesis, are introduced in this section.

Oyster data

At the heart of the thesis are the data collected about passenger journeys using automatic fare collection (AFC) devices. In London, an AFC system was introduced in 2003 under the name Oyster, which provides customers smartcards, which they then use to tap in and out of the transit system. All stations on the Piccadilly Line, and the vast majority of LU stations, require customers to pass through fare gates to enter and exit stations, and most riders now use Oyster cards to do so, though a magnetic strip card (for which TfL cannot access individual origin-destination journey data) is also used for some trips. Because of the widespread availability and use of Oyster, the resulting data can be viewed as representative of the overall users of the line, even including those that still use magnetic cards. Oyster allows TfL to account for journey distance and time of day, both of which affect customer fares.

Table 3-1 provides a sample of records produced by Oyster. In this case, a query was made for journeys between Cockfosters and Holborn stations on September 23, 2011 (this table shows just a portion of all of the trips between those stations on that day). Only completed Oyster journeys—those where customers tap in and out—are included. There is a small number of cases where customers either did not tap in or tap out, but those do not have any evaluative potential since their origins, destinations, and journey times cannot be fully accounted for.

Oyster provides a wide range of data, mostly about fare payment, but for this research only a sub-sample is used, including those noted in Table 3-1:

Table 3-1: Sample Oyster records.

DAYKEY	PID_ENCRYPT	STARTLOC	ENDLOC	T_CEN	T_CEX	JNYTYP
11588	30164031	580	607	430	447	PPY
11588	41537433	580	607	432	449	TKT
11588	18969761	580	607	440	459	TKT
11588	38317266	580	607	454	473	PPY
11588	47107984	580	607	465	480	TKT
11588	41513646	580	607	463	481	PPY
11588	44759809	580	607	465	483	PPY
11588	45585510	580	607	466	483	PPY
11588	20458144	580	607	470	486	TKT
11588	45962166	580	607	470	486	TKT

- **DAYKEY**: The date;
- **PID_ENCRYPT**: An encrypted version of each individual's Oyster smartcard number;
- **STARTLOC**: The starting station of the customer's journey (a 3-digit code is assigned to each station, including Network Rail stations that accept Oyster cards);
- **ENDLOC**: The ending station of the customer's journey;
- **T_CEN**: The time that the customer passes through the fare gate at the origin station (measured in minutes after midnight of the travel day; journeys after midnight but before the close of service retain the same travel day designation and thus date or DAYKEY);
- **T_CEX**: The time that the customer passes through the fare gate at the destination station; and
- **JNYTYP**: Ticket type, where TKT is a pay-as-you-go ticket where the customer is charged per trip and PPY is a pass where the customer has unlimited use of the system for a day, week, or month.

It is straight forward to measure a customer's journey time by subtracting the entry time from the exit time. It needs to be emphasized that the "journey time" discussed here is defined as the time between passing through the faregate at the entry station and passing through the faregate at the exit station, *not* the travel time on the train itself. In the case of the sample data shown in Table 3-1, the median journey time for customers traveling between the Cockfosters and Holborn was 17 minutes; the average was 17.2 minutes; the minimum was 15 minutes; and the maximum was 19 minutes (times are truncated to the minute, as described further below). This is a small sample, but it demonstrates both

the reliability of the Oyster data in documenting journeys and the variability in customer journey times.

Figure 3-6 provides a clear illustration of the variability in Oyster journey times. Though train travel time accounts for the majority of the average customer's journey time, platform wait time (the time spent waiting for the train) and access and egress times between the faregates and the platforms collectively account for a third of customer journey time. In the example presented in Table 3-1, customers 44759809 and 45585510 almost certainly took the same train, as they exited the station at the same time (483, or 8:03). But the first customer arrived at the station one minute earlier and thus had a journey time that was one minute longer (18 vs. 17 minutes, respectively, though these figures are truncated to the minute, as described below). This customer either took longer to get from the gate to the platform and/or waited longer for the train. The Oyster journey time therefore does not directly measure the time for a train to travel from one station to the next.

It is reasonable to assume that certain components of the Oyster journey are more relevant to understanding journey time than others. Access, egress, and interchange time account for an average of 21% of customer journey times, but they are relatively invariable *en masse*. While certain customers may take longer to get from the faregate to the platform, the average customer given a large enough
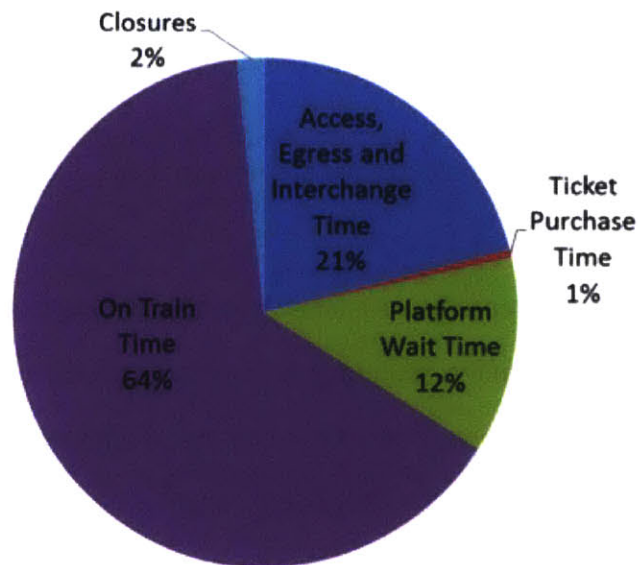
Figure 3-6: Components of an Oyster journey record. (Source: Schil, 2012, based on TfL JTM reports)

sample size will take a consistent amount of time, no matter the train service being provided. This is not true of the on train time or platform wait time, which are heavily reliant on the train service itself. From the point of view of Oyster, a longer journey generally means slower trains (more on train time) or fewer trains (longer platform wait time). In disrupted conditions, these two features frequently come together, but they cannot be differentiated without evaluating train movement data.

Oyster records suffer from an additional issue: They are recorded only to the minute, providing no information about the *second* at which customers enter and exit faregates. As a result, customers 44759809 and 45585510 may in fact have entered Holborn only one second apart, not one minute as reported. As Figure 3-7 shows, this can inflate or deflate journey times significantly. An Oyster record showing passengers entering at 7:00 and exiting at 7:30, a 30-minute trip, may in fact refer to a customer taking just over a 29-minute trip (the shortest-possible journey time) or a customer taking just less than a 31-minute trip (the longest-possible journey time). This almost two-minute difference must be recognized in analyzing Oyster data.
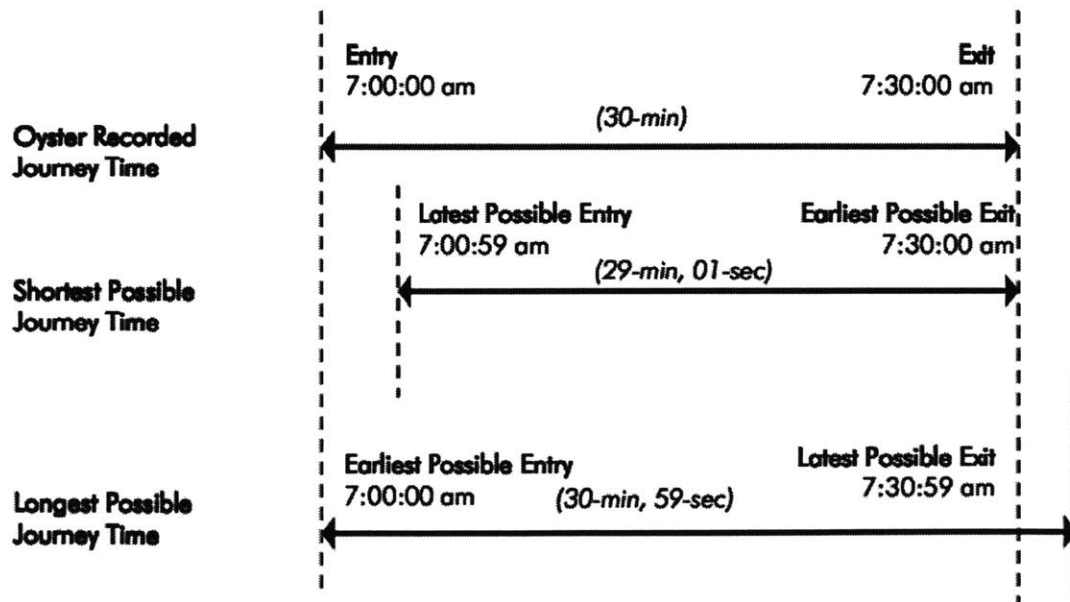
Figure 3-7: Variations in Oyster journey times. (Source: Uniman, 2009, Figure 3-5)

Finally, Oyster provides no accounting for unusual passenger behavior. Certain customers may choose, for instance, to take a roundabout route between two stations or not board the first train to arrive. Their longer journey times may indicate to an analyst a disruption in train service when there is none. Fortunately, these customers represent only a small percentage of the overall ridership, since most people want to complete their journeys as quickly as possible. With a large-enough sample, the effects of these unusual passengers can be minimized. It is necessary to produce an adequate sample when analyzing a dataset of Oyster records. In previous research, such as Uniman (2009), a set of 20 riders with the same origin and destination has been used as the minimum sample size to estimate the mean journey time between two stations.

Train-movement data

At LU, information about the location of trains is provided through a system called NetMIS, whose data can be interpreted as a reasonable proxy of an automated vehicle location (AVL) system. This data is derived in the following sequence. First, track circuits indicate the presence of a train within a specific track circuit. This data is sent to a local signaling computer that in turn sends the train location information to the TrackerNet database, which indicates the location of a train. Because track circuit locations do not directly match the locations of stations, adjustments are made automatically to the TrackerNet data before it is communicated to the NetMIS system. NetMIS can then be used to determine the status of service on a line.

When queried for a particular station, NetMIS produces a dataset similar to that shown in Table 3-2, which shows the arrival and departure times for eastbound Piccadilly Line trains at Acton Town on September 19, 2011. NetMIS can also be used to track specific trains along their journeys, but this thesis concentrates on station-based information such as that found in Table 3-2.

Table 3-2 illustrates what information can be gleaned from NetMIS data. Train numbers, track locations, and destinations are indicated. The train number indicated in the leftmost column is the scheduled trip number and the LCN in the rightmost column is the number assigned to the actual vehicle. By subtracting the arrival or destination time of one train from the previous one, the headway

Table 3-2: Sample NetMIS record.

| Train/Trip | Dwell | | Event Time | | | | | Track Info | | | | | |
	Actual	Scheduled	Arrival	Departure	Scheduled Dep	Last Green Signal Time	Passenger Destination	Track Name (Arrival)	Track Name (Departure)	Loading	In Service	Train Ref No.	LCN (Leading Car Number)
252 / 0	01:05	00:00	07:13:28	07:14:33		07:13:17	Cockfosters	CXSEL_2	CSACT			1251068	250
000 / 0	00:35	00:00	07:15:41	07:16:16		07:15:30	Cockfosters	CXSEL_2	CSACT			1252066	0
275 / 0	00:41	00:00	07:17:36	07:18:17		07:17:16	Cockfosters	CXSEL_1	CSACT			1250915	147
343 / 0	01:13	00:00	07:19:58	07:21:11		07:19:39	Amos Grove	CXSEL_2	CSACT			1249105	128
000 / 0	01:16	00:00	07:22:25	07:23:41		07:22:07	Cockfosters	CXSEL_2	CSACT			1239270	0
260 / 0	00:36	00:00	07:25:09	07:25:45		07:24:47	Cockfosters	CXSEL_1	CSACT			1250281	891
272 / 0	00:33	00:00	07:27:07	07:27:40		07:28:44	Cockfosters	CXSEL_2	CSACT			1252217	245
000 / 0	:	00:00		07:27:31		07:28:55			DMACT_1			1252500	0
345 / 0	00:49	00:00	07:29:54	07:30:43		07:29:38	Amos Grove	CXSEL_2	CSACT			1249474	240

between trains at a specific location can be determined. Counting the number of trains arriving during a specific period indicates train frequency.

In some important respects, NetMIS data are less complete than Oyster. In Table 3-2, the trains that departed the station at 7:16:16, 7:23:41, and 7:27:31 have no train number or LCN associated with them. In addition, the train that departed at 7:27:31 does not have an arrival time, passenger destination, or track name associated with it. These gaps in data are a result of aging physical infrastructure that was not designed for modern computing. On the Piccadilly Line, track circuit data are transmitted to small servers along the line in seven separate segments. At the junction of these segments, information about trains is passed between servers, but small differences in information between two servers can lead to information loss. As a result, certain train IDs are erased (leading to "000" train numbers). This makes tracking trains more difficult (Paul 2010). Similarly, there are circumstances in which the status of a train is unclear. The train that departed at 7:27:31 did not "arrive," which may indicate that it is a "ghost" train, or that it is the same as another train on the track. These

conditions make relying solely on NetMIS to understand the extent and impact of disruptions on the line somewhat problematic.

Train controller log

Combined with other sources of data, however, NetMIS' utility is expanded. For instance, the log of incidents recorded by service controllers offers important information about issues occurring on the line. Throughout the day, controllers write short summaries of the condition of train service on the line. This information generally includes notes about the start time and location of incidents. In addition, they list the (reported) cause of each incident. A sample from a train controller log, from the morning of May 16, 2012, is shown in Figure 3-8.

Controllers, who are the most knowledgeable staff in terms of conditions on the line at any given time, communicate information about delays to their staff and the Network Operations Center, which informs customers about disruptions on the line (this is further described in Chapter 6). The controller

Daily Commentary
05:00 SOT No issues to report at start of traffic. (Pat Horkan)
08:15 PTI Incident. **Hillingdon:-** A report was received of a group of four children becoming seperated from their mother. It is alleged the Train Operator of eastbound train 314 closed the saloon doors too quickly, thus seperating the group. In the initial panic, it is believed the mother called the police however, local station staff and those at Ruislip, liased to allow everyone to be reunited at Ruislip. The east-end Duty Reliability Manager will be investigating this matter further. **(Delay: 4 Mins.)** (Mark Cherrie)
09:13 Signal Remaining at Danger. **Hammersmith:-** Eastbound train 317 was delayed departing the platform due to signal number WD10a remaining at danger. This was due to the train's description becoming lost. The programme machine concerned was cleared in manual mode. Job number 1195516 refers. **(Delay: 4 Mins.)** Further, the same situation arose at 09.30Hrs, with eastbound train 273 being delayed departing the platform. **(Delay: 3 Mins.)** (Mark Cherrie)

Figure 3-8: Sample train controller log.

log information is available on the TfL server the following day, when it is reviewed by other LU staff. They use the controller logs to input information into the CuPID disruptions database, which is used to monitor impacts on train service resulting from incidents (see Chapter 5).

3.4 LU SERVICE PLANNING AND ANALYSIS METHODS

In order to plan for and review operations on the network, LU uses several standard methods. The Route Choice Model (RCM), the Train Service Model (TSM), and the Journey Time Metric (JTM) are three of the most important (see Uniman, 2009; Hickey, 2011; and Schil, 2012). They are discussed in this section.

The RCM estimates route choices for customers travelling between stations. Using passenger responses from TfL's Rolling Origin-Destination Survey (RODS), RCM provides probabilities that customers will take each alternative route. For example, customers boarding at Chancery Lane on the Central Line travelling to Leicester Square, on the Piccadilly and Northern Lines, have two similar-length options: Take the Central Line and transfer to the Piccadilly Line at Holborn; or take the Central Line and transfer to the Northern Line at Tottenham Court Road. Since Oyster does not provide any information about the route taken (since it only provides entry and exit information), the RCM can be used to estimate likely rider route patterns.

The TSM takes advantage of the data provided by the RCM, in addition to historic and projected Oyster and gateline data, to estimate how passengers move throughout the system as a whole by simulating train movement. This information can then be used to determine train load density at different points of the service day and at various locations; these estimates are used as an input into train scheduling. TSM is used principally when service plans are being developed and is not frequently updated.

Finally, the JTM, which has been used by TfL since 1997, is LU's most-frequently used assessment tool for estimating the quality of service for passengers on the line. The JTM is an analysis tool used by line managers to determine what aspects of the service need to be improved, and is updated every four weeks (this is LU's standard "period" of analysis). The JTM uses NetMIS and maintenance records

to compare actual and desired service quality. In the JTM method, the passengers' travel times are made up of five components:

- Ticket purchase time;
- Access, egress, and interchange time;
- Platform wait time;
- On train time; and
- Unplanned and planned closures.

JTM uses NetMIS data about train headways and combines it with other data, including stations closures, elevator outages, and track interruptions. Longer-than-expected (or more-crowded-than-expected) travel time data is weighted by a value of time, so the JTM is not a literal representation of extra time on the system experienced by customers. As shown in Figure 3-9, the JTM total travel time for an individual moving through the system is similar to the Oyster journey time, though the latter includes less access and egress time (because it offers no information about customer in-station time outside the faregates).
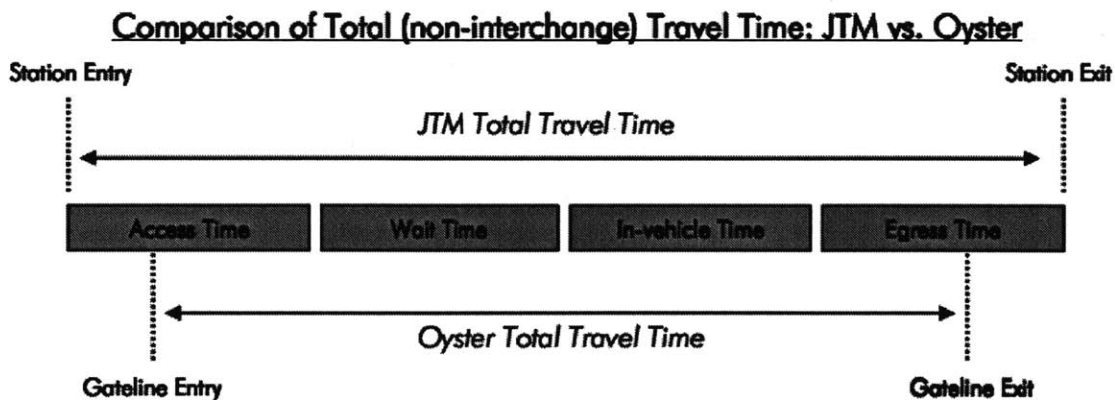
Figure 3-9: Comparison of total (non-interchange) travel time: JTM vs. Oyster. (Source: Uniman, 2009, pp. 53)

JTM is an important tool, but it has several limitations. A JTM assessment cannot be calculated for an individual journey (unlike Oyster), meaning that it does not directly represent the impacts of a disruption on any rider; the metric attributes station performance to individual lines, so the performance of an individual station, even if shared by multiple lines, is attributed to just one of them; it assumes riders take the first train that arrives, which may not be true if trains are particularly crowded; it provides no information about variation in travel times; and relies on a fixed demand matrix (Chan, 2007). Moreover, because it relies on surveys and models, it suffers from small sample sizes, infrequent updates, and costly data collection (Schil, 2012).

4

Categorizing incidents on the Piccadilly Line

This chapter defines the study period for this research and provides an overview of the types of disruptions that occur on the Piccadilly Line. A sample of 29 disrupted days in a normal operations period were selected for in-depth analysis using LU's existing method for identifying poor-performing days. Incidents during the study period were then categorized by location, duration, and effects on passengers. In total, the analysis offers greater insight into incident impacts, but does not attempt to provide further information about incident *cause*.

This view is designed to supplement the organization's existing CuPID disruptions database, with which a comparison is made. The goal of the chapter is to highlight the sources of information currently available within LU. This approach could be replicated for any period on any Underground line to develop a more comprehensive understanding of the characteristics of incidents. In the context of this thesis, the motivation is to lay the groundwork for further use of Oyster data in Chapters 5 and 6 for the purposes of improving our understanding of the impacts of disruptions on passengers and developing potential improvements to the customer notifications system.

4.1 LONDON UNDERGROUND APPROACH TO INCIDENT CHARACTERIZATION

The provision of information about what services are available or are expected to be available at any given time by public transport agencies is an important function, as transit is a customer service. Passengers should be provided clear relevant information about frequency of service and travel times from one station to another, at all times of the day, on both weekdays and weekends. Based on this standard information, customers might expect that service on weekday mornings should be similar across days, weeks, and months.

Of course, variation in service is endemic to transit operations. While in a perfect world trains would operate like clockwork, the reality is that a typical transit system incorporates a wide variety of moving parts, all of which must function normally for service to operate according to plan on any given day. For a perfect service day to occur, the number of passengers arriving at stations must be lower than platform capacity, and riders must enter and exit trains quickly, do nothing to block the doors, and remain healthy. Trains must never fail, switches must always operate correctly, and track fires never happen. Drivers must show up to work on time and react precisely to any changes in information provided to them.

It is not surprising that a "perfect" operational day rarely occurs on most heavily used rail transit lines. *All* days on most transit lines, such as the Piccadilly Line, will experience some degree of disruption, which in this context may be described as divergence from the operations plan. Carrel (2009) defines a disruption as "a single, unforeseen event that causes one or more trains to be unable to complete their trip as scheduled." To be as precise as possible, two separate terms will be used in this document in referring to delay on the transit line. The singular "event" that Carrel describes will be referred to as an **incident**. The primary emphasis of this document, however, is on the larger impacts of that event, particularly on customers. Thus, when describing the cumulative problems that result from the incident, the word **disruption** is used. In this case, the definition of disruptions is extended to any circumstances on the transit line that result in a *customer* being unable to complete his or her trip as

planned.[1] It is almost always possible to look at data representing a variety of aspects of service quality on any individual day and identify times and places where service was less effective than it "should" have been.

Some disruptions, without a doubt, matter more than others. A train arriving one minute later than scheduled may not be noticed by any customer, but a train arriving fifteen minutes late will impact many customers. The second situation would produce significantly longer journey times for travelers and likely require a reworking of crew schedules, while the first may not even be apparent in a statistical method that allows for normal variability in operations. LU currently uses an incident identification and tracking system called the Contract Performance Information Database (CuPID) to record the effects of incidents on the line.[2]

The principal measure used by TfL to characterize disruption severity is the Lost Customer Hours (LCH) metric.[3] After a service manager reports an incident of more than two minutes in duration, staff use a set of estimates of passenger flow based on the incident type; the specific time (15-minute intervals) and date it occurred; and the affected section of the line. These are used to estimate how many riders were inconvenienced by the problem, and for how long. The LCH figure includes a calculation of estimated "knock-on" effects, or delays that last longer than the incident itself, though it is not adjusted based on controller actions. The LCH figure may also be multiplied by a value of time (VoT) in UK Pounds to determine the "social" impact of an incident and incorporate other measures that also affect the experience of passengers on the line. If an incident results in a passenger having to stand instead of sit as he or she normally would, the LCH reflects this, even if the travel time of that pas-

---

1   Why differentiate between trains and customers? The evidence suggests that problems affecting train service and customer service may occur at different times. As such, what is a disruption to the train service may not be one for customers if other trains fill the gap of trains that were supposed to be completing trips then. In other words, if trains are off schedule but nonetheless maintain even headways, passenger journey times would remain normal but crew would have to be reassigned to deal with changed train arrival times.

2   CuPID was originally developed to monitor the performance of the public-private partnership (PPP) entities that maintained the LU network between 2000 and 2010. Though the PPP system has been replaced with an entirely in-house system, CuPID continues to be used to monitor incidents.

3   Lost Customer Hours are known internally as Nominally Accumulated Customer Hours (NaCHs). The unit of passenger impact is one NAX, which equals 100 LCHs.

senger does not change. Other categories of inconvenience include having to take stairs instead of an escalator or elevator.

The LCH measure is thus a representation of not only the total increased journey time customers are presumed to have experienced because of the disruption, but also the increased inconvenience they encounter—even if their travel times are the same. Two incidents with the same effects on train service could result in very different LCH scores, since an incident occurring at rush hour will affect far more people than one occurring in the late evening, for example. In addition, the LCH score for a signal failure is different than that for a train failure, for instance, because they have different estimated "knock-on" effects. If trains begin to queue behind a disabled train—producing secondary effects—the delays that they cause are in theory calculated directly into the LCH score for the primary disruption, and are not logged separately in the CuPID system.

LCH estimates are recalibrated roughly every decade to reflect current demand and based on the LU Train Service Model (TSM). The most recent update was completed in 2011 and is referred to as the "2014" sample (the previous version was modeled in 1999 for "2006," and remained in effect until 2011) (TfL, 2011b). Disruptions caused by planned maintenance work are evaluated separately and are not considered in this analysis.

Each incident and its associated LCH score are logged in the CuPID database. The categorization system used by London Underground to record disrupted events features more than 300 categories, or "cause codes." As shown in Table 4-1 and Figure 4-1, London Underground staff currently rely on a comprehensive set of disruption categories in assigning delays to specific causes. The specificity of this data makes it possible to focus on the most significant causes of problems on the line. There are ten broad groups (on the left side of Figure 4-1). For each category a number of sub-categories and sub-sub-categories are defined. Further detail is provided for two example categories, Customers and Signals.

Table 4-1: TfL disruption cause code categories. (Source: TfL, 2011b)

Principal grouping	Description
Customers & Public	Setting off alarms (malicious or spurious), robbery/theft, accidental track obstruction, suicide attempt, trespass on railway, vandalism, drunks, lost children, crowding, altercations/fights between customers, customers holding doors open, customer injury, boarding/alighting trains, person ill on train.
External Factors	Incidents caused by external parties or factors outside LU⊠s control, including for example bridge strikes, fire in premises adjacent to the railway, local power supply failure, police request etc.
Fleet	Trains defective in service or which are not available to enter service from the depot when required.
Safety & Security	Fire and security alerts (e.g. unattended bags, fire alarms set off) station evacuations including exercises. Note that defects of Fire/Safety equipment are included under Stations Infrastructure.
Signals	Failures of signal equipment, control systems for signals and points, tunnel telephones, and tripcock testers, including overruns of works to these assets.
Staff	Absence, shortage, illness, accident, errors, staff taxi delays, refusal to work on health and safety grounds, industrial action.
Stations Infrastructure	Failure of stations equipment (e.g. escalators, lifts, lighting, fire equipment etc), also overrun of stations project works.
Track & Civils	Defects of Track, Bridges and Earth Structures and Drainage, Track fires, Track obstructions and vegetation impacts, overruns of Track/Civils work.
Other	Any other causes not included above.

CuPID is an assessment tool that provides extensive information about the disruptions that occur on the Underground. It is the primary method by which TfL staff estimate the effects of disruptions. Yet there are several ways in which it could be supplemented to improve its accuracy and value. LCH values are based on a number of projections, including how many people are using the line at any time; how long delays will continue after the incident has ended; and what kind of inconveniences riders will experience when service is delayed. In addition, the estimates do not reflect any live information about service control actions, which could play an important role in moderating the effects of an incident, though previous actions are used to estimate the likely knock-on effects of disruptions (Carrel, 2009). Those projections are based on a detailed analysis of past incidents and are likely to come close to reflecting actual customer effects of disruptions.
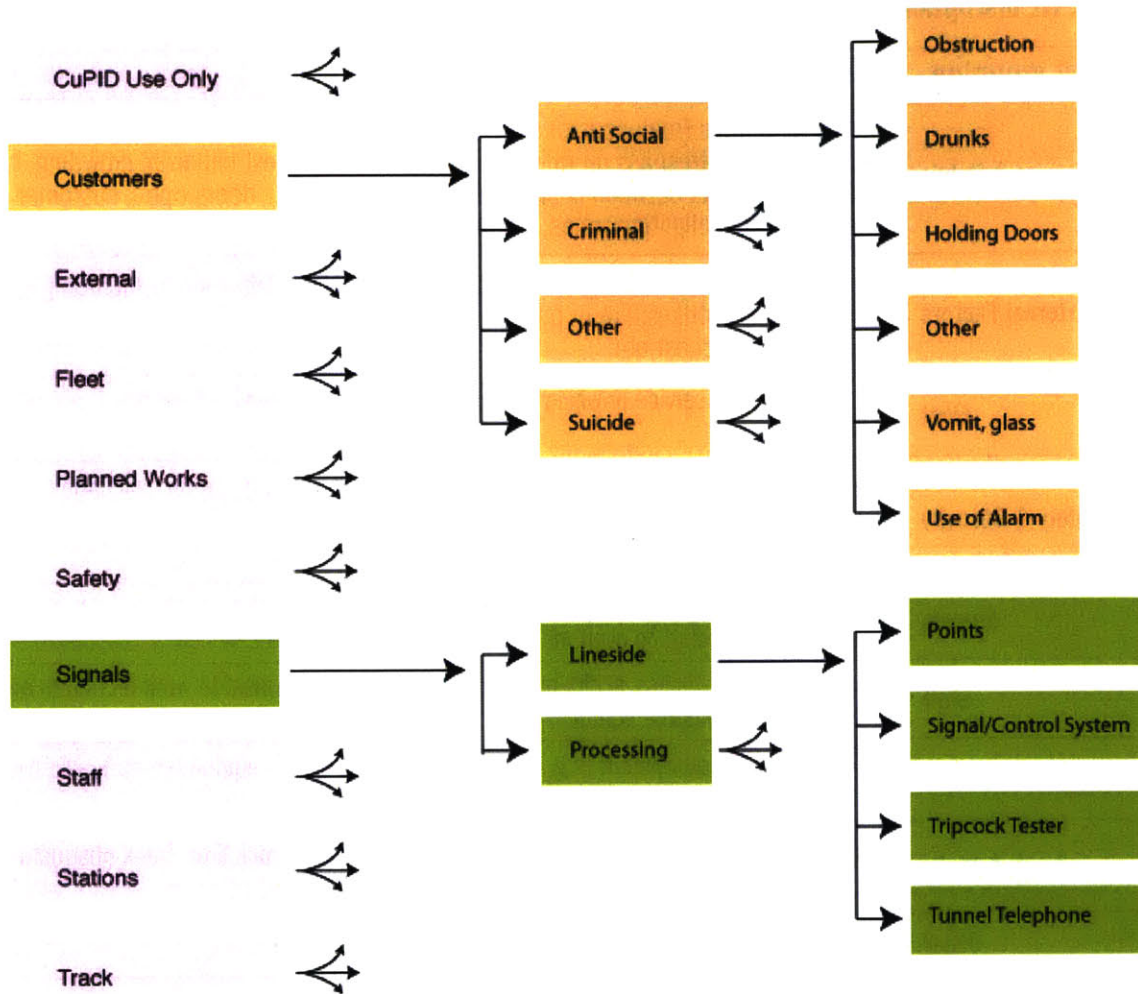
Figure 4-1: London Underground's disruption categorization system.

But TfL does have additional information available that can improve our understanding of the effects of incidents on passengers. This chapter evaluates the information recorded by line service managers to aid in the interpretation of disruption effects. By examining a period of time and classifying incidents in terms of their effects on operations and customers, the information currently provided through the CuPID system can be augmented. Chapter 5 extends this process further, describing how Oyster AFC data can be used to offer additional insight into disruption effects by examining actual variation in passenger journey times.

4.2 SELECTING CASE STUDY SAMPLE DAYS

A three-month period, from April to June 2012, was selected for analysis of incidents. It was limited to three months of retrospective analysis because of limited data availability and quality for previous months. With the exception of the Queen's Jubilee weekend (June 2-5), these months experienced no major abnormalities compared to typical years. These months had normal temperatures and precipitation; by selecting from days in the spring, the best-possible service environment was available for study and the number of incidents caused by issues external to the transit system was minimized. Ridership on the Transport for London network grew during this period compared to previous years, but the system's passenger load has been growing for many years, as described in Chapter 3, and 2012 proved to be no exception.

Perhaps most importantly, the Piccadilly Line did not undergo any major upgrades or renovations during the study period. No sections of the line were shut for any extended period; train availability was normal; and staffing was regular. As a result, the line's operations were not affected by any abnormal service patterns. Thus this period provides a reasonable basis to draw conclusions about how incidents occur on the line in a generally "normal" service environment. From the start, it is worth noting that this study makes no attempt to identify *all* incidents, nor does it claim to provide a universal accounting of the workings of the transit line.

To narrow the focus to the most disrupted situations, not all days within the study period were analyzed. London Underground has a general policy of determining the degree of disruption on a service over a full day. This is measured by the Excess Platform Wait Time (PWT), which attempts to calculate the degree of extra wait on the platform (as compared to schedule) each passenger experienced, in seconds. PWT is roughly, though not exactly, inversely related to the London Underground's headway proxy score, which measures the percent of scheduled headways achieved. The PWT is calculated as shown in Equation 4-1.

$$Excess\ \overline{PWT} = \left\{\frac{\Sigma_i H_i^2}{2 \cdot \Sigma_i H_i}\right\} - \left\{\frac{\Sigma_i E_i^2}{2 \cdot \Sigma_i E_i}\right\}$$

(4-1)

Where,

PWT is the average passenger wait time;

H is the observed individual vehicle headway during a specific period; and

E is the expected individual vehicle headway based on scheduled service during the same period.

Adapted from Uniman (2009), 54.

The headway proxy calculation is based on Equation 4-2.

$$Headway\ Proxy\ \% = \frac{W - M}{W}$$

(4-2)

Where,

W is the scheduled count of headways during a particular period; and

M is the count of missed headways during that period.

According to TfL, "If a gap [in arriving trains] is between 2 and 3 times the planned headway, one headway is counted as missed; if between 3 and 4 times two headways are counted as missed," and so on. (Source: TfL Heartbeat)

These figures are derived using data from NetMIS, which records the arrival and departure of trains from select stations and thus can be used to calculate headways and frequencies. These data are compared to the expected, or scheduled, service. PWT and Headway Proxy can be defined for a period of any length, but the scores are generally used for an entire day's worth of service.

Based on interviews with London Underground personnel, the PWT provides a better description of service quality, generalized across a day, than the Headways Proxy Score. Operators use PWT to examine performance on the line retrospectively and consider the figure as the single most important indicator of the status of the line. Because of different sampling protocols (sampling points, for instance, may not be on all sections on the line), PWT data is not consistent across lines; a "good" score on one line might be mediocre on another. In general, however, scores under 20 seconds are normal and under 10 seconds are excellent. Based on interviews with staff, the Piccadilly Line's operators use a standard of 26 seconds as the nominal target. Days with higher PWT values are considered disrupted.

Twenty-nine weekdays were selected as the case study sample days over the three-month study period with PWT values of 20 seconds or above (there were 65 total weekdays during the period). This allowed the study to focus on the top half of the distribution of disrupted days. Of these days, 18 (62%)

had PWT scores of 26 seconds or above, making them definitively "poor" service days according to the service managers. Table 4-2 summarizes the basic statistics for these study days, and Figure 4-2 illustrates the corresponding PWT scores.

Table 4-2: Basic statistics for selected study days.

Median Excess PWT	30 seconds
Maximum Excess PWT	111 seconds (April 3, 2012)
Median Headway Proxy Score	95.8%
Minimum Headway Proxy Score	80.5% (April 3, 2012)
Maximum Headway Proxy Score	97.6% (June 12, 2012)
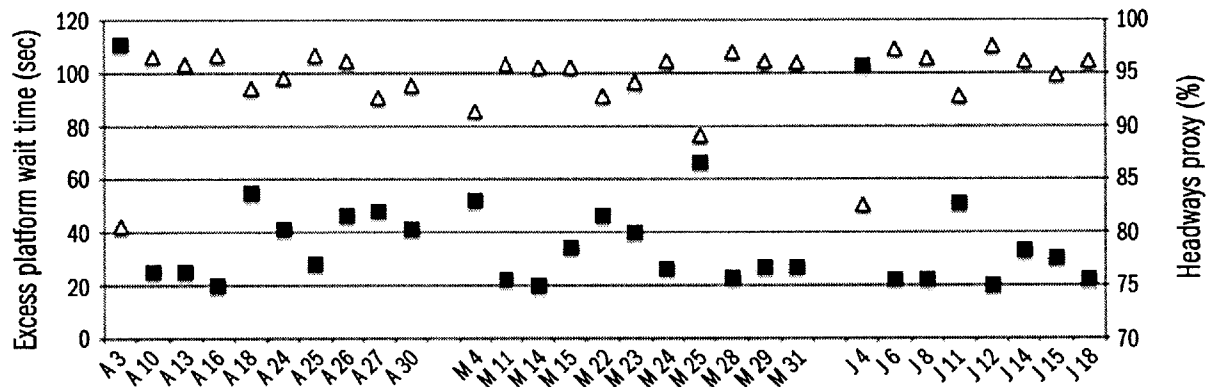
A. Ordered by date.
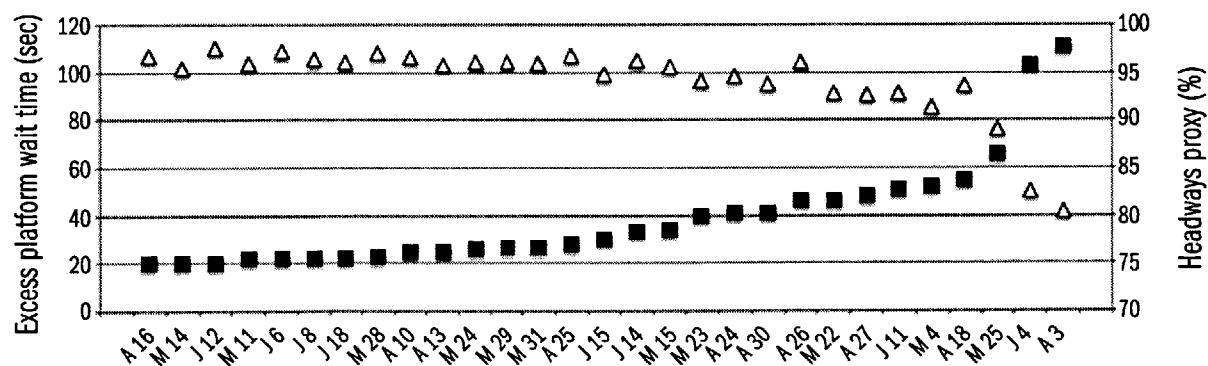
B. Ordered by excess platform wait time score.

■ Excess platform wait time △ Headways proxy

Figure 4-2: Excess platform wait time and headway proxy percentage on study days.

As Table 4-2 and Figure 4-2 show, the case study days performed poorly compared to a "normal" day according to their PWT scores. However, the Headway Proxy Scores provide somewhat different information, in some cases showing quite good performance (as in the case of June 12). As noted, however, service managers on the line do not usually rely on the Headway Proxy Score as an overall measure of performance, so it is not discussed further here and the selection of days was made using the PWT alone.

However, the PWT calculation is not representative of any *single* incident. If there are multiple incidents in one day, the PWT fails to show how much of the overall delay was caused by each of those problems. Moreover, the fact that the PWT is computed across an entire day means that it has a tendency to wash out the effects of shorter, but still severe, problems on the line. As such, this study endeavors to analyze these incidents in more detail to ascertain whether there are other indicators that show their effects more clearly. This does not mean that the use of the PWT score should be eliminated, but rather that its results could be supplemented with other information to form a fuller picture of incidents on the line.

4.3 INCIDENTS DURING THE ANALYSIS PERIOD

As noted at the beginning of this chapter, incidents come in all shapes and forms—some minor and some very significant. How can we identify where and when problems occurred? One approach is to use automated data on transit system performance (in the case of London, Oyster and NetMIS) to identify and analyze incidents and one potential method to examine their impacts—the resulting disruptions—is discussed in Chapters 5 and 6. But we can also use the information already provided by service controllers to begin analyzing these incidents.

As described in detail in Chapter 3, the service controllers on the Piccadilly Line write summaries at the conclusion of each duty period that describe what incidents occurred when, and where. These summaries are based on information about incidents noted as they were occurring. Though this data entry is open-ended, these logs generally provide the following information about incidents:

- Time of commencement;
- Time of termination (problem solved, or at least remediated enough for service to resume);
- Duration (usually consistent with the above two figures, sometimes differs by a few minutes once event is reexamined at the end of the day);
- Location, generally expressed in terms of nearest station or switch; and
- Presumed problem.

During the 29 case study days, controllers noted 116 incidents that had disrupted periods following them of at least 5 minutes in duration. These totaled 2,901 minutes of incident duration (of 34,800 total minutes of service during the period). This amounts to an average of about 4 events and 100 minutes of incidents per day on the Piccadilly Line.

Several points should be noted about these data. First, these data overestimate incidents compared to a "normal" day on the Piccadilly Line. That is because the days for which controller notes were reviewed were those in the top half of the incidents distribution, using the PWT metric. Second, these data do not include incidents whose duration was less than 5 minutes. No doubt there were many such cases, but they were less disruptive to passengers and all were very short. Third, the "minutes" recorded per incident are not weighted by passengers affected, unlike the CuPID system used by LU to describe problems, as will be discussed in detail in Section 4.5. These are the actual number of minutes during which an incident was recorded by controllers on the line.

4.4 INCIDENT CHARACTERIZATION

Using the comments by service controllers, the incidents were characterized into 11 basic types. These reflect the controller's understanding of what occurred. These categories were compiled using a qualitative review of the comments, not a formal system, since controllers are not required to input problem characteristics into any specific pre-defined categories.[4]

The eleven categories chosen for this analysis are documented in Table 4-3.

---

4 It should be noted, however, that discussions with staff members at LU suggest that there is room for improvement in this approach. In interviews, some staff argued that it would be more efficient for controllers to be provided a set list of possible types of incidents from which they could choose when writing their reports. This would reduce analysis time and ensure that similar incidents are categorized consistently.

Table 4-3: Service disruption categories.

Category	Description or examples
Customer problem	A sick passenger or a person committing suicide
Station problem*	Escalator malfunctions or a lack of station staff
Train operator issue	Generally resulting from a staff member not being available on time or needing to be relieved from duty
Heavy traffic	Resulting from higher than expected passenger loads, which may then mean blocked doors and boarding or alighting delays
Signal or track failure	A mechanical failure of the signal or track system
Points failure	A mechanical failure of a track point, where tracks join
SPAD	The acronym for a signal passed at danger, or a train running a red light
Trip	In which a change in operations plans requires a train to be turned around before originally planned
Fire	A fire on the tracks often caused by trash
Train problem	A mechanical problem with the train itself
Other	

This category is not likely to affect train service. Nonetheless, station problems can be associated with significant recorded incidents by controllers and thus are included in this analysis.

These categories are similar to, but not identical, to the cause codes used by TfL to characterize incidents, as described in Section 4.1. The causes described here were determined based solely on controller comments, not on later analysis. This method was used to emphasize the information *known by the service manager* at the time of the incident, or just after, rather than in retrospective review. A comparison with results provided by CuPID is provided in Section 4.5.

The categories identified in Table 4-3 were plotted by time and date, as shown in Figure 4-3. This plot shows no clear pattern of when incidents occurred. Incidents are scattered throughout the day, with no period of the day appearing to have significantly more incidents than others. It should be noted, however, that this chart does show a somewhat lower number of incidents between 8:00 and 10:00 and 20:00 and 22:30, compared to other periods of the day.
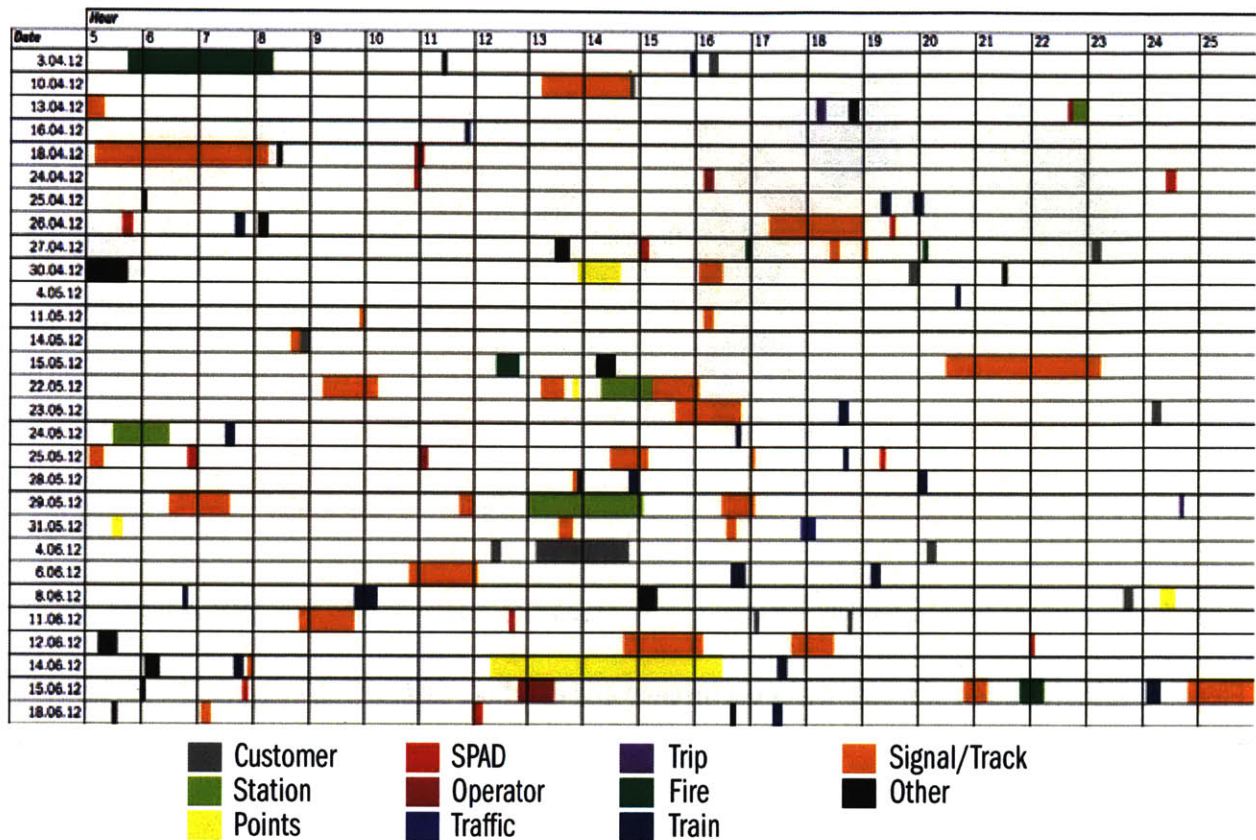
Figure 4-3: Distribution of incidents by time and date, categorized by type.

Different types of incidents had significantly different impacts on passenger service. In total, as Figure 4-4 shows, signal malfunctions were associated with 28% of all incidents and 45% of the total incident duration time. Signal failure was the most common type of incident on the Piccadilly Line during the study period, according to service managers. Other issues, such as points failures, station problems, customer problems, fires, and train problems all contributed to more than 5% of total incident duration time on the line.

Figure 4.44 shows that the distribution of duration by incident type varied tremendously. Train malfunctions were all resolved within 23 minutes and on average lasted only 9 minutes. On the other hand, signal problems lasted up to 187 minutes and took 41 minutes to resolve on average. These data indicate that a large percentage of the signal problems on the Piccadilly Line result in long periods of service delays. Other incident categories are less common, and if they do, often have shorter durations.
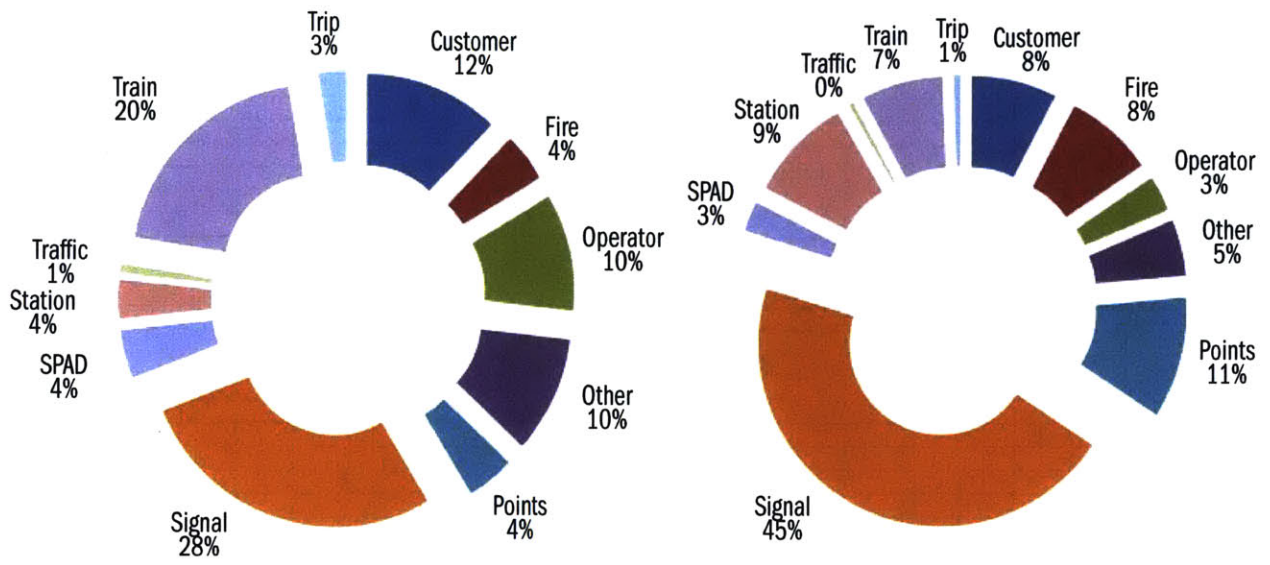
A. Incident occurrence (of 116 total incidents).

B. Incident duration (of 2,901 total minutes).

Figure 4-4: Distribution of incidents by occurrence and duration, categorized by type.
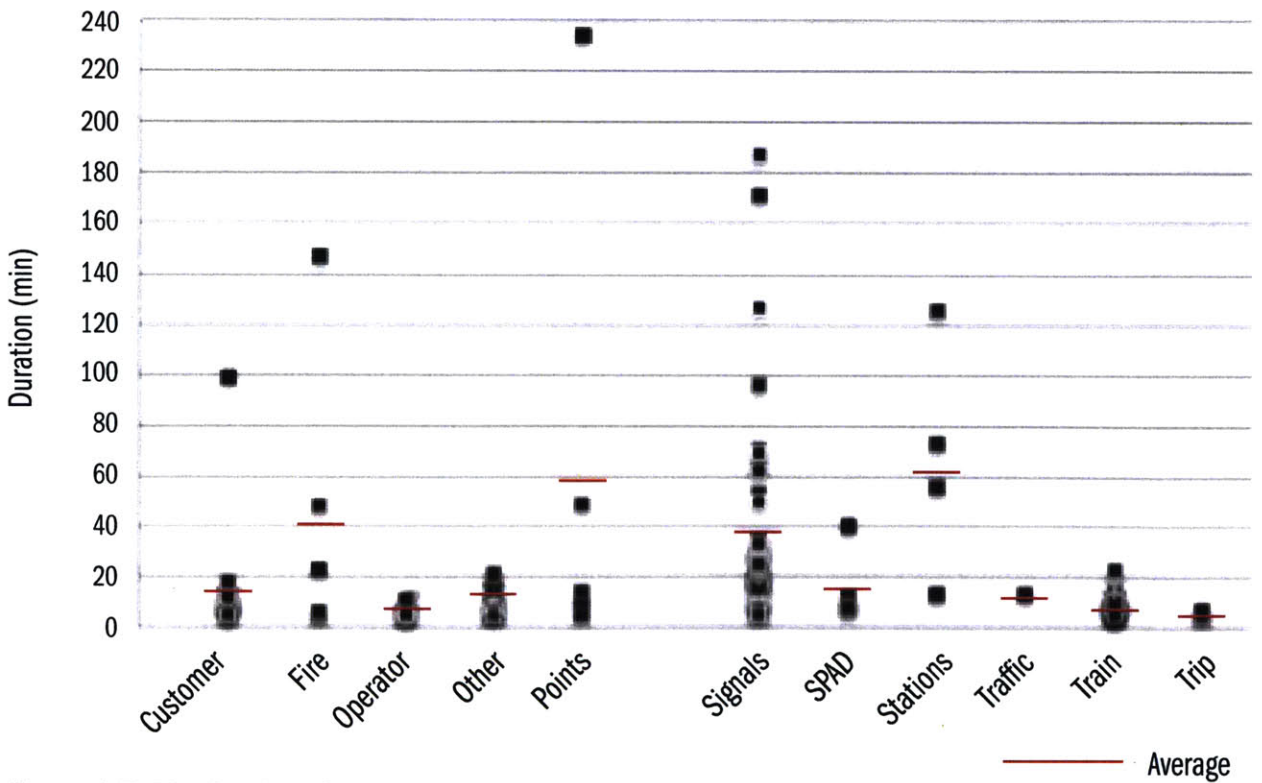
Figure 4-5: Distribution of incident duration, categorized by type.

Incidents are distributed across the day on the Piccadilly Line as indicated in Figure 4-3, but total incident duration is concentrated outside the peak periods, as shown in Table 4-4. The periods in the table were defined by identifying time periods of most concentrated passenger travel. This is good news from the perspective of delaying smaller numbers of riders, as ridership is obviously heaviest in the AM and PM peak periods.

Table 4-4: Operational periods and incidents for weekdays.

Period	Time	Percent of total daily ridership*	Hourly percentage of total daily ridership in period (flow)	Percent of total disrupted time
Early Morning	5:00-7:29	3.0	1.2	20.7
AM Peak	7:30-9:29	17.7	8.9	5.2
Midday	9:30-15:29	30.1	5.0	42.7
PM Peak	15:30-18:59	26.7	7.6	12.7
Evening	19:00-24:30	22.6	4.1	18.8

Defined as the percentage of trips occurring per period at the mid-point of the trip. For example, if a customer taps his or her card at 7:15 and the trip takes one hour, the trip is counted as occurring during the AM Peak because its mid-point, 7:45, is during the AM Peak. Based on an analysis of Oyster entry data for trips on the Piccadilly Line on three weekdays in April. This figure is not a complete reflection of ridership since it only includes trips that begin and end on the Piccadilly Line and only accounts for riders using Oyster cards (who nonetheless account for the vast majority of ridership).

As shown in Figure 4-6, when split into the five periods of analysis, the highest levels of incident duration (total incident duration over the course of the 29 study days) were in the early morning and midday periods. On the study days, there was on average about 8 minutes of incident duration per hour in the early morning period and about 7 minutes of incident duration per hour during the midday period. There were far fewer incidents and far shorter disrupted periods per hour during the AM and PM peak periods. Figure 4-6 also shows the average duration time for the individual incidents that occurred in each period. This chart indicates that incidents were resolved most quickly—or were most minor—in the AM and PM peak periods, compared with the early morning and midday periods.

Total disruption duration   □ Averaged by hour

A. Incident duration.   *B. Average incident duration.*
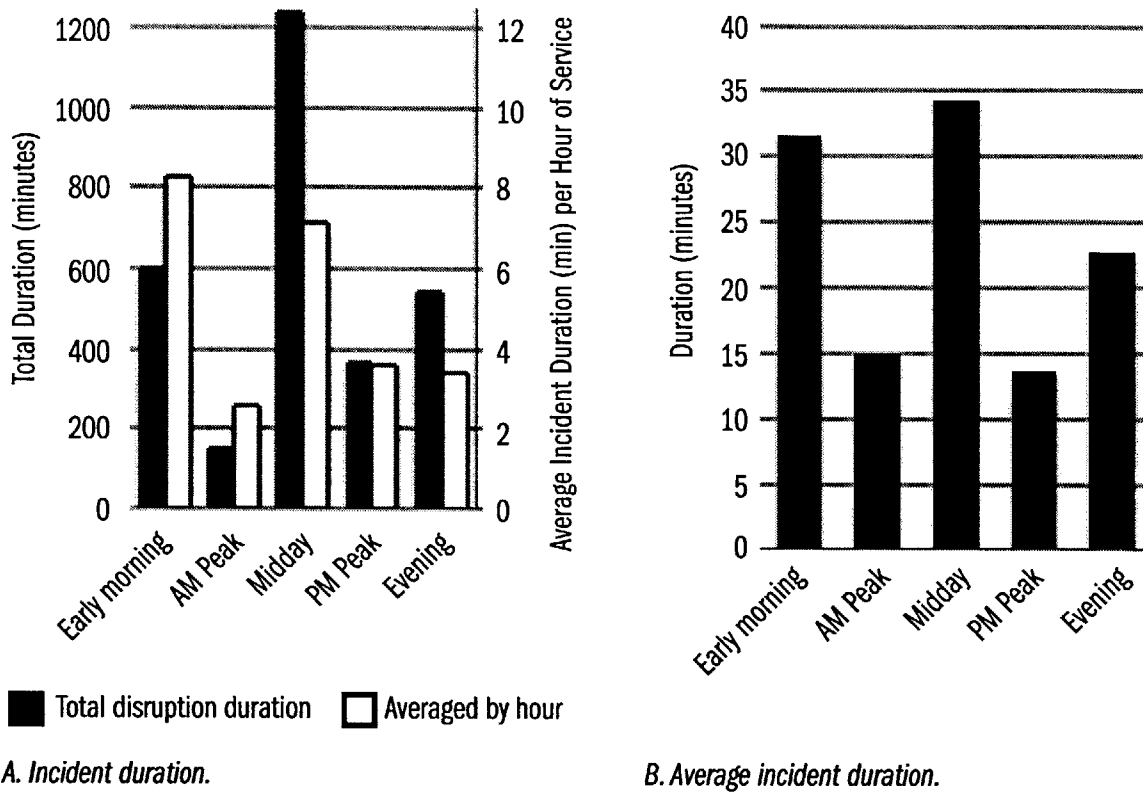
Figure 4-6: Incident duration by time period.

As shown in Table 4-4, passenger loads are concentrated in the AM Peak, midday, and PM peak periods. As a result, despite the fact that a significant percentage of the incident duration occurred in the early morning, these incidents affected relatively few passengers. In fact, despite the lower average incident duration during the AM and PM peaks these incidents likely affected more riders. On the other hand, the high level of incident duration during the midday, when there are many riders using the system, affected many more people.

The incident periods, aggregated over the entire 29-day study period, were concentrated in the Piccadilly Line's North section (north of King's Cross) and the Trunk section (from Acton Town to King's Cross), as shown in Figure 4-7. Data for incidents on each of the branches include incidents that occur at the corresponding stations. For instance, if a problem occurs at Arnos Grove, the incident delay that results from it is included in the total North Branch incident duration period. This does not apply

to Rayners Lane and Acton Town, which both accounted for significant incident durations but are at the confluence of two or more branches (Uxbridge and Rayners Lane branches for the first; Heathrow, Rayners Lane, and Trunk for the second). King's Cross-St. Pancras station, also at the confluence of two branches, did not account for any specific incidents.
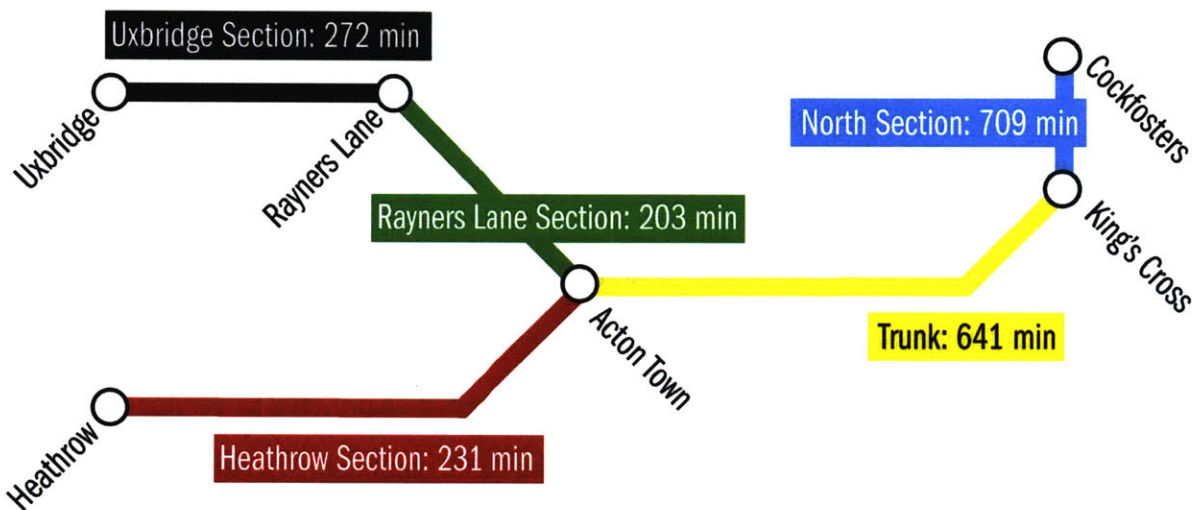
Figure 4-7: Spatial distribution of incident duration across the Piccadilly Line.

When computed by station, as shown in Table 4-5, Acton Town, Ealing Common, Arnos Grove, Uxbridge, and Rayners Lane were the most disrupted locations in terms of total incident duration. This is likely a consequence of the fact that these stations are either crew bases (Acton Town and Arnos Grove), or train turn-around points (the other three stations). Though certain types of incidents can occur any-where in the system (such as a sick passenger), other problems, such as signal or points failure, are most likely to occur where trains switch tracks, which occur at crew bases and at train reversal loca-tions.

Table 4-5: Top stations for cumulative incident duration.*

Station	Cumulative disruption duration (total minutes)
Acton Town	333
Ealing Common	270
Arnos Grove	263
Uxbridge	234
Rayners Land	203
Leicester Square	159
South Kensington	156
South Harrow	114
Cockfosters	114

* All stations not included saw less than 100 minutes in cumulative incident duration during the study period.

4.5 COMPARISON WITH THE CUPID DISRUPTIONS DATABASE

The manner in which the incidents were characterized in Section 4.4 is not completely consistent with LU's current approach. As shown previously in Table 4-6, information provided by service controllers is not always consistent with the information included in the CuPID database. Of a sample of 20 incidents for which both controller and CuPID information is available, the service controller reported different types of incidents compared with those reported by CuPID twice (these cases are highlighted). This is a small proportion of the total, indicating that generally, service controllers have a good sense of the type of incident when it occurs. However, in a small percentage of situations, later investigations show that the understanding that service controllers had about the incident was not completely accurate.

As shown in Table 4-6, incident duration in minutes is distinct from lost customer hours (LCHs), since the LCH figure is calculated by multiplying the incident duration by the estimated number of passengers using the affected part of the line and by the estimated knock-on effects of that type of incident. The rank of the incident durations and LCH figures noted in Table 4-6 demonstrates the lack of direct correlation between these two figures. But these estimates may not reflect *actual* passenger impacts, since the number of riders inconvenienced and the length of the disruption have not been recorded using passenger data. Nor do they reflect different types of controller response strategies, which may affect the length and severity of disruptions significantly. Chapter 5 explores the possibility

Table 4-6: Comparing service controller information with CuPID data.

Date	Time	Incident duration (min), with rank	Service controller cause	CuPID cause	LCH (hours), with rank
13.04.12	5:04	18 (8)	Signal	Signals-Lineside	195 (13)
16.04.12	11:49	6 (15)	Train	Fleet-Defective in Service-Traction	557 (10)
18.04.12	5:10	171 (1)	Signal	Signals-Lineside-Points	2175 (6)
24.04.12	10:56	7 (14)	Operator	Staff-Error-Staff Not in Position	248 (11)
26.04.12	5:40	10 (12)	Operator	Staff-Absence or Shortage	73 (15)
4.05.12	20:39	5 (19)	Train	Fleet-Defective in Service-Batteries	240 (12)
22.05.12	9:14	28 (6)	Signal	Signals-Lineside	67 (17)
24.05.12	5:32	56 (5)	Station	Staff-Absence or Shortage-Stations Staff	23 (20)
25.05.12	5:05	16 (9)	Signal	Signals-Lineside	27 (19)
4.06.12	12:19	10 (12)	Customer	Customers & Public-Criminal	2675 (4)
6.06.12	10:53	71 (3)	Signal	Signals-Lineside	4176 (2)
8.06.12	6:46	6 (15)	Train	Fleet-Defective in Service-Traction	69 (16)
11.06.12	8:47	62 (4)	Signal	Signals-Lineside	14573 (1)
15.06.12	7:49	6 (15)	Operator	Staff-Error-Failure to Shut Rear Cab	1193 (7)
15.06.12	21:52	23 (7)	Fire	Track & Civils-Track-Fire	2529 (5)
15.06.12	0:51	127 (2)	Signal	Signals-Lineside	156 (14)
18.06.12	5:30	6 (15)	Other	Signals-Lineside	66 (18)
18.06.12	7:03	11 (10)	Signal	Signals-Lineside	3262 (3)
18.06.12	11:57	11 (10)	Operator	Staff-Working/Error-Manager Error	789 (9)
18.06.12	16:39	5 (19)	Customer	Fleet-Defective in Service-Saloon Equipment	859 (8)

of using Oyster data to analyze those impacts directly. With increased use of Oyster data to demonstrate actual passenger impacts of a line disruption, the LCH estimates may be significantly improved.

5 | Disruption impact assessment

As described in Chapter 4, TfL maintains an extensive database of incidents categorized by type in the CuPID system. This database includes information about the severity of disruptions associated with incidents, as estimated in terms of LCHs, which are estimates of passenger delay and inconvenience. These methods, however, do not use Oyster data to examine the actual experience of riders in the time following any incident on the line. The goal of this chapter is to use Oyster data to provide a new method to assess the actual impact of disruptions on passengers, with the intention of providing supplementary information to that already provided by CuPID. This chapter proposes one method by which variations in journey times recorded by Oyster could be used retrospectively to examine the effects of disruptions on customers, and it describes how passenger accumulation can be monitored with Oyster tap data.

The chapter begins with a review of TfL's existing methods to analyze the impacts of disruptions in Section 5.1. It then describes how an index of disruption based on journey time variation can be constructed in Section 5.2 and offers a series of case studies on the Piccadilly Line to show how that index might be implemented to measure the passenger impact of a disruption in Section 5.3. Section 5.4 examines how Oyster-based information about passenger accumulation could be used to identify periods of disruption on the line. Section 5.5 compares the Oyster-based methods with existing mea-

sures of disruption severity. Finally, Section 5.6 explores the potential use of Oyster journey time data during disruptions to improve the agency's understanding of customer route choice.

5.1 CURRENT APPROACHES TO ANALYZING DISRUPTION IMPACTS

Service analysts at Transport for London currently use information provided by the agency's smartcard system, Oyster, to understand customer journey patterns on the Underground. By examining records of Oyster taps—both into and out of the system—the performance of the system can be analyzed. Specifically, the time it takes a passenger to move from one station to another can be computed. When there is a disruption, Oyster provides one means, at least for ex post analysis, to examine how customers were affected.

Figure 5-1 shows how journey times can change significantly during disruptions. This chart shows the journey times for all Underground journeys between Knightsbridge and Covent Garden on September 22, 2011. The vast majority of, if not all, of these journeys, can be assumed to have been made on the eastbound Piccadilly Line. The results show that journey times between 6:00 and 16:30 were stable in a range of 10-20 minutes, but by 18:00 travel times had increased to a range of 30-45 minutes. The effects of the disruption are very clear.
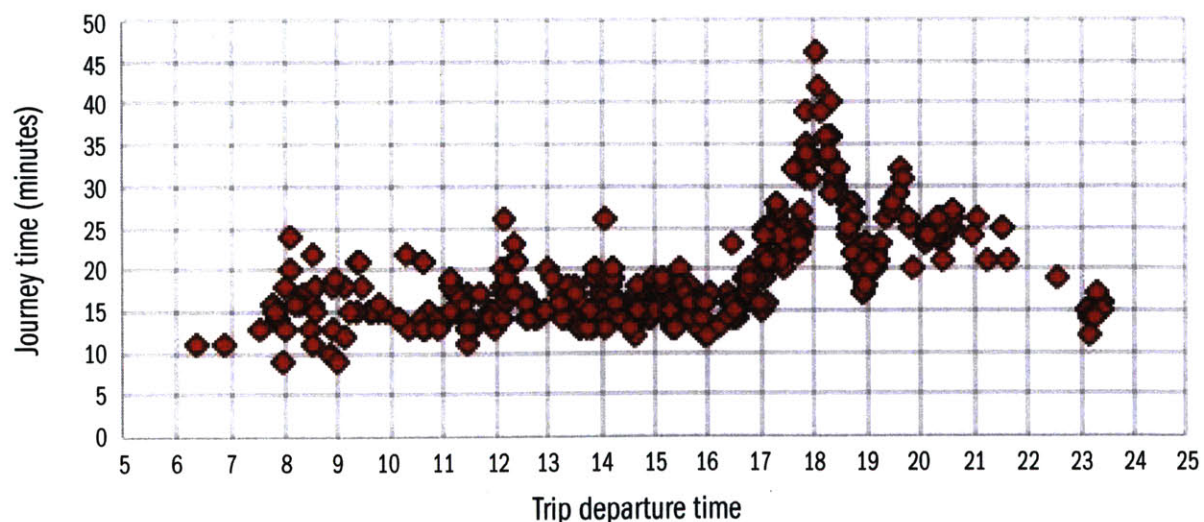
Figure 5-1: Example of Oyster journey times from Knightsbridge to Covent Garden.

Figure 5-2 puts the disruptions experienced on September 22 in the context of travel in the PM peak by passengers on other days in the same week. The mode of the distribution of journey time for customers traveling from Covent Garden to Finsbury Park (eastbound) was far higher than for any of the other days of the week, which recorded normal journey times and did not experience any disruptions. The riders from Covent Garden to Finsbury Park were essentially continuing the path of those who had taken the train from Knightsbridge to Covent Garden, as shown in Figure 5-1, so this distribution is not unexpected.

These increased journey times, however, did not apply to the entire Piccadilly Line, as is also illustrated in Figure 5-2. Riders from Acton Town to Hounslow West, traveling westbound on the Heathrow Branch, had a generally similar distribution of travel times between those destinations during the PM Peak on September 22, as on other days of that week. This reflects the fact that disruptions resulting from the earlier incident were confined to only parts of the line and one direction of service. Comparative analyses such as these offer an effective retrospective look into the circumstances affecting particular travelers on specific route segments during specific periods.

A. Covent Garden-Finsbury Park   B. Acton Town-Hounslow West

■ September 19   ■ September 21   ■ September 23
■ September 20   ■ September 22
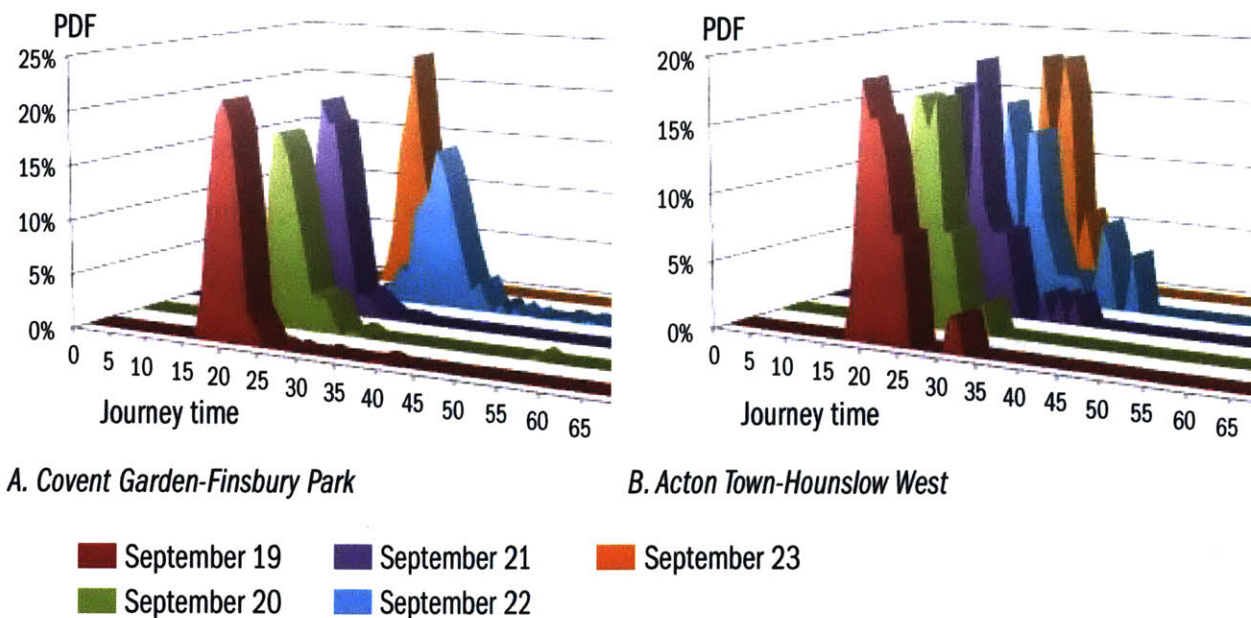
Figure 5-2: Journey time distributions during PM Peak (September 2011).

This analysis of Oyster journeys, however, has several limitations. First, Oyster data for the Underground do not track customers *as* they travel, but rather show only start and end points. On journeys where there is only one option for direct travel (only one line connects two stations, for example), it is reasonable to conclude that virtually all customers chose the direct route. On the other hand, for many journeys through Central London, customers have several options. Customers travelling from St. Paul's Station to Green Park, for example, have three roughly equivalent ways to make the journey, and Oyster offers little insight into which route they took.[1] As such, in terms of considering impacts on a specific line, any analysis of Oyster data must be limited to journeys that can be safely assigned to that route. This eliminates millions of journeys daily from potential evaluation. TfL's Rolling Origin-Destination Survey (RODS) provides one way to address this issue, which will be addressed in Section 5.6.

Second, Oyster data provide journey time from *faregate to faregate*, not actual time riding a train (the latter can be estimated by examining NetMIS track-circuit information). If journeys recorded by farecard data take longer than typical, there are a number of possible explanations. Perhaps a group of tourists heading from one station to another got lost in one of the major stations; perhaps a broken elevator left a person with limited mobility with fewer options; or perhaps the train simply did not show up as expected. Moreover, Oyster journey time data does not clearly differentiate between segments of the journey unless carefully examined. For example, a group of customers who take a significantly longer-than-normal time to travel between Manor House and Northfields may have been affected by a disruption in the northern section of the Piccadilly Line, in the trunk through the center of the city, in the Heathrow Branch, or in all three. Whatever the cause of the delay, Oyster records by themselves are not enough to make a full diagnosis.

Third, while sampling origin-destination pairs may provide some insight into line conditions, examining every origin-destination pair to attempt to identify problems is at best a laborious task. There are 53 stations on the Piccadilly Line and a total of 1,378 potential origin-destination pairs *in each*

---

1   Between those stations, customers can take the Central Line heading west from St. Paul's and connect to the Piccadilly, Victoria, or Jubilee Lines, all of which provide one-transfer service to Green Park.

direction. It would be difficult for any analysis of a single origin-destination pair to represent the overall impact of a disruption. Nor is it clear that examining one origin-destination pair necessarily describes performance in the section of the line being considered, let alone the line as a whole. Moreover, examining Oyster data in terms of individual origin-destination pairs can be effective for heavily traveled routes because of the large number of passengers getting on and off trains throughout the day, but for parts of the system there may be only a few riders and Oyster data can be heavily affected by outliers if the sample size is too small. This is particularly true for periods of the day with low ridership. In sum, this suggests that to use Oyster to help understand disruption impacts retrospectively requires a new method, and that is what is proposed in this chapter.

The limitations of Oyster data underscore the fact that the most valuable type of analysis is one that incorporates different types of data, including Oyster data, train arrivals, and controller reports. The automated fare data by itself cannot fully describe what occurred during a disruption. Indeed, analysis of train movement data available through NetMIS can also provide insight into the state of passenger travel during disruptions.

Figure 5-1 shows that customers took much longer than normal to travel between Knightsbridge and Covent Garden, but it does not provide information about other parts of the line since it is based on Oyster journeys for just that single origin-destination pair. Rather than examine a whole host of potential travel pairs, train travel times can provide a bigger picture of events, as shown in Figure 5-3. They can be tracked live through LU's TrackerNet system (which is accessible in the control room) and recorded in NetMIS. Figure 5-3 confirms the finding from the Oyster origin-destination analysis, showing that there was a significant increase in travel time in the central section of the Piccadilly line on the disrupted day (September 22, 2011, from 18:00 to 19:00). The slope of the lines (each of which represents a train) on the space-time diagram indicates the speed of trains; the faster the train, the flatter the slope. Thus a review of this figure indicates clearly that the travel times between Knightsbridge and King's Cross, and between all the intermediate stations, increased significantly during this period on the disrupted day compared to normal.
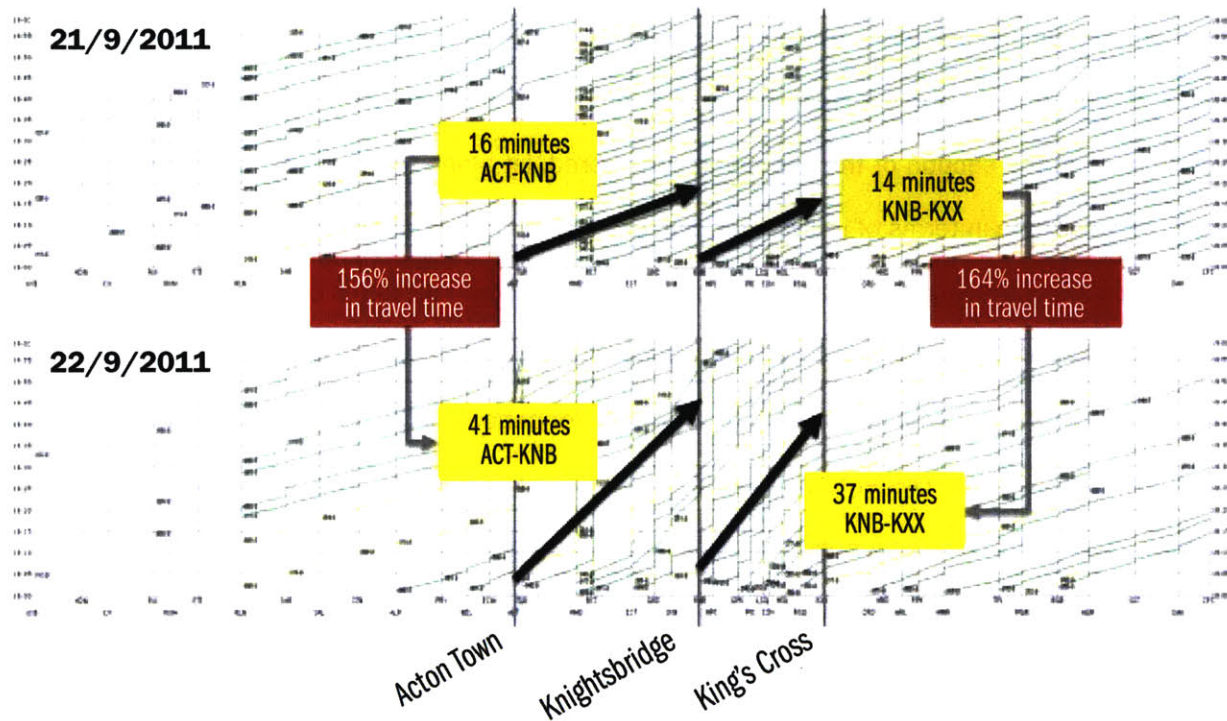
Figure 5-3: Space-time ("waterfall") diagram for eastbound Piccadilly Line trains.

Figure 5-4 demonstrates how NetMIS data provide an indication of changes in train travel time. This figure shows the travel time for all eastbound Piccadilly Line trains running from Hammersmith to King's Cross on September 22, 2011, demonstrating how travel times increased significantly beginning at around 17:30.

NetMIS also provides information about the headways of arriving trains. For example, Figure 5-5 shows how headways at Leicester Square, a representative station, changed over the course of September 22, 2011. These data show increases similar to the travel times data in Figure 5-4, beginning at 18:00, but by 19:00, headways at Leicester Square have largely returned to normal. The discrepancies between the information available from Figures 5-4 and 5-5 provides a clear indication that headways only provide a partial view of the overall impact of a disruption on passengers.

As described in Chapter 3, analysts at London Underground take advantage of this information about train movements to determine the extent of disruptions on the line. However, the reliance on train-tracking information has its limitations. First, while it keeps service managers aware of when
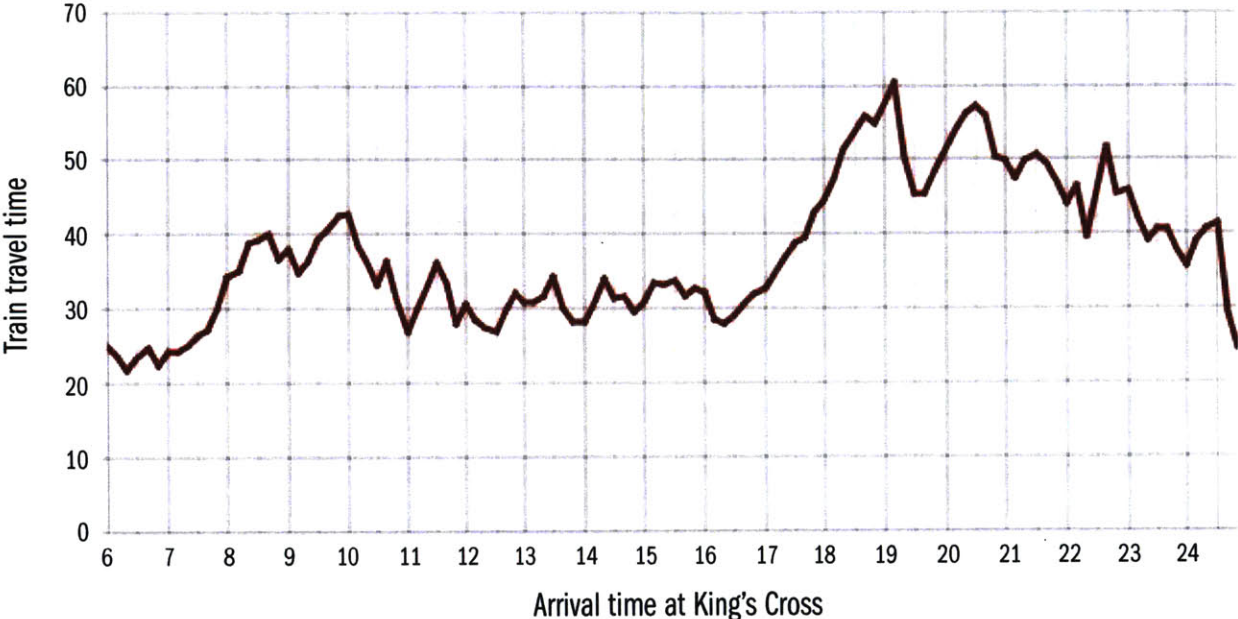
Figure 5-4: Example train travel times graph.
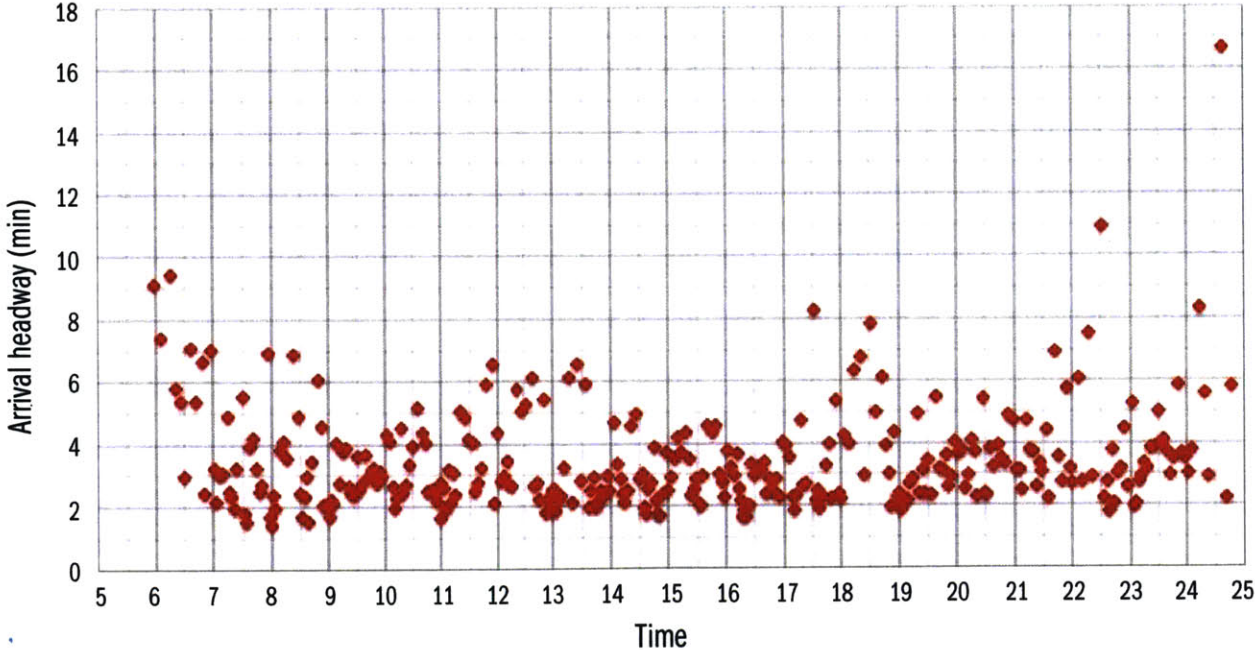
Figure 5-5: Example arrival headways.

trains arrive (compared with their schedule), it provides no direct information on customer impacts. While it may be intuitive to posit that increased train headways result in increased passenger delays, a doubling of headway in a heavily served section of the line may have little effect on the journey time of the average passenger using the line. As such, even when trains are quite delayed, a disrupted situation may have been "resolved" in terms of passenger journey impacts. This is particularly problematic because the LU's primary instrument for measuring disruption impact, the LCH metric, relies on average headways over specific period to be calculated.[2]

Transport for London is developing a new system to describe disruption impacts with the goal of determining the recovery time and recovery rate for incidents. Agency staff have defined recovery rate as the amount of time it takes trains to return to normal service after an incident has been resolved. Analysts have calculated Excess Platform Wait Time (PWT) figures in 15-minute increments and then tracked how long it takes for service to return to normal (or the PWT proxy target), using a 7-period moving average of the PWT (7 15-minute periods). This method is illustrated in Figure 5-6. It shows the PWT figure for each 15-minute period as a blue bar and the moving PWT figure (which averages the seven previous PWT scores, or 1h45 of service), as the red line. The method calculates the time between the end of the incident (shown in red at the bottom of the graph) and the end of the recovery, when the 7-point moving PWT average falls below the daily target. This method makes it possible to compare different types of disruptions and the recovery rate.

This approach is promising, but it also relies on the NetMIS data to determine the disruption impacts. This makes sense when the primary goal is to analyze how to get trains back to the operations plan as quickly as possible. Yet it may be is less effective in terms of determining the true effects on passengers. For that type of analysis, it is necessary to compare train recovery data jointly with passenger journey time information. This chapter posits that it is feasible to use Oyster data to develop

2   It is true, of course, that a service with longer headways (and thus lower frequencies) necessarily results in more crowded trains, assuming constant passenger demand. If trains become so crowded as to force passengers to wait for the next train to board, Oyster travel time variations will reflect the problem, but if trains are simply *more* crowded than normal, Oyster provides no information. This demonstrates one of the significant benefits of using a system such as LCH and the JTM, whose premises are that effects such as passenger crowding can be estimated during disruptions.
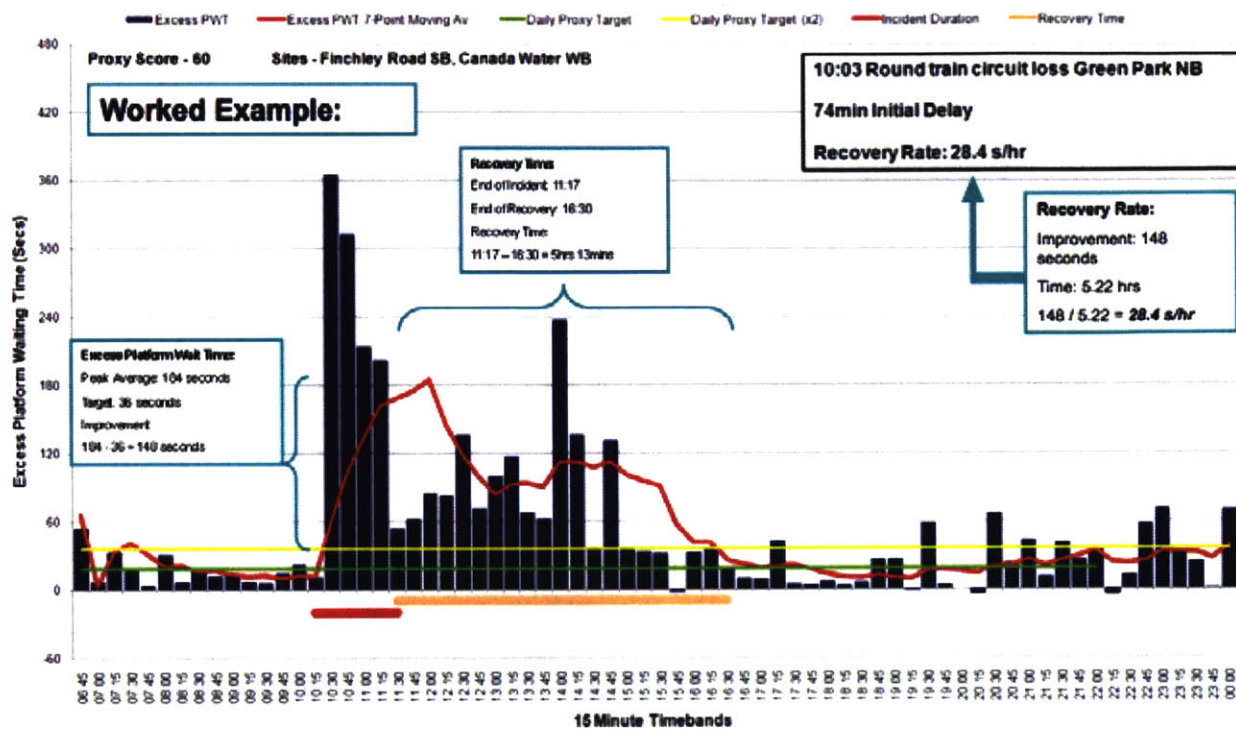
Figure 5-6: Service recovery time and recovery rate chart, showing calculations, for Jubilee Line service on June 6, 2012. Source: Harper (2012)

an aggregate measure of passenger delay over short periods for entire line sections; this method is described in the following section.

5.2 DEVELOPING AN AGGREGATE MEASURE OF JOURNEY TIME VARIATION

As described in Section 5.1, there are significant limitations to the use of individual origin-destination pairs for an analysis of Oyster data. As shown in Figure 5-7, even the most-traveled origin-destination pairs on the London Piccadilly Line saw fewer than 20 individuals in 15-minute periods during much of a sample day (April 24, 2012). An analysis of Oyster records should incorporate a reasonable sample to prevent conclusions from being affected by outliers, but Figure 5-7 shows that this is not possible for short periods of time for individual origin-destination pairs for most of the day.

This research thus proposes to extend the value of the Oyster AFC data source by suggesting a new tool to examine a much larger dataset and combine relevant information in order to assess the impacts of disruptions on passengers. The key insight is to aggregate Oyster smartcard data beyond the
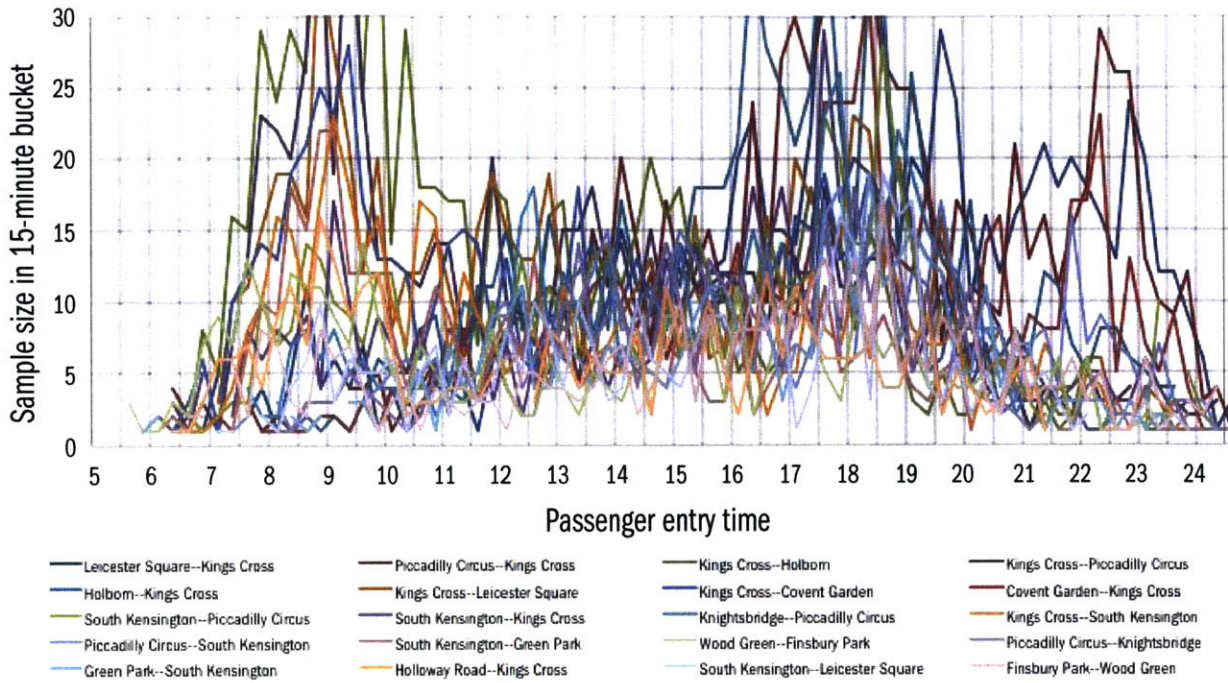
Figure 5-7: Entries by 15-minute bucket on Piccadilly Line top 20 origin-destination pairs.

origin-destination level. This analysis uses a "basket" of journeys across a section of a line to evaluate

the degree to which journey times are varying from the norm. The basket consists of all journeys occur-

ring *within* a specific line section, but excluding journeys that extend beyond the segment of interest.

This method ensures the following:

- A larger sample, even during short time periods, since multiple origin-destination pairs are considered, not just one;
- Reduced sensitivity to outliers through the use of large samples and a focus on the mean variation from the norm; and
- A focus on section-level problems, not difficulties with any specific origin-destination pair.

The first step is the selection of the baskets for meaningful analysis; because the Piccadilly Line is

quite lengthy, it is necessary to divide the line into a number of sections in order to perform an analysis.

The Piccadilly Line was divided into five separate sections, as shown in Figure 5-8. The choice of the

sections was based on two criteria. First, the location of reversing tracks was noted, as shown in this

figure. If sections of the line are to be studied "independently," trains must be able to turn around at

the ends of such sections in the case of a disruption. A reversing track allows controllers to keep ser-

vices running normally on one part of the line even when there are major problems elsewhere. Second, service managers were interviewed in order to determine their typical actions when disruptions occur. Based on those conversations, Acton Town and King's Cross were identified as common locations for train reversal during disruptions. Thus these locations provide natural boundaries for the sections.

The result of this division is illustrated in Figure 5-8 and documented in Tables 5-1 and 5-2. The North section runs from Cockfosters (CFS) to King's Cross (KXX); the Trunk runs from King's Cross to Acton Town; the Heathrow section runs from Acton Town to Heathrow (H123/H4/H5); and the Rayners Lane section runs from Acton Town to Rayners Lane. The section of the line from Rayners Lane to Uxbridge (which we will call here the Uxbridge section) is excluded in this analysis. This section of the corridor is shared with the Metropolitan Line, and so even if the Piccadilly Line suffers significant delays in this area, customers have an easy alternative service available to them. As shown in Table 5-1, the line sections are of similar length; travel time from end to end on each section is between 20 and 32 minutes. Service is provided more frequently for the North and Trunk sections, and significantly less frequently to Rayners Lane.

It should be noted that while these four sections do provide a useful way of understanding performance on the Piccadilly Line, it would be reasonable to increase the number of sections so as to have a finer-grain level of detail about performance in each of the sections. For example, Arnos Grove (AGR) is a crew depot and often the site of train reversals; the same could be said for Northfields (NFS) on the Heathrow Branch. In most cases, increasing the number of sections will decrease the statistical reliability of the results, but as long as there are a reasonable number of origin-destination pairs within the section (to ensure that representative Oyster results can be produced for small time periods), the results will be representative of the conditions within these segments and provide even more accurate information about disruptions than broader section selections. However, smaller line sections are more likely to be influenced by problems occurring at a specific station (such as an escalator malfunction) that would increase *total* journey time but not the time on the train itself, which is what is of most interest here. In a larger sample, such problems would be averaged out.
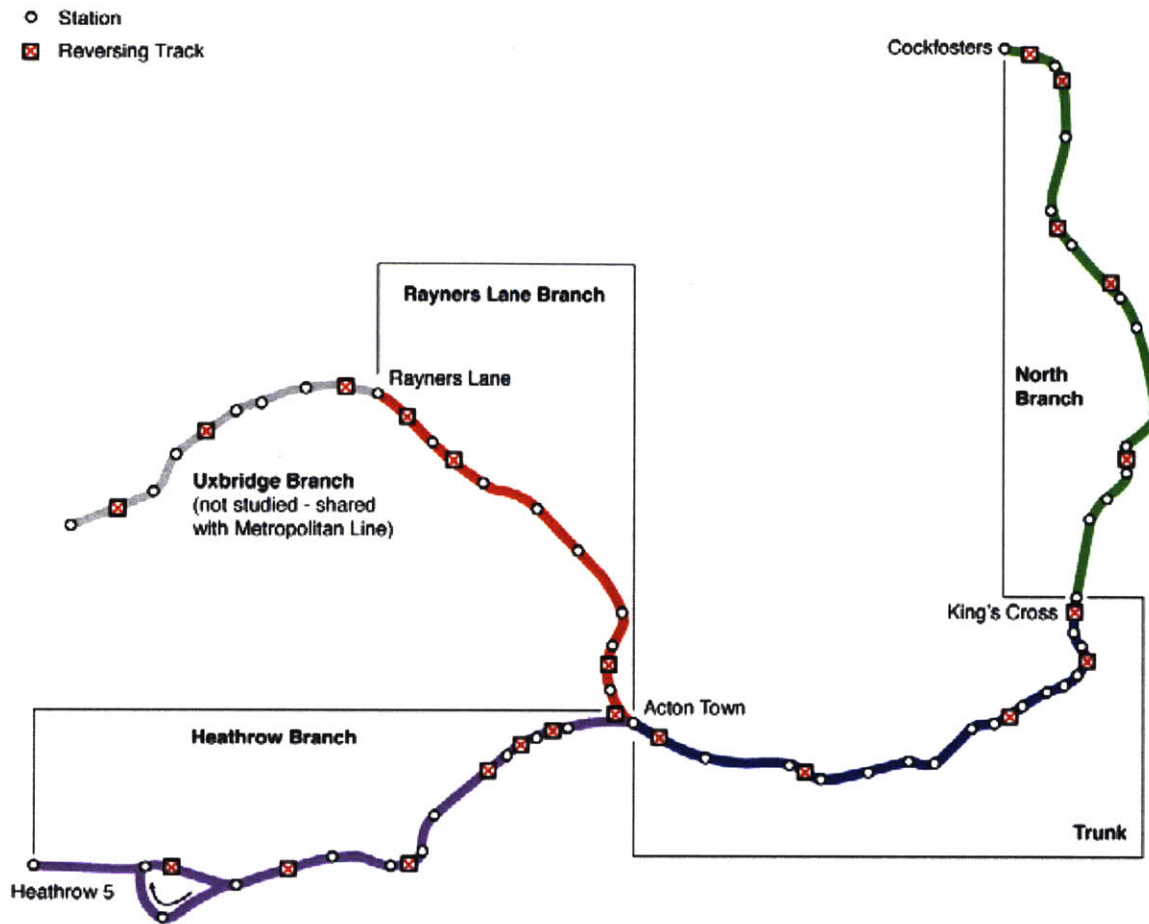
Figure 5-8: Study sections for Piccadilly Line service, showing reversing track locations.

Table 5-1: Piccadilly Line section characteristics.

Section	Length, km (rounded)	End-to-end scheduled train travel time (min)*	Stations (including end points)	Peak trains per hour
North	16	32	13	24 KXX-AGR; 18 AGR-CFS
Trunk	14	32	16	24
Heathrow	19	22	12	18 ACT-NFS; 12-NFS-Heath.
Rayners Lane	11	20	9	6

* Source: TfL Journey Planner, leaving at 8:00 on a weekday morning.

All of the journeys with origins and destinations on the Piccadilly Line within each of the sections were analyzed (about 185,000 trips a day), with several important exceptions. Because the sections are being analyzed independently, trips that span sections (such as from a station in the North section to one in the Trunk) are not considered. Trips where passengers had a direct line alternative for the whole trip (such as between Finsbury Park and Green Park, which is also served by the Victoria Line) were excluded. Trips where only Piccadilly Line service is available directly but easy transfers from other lines are possible (such as from Hammersmith to Leicester Square, where customers may use the District Line from Hammersmith to South Kensington before transferring to the Piccadilly Line) were also excluded. In total, as shown in Table 5-2, the included trips account for about 30% of all daily Oyster trips that begin and end on the Piccadilly Line. Most of the riders considered took trains in the North or Trunk sections, with smaller numbers on the Heathrow and Rayners Lane sections.

We can be relatively confident that these trips represent trips that were taken entirely on the Piccadilly Line and therefore that journey time variation in the Oyster sample is a direct reflection of service on the line. It is possible, however, that certain of these trips occurred at least partially outside the Piccadilly Line if passengers made "irrational" journey decisions. For example, a rider could travel from Green Park to Holborn by taking the Jubilee Line to Oxford Circus and then transfer to the Central

Table 5-2: Piccadilly Line section passenger and OD-pair characteristics.

Section	Sample daily Oyster journeys	OD pairs in analysis per direction*	Number of excluded OD pairs (duplicated by other lines)
North	18,867 (10%)	69	9
Trunk	21,012 (11%)	44	76
Heathrow	9,894 (5%)	63	0
Rayners Lane	5,313 (3%)	34	1
Other sections (including Uxbridge section) and cross-sectional trips	130,242 (70%)	1,082	n/a

* Source: Oyster records for April 16, 2012 (sample day). Only trips with both origins and destinations on fthe Piccadilly Line are included in the analysis.

Line. This, however, is out of the way and quite inconvenient compared to the direct Piccadilly Line journey and thus unlikely.

The validity of the selected sections from an analytical perspective is demonstrated in Figure 5-9, which documents the number of Oyster records for each of the Piccadilly Line sections and directions across the course of the day. The figure demonstrates that, with the exception of the westbound Rayners Lane section, all of the sections feature sample sizes of 20 Oyster records or more for 15-minute sample periods throughout the day, from around 7:00 to 20:00. Higher ridership on the Trunk and North sections of the line produce sample sizes available from Oyster records of greater than 80 per 15 minute bucket for most of the day. The figure indicates that conclusions from a study of Rayners Lane section records will be less reliable than those from a study of records from the Trunk and North sections (and the Heathrow section to a lesser extent). In future analyses that require creating potential line sections, examining the size of the sample of records derived from each is an important step.
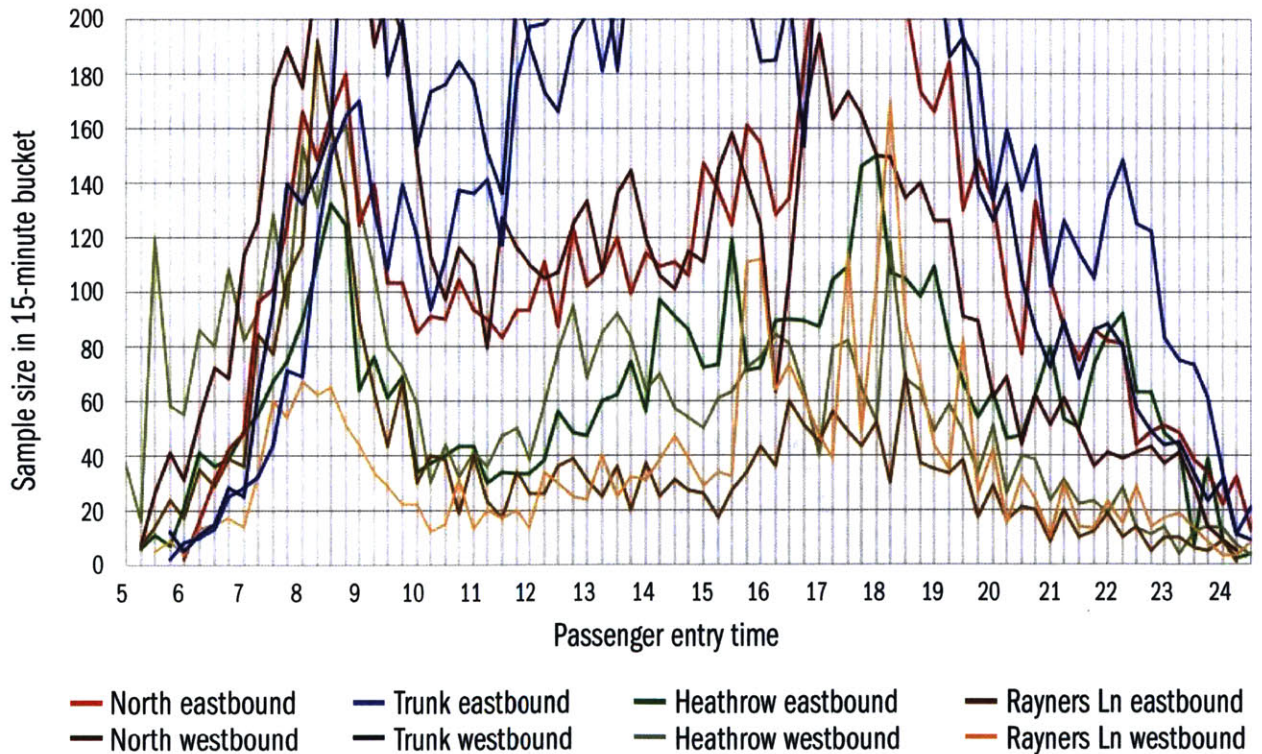
Figure 5-9: Entries by 15-minute bucket on Piccadilly Line sections, by direction.
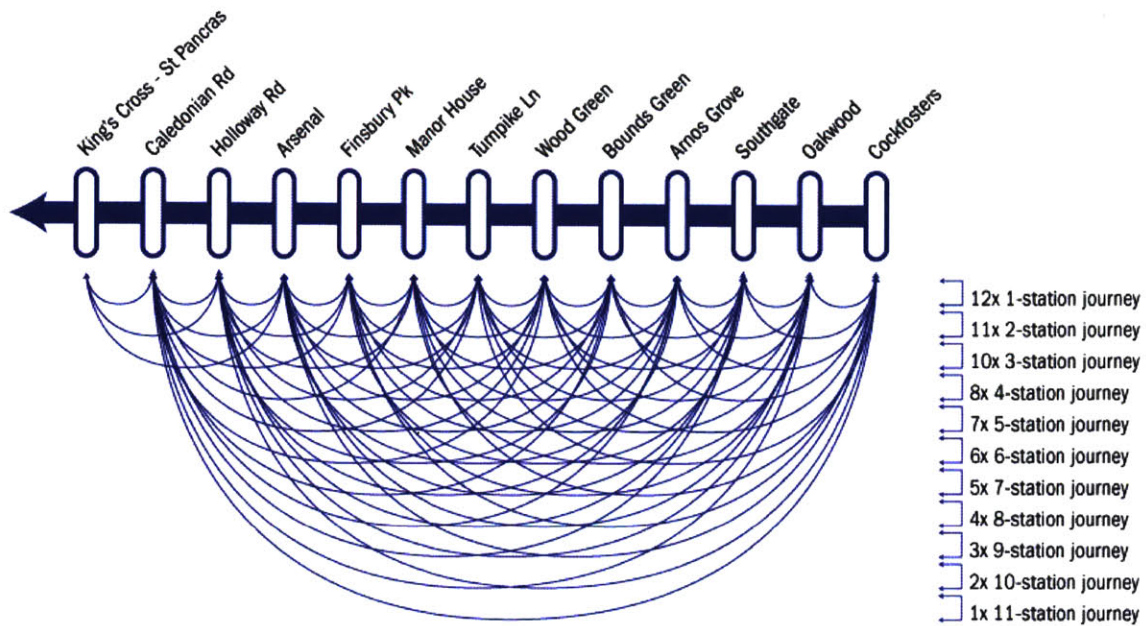
Excluded journeys include all trips from King's Cross to stations from Finsbury Park north, since passengers on those journeys may be taking the Victoria Line for at least part of their journeys

Figure 5-10: Origin-destination journey basket for the North section of the Piccadilly Line.

For each of the sections, all origin-destination pairs where direct Piccadilly Line trips were possible and no alternative route was available (even for a part of the journey) were selected. As shown in Figure 5-10, this produced a total of 69 origin-destination pairs for the North section—in each direction. This very large sample ensures that estimates of journey times on the Piccadilly Line can be reliably derived for even short periods of time.

In order to evaluate the impact of a disruption on passengers moving along a line, individual passenger journey time records were compared with expected journey times. The goal of this analysis is to provide an overall estimate of average change in journey time from what would be expected for all riders on a particular section and direction in order to indicate the severity of the disruption at a specific point on a specific day. Equation 5-1 is used to calculate the relative variability v in passenger journey time for a given section. The equation requires the following information:

- All Oyster records for representative origin-destination pairs in each direction in each line section, including trip start and end times, with the actual trip time for each journey (t_j); and
- "Normal" Oyster travel time for all possible origin-destination pairs, differentiated by hour (T_j). This hour-by-hour differentiation is necessary because of significant changes in normal journey times depending on the period of the day. This data has been compiled from median travel times on origin-destination pairs over the course of 5 weekdays with "normal" service (with PWT scores of less than 20).

Once the actual and normal travel times for all trips in the line section have been compiled, the normal travel time is subtracted from the actual travel time to find the variation from the norm, as shown in Equation 5-1. The variation of *each trip* is calculated separately. These variations are averaged across all trips (n) and across all minutes (m) in the sample period. This equation offers a reliable estimate of the variation in travel time for passengers on a section of the line. The result is a percent difference from the norm; in other words, it indicates average variation in travel time in a specific line section. Note that it should not be assumed that such variations apply to every origin-destination pair within the section and period.

In this analysis, a period (m) of 15 minutes is used because this corresponds with the minimal sample size necessary for analyzing the journey time patterns on the Rayners Lane section of the line. This period could be shortened for a larger sample or expanded for a smaller one. Note, however, that the relative variability is re-calculated every minute and thus provides a rolling estimate of journey time variability, rather than one confined to individual 15-minute periods.

$$v = \textit{Relative variability for trips t in line section} = \left\{ \sum_{i=0}^{m} \left[\frac{\sum_{j=0}^{n} (t_j - T_j)/T_j}{n} \right] \right\} / m \qquad \textbf{(5-1)}$$

Where,
m is the number of minutes during the representative sample period;
n is the number of sample trips;
t_j is the actual journey time for trip j ; and
T_j is the normal journey time for trip j (characteristic of the OD pair of the trip).

Table 5-3 demonstrates how this measure of variability can be implemented. Each individual trip's Oyster journey time is compared with its normal journey time, producing an individual variation. Subsequently, all variations within the time period and line segment are averaged to indicate the overall level of passenger journey time variation on the line section. It is important to note that negative variation is possible because the "normal" reference point does not represent the journey time on a perfect day but rather a normal day, so it is possible for passengers to move through the system faster than "normally." In this hypothetical example, all journeys between stations "A", "B", "C", and "D" are analyzed in the A to D direction. The overall variation v is in effect weighted by the segments with more passengers, so the measure is reflective of the average passenger.

This approach offers the analyst the ability to assess how customer journey times diverge from the norm during disruptions. The information may be useful in determining when disruptions begin and end in terms of customer impacts. However, it should be noted that there are several limitations of this method. First, the only source of data used is Oyster. Currently, the data is provided in 1-minute increments (not seconds), meaning that very short trips may show high variability even if they diverge

Table 5-3: Hypothetical dataset of travel within line section from 8:00 to 8:05.

Passenger	OD pair	Departure time	Actual journey time (min)	Normal journey time (min)	Variation (%)
1	A-B	8:00	9	10	-10
2	A-B	8:01	13	10	30
3	A-B	8:02	15	10	50
4	A-C	8:01	18	15	20
5	B-C	8:00	11	12	-8
6	B-C	8:00	13	12	8
7	B-C	8:05	16	12	33
8	A-D	8:03	27	26	4
9	A-D	8:04	26	26	0
10	C-D	8:00	5	5	0
11	C-D	8:05	5	5	0
All passengers	n/a	n/a	n/a	n/a	$v = +12\%$

by only a minute or two. Second, the method does not highlight performance for particular origin-destination pairs. If there is a specific problem occurring at an individual station or in a subsection of the line section selected, it will not necessarily affect the overall level of variation for the line. Finally, the data become available only at the *end* of passenger journeys, when they tap out. Thus the beginning (and end) of a disruption is delayed in terms of when the data presents it, which is a concern for real time use of this system, as will be discussed in Chapter 6. However, for this analysis, it is reasonable to use this data in the context of retrospective analysis of incidents.

5.3 ANALYSIS RESULTS AND DEVELOPMENT OF A MEASURE OF DISRUPTION IMPACT

Figure 5-11 shows the variability analysis results for two disrupted days, May 14 and 15, 2012. The source of data for these charts is variation in journey time for each section and direction during the sample time period. At time t_j the relative variation v is calculated for the period over the previous fifteen minutes. Using a rolling view of the sample prevents major minute-to-minute shifts in travel time variation and illustrates the trend, not the moment-to-moment performance. The chart provides information in terms of passenger entry time. As was shown in Figure 4-2, the two sample days recorded Excess Platform Wait Times (PWT) of 20 and 34 seconds, respectively, indicating that the first day was mildly disrupted and the second more severely disrupted. These indicators are borne out in Figure 5-11, as May 15 shows significantly higher values of travel time variation than the previous day.

A close look at the figures provides an overview of the troubled line sections during different times of the day. Figure 5-12 shows the westbound progression of journey time variation on different sections of the Piccadilly Line from 8:30 to 11:00 on May 14, 2012. This chart shows that average journey time variations never exceeded 30% in this period. This indicates that while passengers did see higher-than-normal journey times during this period, the impacts never became extreme.

A look at the variation on the line on this day provides some insight into the manner in which disruptions move through the system. The disruption appears to commence in the North section, move into the Trunk, and then shift to the Heathrow and Rayners Lane sections. This is as expected, since

A. 14 May, 2012 westbound.

B. 14 May, 2012 eastbound.

C. 15 May, 2012 westbound.

D. 15 May, 2012 eastbound.

North — Trunk — Heathrow — Rayners Ln
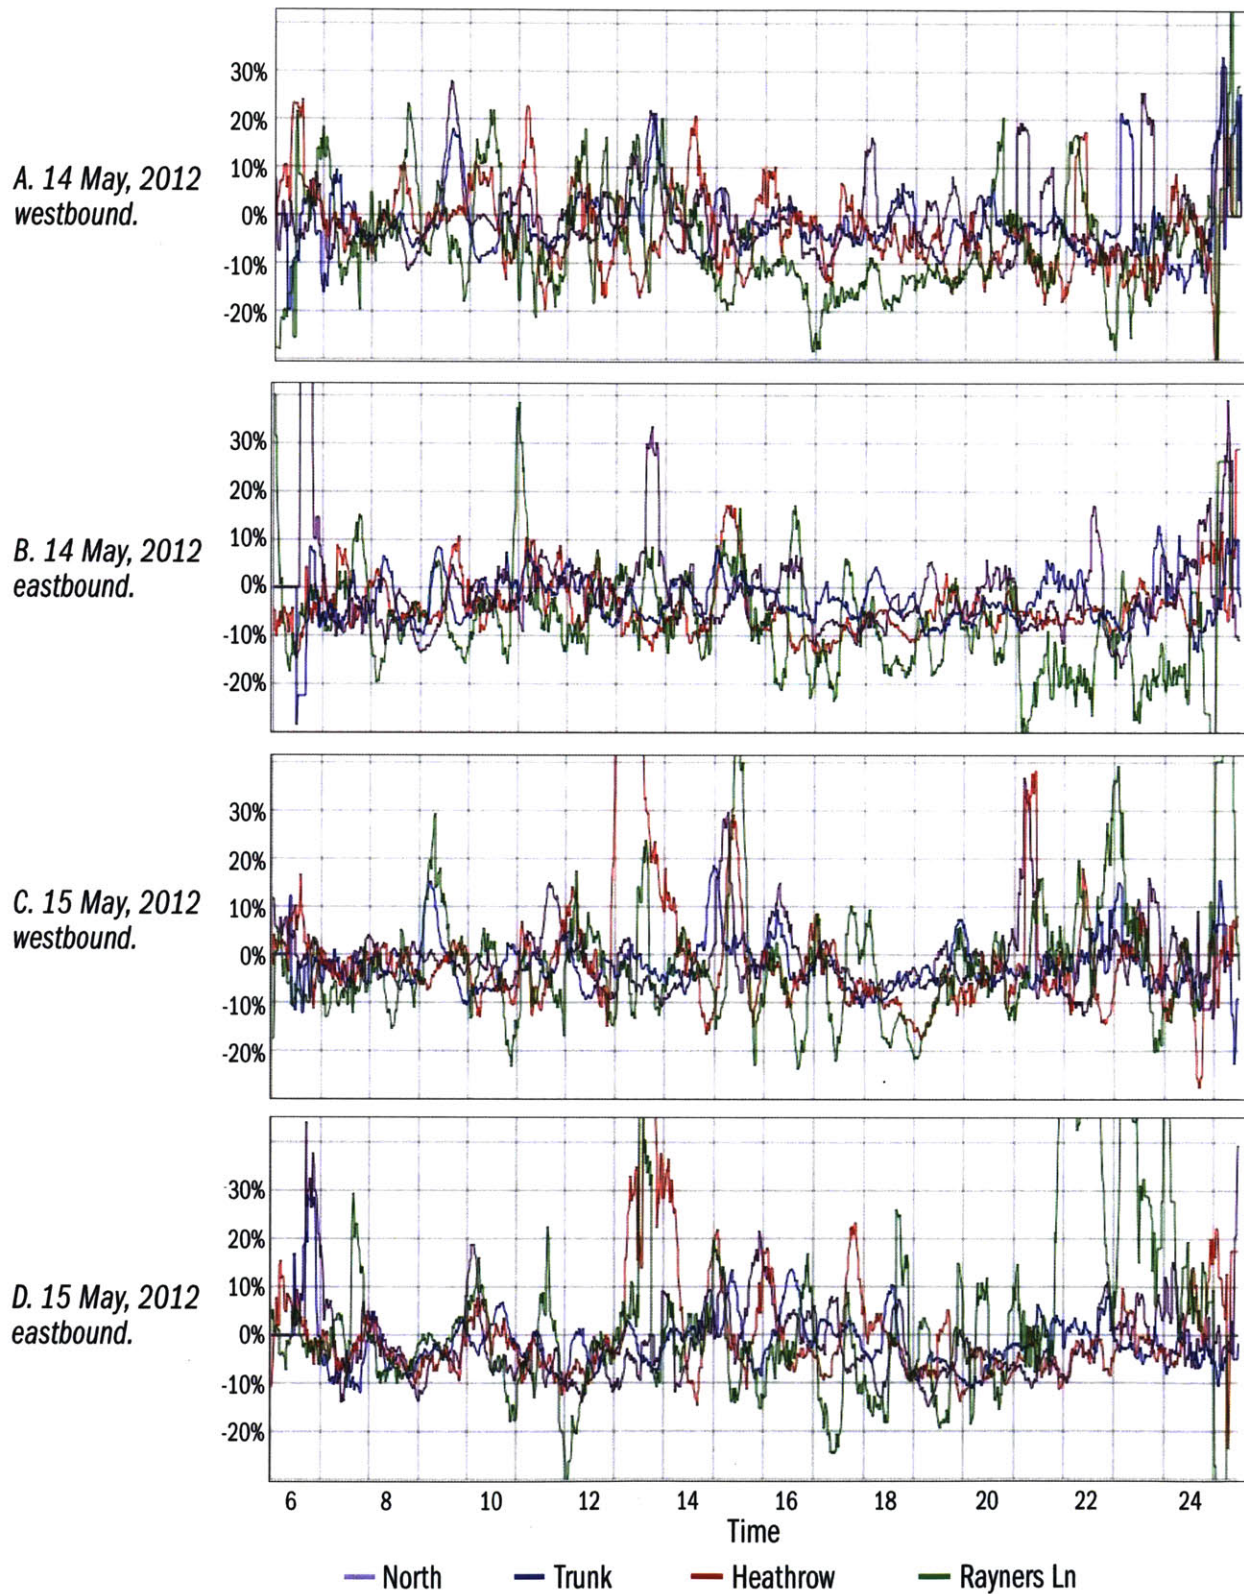
Figure 5-11: Journey time variation index throughout day and line section.

westbound service begins in the North section and then extends throughout the line; what is being shown on the chart is a propagation of passenger delays from one section to another, following the trains.
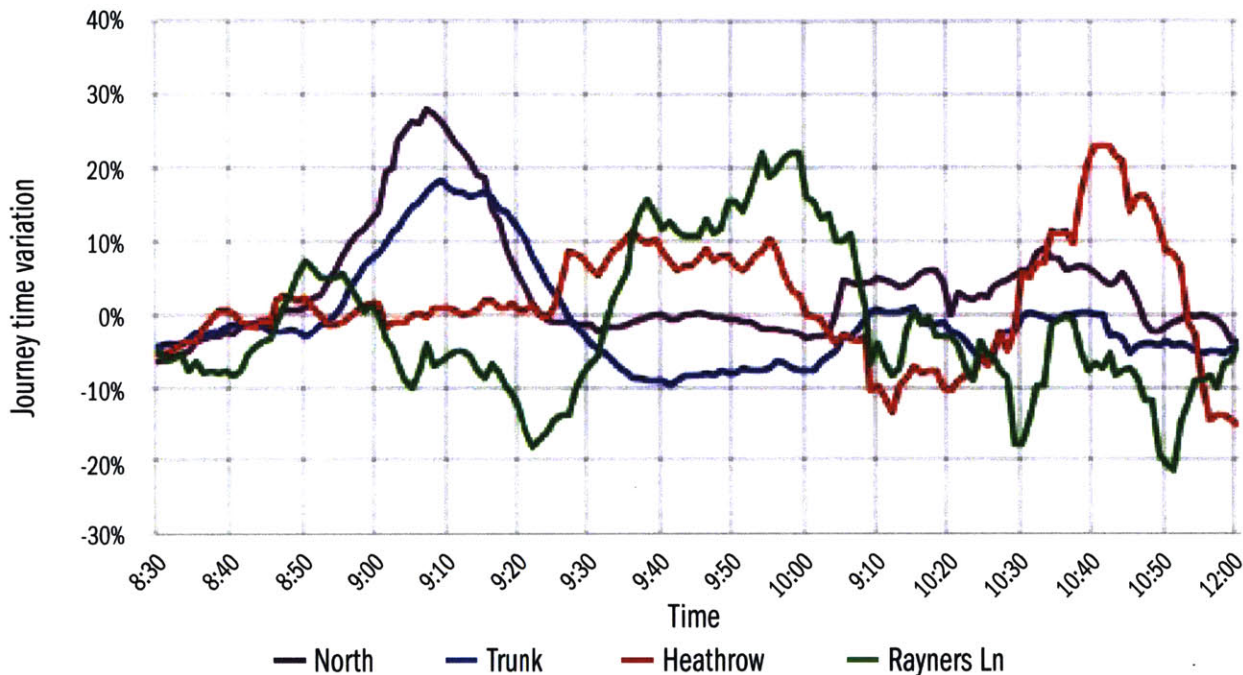
Figure 5-12: Journey time variation westbound from 8:30 to 11:00 (May 14, 2012).

Results from May 15 indicate a far more severe disruption, beginning on the westbound Heathrow Branch at around 12:30, as shown in Figure 5-13. Passenger delays spread to eastbound service by 13:05, when difficulties westbound begin to affect eastbound operations. Westbound passengers saw increases in journey times of around 80% at the height of the disruption; eastbound they saw increases of up to 105%. These problems, however, did not propagate to the Trunk eastbound, indicating that while certain projections might be made about where problems are likely to occur, one cannot simply assume that variation in passenger journey times will propagate from one section of the line to the next. Similar trends are evident in Figure 5-14, which illustrates the occurrence of a major disruption on the Rayners Lane Branch between 21:00 and 24:00.
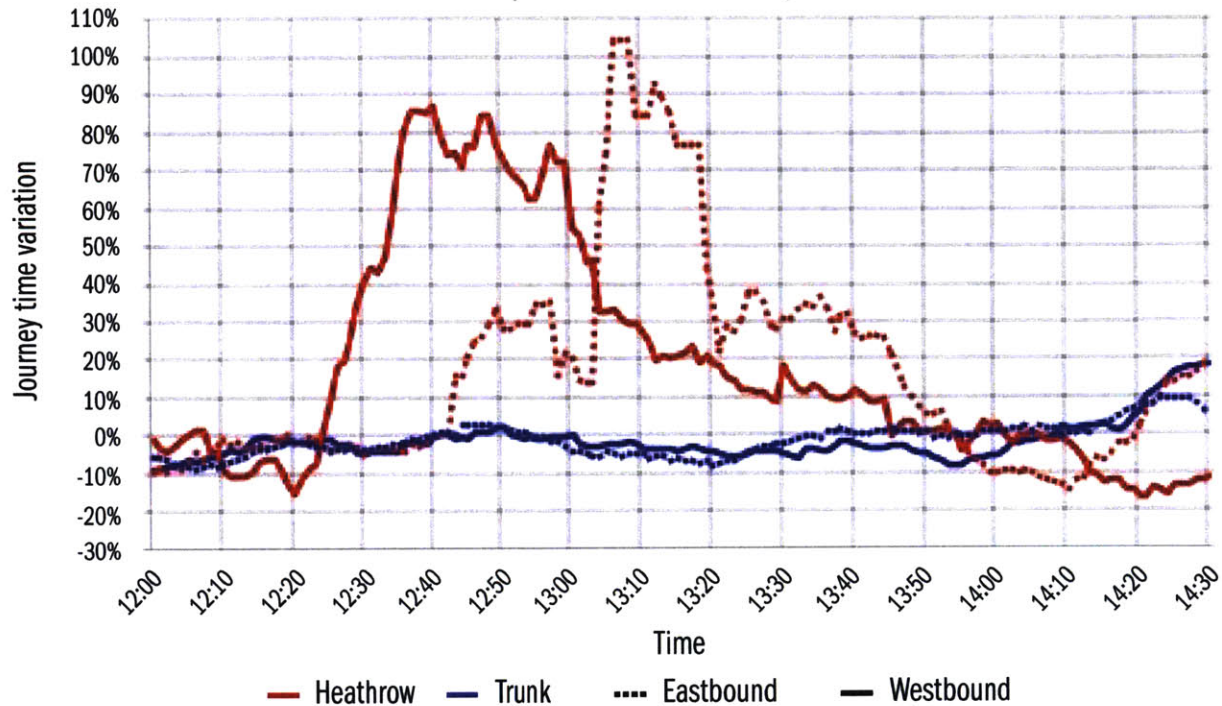
Figure 5-13: Journey time variation by line section from 12:00 to 14:30 (May 15, 2012).

It should be emphasized that the results of the analysis include more than just the direct effects of the incident itself. Rather, the metric *v* computes all passenger delay compared to that occurring on a regular day of service and thus provides insight into the "knock-on" effects that are also estimated in the LCH calculation, as well as the potential impact of controller actions. In addition, the propagation of disruptions depends directly on the decisions controllers have made about how to respond to the incident. However, as both Figure 5-13 and 5-14 illustrate, the propagation of delay throughout the system cannot easily differentiate between "direct" effects and knock-on effects. In the case of eastbound Rayners Lane services, problems on the line continued from 21:20 until 23:40, long after the actual incident had been resolved. No matter whether direct or indirect, however, this mechanism does demonstrate when and in what section of the line passengers experienced increased journey times.
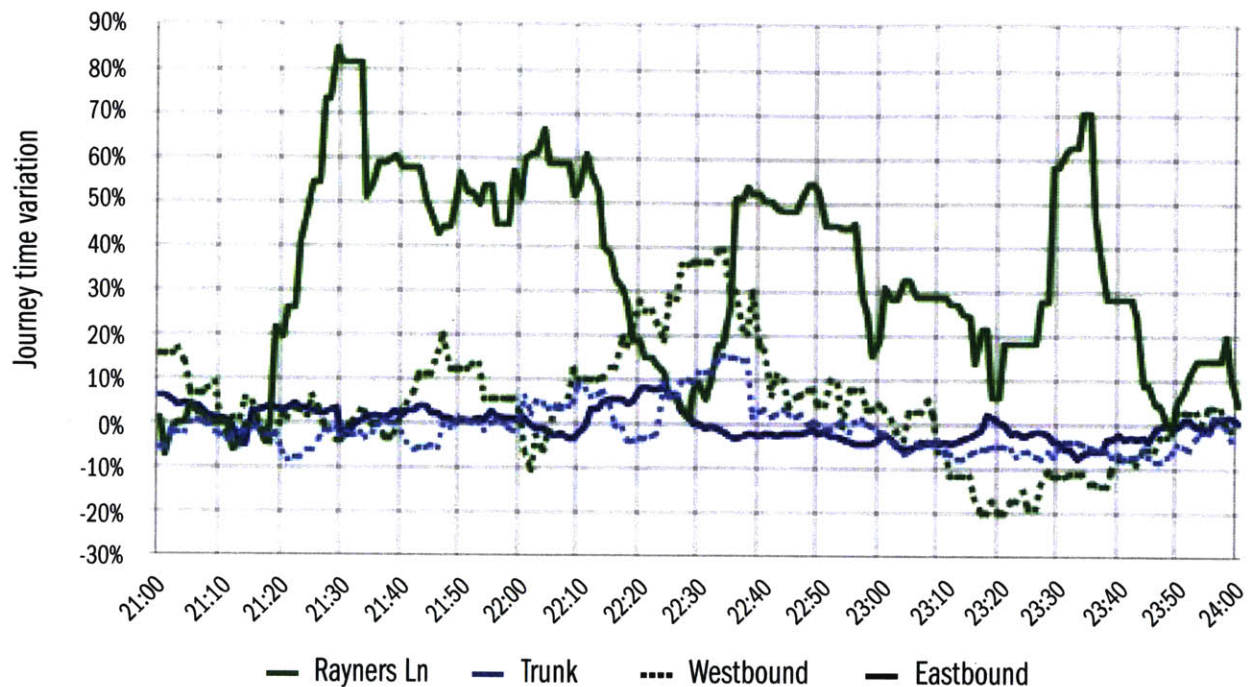
Figure 5-14: Journey time variation by line section from 21:00 to 24:00 (May 15, 2012).

A review of several days of data indicates a ±20% margin of journey time variability for "normal" service. When variability exceeds 20%, conditions on the line are not normal. This margin is used for the calculation of total variability resulting from a disruption (see below) and the discussion of real-time use of Oyster data (see Chapter 6). It should be noted that this baseline may be worth reexamining based on additional data and further analysis.

The minute-by-minute variations in journey times clearly indicate which parts of the day saw the most disrupted service, but they can also be used to provide an overall estimate of disruption impacts on passengers. This thesis proposes a method to incorporate Oyster data into the retrospective analysis of disruptions and monitor the degree to which passengers were affected by a specific disruption.

As shown in Equation 5-2, the total disruption impact V is measured retrospectively by calculating the integral of the variability curve v above the baseline 20% variability, for disruptions where the v curve rises above 20% for five minutes or more. This integral is an indicator of the degree to which passenger service has been disrupted and suggests the severity of the disruption in terms of customer impact. This integral allows a direct comparison between different disruptions. Note that because the

integral is simply the area under the curve, not the height or breadth of the curve, a long, minor disruption could produce a V score similar to that for a short but severe disruption. Dividing V by the affected time period, as documented in Equation 5-3, provides an indication of the intensity I of the problem.

$$V = \textit{Total disruption impact} = \left\{ \int_a^b v_i \, dt \right\} - 0.2(b - a) \text{, for periods where } v > 0.2 \qquad \textbf{(5-2)}$$

Where,
a is the commencement of the disrupted period in minutes (when v first reaches 0.2);
b is the end of the disrupted period in minutes (when v drops below 0.2); and
v is the variability measure as defined in Equation 5-1.

$$I = \textit{Disruption intensity} = \frac{V}{(b - a)} \text{, for periods where } v > 0.2 \qquad \textbf{(5-3)}$$

Where,
a is the commencement of the disrupted period in minutes (when v first reaches 0.2);
b is the end of the disrupted period in minutes (when v drops below 0.2); and
V is the total disruption variability measure as defined in Equation 5-2.

V and I are both relative metrics that allow the effects of disruptions to be compared, but they do not "represent" anything, unlike the LCH metric, for example. Figure 5-15 provides information about the relative scales for the V and I metrics. A disruption with a V value of 5 would be relatively minor while one with a value of 20 would be major. A disruption with an I value of 0.5 would be "sharp," meaning that it would affect passengers significantly over a short period of time, whereas one with a value of 0.1 would be "slow," meaning that it would affect passengers over a long period of time.
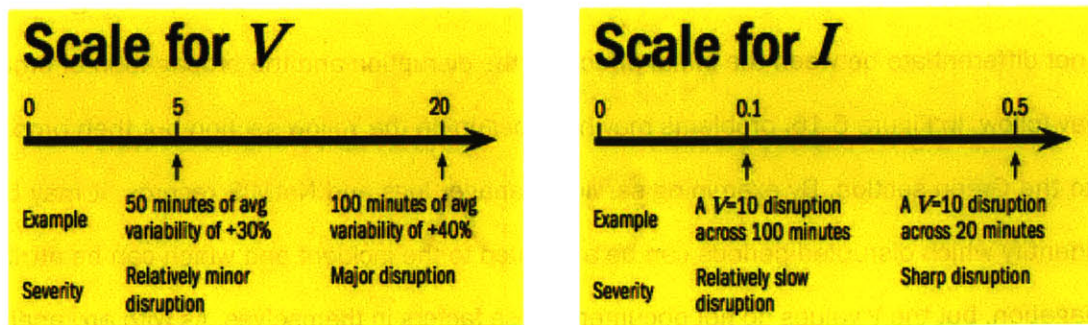
Figure 5-15: Scales for V and I metrics.

The *V* metric is not comparable directly across sections because the number of passengers affected will be different. One way to increase the utility of the *V* scale would be to multiply it by the number of users on the relevant section of the line during the affected period; this would allow for a direct comparison between sections. One problem, however, is that reliable estimates of passenger flow on the Piccadilly Line are not yet fully developed. That said, Ravichandran (2012) has developed a method by which to do so, and once results are available for the line throughout the day, these could be multiplied by the *V* value for useful analysis. In turn, the overall disruption impacts along the line could be derived by summing the values for each section. This would provide a useful indicator of the full impacts of a disruption and provide a comparison with the LCH values calculated from NetMIS data.

In the meantime, the disruption impact metric remains valuable even if the results it produced may not be directly comparable between lines. Figure 5-16 illustrates how *V* may be calculated based on *v*, the relative disruption variability curve, for a hypothetical four-section line. The graph shows the journey time variability for each of those sections on one day, between 15:00 and 17:00. As the figure demonstrates, the Green section of this line had the most significant disruption, occurring between 16:10 and 16:50, and its integral (1) is calculated above the 20% margin to produce the *V* score. The Yellow section had a relatively minor disruption between 15:50 and 16:10, but it stayed above 20% for more than five minutes, so its integral (2) is also calculated. The Green section also had a disruption that reached a 20% variation between 15:45 and 15:50 (3), but because this disruption lasted less than 5 minutes, its integral is not calculated and it is not considered a disruption worth reporting under this metric. These short disruptions can be overlooked as they have minimal effect on passengers.

V does not differentiate between the initial effects of the disruption and the propagation of problems that may follow. In Figure 5-16, problems may have begun on the Yellow section but then propagated over to the Green section. By examining service manager logs and NetMIS records, it may be possible to identify which disrupted periods can be attributed to the incident and which can be attributed to propagation, but the *V* values do not document these factors in themselves. As with any analysis of disruption impacts, differentiating the "direct" impacts and the "knock-on" effects may not be

possible—particularly when other incidents may follow major incidents. In addition, the metric does not account for different control strategies, which can affect passenger delay significantly. Nonetheless, by choosing specific incidents that appear not to have any follow-up incidents, the disruption impacts of those problems could be analyzed by line section using the V metric described here. In addition, the metric has the advantage of offering the opportunity to compare recovery strategies; with similar incidents, the impacts of different controller actions could be evaluated.
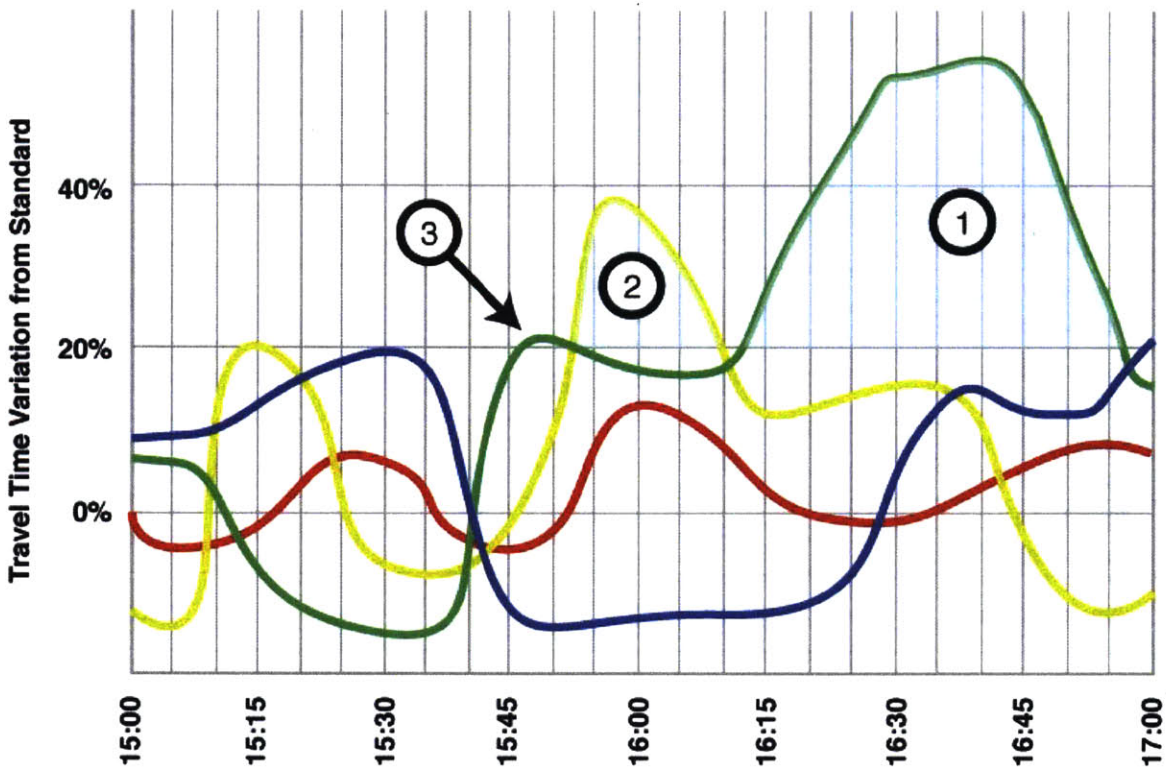
Figure 5-16: Calculating the total disruption impact V.

Altogether, this method provides an indication of when disruptions began and were resolved, and their severity, both in terms of overall disruption impact (*V*) and intensity (*I*). Table 5-4 indicates when and where disruptions occurred, by section, on May 15, 2012. The most significant disruptions by overall journey time variability occurred on the Heathrow westbound section from 12:28 to 13:19, with a *V* of 17.43, and on the Rayners Lane eastbound section from 21:19 to 22:18, with a *V* of 19.78. These two

Table 5-4: Disrupted periods of at least 5 minutes according to variation index.

Affected section	Start	End	Duration (min)	V	I
Heathrow westbound	12:28	13:19	51	17.43	0.34
Rayners Ln eastbound	13:00	13:14	14	3.05	0.22
Heathrow eastbound	12:45	13:45	60	13.92	0.23
Heathrow westbound	14:46	14:58	12	0.58	0.05
Rayners Ln westbound	14:49	15:07	18	3.21	0.18
Heathrow eastbound	17:14	17:21	8	0.17	0.02
Heathrow westbound	20:42	20:56	14	2.17	0.16
Rayners Ln eastbound	21:19	22:18	59	19.78	0.34
Rayners Ln westbound	22:19	22:39	20	2.06	0.10
Rayners Ln eastbound	22:35	23:18	43	7.63	0.18
Rayners Ln eastbound	23:27	23:43	16	4.03	0.25
Whole line	n/a	n/a	n/a	74.03	0.24

section-based disruptions had the same I of 0.34, which was the highest level of intensity for any of the disruptions measured on this day.

It should be reemphasized that both the V and I figures are *relative* ones; they reflect neither the number of customers affected, nor the number of hours individuals were delayed. On May 15, the disruption on the eastbound Rayners Lane service noted above produced a similar V score as that on the westbound Heathrow section, with identical intensity scores. But the *number* of people affected on the Heathrow section (at midday) was likely higher than that on the Rayners Lane section (at night). But V only accounts for variability within each section compared to standard journey times and thus makes no adjustment based on the number of users affected, unlike the LCH measure, which does attempt

to account for this.

Hence, the primary use of the total disruption impact measure V is to provide a way of comparing different disruptions to evaluate their relative severity. It may also be used as a measure of the entire day's disruption variability by summing disruptions, particularly by line segment. As such, it could provide a supplement to the full-day PWT and LCH metrics to provide a summary of the degree to which passengers were affected by problems on the line.

5.4 PASSENGER ACCUMULATION

In addition to providing important information about the time it takes riders to complete journeys between origin-destination pairs, Oyster can be utilized to track rider accumulation in the system. This could prove to be a useful method to explain the effects of incidents on passengers. After all, if a disruption causes increases in journey times for passengers on the system, there will be more riders than normal *in the system* during disruptions—assuming steady entry of passengers into the system. This section describes how Oyster data could be used to measure this passenger accumulation and how this measure could be used to understand disruptions retrospectively. Chapter 6 investigates the potential role of such a measure in tracking disruptions in real time.

On the Underground, information about the origins and destinations of passengers using the system can be used to estimate the number of passengers on the system at any point in time. Though many journeys have only one credible route for a specific origin-to-destination movement, others—usually those that require an interchange—have several possible routings because of the complexity of the LU network in Central London. The RODS surveys conducted by LU (discussed further in Section 5.6) provide some insight into how passengers complete their journeys. For more complicated journeys, such as those that involve bus trips, Gordon's work (2012) allows for projection of actual journeys.

Oyster entries alone also offer *some* insight into how passengers were using the system at a point in time. Oyster can measure the number of people in the system (excluding the small minority of people who use magnetic tickets) at any time simply by summing the number of passengers who have tapped in but not yet tapped out. In other words, subtracting the cumulative number of exiting passengers

from the cumulative number of entering passengers. These estimates of passenger accumulation can also be developed for specific lines by considering stations where service on only that line is provided.

For instance, on the Piccadilly Line, of 53 total stations, 32 have no interchange with any other line and therefore we can be confident that customers who enter there will be riding the Piccadilly Line (eastbound or westbound), at least for the first part of their journeys. This information can, in addition, be divided up into the line sections described in Section 5.2. There are 11, 4, 11, and 6 stations with no interchanges in the North, Trunk, Heathrow, and Rayners Lane sections, respectively, that could be used for such an analysis. There are no stations in the line section from Rayners Lane to Uxbridge with such characteristics since all these stations are shared with the Metropolitan Line. This demonstrates one inherent weakness of this type of analysis: It cannot be performed on a line where multiple lines share stations.

This accumulation method has a number of other deficiencies. Since we do not know the destination of riders who have tapped in but not yet tapped out, their route patterns or direction of travel cannot be projected (with the exception of riders at termini stations). Because of the lack of knowledge about direction of travel, an increase in riders originating from a station compared to normal could mean a disruption in either or both directions—but this method does not show which is being affected. These issues require that this approach be developed further.

Nonetheless, examining passenger accumulation using Oyster tap-ins at Piccadilly Line stations with only Piccadilly Line service may offer some insight into the status of service on the line. Figure 5-17 illustrates passenger accumulation from Piccadilly-line only stations on September 19-23, 2011. As noted in Section 5.1, September 22 saw major disruptions to eastbound service between 17:00 and 20:00.

Figure 5-17 demonstrates that the five days saw very similar levels of passenger accumulation in each of the four analysis sections throughout the day. These charts document passenger accumulation at 10-minute intervals. Recall that the passenger accumulation indicated here is simply the number of passengers who have tapped in at Piccadilly Line-only stations, but not out, in each of the

sections. Thus it must be emphasized that if the chart indicates that there are 1,000 passengers from the Heathrow section currently on the system, this simply means that there are 1,000 passengers who tapped into Piccadilly Line-only stations on the Heathrow section who have not yet tapped out, not that there are a total of 1,000 riders on the Heathrow section at this time. These passengers may or may not still be in the Heathrow section, or even on the Piccadilly Line, and other riders from other parts of the Piccadilly Line or other lines altogether may have entered the section.

However, examining the section of passengers in the Trunk section clearly demonstrates the effect of the service disruption that began at around 17:00 on September 22; then, there were clearly more passengers originating from the Trunk section than on the other days. This was also true of the Heathrow section beginning at around 18:30 on September 22. This evidence suggests that it is possible to use such Oyster information to track increased rider accumulation and thus effects of the disruption.
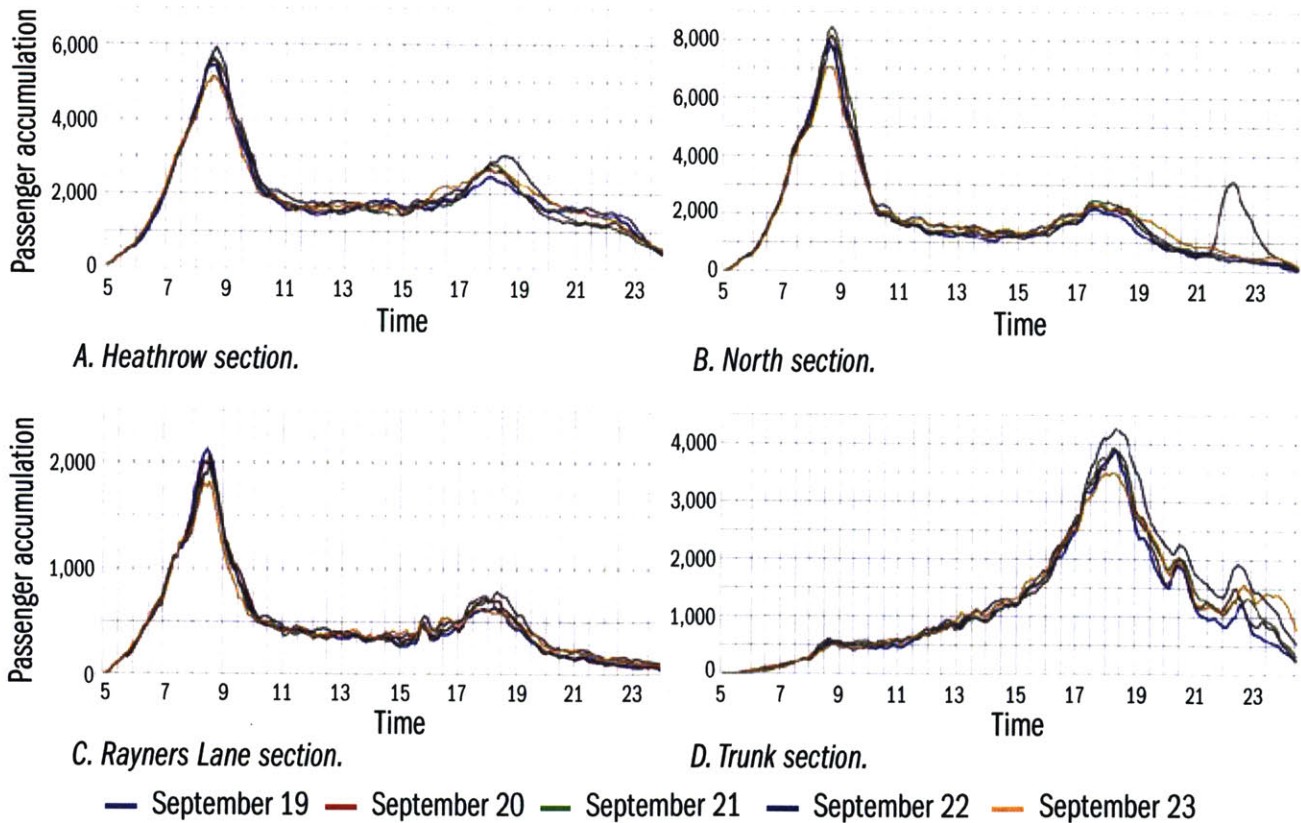
Figure 5-17: Passenger accumulation by Piccadilly Line section.

That said, Figure 5-17 also illustrates a major reason to be cautious in using such passenger accumulation information to examine the effects of disruptions. The chart indicates that there was a significant increase in passenger accumulation on the North section beginning at around 21:30 on September 20. A quick evaluation of these charts might suggest that there had been a major disruption that had backed up travelers very significantly. However, a more in-depth analysis shows that September 20[th] was the date of a major match for the Arsenal football team. Arsenal has a stadium located adjacent to a North section Piccadilly Line station (Arsenal). When the game ended, passengers flooded into the system, therefore causing a major increase in ridership there.

A careful examination of the data, though, can still be quite revealing. Figure 5-18 illustrates ridership on the four Piccadilly Line study sections on September 22 as a percentage of that on September 21, a "normal" service day with no major disruptions. All sections, in particular the Rayners Lane section, saw ridership above normal between 18:30 and 20:00, indicating a disruption during this period. Examining this data from an absolute perspective (B) may actually be more useful in terms of identifying the most challenged sections of service. The data show that ridership on the Trunk section of the line was 200 passengers above normal beginning at around 17:00 and continuing until 22:45 or so, with a peak of about 950 riders above normal at around 18:40. In the context of changes in absolute volume, the changes on the Rayners Lane section, which serves fewer riders at the stations available for analysis, no longer appear nearly as significant.

One way to summarize the information from the graph in Figure 5-18 is to measure the area under the curve of the disrupted passenger flow line, in a similar manner to the one used in Section 5.3. A minimum number of passengers affected by a disruption can be set as a margin (in this case, 200 additional passengers riding than typical will be used to illustrate the concept), and the area of the curve above that level can be calculated using Equation 5-4.
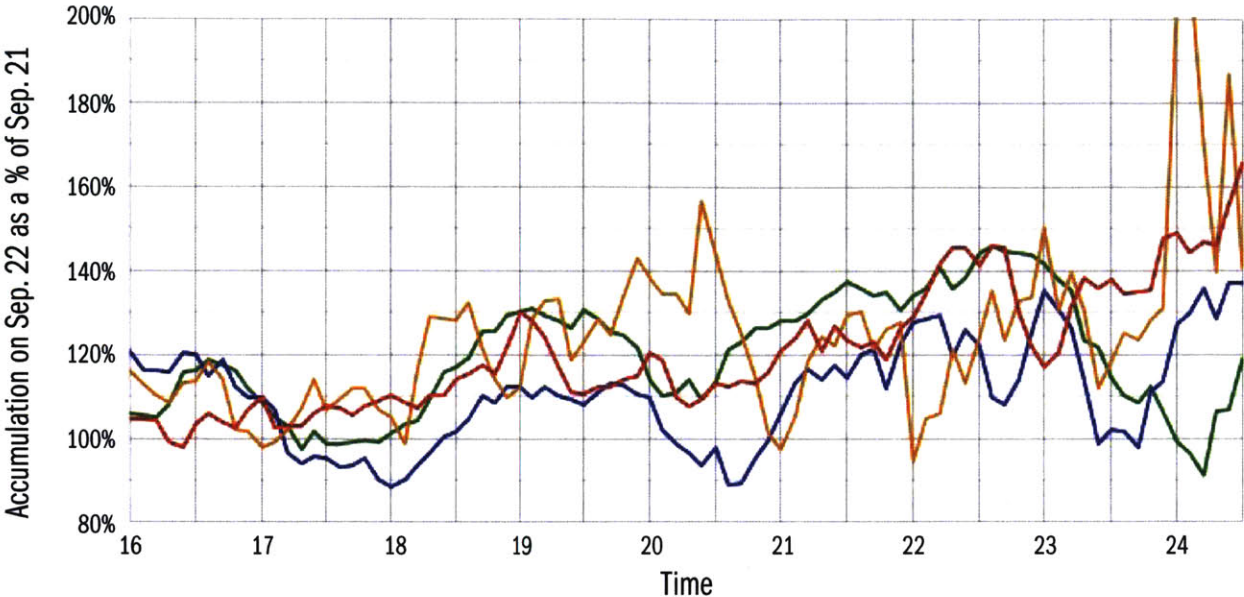
A. Percent variation in passenger accumulation compared to September 21.
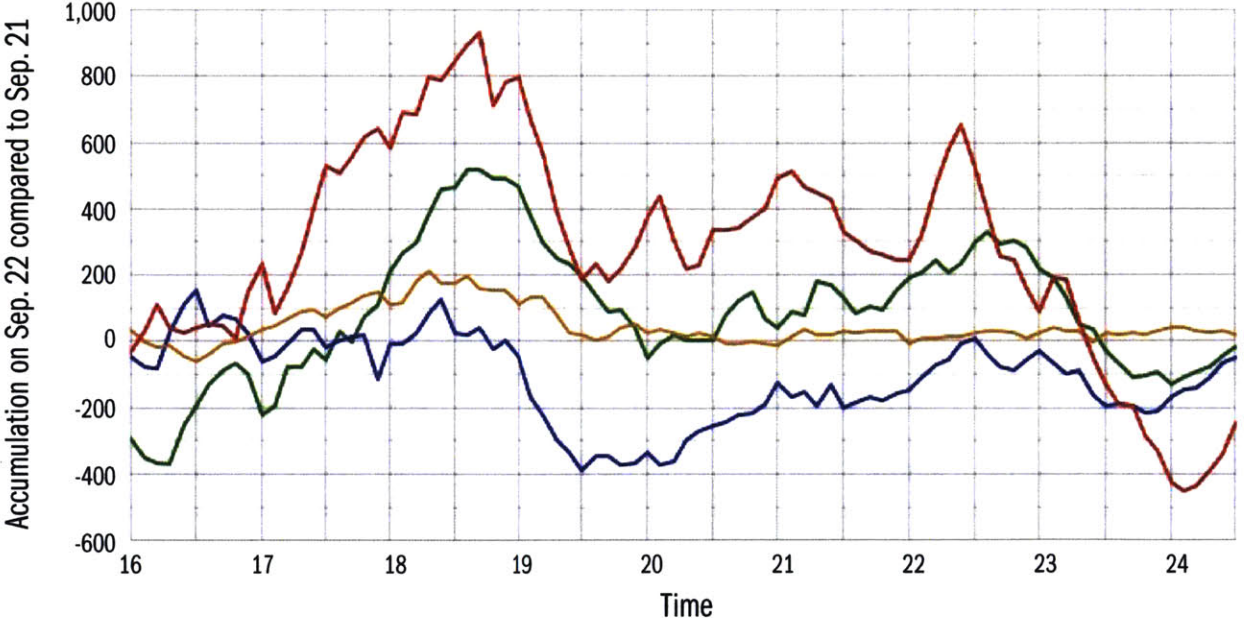
B. Absolute variation in passenger accumulation compared to September 21.

— North　　— Trunk　　— Heathrow　　— Rayners Ln

Figure 5-18: Passenger accumulation by Piccadilly Line section (September 22, 2011).

$$F = \textit{Passenger accumulation variability} = \left\{ \int_a^b (p_i - p_s)\, dt \right\} - 200(b - a) \tag{5-4}$$

Where,

p_i is the accumulation of passengers on the study day on a specific segment;

p_s is the accumulation of passengers on a normal day and time on the same segment;

a is commencement of the disrupted period in minutes (where $p_i - p_s$ first reaches 200); and

b is end of the disrupted period in minutes (where $p_i - p_s$ drops below 200).

The *F* measure, like the previous measures discussed in this chapter, is a relative one that allows different disruptions to be compared, and the scale for the metric is shown in Figure 5-19. However, unlike *V* and *I*, *F* allows comparisons between different line sections, because it represents the absolute increase in accumulated riders over what would be expected for a given period. If one section of the line has a considerably higher *F* score than another, the disruption has likely had a far greater impact in that section in terms of the number of passengers affected. For September 22, this method produces the results shown in Table 5-5.
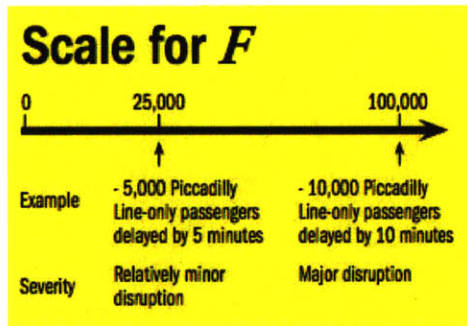
Figure 5-19: Scale for *F* metric.

Table 5-5: Passenger accumulation on September 22.

Line section	Start	End	*F*
Trunk	17:00	17:06	310
Trunk	17:18	19:24	95,090
Trunk	19:36	19:42	340
Trunk	19:48	22:48	50,600
Heathrow	18:00	19:24	27,070
Heathrow	22:06	23:00	6,060

When compared with *V* and the relevant LCH measures from the CuPID database, *F* provides an additional measure to better understand disruption impacts on passengers. These comparisons will be used further in Chapter 6 for the purposes of identifying real-time problems on the line.

5.5 COMPARISONS WITH NETMIS, SERVICE CONTROLLER, AND CUPID DATA

Comparing the results of the travel time variation analysis described previously with NetMIS data on train frequency provides some insight into the differences between data from these two sources. In Figure 5-20, variations in Oyster travel time throughout the day on April 24, 2012 are shown in solid lines, with frequencies of trains running eastbound through four stations (from NetMIS) shown in dotted lines. These four stations are representative of the four sections identified for the travel time variation analysis. By examining the variation from normal frequencies and headways, we can develop a reasonable idea of the state of traffic.

Figure 5-20 demonstrates that the analysis of frequencies provides only a general indication of service quality. While the Oyster variation analysis shows quite significant delays for customers on the Rayners Lane section between 18:00 and 19:00, no change is seen in the frequency, though it is lower than scheduled. Moreover, while frequencies *do* decline on the Trunk and North sections, this decline occurs mostly after customer journey times have been affected on the North section.

A closer look, such as is shown in Figure 5-21, is equally revealing of the differences between the type of information provided by NetMIS and Oyster records. Significant drops in frequency in the North and Trunk sections of the line (from about 24 to 16 trains per hour) occurs almost two hours after major passenger journey time increases have been recorded. This may indicate that passengers have chosen to switch to another line in order to avoid delays. That said, it is also possible that even reduced frequencies by 19:00 do not significantly affect passenger journey times because even frequencies of 16 trains per hour (recorded on the North section at 19:00) provide a train every 3.75 minutes. This latter possibility, though, assumes adequate space on the trains to handle the passenger load, which may be unrealistic considering the loads on central sections of the Piccadilly Line.
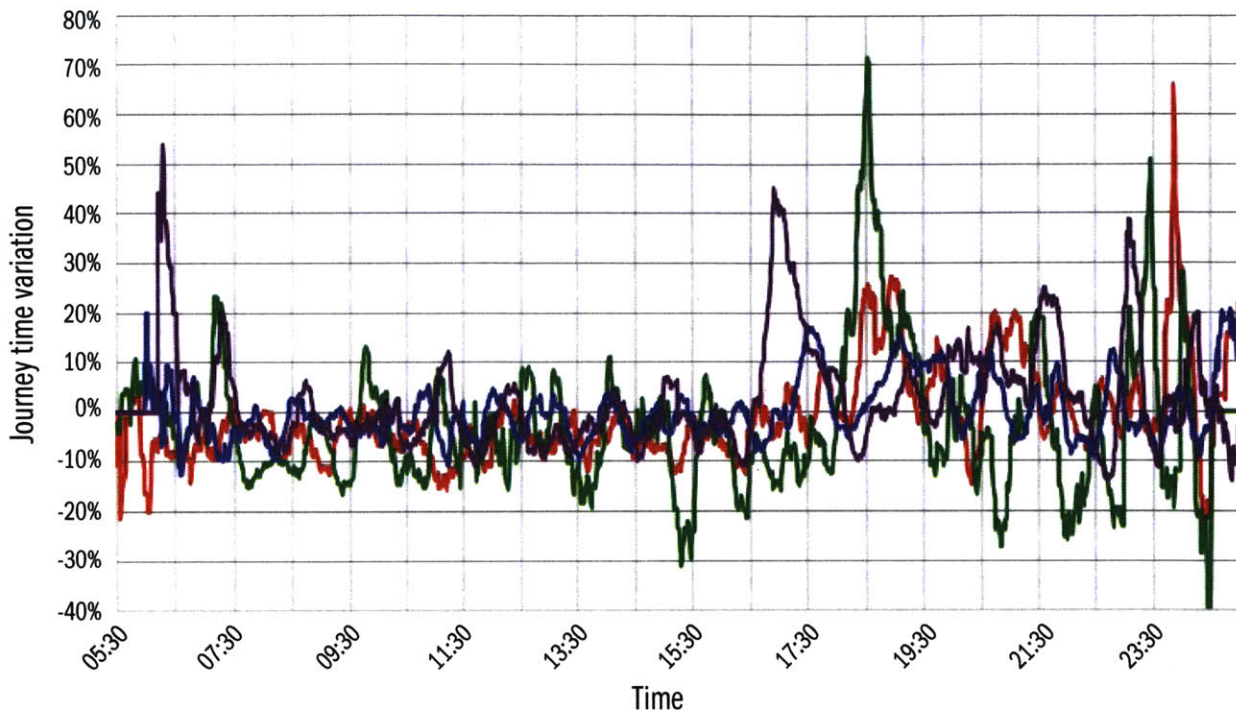
A. Oyster journey time variation.
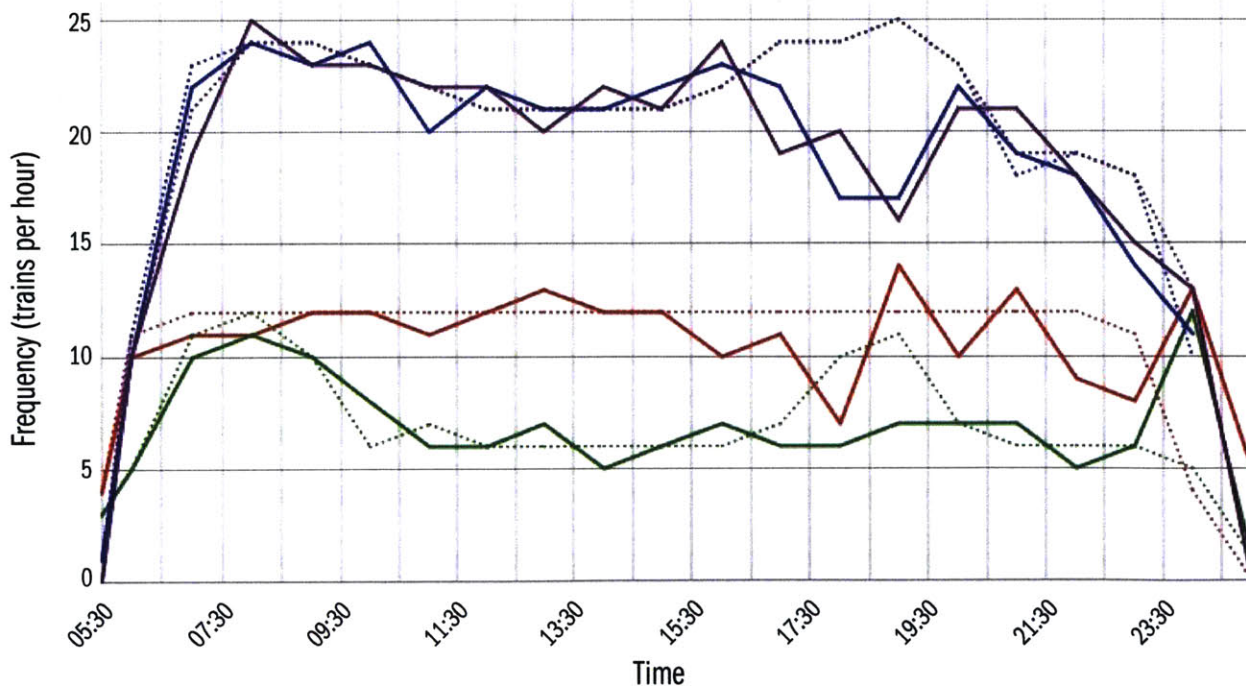
B. Train frequencies per hour at representative stations, compared to schedule.

—— Trunk　　—— North　　—— Heathrow　　—— Rayners Ln　　···· Scheduled

Figure 5-20: Comparing Oyster journey time variation and train frequencies, eastbound (full day).
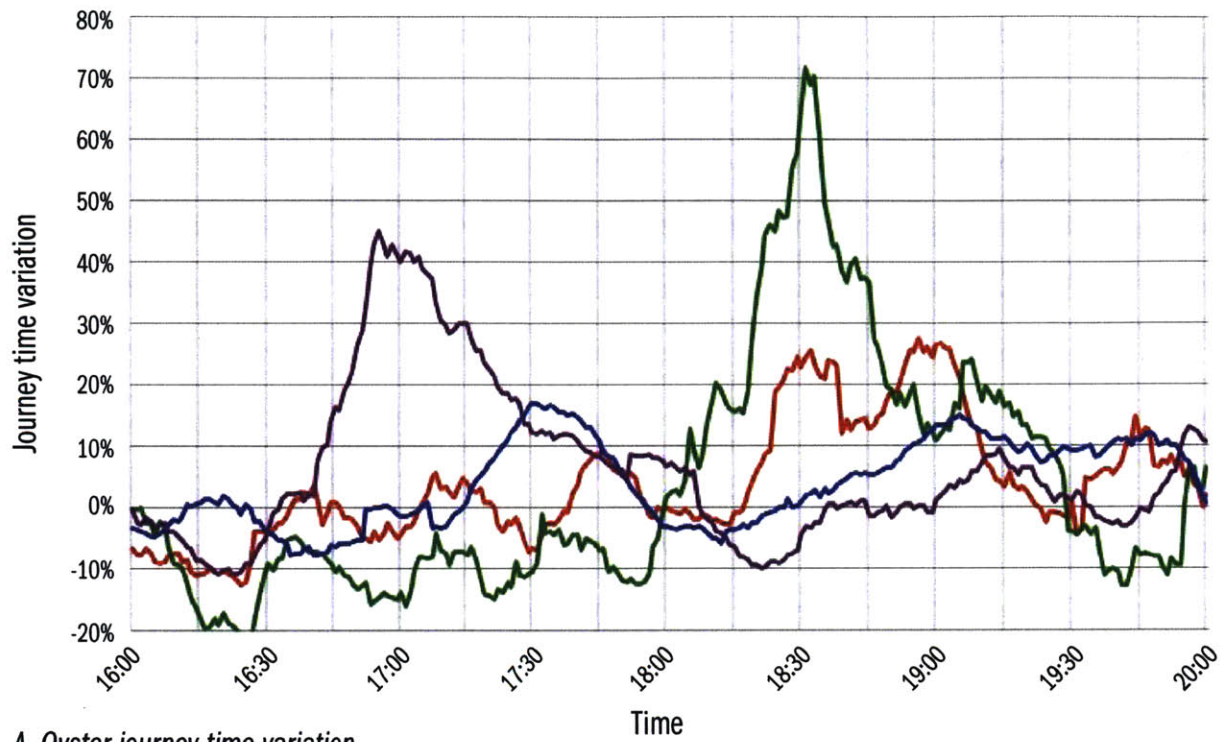
A. Oyster journey time variation.
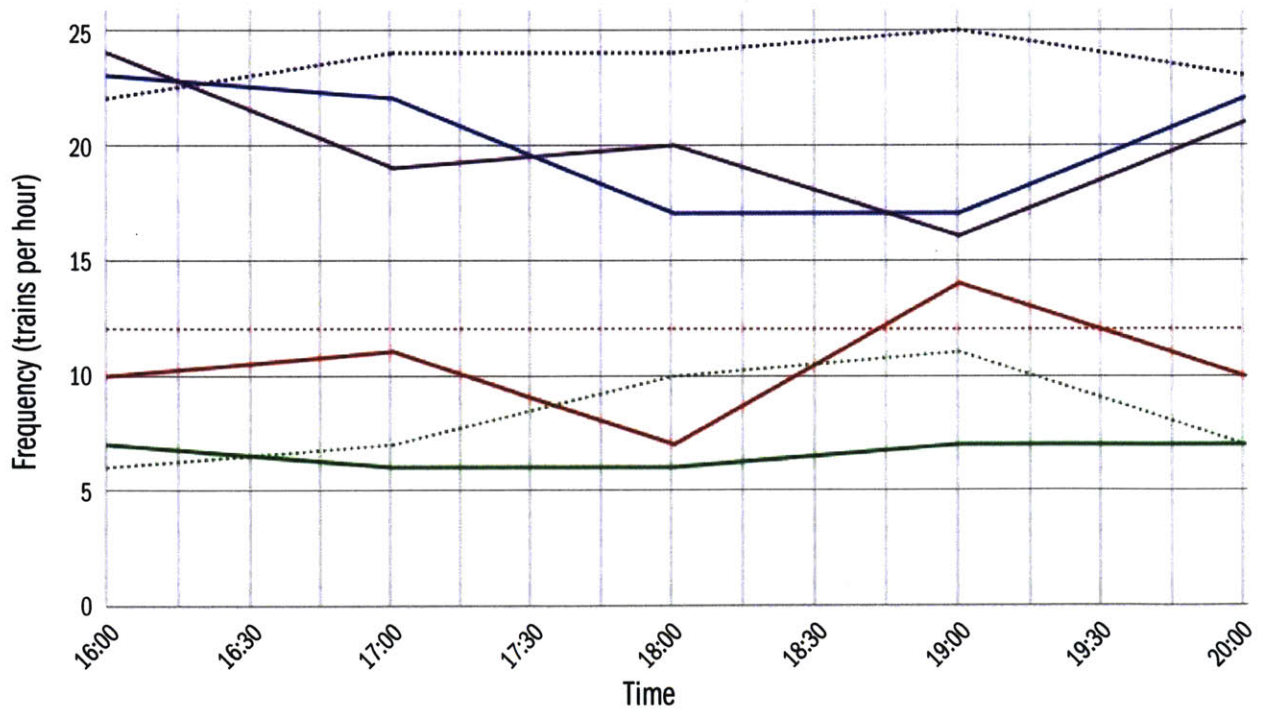
B. Train frequencies per hour at representative stations, compared to schedule.

—— Trunk   —— North   —— Heathrow   —— Rayners Ln   ···· Scheduled

Figure 5-21: Comparing Oyster journey time variation and train frequencies, eastbound (evening).

One way to improve the value of the NetMIS data analysis is to examine all headways during the period studied, not just the frequency of trains over a one-hour period. Figure 5-22 shows the headways at representative stations. This analysis provides a more in-depth view of what problems are occurring by showing specifically where and when the highest headways occurred. One major difference between looking at the headways themselves and the variation from the schedule, is that while with the former, it is difficult to see much disruption for the frequently served part of the system (the North and Trunk sections), the latter compares every section to its normal situation. This discrepancy provides two views of the same data: The former indicates the general effects on passengers (as previously described, they may not be seriously affected by higher headways in high-frequency sections) and the latter shows effects on operations, which likely are more seriously affected by headways in terms of percentage variation.

A point to note here is that while headways are one interesting indicator of service performance, London Underground has developed headway management tools that are intended to improve the regularity of headways even during disruptions. On the Piccadilly Line, train drivers are provided count-down clocks in stations in Central London that give them information about when they should depart; this regulates headways, ensuring relatively evenly spaced service even when there are fewer trains operating on the line. As such, some of the variation in headways documented in Figure 5-22 may be resolved over time.

Figure 5-23 compares journey time variation and headway variation eastbound between 16:00 and 20:00 on April 24[th]. This comparison provides a bit more context for the disruption that occurred. Large headways occur on the Rayners Lane section before the significant increase in journey time variation from the norm, as expected. For the North section, however, major journey time variation appears to occur before the variation of headways. This may be because of train regulation that is instructing drivers to move more slowly through the system, therefore maintaining relatively steady headways but at a lesser pace.
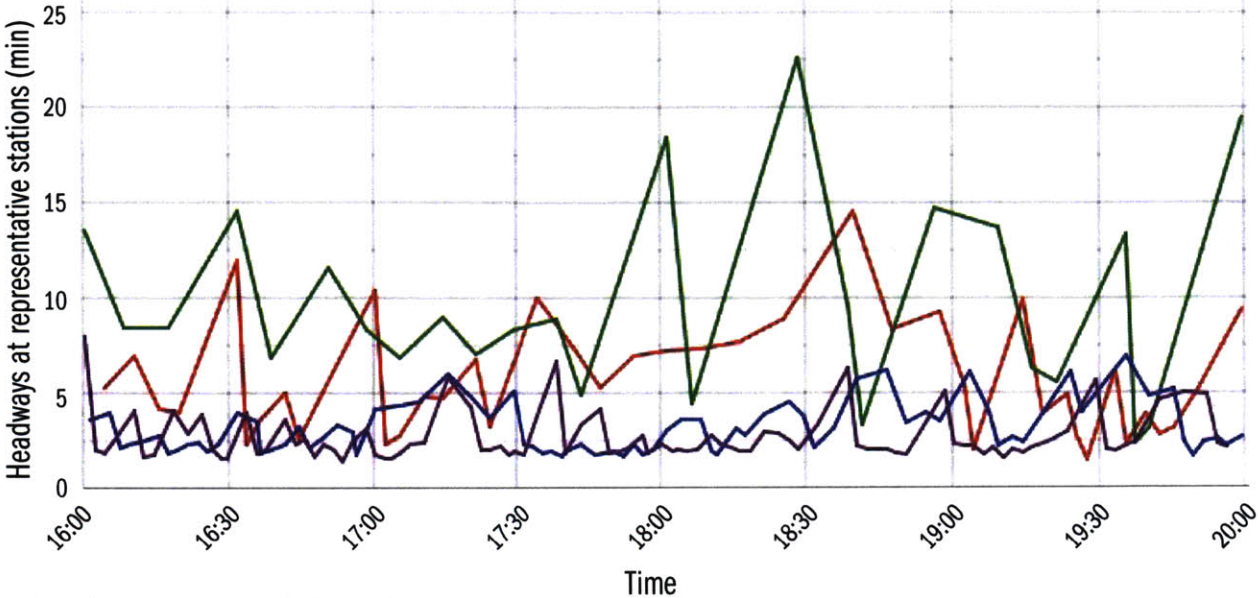
A. Headways at representative stations.
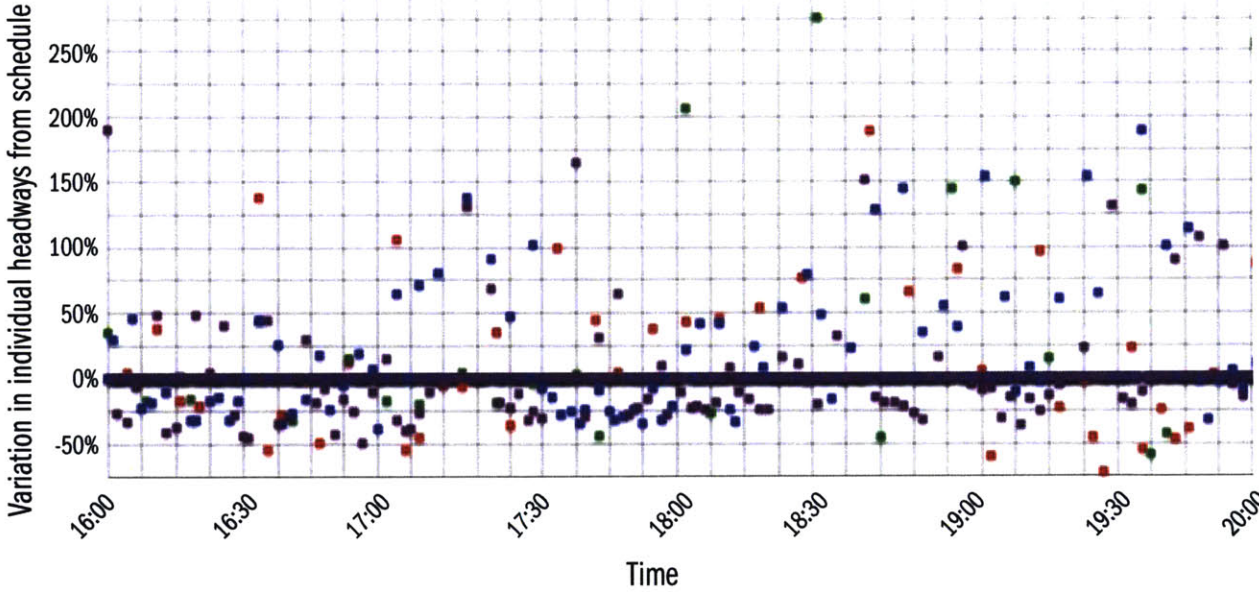
B. Variation in schedule for each headway

— Trunk   — North   — Heathrow   — Rayners Ln

Figure 5-22: Headways and headway variations at representative stations, eastbound (evening).
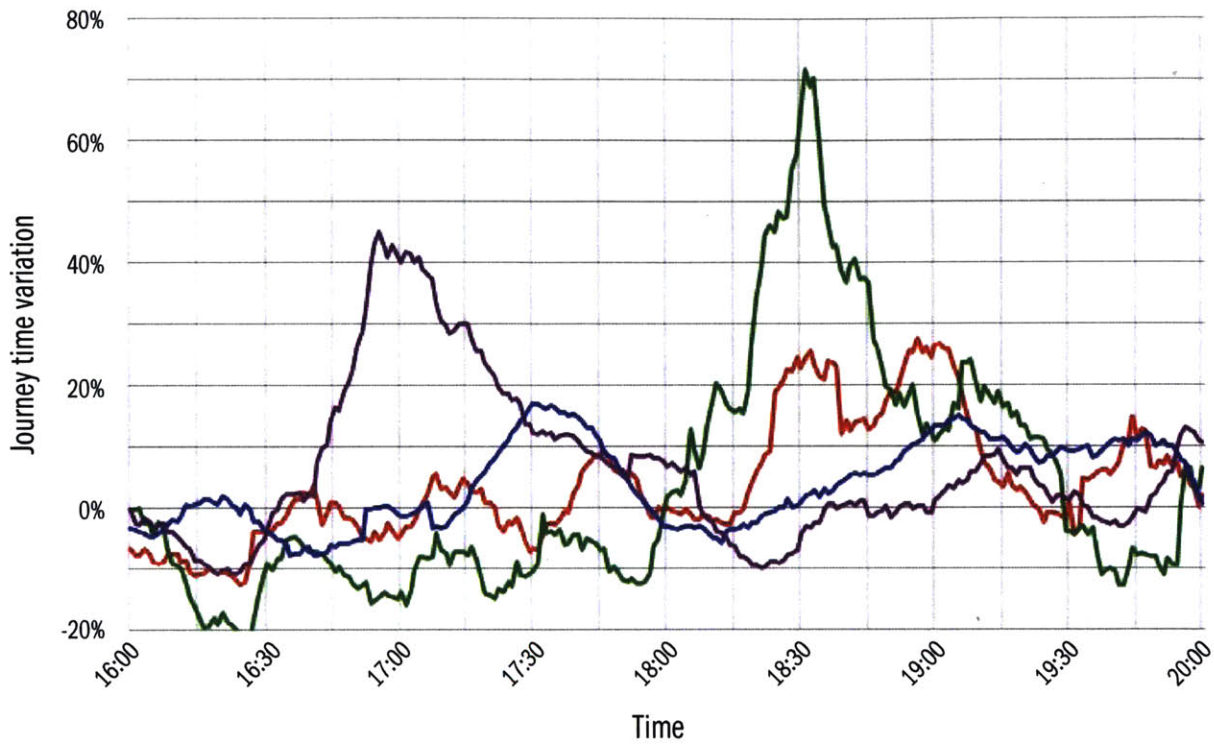
A. Oyster journey time variation.
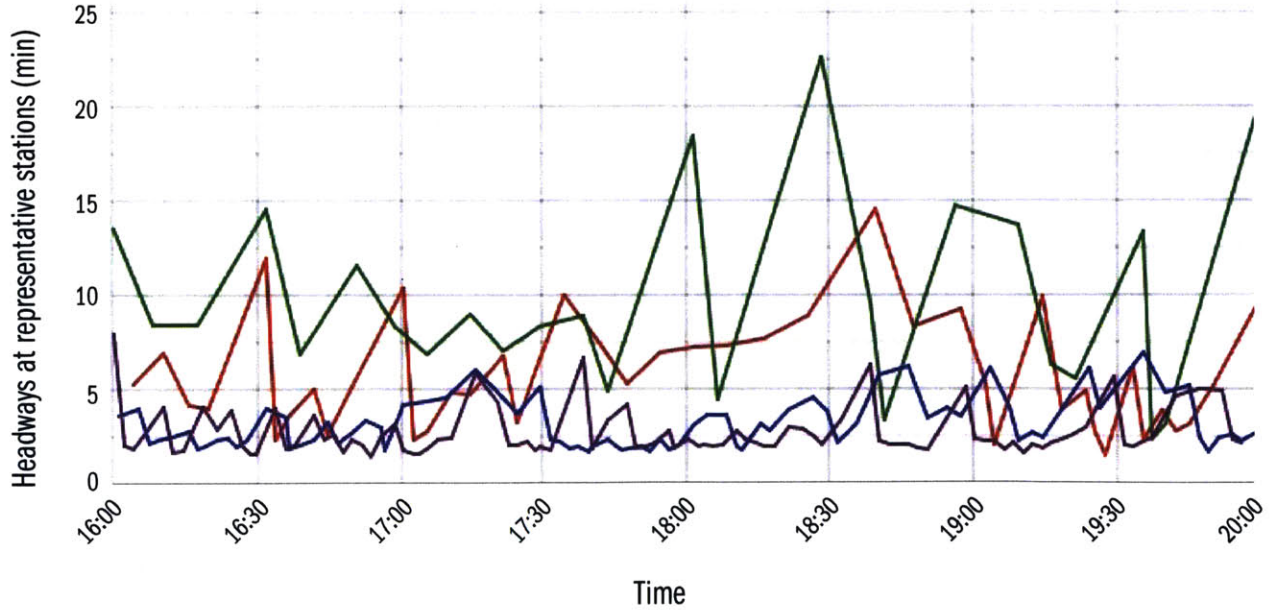
B. Headways at representative stations.

— Trunk   — North   — Heathrow   — Rayners Lane

Figure 5-23: Oyster journey time variation and headways, eastbound (evening).

By comparing NetMIS and Oyster data, it is possible to get a reasonable idea of the impacts on both passengers and service during specific disruptions. The journey time variation approach proposed here offers real potential for describing the effects on the system of an incident by revealing just where and when passenger journeys were affected, without relying on imputed information from train delays to make assumptions about how disruptions affected them.

Comparing journey time variation to service manager logs also provides a useful point of discussion. Figure 5-24 shows the variation in journey time on all four sections from 16:00 to 20:00 on April 24, 2012. These disruptions are quantified using the total disruption impact index V, summarized in Table 5-6, which shows that the most significant delays for passengers occurred westbound in the North and Rayners Lane sections of the line. It is clear from the chart that a problem westbound on the North section of the line propagated to westbound service on the Rayners Lane and Heathrow sections, which then propagated to eastbound services in these sections. Eastbound service in the North section also appears to have been somewhat affected by westbound disruptions.
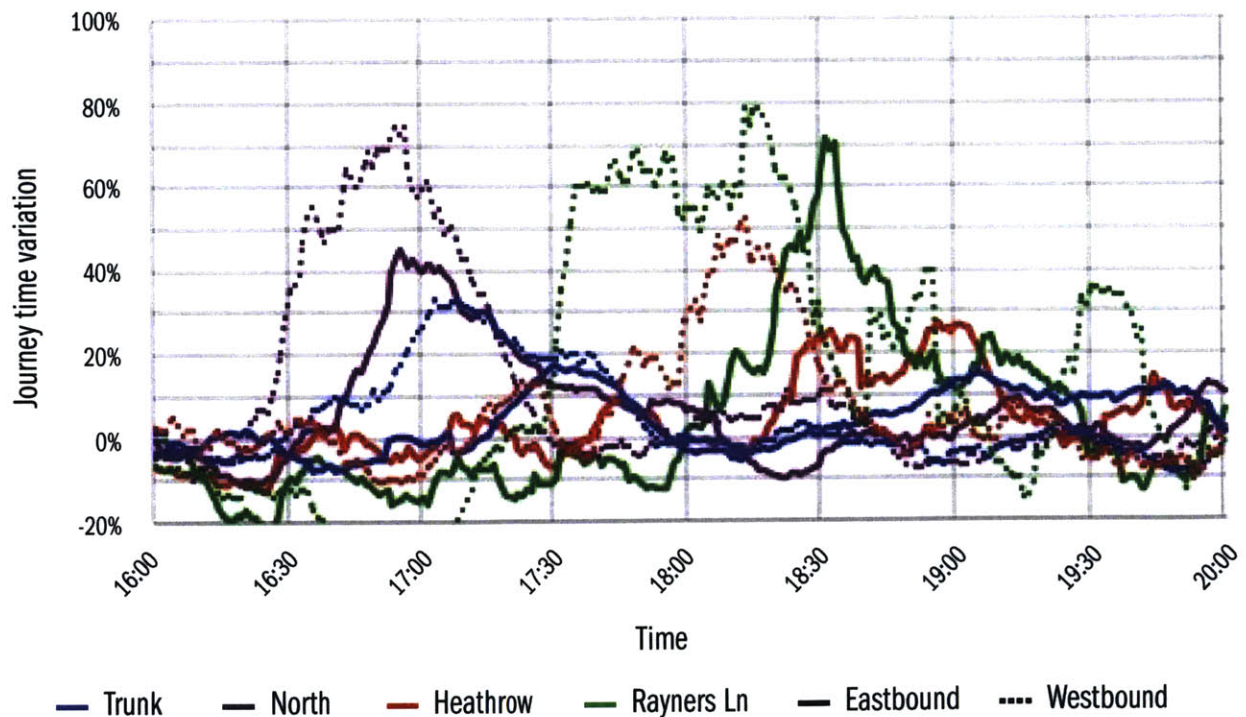
Figure 5-24: Journey time variations, east and westbound journeys.

Table 5-6: Total disruption impact V across sections (April 24, 2012, 16:00-20:00).

Line section	Start	End	V
Heathrow eastbound	18:26	19:05	1.06
Heathrow westbound	17:59	18:28	5.51
Rayners Ln eastbound	18:19	18:48	7.51
Rayners Ln westbound	17:31	19:41	26.26
Trunk eastbound	-	-	0
Trunk westbound	16:56	17:23	2.12
North eastbound	16:48	17:21	4.49
North westbound	16:28	17:18	16.07
Total	-	-	63.02

A review of the service controllers logs provides more detail about what occurred. The train service controller only recorded one incident that lasted more than five minutes on April 24[th]: a SPAD (signal passed at danger) incident at 16:10. According to the controller log, this incident occurred at Arnos Grove westbound and was resolved in 11 minutes. When recorded in the CuPID disruptions database, as shown in Table 5-7, this data was used to establish the NACHs figure for the incident, which was a relatively high 21.82 NACHs (or 2,182 Lost Customer Hours). When visualized in terms of journey times in Figure 5-24, it is clear that the impacts of the disruption were felt quite strongly by passengers entering the system from around 16:30 to 17:15. The Oyster journey time variation analysis, however, provides additional information that CuPID does not—notably that increased travel times also occurred eastbound, and that they spread to the Rayners Lane and Heathrow sections by 17:30.

According to CuPID, at 17:41, a train operator refused to pick up a train because of health concerns at Arnos Grove westbound. This caused service to be delayed as a new operator had to be found. According to Oyster journey time records, however, this incident caused no visible impact on customers; westbound journey times in both the North and Trunk sections between 17:40 and 20:00 remained within the normal fluctuation range of ±20%. Yet CuPID, assuming that the delays affected passengers, assigns LCH values to this incident. This is one illustration of the differences in results between the existing estimations provided by CuPID and the passenger data offered by Oyster.

At 18:46, a westbound train was delayed entering sidings at Barons Court, according to CuPID,

causing a disruption. This was assigned a NACHs score of 8.19 (819 LCHs). The journey time variation

analysis shows an increase in average journey times on the Rayners Lane branch beginning at around

19:20, which might be a result of that incident.

Table 5-7: CuPID incident reports for Piccadilly Line (April 24, 2012, 16:00-20:00).

Location	Time	Duration	Incident	NACHs
Arnos Grove	16:10	11	SPAD	21.82
Arnos Grove westbound	17:41	3	Train operator refused to pick up train	2.60
Arnos Grove westbound	17:44	120	Train operator refused to pick up train	5.94
Barons Court westbound	18:46	6	Delayed departure into sidings	8.19

These examples provide evidence of the importance of considering several types of data in analyzing

the effects of disruptions. While CuPID and NetMIS offer valuable insight into train operations, they

only offer limited information about the passenger impacts in many situations—and that is where the

analysis of Oyster journey time changes may be helpful. Ideally, combining these three data sources

not only for the Piccadilly Line, but all London Underground services, can provide additional informa-

tion to supply Transport for London's disruptions analysis team the most updated and reliable informa-

tion about how efficiently riders are moving around the system.

There is room for future research on the use of Oyster data to get a better handle on the problems

encountered on the line. It also suggests that in order to fully understand the impacts of disruptions, it

may be useful to take advantage of both travel time variation data provided by Oyster *and* passenger

impacts estimates from CuPID. Using a case study from September 22, 2011, the various information

sources are compared below, offering an overall view on the impacts of a disruption that occurred

between 16:00 and 19:00, primarily affecting eastbound Piccadilly Line service in the Trunk section

of the line.

Two NetMIS measures—arrival headways and train travel times—are compared to the two Oyster

measures described in this chapter in Figure 5-25, which shows performance on the Piccadilly Line

eastbound on September 22, 2011. These provide a useful comparison for determining the informa-

tion that is gained from the different approaches.

The headways information provides little evidence of any pattern throughout the period. Though there is some short-lived increase in headways on the Trunk section beginning at around 18:15, that increase is momentary and does not correspond directly with the increase in journey times experienced by passengers according to the journey time variation analysis. It is true that the journey time variability, passenger accumulation, and train travel time information provide similar information about the status of disruption in the Trunk section. But whereas passenger accumulation and journey time variability demonstrate that there were disruptions in the Heathrow and Rayners Lane sections of the line, the train travel times data do not indicate this, illustrating that riders may have been impacted in a way that is not easily conveyed through an examination of train movements. In addition, the journey time variability information identifies the severity of the disruption on each of the sections for the average rider over the study period, while the train travel data does not illustrate this directly. This comparison thus highlights the value of using AFC data to analyze disruption impacts on passengers.

5.6 UNDERSTANDING CUSTOMER ROUTE CHOICE

As described in Chapter 3, Transport for London currently uses the Rolling Origin-Destination Survey (RODS) to estimate the route choices of passengers in the Underground network. RODS plays an important role in understanding passenger travel patterns on the system because it provides information about where customers choose to make connections. Since many passengers in the system take more than one line to reach their destinations, this information makes it possible to estimate passenger flow at the line segment level more accurately. However, while useful, RODS has a number of limitations. Because it is a survey, it only samples a small percentage of overall customers, meaning that it does not represent all travel patterns. Since it is "rolling," it uses data that can be several years old to monitor customer patterns. People surveyed ten years ago may not act the same as people today, given changes in the system. This section suggests that RODS could be improved with supplementary material based on current Oyster information.
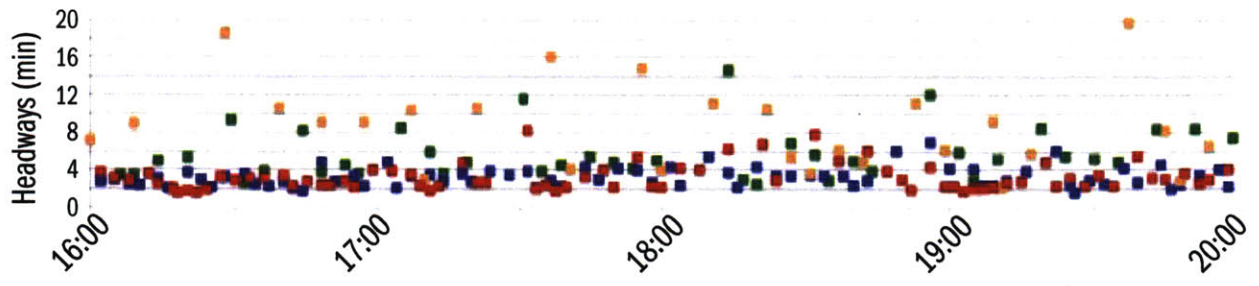
A. Headways at representative stations (NetMIS).
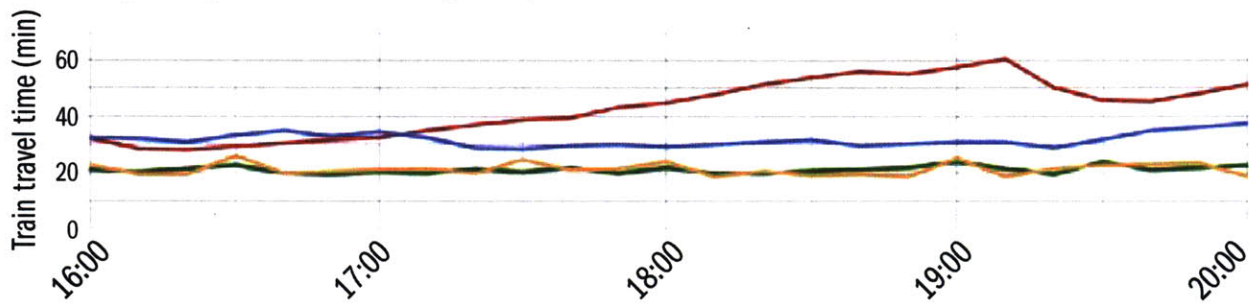
B. Train travel times across representative sections (NetMIS).
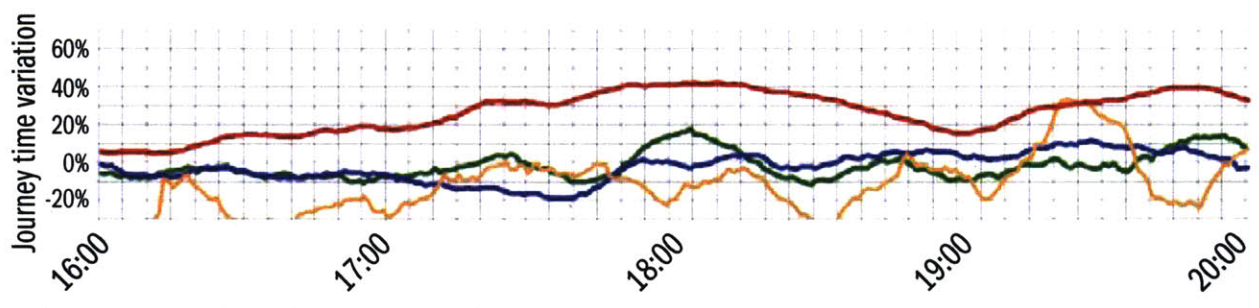
C. Journey time variation in line sections (Oyster).
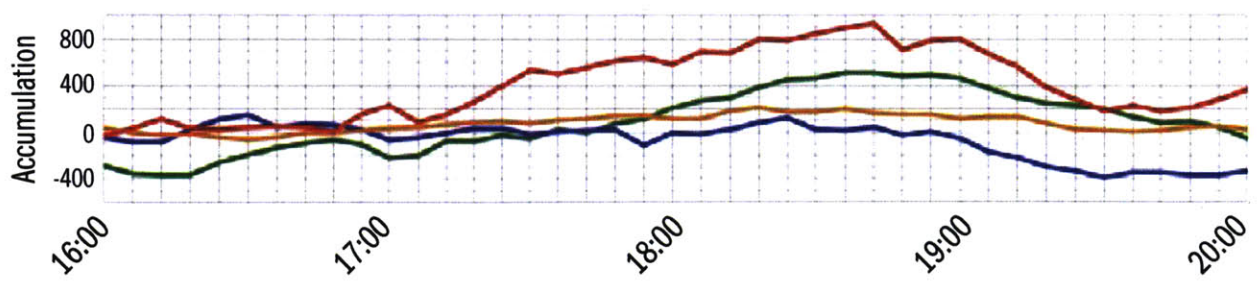
D. Absolute variation in passenger accumulation (Oyster).

— Trunk   — North   — Heathrow   — Rayners Ln

Figure 5-25: Various measures for examining extent of disruption, eastbound (September 22, 2011).

While this research primarily focuses on trips that both begin and end at Piccadilly Line stations, a look at disruptions also provides insight into customer route choice. Indeed, disruptions provide a unique opportunity to analyze what riders do when a disruption occurs on the line. By tracking how customer journey times are affected by delays, we may be able to estimate what route they took.

Customers traveling on the London Underground between Knightsbridge and Finsbury Park have two primary alternatives for travel, as shown in Figure 5-26. Though the route is served directly by the Piccadilly Line, customers making this journey have the option of switching to the Victoria Line at Green Park to complete their trips to Finsbury Park. The Transport for London journey planner does not recommend this path, in part because the transfer at Green Park involves a long walk. But the train trip between Green Park and Finsbury Park is quicker on the Victoria Line than the Piccadilly Line (about 14 minutes versus 18 minutes), so some customers might see an advantage in switching lines. RODS data suggest that 93.5% of passengers take the Piccadilly Line for this journey.
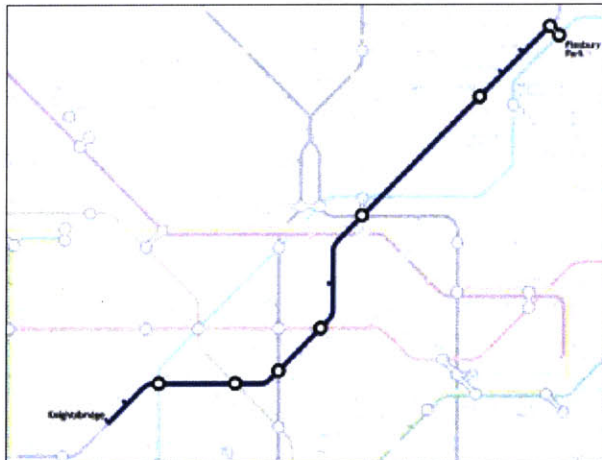
A. Direct on Piccadilly Line.   *B. Interchange at Green Park to Victoria Line.*

Figure 5-26: Routes between Knightsbridge and Finsbury Park.

However, the comparison in Figure 5-27 suggests otherwise. September 21, a "normal" operations day with few disruptions, shows median journey times of about 27 minutes throughout the day between those stations. Because of the similar trip times for customers taking the Piccadilly Line the whole way and those transferring to the Victoria Line, this data provide no information about route choice. How-

ever, on September 22, the Piccadilly Line experienced significant delays beginning at around 17:00 and ending at around 24:00; the Victoria Line had no problems during this period. Figure 5-27 shows very clearly the large increase in travel times that most customers experienced. Yet it also shows a significant number of people with average travel times during the same period; these people *must* have switched to the Victoria Line rather than remain on the Piccadilly Line, as there is no other way they could have had such low travel times. Customers were told about minor delays, which likely did not alter their travel patterns significantly. But these delays were not "minor" in that they increased trip times by almost 100%.
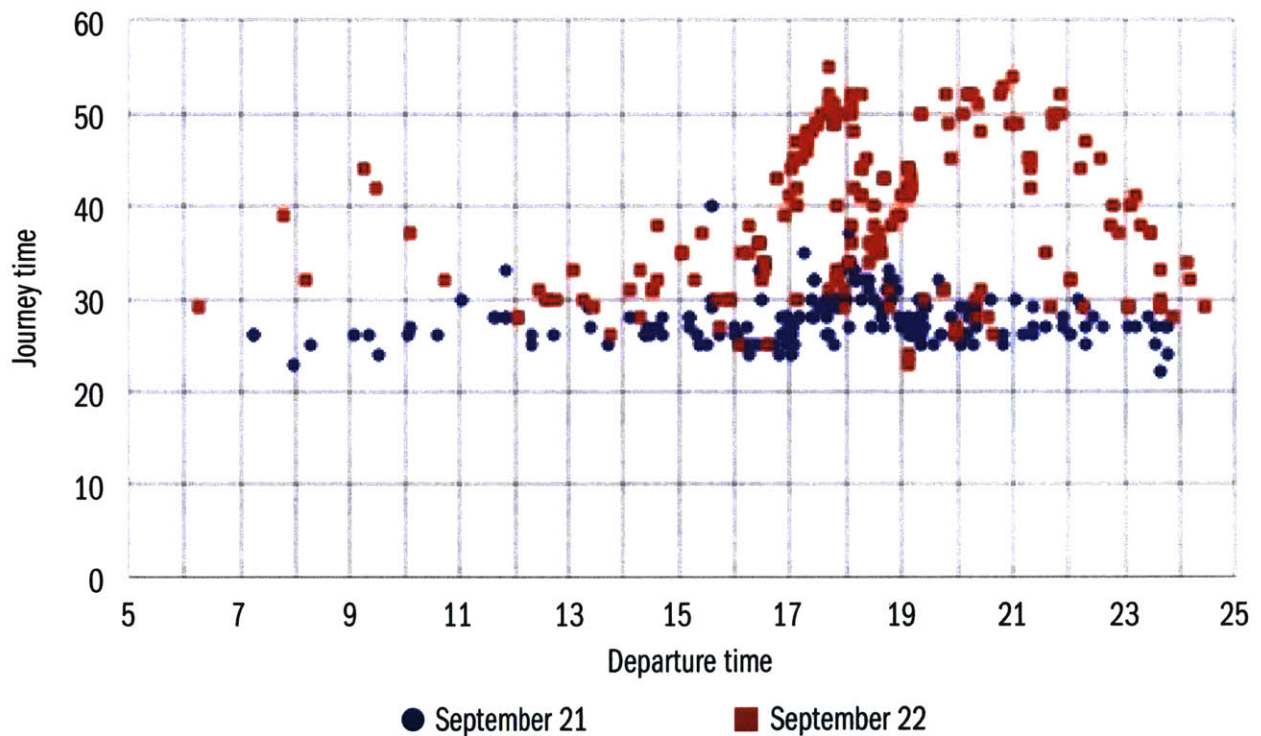
Figure 5-27: Oyster journey times from Knightsbridge to Finsbury Park (September 2011).

The data show that between 17:00 and 24:00, 72% of customers traveling between Knightsbridge and Finsbury Park had journey times of above 35 minutes. We can safely assume that they took the Piccadilly Line. But the remaining 28% likely took the Victoria Line to achieve such low journey times. This is significantly higher than the 6.5% that RODS estimates typically take the Victoria Line for a part of this journey. Some of these passengers were undoubtedly influenced by the announcement of a Minor

delay at around 17:30, or it is possible that this notification made no impact at all. This effect cannot be completely evaluated, however, a thorough study of customer response to passenger announcements has not yet been conducted by a public transport system such as TfL. Thus before making any broad conclusions from information derived from disrupted days, these sociological questions must be addressed.

Nonetheless, analyses of disrupted days can provide unique insights into customer travel patterns when one potential route path is disrupted and the other is not—especially when riders are not thoroughly informed of a problem on the line (which would alter their decision making). Further analysis of origin-destination pairs presented in Table 5-8, using disrupted days in the study period, may be a valuable undertaking for Transport for London to determine the degree to which RODS accurately reflects customer route choice on the Piccadilly Line. Using the journey time data for disrupted days can thus play an important role in informing the development of a more accurate rider flow model for the system as a whole.

Table 5-8: Origin-destination pairs where customers have two clear alternatives (one of which is the Piccadilly Line) for making their journeys (Source: RODS).

Route	RODS OD% Piccadilly	RODS OD % for Alternative Route
Green Park-Finsbury Park (eastbound)	0.0	100.0 (Victoria Line)
Finsbury Park-Green Park (westbound)	2.6	97.4 (Victoria Line)
Green Park=King's Cross (eastbound)	0.0	100.0 (Victoria Line)
King's Cross-Green Park (westbound)	0.0	100.0 (Victoria Line)
King's Cross-Finsbury Park (eastbound)	0.0	100.0 (Victoria Line)
Finsbury Park-King's Cross (westbound)	0.1	99.9 (Victoria Line)
Hammersmith-King's Cross (eastbound)	50.1	49.9 (Circle or Hammersmith and City Lines)
King's Cross-Hammersmith (westbound)	66.5	33.5 (Circle or Hammersmith and City Lines)

6 | Using AFC data to monitor disruptions in real time

Transport for London's Network Operations Centre (NOC) provides customers with the latest information about the status of service across the regional transit network, which is available online and at stations. Updates are communicated to customers to reflect the delays that might be expected on specific lines, with the assumption that customers can benefit from this additional information about their journeys. In theory, accurate information offered to customers may allow them, in response to a delay, to change their travel modes or their journey departure times, alter their routes, communicate their late arrival times to friends or colleagues, or simply be more understanding of the disruptions they face on their journeys.

This chapter summarizes how these announcements are currently made, with a focus on a comparison between the notifications made to customers and the service manager assessments of disruptions during the study period described in Chapter 4. The chapter then evaluates how the Oyster journey time variation method described in Chapter 5 could be used to monitor service performance in real time. It compares this method with information provided by live passenger accumulation and NetMIS-based measures, including large headways, lateness, and train travel times. Given this universe of data available to monitor service conditions in real time, this chapter notes the positive and negative characteristics of several types of information. It particular, it shows that the journey time

variation measure, despite suffering lags in information availability, nonetheless can offer insight into the status of service on the line in real time. A case study is used to show how information about route performance changes over time, and suggests a mechanism through which real-time journey time variation could be used so that announcements to customers better reflect likely delays. Finally, the chapter provides an example of how such Oyster-based announcements could differ from the existing NOC notifications, under certain conditions.

As such, the goal of this chapter is to explore the potential of incorporating several types of real-time information, including AFC data, to refine the notifications given to passengers. This suggests a framework for using Oyster information as part of an updated service status notifications regime.

6.1 THE PASSENGER NOTIFICATIONS SYSTEM

Because the London Underground provides frequent service throughout the day at almost all stations, customers rarely rely on schedules to decide when to arrive at a station to start a journey. As shown in Figure 3-5, on the Piccadilly Line, trains are scheduled to arrive at headways of 10 minutes or less at all times of day at all stations with the exception of those from Rayners Lane to Uxbridge, where Metropolitan Line trains also provide service. These frequencies make walk-up service the norm for riders and allow them to make travel decisions without first consulting schedules.

Yet the reliability of train service can affect passengers significantly, sometimes forcing customers to wait far longer than just the few minutes they expect. Transport for London, like many public transport agencies, is working to address the occurrence of delays on the line. At the same time it has developed a standard approach to informing customers about problems when they do occur, such that customers can make more informed decisions about their travel paths when they are using, or thinking of using, the system. The NOC, operated by the London Underground, has established criteria that are used to inform customers about the status of service; each LU line is evaluated independently. There are four operations status indicators communicated to customers, ordered by magnitude of disruption:

- **Good service**, indicating no problems on the line;
- **Minor delays**, indicating small delays on the line;
- **Severe delays**, indicating significant delays on the line;
- **Line closure**, **Part suspended**, or **Suspended**, indicating no train service on part or the entire line.

Good service and line suspension are self-explanatory, but customers are provided no additional information about what distinguishes a "minor" from a "severe" delay. This leaves the onus of decision making on the rider during such disruptions, though LU staff suggest that "minor" delays are meant to encourage customers to stay on the route they had planned (with the expectation of slower-than-normal service) and "severe" delays encourage customers to change journey plans. If there are minor delays, customers are expected to assume that travel times are likely to be longer than normal, but the magnitude of the delays during this disrupted period is not made clear to riders. This uncertainty also exists during periods characterized by severe delays. For example, TfL does not provide customers estimates of when normal operations are expected to resume. The agency, though, is working to develop measures of estimated resumption of service (EROS) that may be implemented in the next several years.

Nevertheless, the NOC has clear guidelines to determine whether a delay is "minor" or "severe." Four metrics—headways, slow trains, stoppages, and percentage of scheduled trains in service—are used to establish the service quality on a line. The criteria for each metric differ by line; those for the Piccadilly Line are shown in Table 6-1. It should be noted that the criteria are different during "core times" (from 7:00 to 9:30 and 16:30 to 19:00) than during the rest of the day, because of the importance of these peak periods. In addition, TfL defines suspended service as existing when there has been no train movement for at least 15 minutes.

Table 6-1: Criteria for determining service status on the Piccadilly Line. Source: TfL NOC.

A. Minor Delays.

Station	Trains between Heathrow Airport-Arnos Grove		Trains between Acton Town-Uxbridge and Arnos Grove-Cockfosters	
	Core Times	Other Times	Core Times	Other Times
Headways	3x normal lasting > 10 mins for more than one train	4x normal lasting > 15 mins for more than one train	2x normal (not consecutive cancellations)	2x normal (not consecutive cancellations)
Trains moving slowly	> 10 mins of blocking back with 3x normal headway for more than one train	> 10 mins of blocking back with 3x normal headway for more than one train	> 20 mins of trains moving slowly in an area but trains are not being terminated short of scheduled destination	> 20 mins of trains moving slowly in an area but trains are not being terminated short of scheduled destination
Stoppage/Sit down	Between 10-15 mins	Between 15-20 mins	Sit-down < 15 mins; advise minor delays until gap regulated	Sit-down < 15 mins; advise minor delays until gap regulated
% of Scheduled Trains in Service	Between 75-85%	Between 70-85%	Between 75-85%	Between 70-85%

B. Severe Delays.

Station	Trains between Heathrow Airport-Arnos Grove		Trains between Acton Town-Uxbridge and Arnos Grove-Cockfosters
	Core Times	Other Times	All Times
Headways	4x normal lasting > 15 mins for more than one train	5x normal lasting > 20 mins for more than one train	3x normal or there are > 2 consecutive cancellations
Trains moving slowly	> 15 mins of blocking back and/or trains being terminated early	> 20 mins of blocking back and/or trains being terminated early	> 20 mins of blocking back and/or trains being terminated early and/or trains working platform-to-platform
Stoppage/Sit down	> 15 mins gap for more than one station	> 20 mins gap for more than one station	> 15 mins
% of Scheduled Trains in Service	< 75%	< 70%	< 70%

The criteria described in Table 6-1 are evaluated based on live NetMIS data, which is communicated to service managers using the TrackerNet live train-tracking system. The data have been refined significantly over the years to best reflect TfL's understanding of the manner in which train travel conditions affect passenger flow, and are unique to the Piccadilly Line (the information for each service reflects its particular characteristics). As the table demonstrates, trains must be experiencing significant delays, with more than 10 minutes of blocking back or headways that are three times longer than normal (at peak hours), for even a minor delay announcement to be made. These service metrics are relatively conservative, requiring service to be significantly affected before passengers are informed of any problems. This is a reasonable approach, since it would be inappropriate to warn customers about service problems that are very minor or that may last only a short time. Too many warnings may be even less useful than too few.

The NOC uses NetMIS data to determine whether any portion of the London Underground system fits either the minor or severe delay status criteria and then updates the notifications provided to customers accordingly. That said, based on discussions with TfL and LU staff in January 2013, NOC personnel and train service controllers base announcements to customers only partially on the criteria used in Table 6-1. Often more relevant are subjective assessments of service based on observed train movements; in other words, notifications do not always follow the guidelines and differ depending on who is making the decision. Nonetheless, notifications are distributed through the following media, which offer information about all of the system's lines:

- Screens located in stations;
- The TfL Website;
- Twitter (TfL, London Underground, and Line-specific accounts);
- Email and text message alerts for which customers have signed up; and
- Third-party applications that receive TfL's data feed.

All of these content providers also offer additional information about issues on the line that do not affect train service directly such as elevator outages. In addition, customers have access to next train information via LED indicator screens on train platforms. This data can also be accessed online for most stations. All in all, customers are provided access to a variety of sources with information about

many aspects of their trips.

The study period discussed in Chapter 4, consisting of 29 somewhat (or very) disrupted days in April, May, and June 2012, offers a starting point for analyzing the information provided to customers. Figure 6-1 shows the passenger announcements made about the Piccadilly Line on each of the study days. This chart demonstrates that a number of the study days, despite their elevated PWT scores (as shown in the rightmost column), had no delay announcements; May 31, for example, had a PWT score of 27 seconds but no announcement. Figure 6-1 provides information only about the *most severe* announcements made during any specific period, meaning that if there is a part line suspension on one branch but minor delays on another, only the former is shown.

Date	Hour	PWT
3.04.12		111
10.04.12		25
13.04.12		25
16.04.12		20
18.04.12		55
24.04.12		41
25.04.12		28
26.04.12		46
27.04.12		48
30.04.12		41
4.05.12		52
11.05.12		22
14.05.12		20
15.05.12		34
22.05.12		46
23.05.12		40
24.05.12		26
25.05.12		66
28.05.12		23
29.05.12		27
31.05.12		27
4.06.12		103
6.06.12		22
8.06.12		22
11.06.12		51
12.06.12		20
14.06.12		33
15.06.12		30
18.06.12		22

Legend: ■ Part suspension   ■ Severe delays   ▨ Minor delays

Figure 6-1: Announced delays in study period.

Figure 6-2 provides the same information, but separately for each section of the Piccadilly Line (as defined in Chapter 5). Table 6-2 summarizes the data presented in the figure. In both cases, only the most serious problem is included when multiple announcements overlap when tabulating the overall

conditions. In other words, if there is a part suspension in one part of the North section at the same time as a minor delay on another part, only the former is counted. If different types of delays follow one another with no break, they are counted as one continuous delay.

A. Trunk section

B. North section

C. Heathrow section

D. Uxbridge section

■ Part suspension   ■ Severe delays   ■ Minor delays
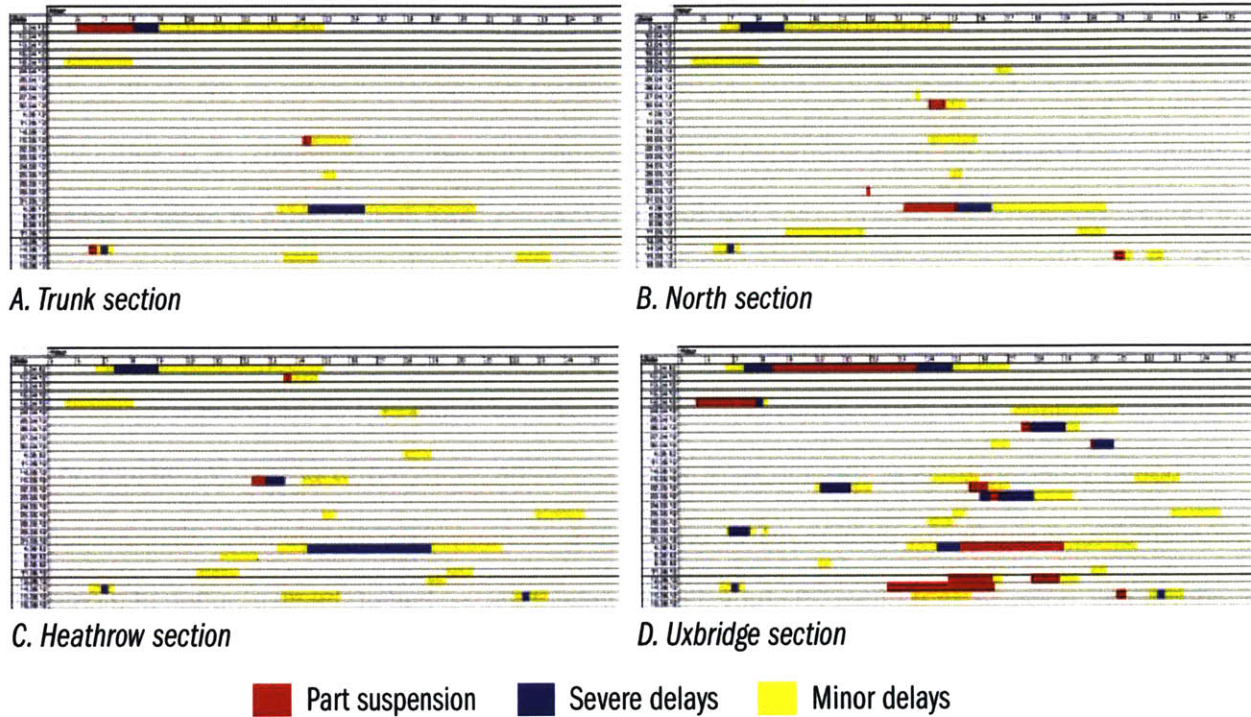
Figure 6-2: Announced delays, by line section.

Table 6-2: Delays announced in study period, by line section.

Line section	Number of announced delays	Minutes of suspended service (rounded to 5 mins)	Minutes of severe delays	Minutes of minor delays	% of service time with notifications*
Trunk	8	155	190	1,095	4.1
North	14	175	185	1,305	4.8
Heathrow	17	45	440	1,650	6.1
Uxbridge	26	1,165	590	1,605	9.7

* Of 34,800 total minutes of service on 29 study days.

Both Figure 6-2 and Table 6-2 indicate that the Uxbridge section is the branch of service that was most likely to receive passenger notifications. In fact, during the 29 study days, passengers on that section were informed of some delay in the service—whether minor, severe, or suspended—9.7% of the total service time. Discussions with Piccadilly Line staff reveal why the Uxbridge section (from Acton Town to Uxbridge) experienced far more passenger announcements than other sections of the line (notably the Trunk), despite it being the source of roughly an equivalent number of incidents as the North or Trunk sections (see Chapter 4). The explanation is that Piccadilly Line service controllers, as a matter of policy, reduce service to the section of the line between Rayners Lane and Uxbridge in order to preserve service on other parts of the line. Because this outer branch is also served by the Metropolitan Line, customers have an alternative. As a result, it is misleading to examine only the passenger notifications information to determine the presence of incidents on any individual part of the system.

Though other sections of the line were less likely to be subjected to passenger disruption announcements, a significant percentage of passengers on many days experienced such announcements (recall, however, that these 29 study days represent the top half of the distribution of disrupted days from April to June 2012). The frequency of passenger announcements suggests that the quality of the information provided is of utmost importance. Riders with the ability to check the status of service in advance may choose to change mode or route, alter the timing of their trips, or not travel at all if they understand that service on the line is disrupted. Others who are already in the system may decide to change their planned routes if they learn that there are major problems in one specific area. This is particularly true for travelers in the trunk section of the line, where there are several alternatives, both on the Underground and on local buses, for many journeys.

It is clear, however, that the existing passenger notifications system, while useful, is not always fully representative of the disruptions occurring on the London Underground. Figure 6-3 compares the most severe passenger notifications in the four sections of the line at all times during the study period with all incidents that service managers recorded as lasting at least 5 minutes (a short time, but one that allows virtually all incidents to be identified clearly). The illustration shows that a large number of

incidents recorded by service controllers are never announced to customers. In addition, many incidents that *are* announced to customers are only noted publicly fifteen minutes to as much as an hour after the controller indicates that the incident occurred.

As shown in Figure 6-3, the majority of long-duration incidents recorded by service controllers *are* eventually passed on to riders using the system. However, in many cases, the notifications system does not always accurately indicate where and when disruptions are occurring (see also Section 6.4 for an example).
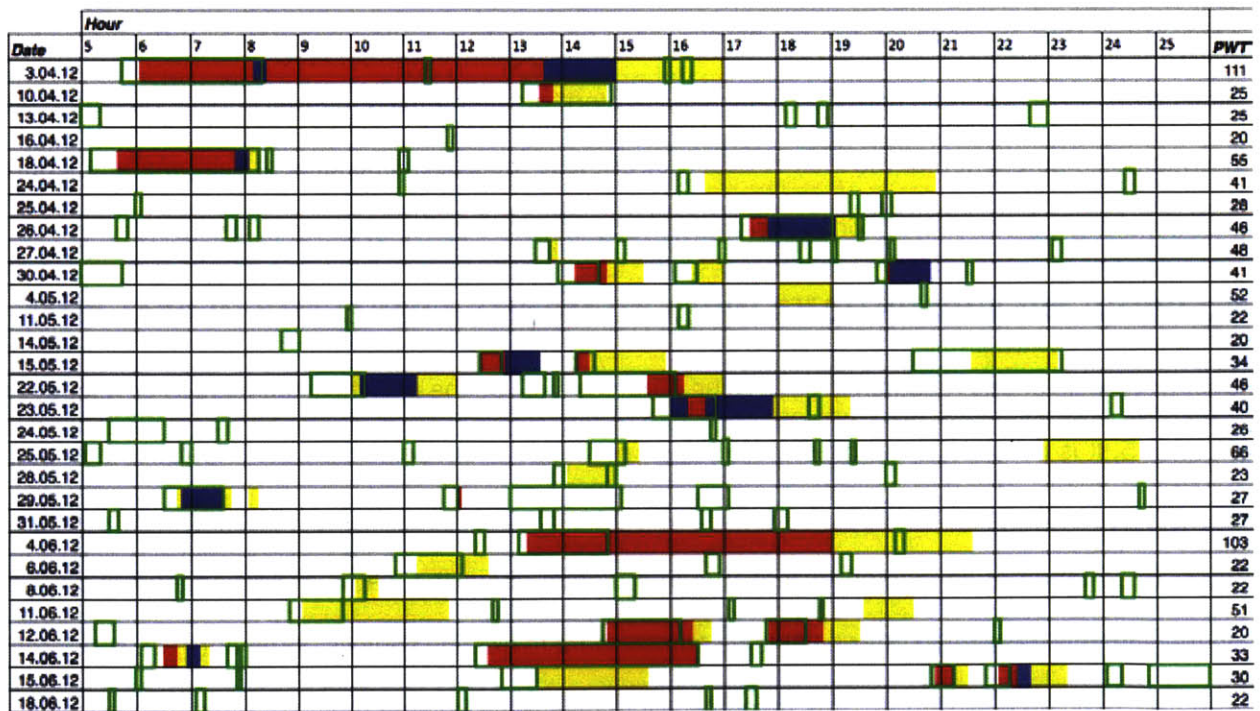
Controller comment — Incident with a duration of ≥5 min
Customer notification* — Part suspension — Severe delays — Minor delays

* Chart shows most severe disruption notification announced to customers at any point on the line.

Figure 6-3: Comparing controller assessments of problems and delay announcements to customers through the passenger notifications system.

The clearest explanation for the discrepancy between controller comments and passenger notifications is that the two use different metrics to evaluate disruptions. While service controllers record information that is relevant to every part of the service—including items only marginally related to passenger travel times such as delays in crew duties—the NOC's TrackerNet-based system is designed specifically to show train movements. In addition, not all disruptions are perceived to be "bad enough" to communicate to customers; if the problem is expected to be resolved quickly, why concern customers needlessly?

Nevertheless, there would be great value in working to improve the customer communications system for the most basic of reasons: not all customers trust the existing messages. Anecdotal evidence based on conversations with LU passengers, and LU and TfL staff suggest that the disruptions notifications that are provided are not always relied upon because passengers frequently experience either slow service with no corresponding service alert, or, alternatively, normal service despite a notification. The perception of the lack of accuracy of these messages encourages customers to ignore them, making it more difficult for agencies like TfL to influence passengers when desired.

6.2 USING THE DISRUPTIONS INDEX TO MONITOR PROBLEMS IN REAL TIME

One hypothesis of this research is that, using methods such as the analytical techniques developed in Chapter 5, it is possible to supplement existing knowledge about service on the line—based on train movements—with additional data derived from Oyster passenger taps to improve the provision of customer information. If connected to real-time information about the status of passenger movements on the line, this data could offer a valuable complement to existing sources. If the data are consistent with that suggested by NetMIS, this resource will reinforce the understanding of the service status. If they are different, however, Oyster may offer additional insights into the degree to which passenger movements are being affected by service. The in-depth nature of the Oyster information may also be beneficial in terms of targeting announcements to make them as specific and relevant as possible.

Before proposing how the Oyster data could be used to monitor passenger movements in real time, several caveats must be noted so as to emphasize the limitations of this technique. First, the

analysis conducted here has been performed retrospectively, with a complete set of data, *as if* the data were available in real time. However, all fare gates on the London Underground system cannot currently automatically transmit tap-ins and -outs by customers in real time to the NOC or to other TfL locations, a necessary step for conducting the analysis to support real-time decision making, as suggested here. While certain stations, such as several on the Jubilee Line, are capable of transmitting such information within minutes, most stations on the Piccadilly Line do not currently have this capability. Planned upgrades to the faregates, the establishment of a data warehouse within TfL, and the overall improvement in information technologies, are expected to make these resources available within the next decade, however. Because of the need to receive the most up-to-date data as possible, the lack of such real-time information is the major stumbling block that would have to be overcome before this proposal could be fully implemented. The computations described here can be conducted quickly on a low-end computer, so nothing about the analysis is technically difficult, but any implementation of this tool would require software development to automate the analysis.

Second, while a metric based on passenger travel time can provide useful information about the passenger experience, it cannot be the only tool used to determine the announcements to make to riders. Service managers who are attempting to restore service after a major incident may prefer having customers enter the system more slowly than normal in order to reduce crowding and improve the chances of a faster recovery. As such, it might be useful to maintain a severe delay announcement after the incident has ended. The decision on whether to do so must remain in the hands of a service manager, even if supported by data tools such as the one described here.

Third, unlike the NetMIS-sourced data currently used to provide information about disruptions, data provided by smartcards suffer from a significant lag between when a problem occurs *on the line* and when affected customers exit the system. The method described in Chapter 5 requires customers to leave the system before their journey times can be recorded and the extent of their delay can be understood. As such, the exact time period for the disruption's occurrence is never known precisely, although it may be estimated.

For example, Figure 6-4 compares travel time variation for customers based on their entry times (in purple), and their exit times (in gold). The entry timeline (which is used for the analysis in Chapter 5 and this chapter) shows how significant the disruption was for customers, based on the time they *entered* the system; the exit time line shows the disruption for customers, based on the time they *exited* the system. Naturally, for the day shown (April 24, 2012), the increase in travel time appears "earlier" for entering customers than for exiting customers. For example, customers who entered at 16:35 experienced journey time variations of more than +50%, but customers who *exited* at the same time experienced variations of less than 10%. Only at 16:55 did the variations in travel time for exiting passengers reach +50%.

Thus the time at which a problem on the line (such as slow trains) caused increased journey time occurred sometime between 16:35 and 16:55. The bimodal distribution of exit times for this group of passengers is also interesting, as it suggests that of passengers who entered between 16:30 and 17:00 (those who were most disrupted), about half exited between 16:45 and 17:00 and half between 17:30 and 17:45. This may be a reflection of the particular OD pairs sampled in the North section of the line profiled here. The lower variability of exiting passengers is probably a manifestation of the fact that while the delays on the line affected a select cohort of people entering the system, their journey times varied and their exits overlapped with people entering at other times who did not experience the service delay.

It is on the system, between tap-in and tap-out, that delays are manifested. As a result, performing an Oyster journey time analysis can offer some information on when the disruption occurred, but is not definitive. There is an inherent lag between when passengers enter the system and when information about their journey times is known. When real-time information is plotted in this section, it is graphed showing the travel time variation only for *exiting* passengers, but based on their *entry* times. Nonetheless, as will be shown in Section 6.3, this method still has the potential to provide insight into the severity of disruptions experienced by passengers, even compared with data sources without such a lag (such as NetMIS or the passenger accumulation measure discussed in Chapter 5).
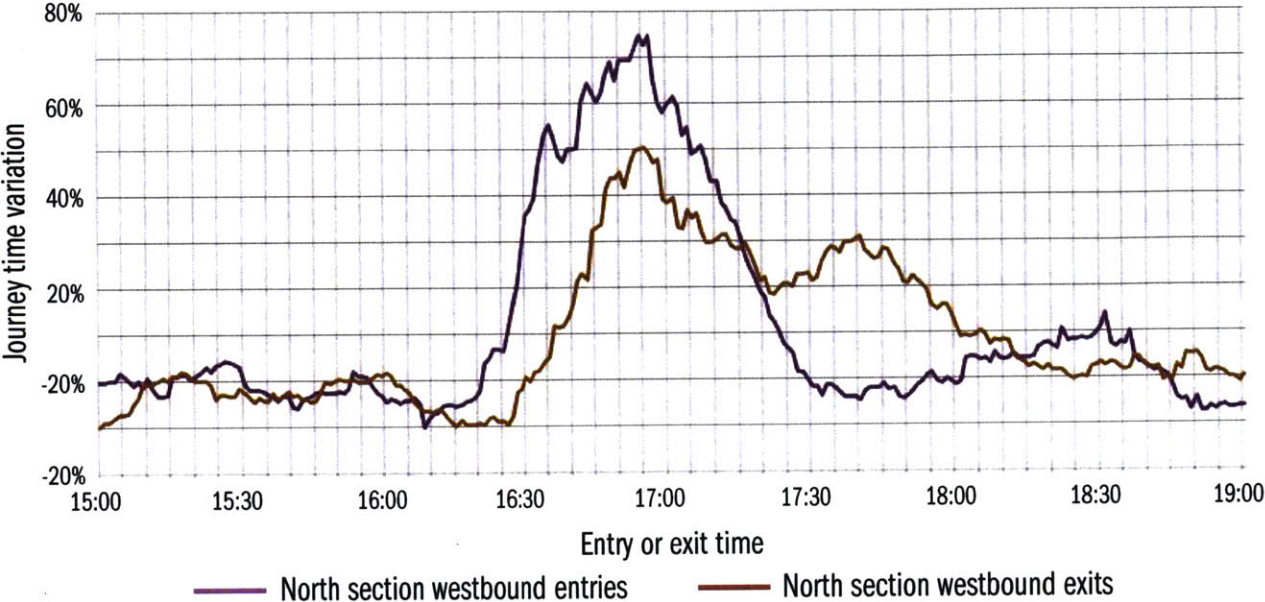
Figure 6-4: Journey time variation on the Piccadilly Line by Oyster entry and exit time.

Having noted these reservations about the journey time comparison method, it is also appropriate to describe the benefits to providing Oyster input into the real-time passenger notifications system. Unlike NetMIS, Oyster provides actual information about when a customer's journey begins and ends. As such, it offers an alternative perspective that identifies exactly which group of customers is affected by train travel problems. By segmenting the Oyster journeys recorded into baskets by section and direction, the method allows a quick way to examine the overall status of customer travel patterns on the network.

Using the same technique as described in Chapter 5, Oyster data for a large number of travelers moving throughout the London Underground system can be compiled and compared to standard performance. The method is summarized in Figure 6-5.
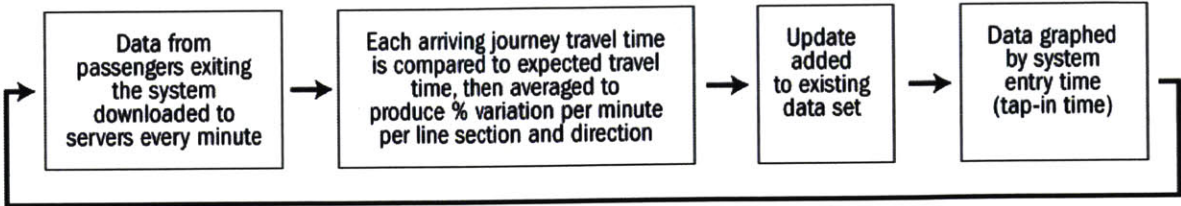
Figure 6-5: Method for acquiring and graphing real-time Oyster data.

This data could be charted on a live screen with constantly updated data in graphed form, as shown in Figure 6-6. This graph shows the travel time variation of customers on the Piccadilly Line on April 24, 2012. The chart is graphed as if it is being generated at 17:20. As such, the chart *only includes* information about customers whose journeys have ended by 17:20 and does not represent the full extent of customer data that will be available later when more customers have exited the system. At 17:20 there are many customers still on the system who entered it many minutes earlier. If they have not exited yet, however, their information cannot be included since it is not yet known. This real-time data, therefore, is only a partial representation of the true state of the line at this time.
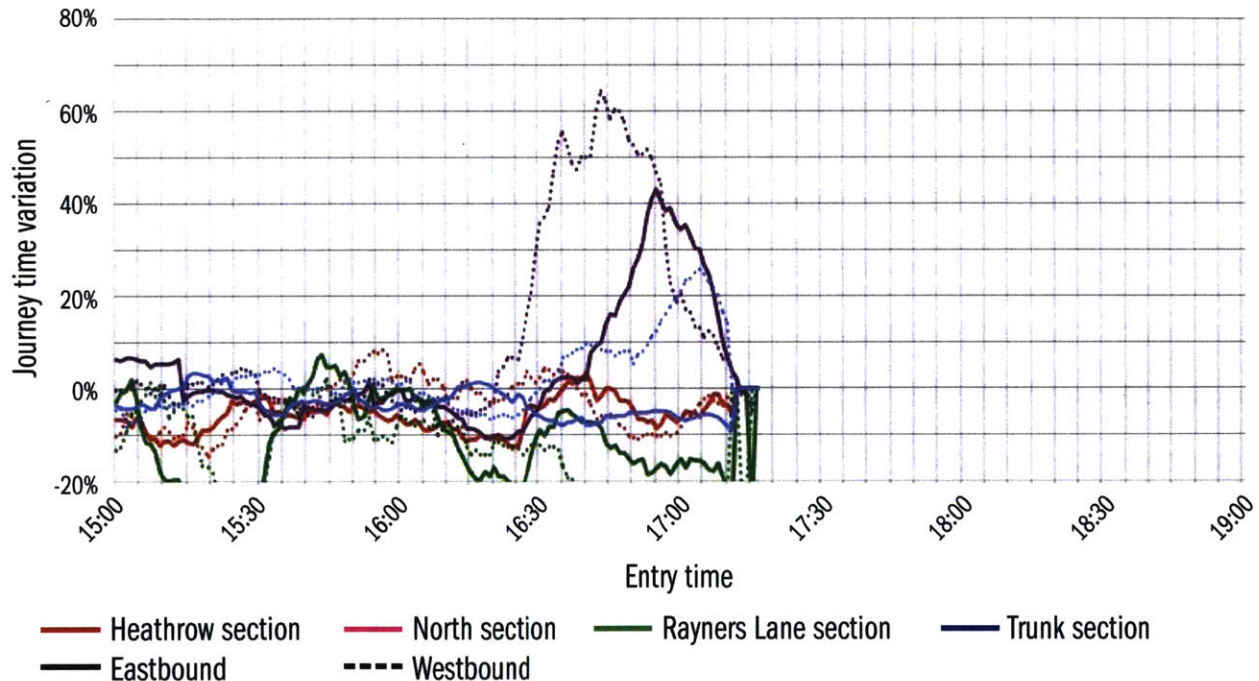
Figure 6-6: Example Oyster-based journey time indicator screen.

This real-time data is plotted using a rolling horizon technique. This means that, as in Chapter 5, the variation v presented at a specific time represents the average of v over the previous fifteen minutes. This provides a smoother curve for presentation and makes trends easier to read. However, if the data were presented in chart form as shown in Table 6-3, it may be easier to understand. This table shows the highest journey time variation for passengers entering within 15 minutes of the enquiry time from

16:30 to 18:00 on April 24, 2012. Information for each of the Piccadilly Line sections, in both direc-

tions, is presented for each ten-minute enquiry point. The table does not reflect high journey time vari-

ability for riders boarding more than 15 minutes before the enquiry point. For example, if a real-time

enquiry is made at 17:00, only entries occurring after 16:45 are considered. The sample of riders is

limited to recent boardings because the real-time data is meant to reflect current, rather than past,

conditions. Including passengers entering long before the enquiry time could be a poor reflection of

conditions on the system at the enquiry time.

As the table shows, significant disruptions for passengers (highlighted for travel time variations

of ≥ +20%) are recorded at 16:50, 17:00, and 17:10 in the westbound North section; at 17:10 and

17:20 in the eastbound North section; at 17:20 in the westbound Trunk section; and at 18:00 in the

westbound Rayners Lane section.

Table 6-3: Highest journey time variations for passengers entering within previous 15 minutes.

Line section	Highest journey time % variation v within previous 15 minutes; Enquiry time									
	16:30	16:40	16:50	17:00	17:10	17:20	17:30	17:40	17:50	18:00
Heathrow eastbound	0	0	0.3	0	0	0	0.8	0	0	5.8
Heathrow westbound	0	0	2.1	0	0	0	0	2.2	0	0
Rayners Ln eastboudn	0	0	0	0	0	0	0	0	0	0
Rayners Ln westbound	0	0	0	0	0	0	0	0	0	29.2
Trunk eastbound	0	0	0	0	0	0	0.9	7.0	10.1	2.6
Trunk westbound	0	0	0	0	8.6	23.0	16.6	10.1	7.6	0.1
North eastbound	0	0	0.6	9.7	29.4	24.8	10.9	4.3	5.9	2.9
North westbound	0	14.2	40.5	33.8	32.7	12.3	5.8	0	0	0

Table 6-3's data is illustrated in Figure 6-7, which charts journey time variability between 16:30 and 18:00 based on the information that service controllers would have every ten minutes during this period. This sequence of points documents what a live feed of Oyster data would show using the method described previously. The figure demonstrates how the information provided by a real-time analysis of Oyster journey time variation data changes as more data becomes available (as customers exit the system). It also allows controllers to examine data even for riders who boarded more than 15 minutes before the enquiry point, whose data are excluded from Table 6-3.

Consider the cohort of North section riders who entered the Piccadilly Line westbound at 16:35. At 16:42 (not shown), their travel time variation reaches roughly +15%; at 16:45, the variation reaches +23%. By 16:50, the last point when this cohort would be included in the analysis for Table 6-3, their variation increases above +40%. By 17:00, this cohort reaches its maximum variation of +58%. In this example, an indication of a disruption in the form of an increasing height of the curve for one direction and section of service begins to occur 10 minutes after the entry of a cohort of delayed riders. The information lag between when customers affected by a disruption board and when it becomes apparent that their travel is slowed (thanks to tap outs) means that the full severity of the disruption is not obvious for quite a while. Yet the trend is clear very soon—a significant degree of disruption even just 10 minutes after riders enter, suggesting that this method has potential for real-time application.

For comparison, the source of this disruption was recognized by service controllers at 16:10 as a SPAD for a train heading westbound at Arnos Grove. According to the log, the incident was resolved in 12 minutes. Examining the Oyster data suggests that any impacts of the disruption on passengers, however, occurred only later in the day. In other words, passenger delays cannot be projected directly from the information about the incident itself, but rather must be evaluated through an analysis of other data, such as that available from Oyster or NetMIS.

The example presented in Figure 6-7 demonstrates some of the benefits and failings of the travel-time-based disruption analysis. The magnitude of a problem does not become fully apparent until some time after many of those who are affected by the delays have boarded; indeed, at least some of
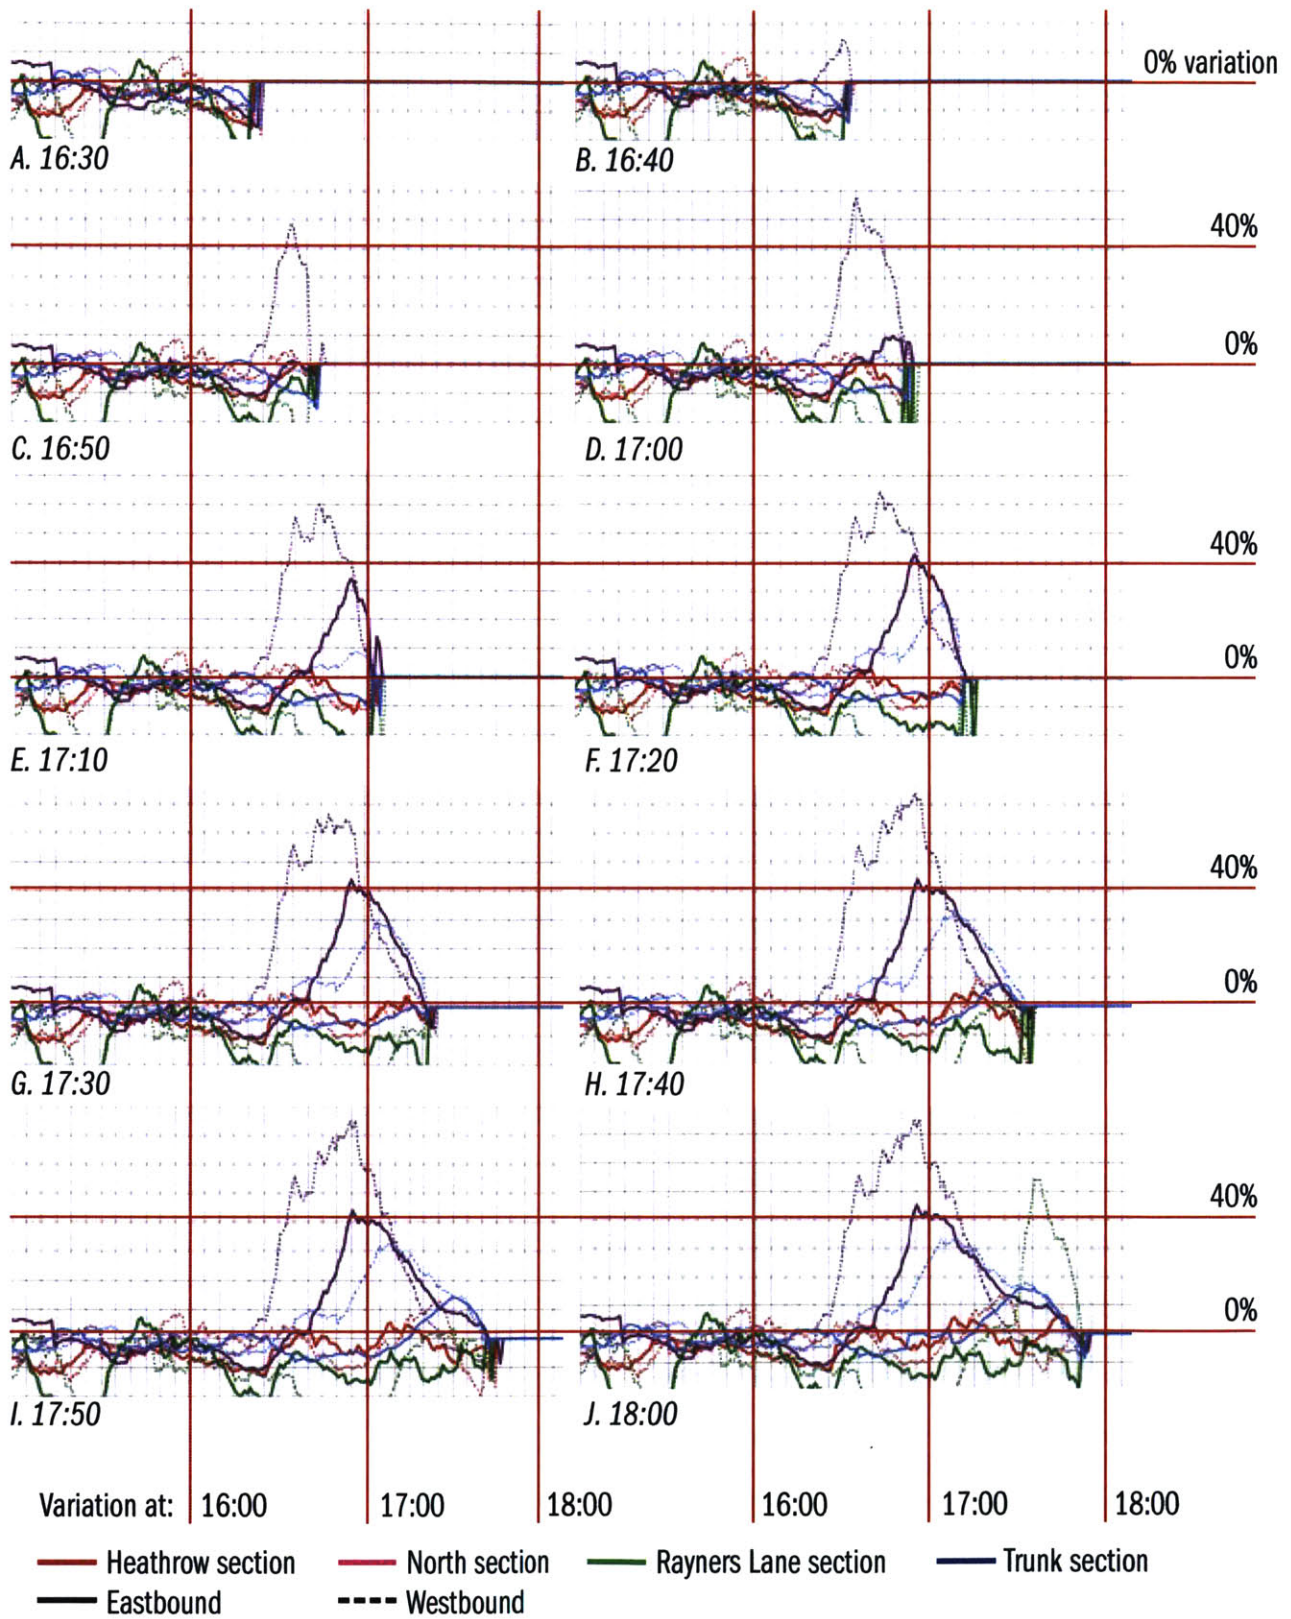
Figure 6-7: Journey time variation based on "real-time" information, graphed at enquiry time.

them must complete their journeys before their extended travel times are included. Though at 17:50, there appear to have been no significant delays affecting westbound Rayners Lane section passengers who boarded at 17:35, by 17:55 (not shown), it becomes apparent that those riders were affected by travel times that varied by +45% above normal, and by 18:00 by +55%. It should be noted that the method described in this section is significantly less effective for parts of the system with less rider-ship, such as the Rayners Lane section; the gap in information flow in this example is likely a conse-quence of the fact that there were few travelers in the period taking rides short enough to be evaluated quickly, unlike on the North and Trunk sections, which have enough customers taking short trips at any given point during the day to ensure faster availability of information about disruptions.

In addition, the full extent of the problems riders are experiencing takes longer to appear, since that requires a majority of the people boarding at any particular time to have completed their trips. Be-cause they are experiencing a disruption, their travel times take longer than normal and thus will take longer to manifest in the calculations. Even so, the existence of an increasing problem on the line for riders who boarded ten minutes before probably indicates that there are similar travel time variations for riders boarding at the analysis time. At 16:45, the information manager knows based on Oyster data that there were increasing problems on the westbound North section for riders boarding at 16:35 (though of course NetMIS data is likely also providing similar information about the existence of prob-lems on the line, even earlier). By 17:00, however, the analysis makes clear that riders who boarded at 16:45 had similar travel time increases and by 17:30, it becomes evident that riders who boarded at 17:00 *also* saw the same travel time increases. The gap between the time when we know that travel times are significantly longer than the norm (+20%) and the time when customers who experience these longer trips board is summarized in Table 6-4 for riders on April 24 between 16:30 and 17:30, for cohorts of entering riders every 10 minutes.

The evidence from Table 6-4 indicates that the Oyster journey time analysis is quickest in estab-lishing the cohorts of riders that are seriously affected by disrupted conditions for those cohorts in the middle of the disrupted period. The highlighted cohorts feature the shortest delay between the entry

Table 6-4: Delay between entry time of disrupted cohort and real-time disruption identification.

Line section	Entry time for rider cohort experiencing delays of ≥ +20%	"Real" time at which analysis tool identifies disruption occurrence	Time delay (min)
North eestbound	16:30	16:42	12
	16:40	16:46	6
	16:50	17:01	11
	17:00	17:10	10
	17:10	17:38	28
North eastbound	16:50	17:10	20
	17:00	17:10	10
	17:10	17:26	16
	17:20	18:00	40
Trunk westbound	17:00	17:20	20
	17:10	17:22	12
	17:20	17:40	20

time of a cohort that is severely disrupted and the moment at which the real-time analysis tool can identify that that cohort is experiencing journey times that raise concern. In the case of travelers on the North section travelling westbound during this period, it would have only taken the analysis tool six minutes to show that customers boarding at 16:40 were experiencing higher-than-normal variations in their journey times.

Section 6.2 has shown how Oyster journey time information might be used to identify the occurrence of disruptions, but it does not establish the degree to which this method is useful in a real-time environment when compared to other information provided by Oyster, such as passenger accumulation (riders who have not yet tapped off), and NetMIS, such as train headways or lateness. As such, Section 6.3 compares the results of the real-time Oyster-based analysis with four other measures of disrupted conditions.

6.3 COMPARING DATA SOURCES

Section 6.1 presented the criteria used by TfL to determine when to make announcements to customers on the status of service on Underground lines. NetMIS headway data provide the most important source of information with regards to operations. When trains arrive at reduced frequencies, carriages become more crowded, customers sometimes are required to wait for another train to arrive in order to board, and service generally slows down.

Chapter 5's study of the disruption in the afternoon of September 22, 2011 shows that while passenger journey time variations increased above the +20% level for customers boarding in the Trunk section of the line between 17:10 and 18:50, most eastbound train arrival headways at a station midway on the Trunk remained within the standard variation until about 18:15. This indicates that the headways measure was failing to provide a full representation of the consequence of the incident. As such, passengers who might have been warned in advance to choose a different path or wait to make their journeys remained uninformed. As shown in Table 6-5, a minor delay in the service was announced only at 18:21, more than an hour after the first customers entering the Piccadilly Line saw delays of +20% or above.

Table 6-5: Passenger notifications on September 22, 2011.

Start time	End time	Notification	Location
16:00	18:21	None	n/a
18:21	19:00 on	Minor Delays	Between Hammersmith and Arnos Grove eastbound only. Good service on the rest of the line.

Of course, TfL has adopted a conservative approach to informing customers about the impact of disruptions on the line, as stated in Section 6.1 and reinforced by the fact that the NOC does not always follow its official guidelines. Informing customers about minor problems or giving customers information about a problem that does not exist would lead to decreasing customer satisfaction. The agency has spent years perfecting this approach and has developed a tried and trusted method. Nonetheless, it is worth evaluating whether Oyster data can be used to identify the full extent of a disruption early enough to supplement the information provided by NetMIS headway and other sources.

Table 6-6 offers a comparison between five types of data all potentially available in real-time for Piccadilly Line service controllers. The data shown compare the live information available to controllers at ten-minute intervals between 16:00 and 19:00. The NetMIS and passenger accumulation data is "final" (barring technical difficulties and lost data, there is little delay in transmission and the initial data are the same as they would be in a retrospective analysis); it should be emphasized that the journey time variations presented in the leftmost data column is incomplete, as discussed in Section 6.2. The information provided by the different types of data is not directly comparable, but across all five datasets, the trend is clear: relatively normal service between 16:00 and 16:30 gradually becomes more and more disrupted.

Table 6-6: Comparisons between Oyster-sourced and NetMIS-sourced real-time information.

Enquiry time	Real-time Oyster		NetMIS			Retrospective Oyster
	Highest journey time % variation v within 15 min	Passenger accumulation F	Average train lateness at LSQ (10-min interval)	Average train headway at LSQ (10-min interval)	Running times % variation from schedule HMS-KXX (10-min interval)	Highest journey time % variation v within 15 min
16:00	7.1	-36	2.8	2.7	23.6	11.6
16:10	5.0	+110	1.1	3.3	9.2	7.2
16:20	3.0	+38	5.6	2.4	7.9	8.8
16:30	1.5	+40	0.4	2.5	11.9	14.6
16:40	8.0	+44	7.0	3.3	17.4	14.9
16:50	9.2	+9	3.5	2.6	22.6	17.1
17:00	9.7	+231	4.4	3.1	25.5	19.3
17:10	11.3	+156	4.5	3.3	34.1	21.2
17:20	7.4	+264	5.2	3.0	41.9	31.3
17:30	12.4	+529	7.4	2.7	49.3	32.3
17:40	22.5	+556	11.0	3.4	51.8	33.4
17:50	17.0	+616	12.6	3.2	65.4	40.3
18:00	26.5	+584	12.6	3.9	71.8	41.0
18:10	23.3	+682	16.2	3.5	82.8	42.1
18:20	24.8	+796	23.2	6.3	97.5	42.1
18:30	23.9	+844	25.6	4.9	107.3	40.2
18:40	21.6	+932	23.8	6.4	114.6	36.2
18:50	23.9	+713	No data*	4.4	111.2	31.9
19:00	16.4	+799	27.9	2.9	121.5	24.7

Lateness is relative to schedule; because of gaps in data, not all trains are identified by train number and therefore are not comparable to their schedules.

Table 6-6 demonstrates that train headways at Leicester Square (LSQ)—a station in the middle of the Trunk and with service characteristics indicative of the service throughout the line section—provide relatively weak information about customer service. When calculated in terms of average train headways over the preceding ten-minute interval, major disruptions only appear beginning at 18:20, when headways exceed 4 minutes for the first time; this condition lasts until 18:50, as highlighted. Retrospective information shows that disrupted conditions started affecting passengers boarding at 17:10, but there is little in the headways data to confirm this until 18:20.

Other NetMIS data points are, in fact, more informative. Average train lateness, calculated by determining the difference in arrival time at Leicester Square of trains over a ten-minute period and the times when they were supposed to arrive, shows lateness of 10 minutes or more from 17:40 on, as highlighted. This calculation, however, may be flawed: For example, it could be affected by a service change in which a significant number of trains are renumbered. Because the Piccadilly Line's train identification system suffers from significant information gaps (as described in Chapter 3), the accuracy of the lateness information cannot be guaranteed. In addition, the lateness of trains as compared to schedule may not greatly affect riders (though it would certainly affect crew assignments): If all trains are thirty minutes late, but with steady frequencies and normal travel times, riders will not perceive any change in service.

The variation in running times for trains in the Trunk section of the line, from Hammersmith to King's Cross, appears to correlate closely with the disruptions actually experienced by customers (see the two rightmost columns in Table 6-6, with highlighted areas showing most severely disrupted times). This makes intuitive sense, as monitoring train travel times between two stations is a direct reflection of customer journey times, though these data would also suffer from a delay since a train has to arrive at the destination station before its travel time can be calculated. As such, it could be useful to assess train travel times in an aggregate method much as has been done with Oyster data here, examining a set of ODs across a section of a line to determine the state of service in order to evaluate travel times on both short and long segments simultaneously. That said, the sometimes-limited accuracy of Net-

MIS data, as shown in Chapter 3, suggest that having alternative sources of data reaffirm conclusions derived from tracking trains would be useful. In this example, it appears to be correct, but this may not be true in all circumstances.

This leads to the question of whether the onset of a disruption and its severity can be identified using real-time Oyster data—using the passenger accumulation information described in Section 5.4 and/or the real-time variability method described in Section 6.2. As Table 6-6 shows, passenger accumulation on the Piccadilly Line exceeds 500 passengers above normal beginning at 17:30, as highlighted, significantly sooner than the headway measure implies a reason to be concerned, and ten minutes before the lateness measure alert. As with the running times information, then, it potentially provides a very useful source of information about the seriousness of a delay.

The final method is the journey time variation metric discussed in Chapter 5 and refined for real-time purposes in Section 6.2. This variation measure, like the train travel time information, suffers from a potentially fatal flaw: A lag in information between when the disruption is likely to have occurred and when the data is provided. This is a consequence of the fact that the journey time variation can only be calculated at the *end* of a trip, when a customer taps out of the system. This is likely to cause significant delays in understanding the severity of a disruption. Nevertheless, Table 6-6 shows that a "significant" level of disruption becomes evident using this method beginning at 17:40, which is the same moment at which the NetMIS-based train lateness information shows a serious disruption occurring and just ten minutes after the passenger accumulation information begins to indicate a problem on the line. That said, it comes a full half-hour after the running times information shows a service delay on the line. Even so, this measure of customer delays appears to provide useful real-time information about the status of service.

These five measures are compared with the actual (retrospective) passenger journey time variability in Figure 6-8. The measures are charted at a relative scale to allow for a direct comparison between them. As the graph demonstrates, of all the measures, the real-time journey time variation (in this case, when multiplied by two to be charted at a relative scale) offers the closest approximation of

the *actual* customer journey time variation. This indicates that the real-time Oyster data—even with its inherent delay—can provide a reasonable proxy for the actual journey time variation and the severity of the disruption. In this example, it appears to provide the most accurate real-time estimate of rider disruptions of the five measures compared here. This is not to say that this will always be true, but rather that the real-time estimate of Oyster journey time variation offers a valuable supplement to the existing information.
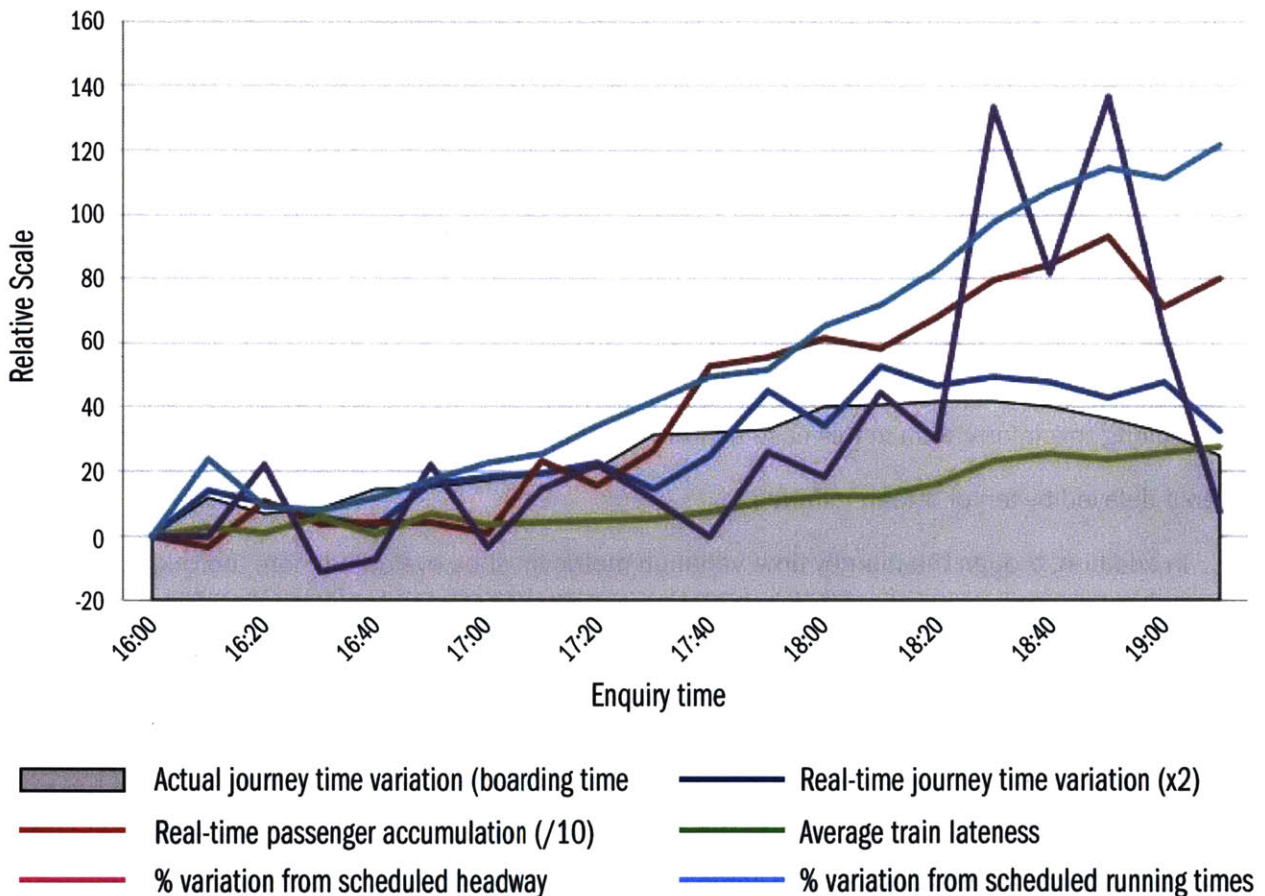
Figure 6-8: Comparing real-time information sources with actual journey-time variation, on a relative scale.

Despite the fact that real-time journey time variation *v* suffers from a lag in data availability, when plotted in Figure 6-8 *relative* to the four other measures discussed above, it comes closest to matching the actual journey time variation curve in terms of slope. All of these information sources are likely to be useful in determining in real-time the likely impacts of an incident. A complete controller information

system would provide information about all of these indicators. Future research should investigate the degree to which each of these measures provides the most reliable estimates of disruption severity. What is particularly worth examining is at what point each of the measures indicates when a problem is "bad enough" to inform customers. Though the real-time Oyster method is but one within a universe of potential data sources to measure real-time impacts of a disruption, it is discussed in the following section more extensively to evaluate how it might be used to inform customers of delays.

Compared to other metrics, particularly the passenger accumulation and certain of the NetMIS data, the journey time variation measure is likely to be less useful in determining the onset of a disturbance for passengers on the line. The inherent lag between when passengers experience delays and when they tap out make using it as a tool for making the first announcement problematic. On the other hand, tracking Oyster data may be more useful for determining the conclusion of a particular disturbance for passengers on a line. By using actual passenger data, the journey time variation metric offers the opportunity to clarify when journey times have returned to normal. Though there would be a lag in acquiring this information in this case as well, it still would offer insight into customer performance, even if delayed by ten or fifteen minutes.

In addition, though the journey time variation metric must be evaluated more thoroughly to be understood completely, the results of this comparison suggest that, when scaled correctly, it may provide an accurate assessment of the eventual severity of a disruption for passengers. If the metric defined in Chapter 5 is accepted as a reasonable estimate of actual passenger disruption, then the use of it in real-time circumstances may point to the severity of rider delays, as shown here. This technique, however, requires additional study and comparisons for other disrupted events before it can be recommended.

6.4 COMPARING THE DISRUPTIONS INDEX WITH CURRENT PASSENGER NOTIFICATIONS

The method used to take advantage of real-time Oyster data described in Section 6.2 does not specify how the results can be used to indicate a degree of disruption to passengers. It is not clear that there

is a direct relationship between the NetMIS-based announcements given to customers today and the Oyster-derived real-time data. But it is possible to envision using the Oyster information to assign similar levels of disruptions information. In this context, two types of delays are described—minor and severe delays—with the assumption that anything more serious (a part or full line suspension) would be immediately communicated by service controllers in touch with operators.

The Oyster data suggest that any variation in travel time of above about 20% reflects greater-than-normal delays for customers, and any variation that is twice as significant (more than 40% above standard travel times) implies significant rider delays. It is suggested therefore that the following criteria could be used to inform announcements using the Oyster data:

- < +20% variation in travel time: **Good Service**;
- Between +20% and +40% variation in travel time: **Minor Delays**;
- > +40% variation in travel time: **Severe Delays**.

These represent one simple way in which the Oyster data could be used, keeping in mind that it is intended to be supplementary to the existing NetMIS-based system, not a replacement. Note that this is a hypothetical use of the Oyster data to make notifications, and any reform of the existing system should incorporate Oyster, NetMIS, and controller data simultaneously. The exact specifications for when the appropriate moment is to make an announcement should be based on further analysis. Nonetheless, a diagram of how an Oyster-based announcement system could be used is presented in Figures 6-9, 6-10, and 6-11.
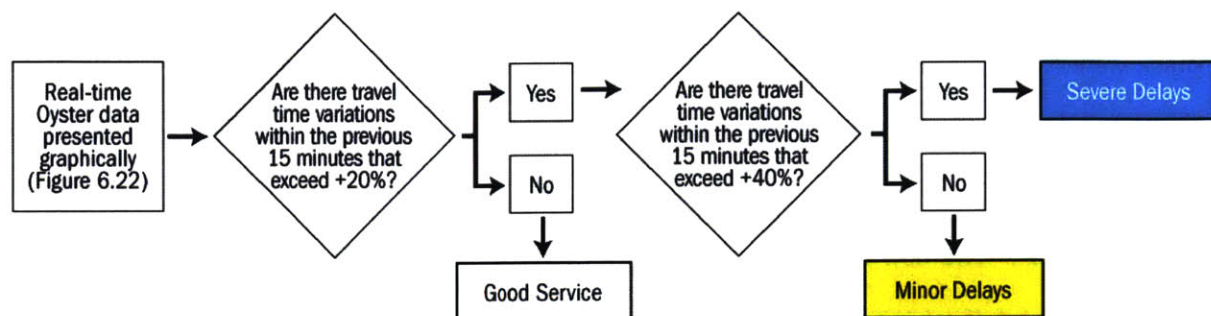
Figure 6-9: Diagram of use of real-time Oyster data to update passenger notifications.
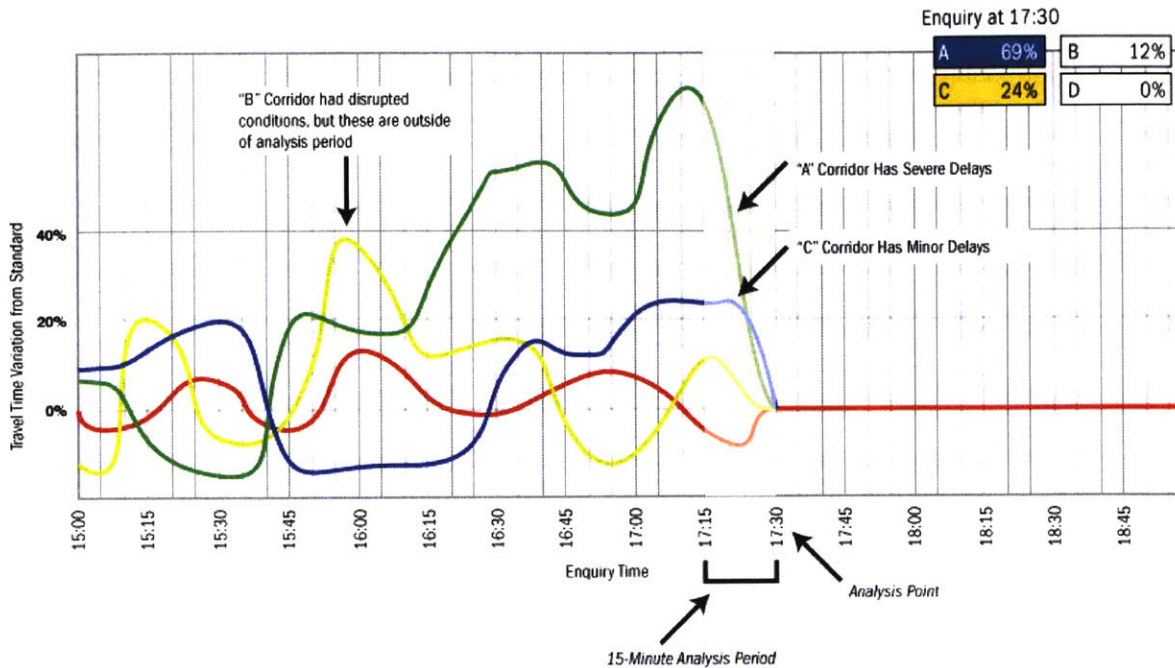
Figure 6-10: Hypothetical use of real-time Oyster data to update passenger announcements.

Figure 6-10 shows a hypothetical real-time Oyster-based travel time variation chart based on information "known" at 17:30. The four imaginary sections of the line presented in the diagram show varying performance over the course of the day, but only information presented in the previous 15 minutes is used as an input to determine the status of service on the line. As a result, the hypothetical "A" and "C" Corridors are determined to have delays significant enough to inform customers, whereas the rest of the line is not.

The information at the top right of the chart indicates the degree of severity by noting the highest level of journey time variation v in each section of the line, as first documented in Table 6-3. This offers a quick view of the status of service. This could be used in association with the other real-time measures defined in Section 6.3 to determine the appropriate service status level presented to customers at any given time.

Figure 6-11 shows how this information might be presented to customers on TfL's standard Service updates chart. If relying on Oyster data alone, customers could be informed of problems on individual sections of the line; in this case, because the "A" and "C" Corridors are experiencing significant

Service updates at 17:30

	Now Later This weekend
Piccadilly	Severe Delays
A Corridor	Severe Delays
C Corridor	Minor Delays
Rest of Line	Good Service
District	Good Service
H'smith & City	Good Service
Jubilee	Good Service
Metropolitan	Good Service
Northern	Good Service
Overground	Good Service
Circle	Good Service
Victoria	Good Service
Waterloo & City	Good Service

Figure 6-11: Hypothetical service updates using real-time Oyster data.

passenger delays, they are called out specifically as having "Severe" and "Minor" Delays, with the rest of the Piccadilly Line (the hypothetical "B" and "D" Corridors) described as having Good Service, in line with what is presented in Figure 6-10. That said, as discussed in Section 6.3, a comprehensive use of all available data sources would be preferable.

As shown in Figure 6-7, the severity of a disruption is often not clear for many minutes after the incident has begun. As such, real-time data must take into account a broad range of input. One way to do that is to provide customer alerts if there are any measurements of travel time variation of 20% or above the norm within the previous fifteen minutes, under the assumption that riders entering the system as far back as fifteen minutes before the analysis time provide some information about service on the line now. In other words, if controllers receive information at 15:30, they should make an announcement if *any* section of the chart exceeds the +20% baseline. Each section and direction of the line should be examined independently.

In Figure 6-12, the passenger notifications provided by the NOC are compared with potential notifications using the Oyster system, between 15:00 and 19:00 on April 24, 2012, a day with a PWT score of 41 seconds. A signal problem on the North section caused a significant delay for customers at around 16:25, producing repercussions in the line's service for the remainder of the afternoon.
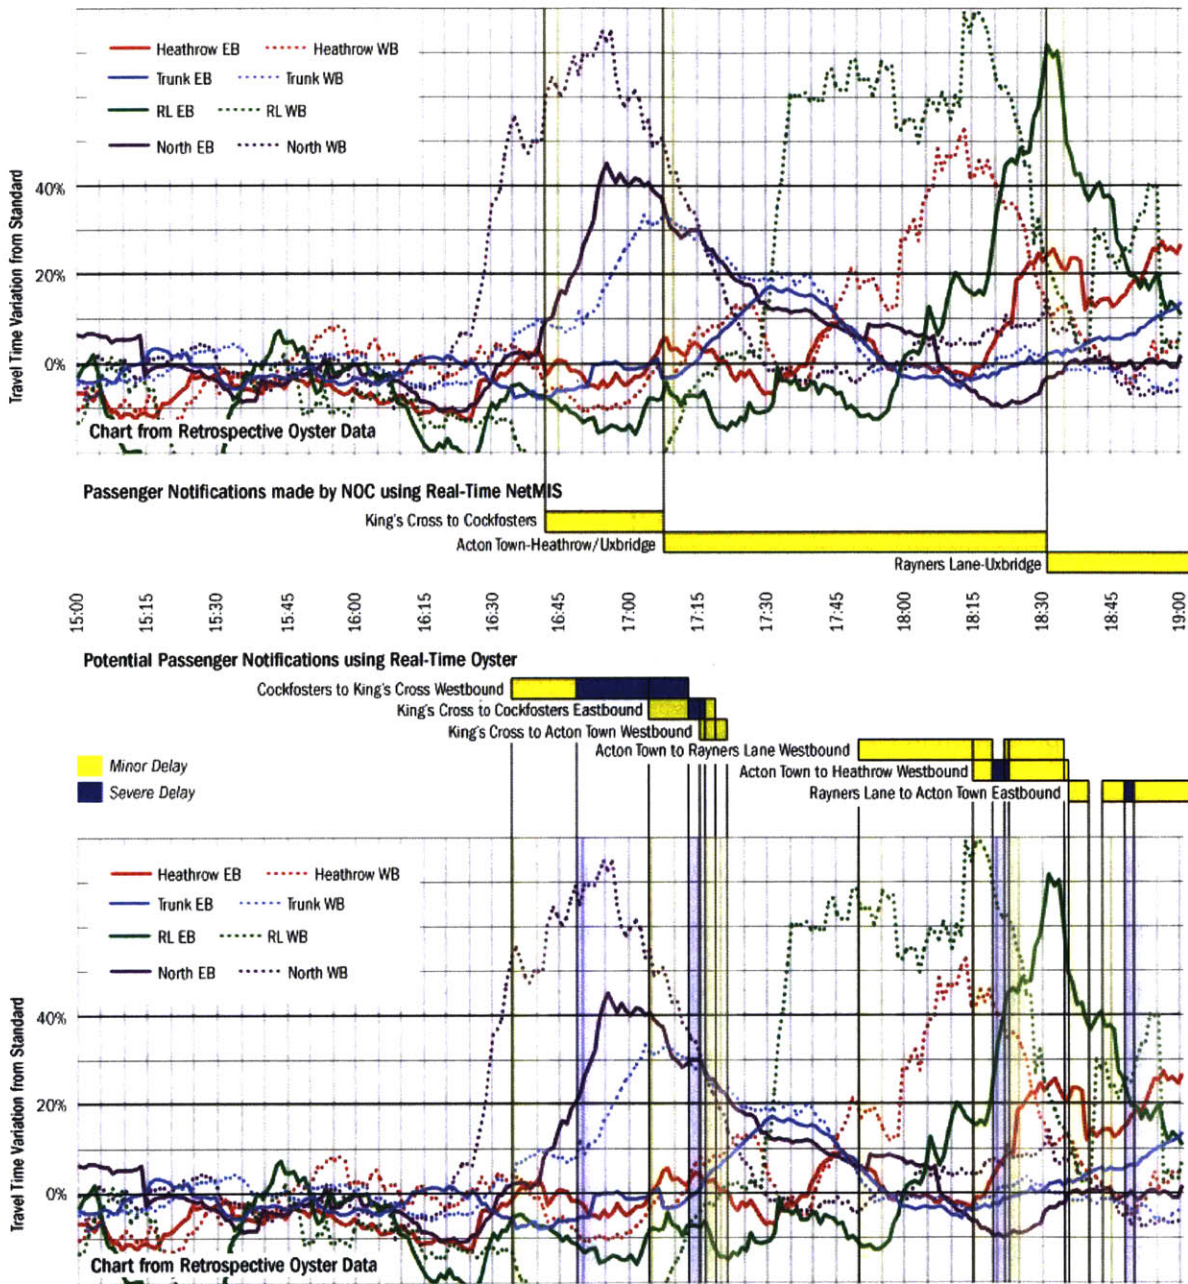
Figure 6-12: Comparing existing passenger notifications with proposed Oyster-based announcements.

The graphs at the top and bottom of Figure 6-12 show the actual average passenger travel time variation from Oyster over the period, based on a retrospective analysis. The charts show that passengers who entered the North section westbound at 16:25 were the first to experience disrupted travel conditions. These problems propagated first to riders on the eastbound North section, then the westbound Trunk, Rayners Lane section, and Heathrow section and finally the eastbound Rayners Lane section of the line.

Passenger notifications consisted entirely of Minor Delay announcements, as indicated at the top of the figure. Customers traveling from King's Cross to Cockfosters (the North section) were alerted first, from 16:43 to 17:08; at 17:08, customers traveling from Acton Town to Heathrow and Uxbridge (the Heathrow and Rayners Lane sections) were informed of problems until 18:31; then, for the rest of the evening (until 20:54, off the chart), only customers on the section from Rayners Lane to Uxbridge were informed of delays.

This record of passenger notifications points to the shortcomings of the existing passenger announcement system. Customers are generally not told whether the delays are in one direction or the other.[1] This has the effect of potentially informing riders about a problem they are unlikely to experience—though it is worth pointing out that one basic issue is that even in cases where TrackerNet data may show specific problems, these are not always communicated by the NOC. For example, riders on the Heathrow branch were informed that they were likely to experience minor delays beginning at 17:08. Yet riders on the eastbound Heathrow branch did not see delays as large as 20% until 18:25 at the earliest—more than an hour later. On the other hand, customers traveling in the line's Trunk section (from Acton Town to King's Cross) were never told about any potential delays, despite the fact that entering westbound Trunk customers experienced delays of 20% or more on average from 16:55 to 17:25. This was a full half hour during the afternoon rush period, when thousands of riders who likely had other travel route options were not informed of potential delays on the line.

---

1   This is not always true: There are specific circumstances in which the notification does indicate the direction of travel in which customers are affected. But most announcements do not differentiate by direction, despite the fact that most customers are interested in only one direction at a time.

The retrospective Oyster data presented in Figure 6-12 point to the advantage of using the analysis tool proposed in this thesis. Based on the method presented in Section 6.2 and considering the highest variation during the previous fifteen minutes, the bottom chart indicates when either minor or major delays would hypothetically have been announced for the line if Oyster real-time information were used and compares it to the actual travel time variation recorded retrospectively. Comparing these theoretical passenger announcements with the actual notifications made by the NOC shows the advantages and disadvantages of the system. Tables 6-7 and 6-8 indicate how several "issues" (part of the overall pattern of disruption) that occurred in the afternoon of April 24 were responded to by NOC and how they theoretically would have been addressed using the Oyster real-time method.

Table 6-7: Comparing real-time Oyster and NetMIS-based passenger notifications at commencement of passenger delay.

Section of the line where travel time variations are increasing	Time that journey time variation passes +20% for entering riders (retrospective information)	Actual NetMIS-based (NOC) response	Theoretical Oyster-based real-time response
North section westbound	16:27	16:42 Minor delays for North both directions	16:35 Minor delays for westbound North
North section eastbound	16:47	16:42 Minor delays for North both directions (same as above)	17:05 Minor delays for eastbound North
Trunk section westbound	16:56	None	17:15 Minor delays for westbound Trunk
Rayners Lane section westbound	17:32	17:08 Minor delays for Rayners Lane both directions	17:50 Minor delays for westbound Rayners Lane
Heathrow section westbound	17:59	17:08 Minor delays for Heathrow both directions	18:15 Minor delays for westbound Heathrow
Rayners Lane section eastbound	18:18	17:08 Minor delays for Rayners Lane both directions	18:36 Minor delays for eastbound Rayners Lane
Heathrow section eastbound	18:25	17:08 Minor delays for Heathrow both directions	None

Table 6-8: Comparing real-time Oyster and NetMIS-based passenger notifications at end of passenger delay.

Section of the line where travel time variations are decreasing	Time that journey time variation falls below +20% for entering riders (retrospective information)	Actual NetMIS-based (NOC) response	Theoretical Oyster-based real-time response
North section westbound	17:19	17:08 Good service	17:13 Good service
North section eastbound	17:23	17:08 Good service	17:19 Good service
Trunk section westbound	17:24	None	17:22 Good service
Rayners Lane section westbound	18:31	Past 19:00	18:35 Good service
Heathrow section westbound	18:28	18:31 Good service	18:36 Good service
Rayners Lane section eastbound	18:50	Past 19:00	Past 19:00
Heathrow section eastbound	18:38	18:31 Good service	None

As indicated by Tables 6-7 and 6-8, the Oyster real-time analysis method appears to be more accurate (the highlighted cases) than the current NOC method on about two-thirds of events, both in terms of when disruptions begin (Table 6-7) and when they end (Table 6-8). This result must be viewed with some skepticism, however, as the result is being measured using the criterion (travel time variation) that itself is used to develop the Oyster delay-based announcements. But if it is accepted that Oyster travel time variations do accurately represent the experience of riders on the line, this comparison provides useful insight into the ways in which the existing announcements could potentially be improved by incorporating Oyster data.

The evidence from this comparison *does* show that there are several cases where the NOC information is more accurate than the Oyster-derived information. Because the NOC's announcements are skewed towards telling customers in both directions about delays (even when the problem is initially in only one direction), the NOC alerted riders to problems on the Heathrow Line in both directions at

17:08. Though the eastbound delays did not begin until 18:25, the Oyster-based system never revealed this to be an issue worth announcing to customers, so it was ultimately less useful. This was also true for the North branch eastbound service; there, delays began at 16:47 but the NOC made an announcement there at 16:42, whereas Oyster only recognized the problem at 17:05.

The example of April 24, 2012 documented in this section may provide an idealized case for the use of real-time Oyster data to establish the presence of and severity of disruptions. There are several potential problems with this approach that deserve repeating. The first and most important of these difficulties is the inherent lag in the farecard data between when customers enter the system and when it becomes clear that there are delays on the line. As shown in Section 6.2, these delays nonetheless can be caught within ten minutes, though that is less likely for less-utilized sections of the service. As shown in Section 6.3, despite the lag, the real-time Oyster data is nonetheless more closely correlated with the ultimate severity of journey time variation of customers than other measures such as passenger accumulation and headways. Future research should investigate the ways in which these measures can be integrated most effectively to minimize the impact of lag.

Oyster data has also the potential to produce "false alarm" events in which a problem appears to be particularly bad at one point in time, thus inducing a passenger notification, but which disappear once more data becomes available. In theory, this could occur if a number of riders took short but longer-than-usual trips in the first fifteen minutes after they boarded, but most riders (exiting after the fifteen minute initial period) took trips that lasted a normal amount of time. This is unlikely to occur because it would require the short trips to be longer than usual even as long trips are normal; unless a group of passengers is stranded on a platform for some reason, this should not occur (since people taking longer trips share the train travel segment with those taking the short trips). In addition, the sections of the Piccadilly Line described in this thesis are long enough and with enough passenger data available that outliers are unlikely to affect the average significantly. That said, if the line were divided into smaller analysis sections, data error resulting in "false alarms" would be more likely.

The deficiencies in the Oyster data reaffirm the argument that it would be inappropriate to rely

on this information alone to make determinations about the status of service on a particular line. Yet there are also deficiencies in the information provided by the current NetMIS-based system, notably thanks to incomplete data provision (which is less of a problem with Oyster). Thus the evidence suggests that both the NOC's NetMIS system and the proposed Oyster real-time method offer insight into the travel time variations affecting passengers on the line. However, neither tool is reliable enough on its own to ensure that riders are receiving the most accurate possible information about London Underground services. If NOC decision makers are presented with both types of information, each providing proposed service announcements, they may be able to make qualitative assessments based on accumulated knowledge about the appropriate message to provide customers. Though there are guidelines in place by NOC to use NetMIS to automatically determine how customers are provided information, these are merely suggestive; NOC has the ability to make subjective judgments, and for good reason. For instance, if a disrupted service is moving along the line from east to west, it may be appropriate to announce to customers that there is a delay before such a delay has actually occurred. For instance, the major disruption that began in the westbound direction on the North section of the Piccadilly Line on April 24th affected service on the Trunk, then the Rayners Lane section, and finally the Heathrow section. It may be useful to inform customers in advance that they are likely to experience problems for a section that is likely to, but has not yet, experienced increasing travel times.

6.5 IMPROVING INFORMATION FLOW WITHIN THE AGENCY

None of the analysis presented here or in Chapter 5 can be completed without changes in the way in which Transport for London approaches information sharing between different departments in the agency. At the heart of the problem is that while TfL has several very useful datasets available, these are not universally accessible by agency staff, according to interviews with employees. For example, whereas NetMIS information is used to analyze service at the London Underground, analysts at TfL headquarters are generally unfamiliar with how this source can be interpreted and do not use it in their daily work. The inverse is often true for Oyster data, though staff at TfL have begun investigating the possibilities of sharing more data. This lack of comparative assessment inhibits a thorough under-

standing of how the system is functioning and what improvements may be feasible.

If upgrades are made as planned to faregates that would allow live Oyster data to be streamed to TfL, this data should be shared directly with NOC personnel so that they can incorporate it into the London Underground's passenger disruptions notifications. This will require understanding of the potential use of AFC data currently confined to TfL headquarters to be passed to staff at the Underground, an important but likely time-consuming process.

7 | Conclusion

This chapter summarizes the research conducted in this thesis and suggests avenues for future research on the impact of disruptions on customers using the system. Section 7.1 provides an overview of the insights provided by the thesis as a whole, offering a chapter-by-chapter summary of the research undertaken and the methods developed to describe disruptions on the London Underground Piccadilly Line. Section 7.2 offers a note of caution, identifying significant limitations of this research and additional work needed to be fully implementable. Much of the research has high potential for application on the London Underground, however, and corresponding recommendations are described in Section 7.3. Section 7.4 notes how the research might be expanded in other directions that take advantage of AFC data in the analysis of the effects of disruptions.

7.1 RESEARCH SUMMARY

This thesis reviews disruptions on the London Underground Piccadilly Line, analyzing not only the character of incidents that occur on the line today, but also offering improved methods to understand their impacts on passengers. The thesis has three primary components: First, a description and categorization of incidents that occur on the Piccadilly Line; second, the development of a new measure of impacts on passengers resulting from those incidents using AFC data; and third, an exploration of the

potential use of AFC data in real-time applications to monitor the service on the line. As an ensemble, the thesis offers new insight into the impact of disruptions than was hitherto available to the London Underground.

Chapter 4 characterizes and categorizes incidents that occurred on the Piccadilly Line during a three-month study period composed of 29 analysis days in spring 2012. The chapter summarizes the tools currently used by the London Underground to evaluate service, including the CuPID disruption database, the excess platform waiting time metric, and the headways proxy metric. By reviewing the occurrence of incidents on the Piccadilly Line including service controllers' logs, the chapter notes the concentration of problems relating to signals, especially when weighed in terms of the duration of incidents. It also shows that incidents are disproportionately clustered in the early morning and midday periods of the day, and in the Trunk and North sections of the line. It finds that stations where trains pull in and out of sidings, such as Acton Town and Arnos Grove, account for a large percentage of incidents on the line.

The thesis emphasizes the difference between "incidents" and "disruptions," describing the first as the specific "problem" on the line that triggered delays, such as a broken rail, a sick passenger, or a stalled train. The duration of an incident is the length of time between the problem being identified and it being "solved." A disruption, on the other hand, is defined as the cumulative impact of the incident on the passengers using the line. This distinction is essential to understanding service provided on the line and allows us to differentiate train and customer impacts.

Chapter 5 proposes a novel approach for investigating disruptions. It first identifies the limitations of approaches based on AFC origin-destination pair analyses and notes that NetMIS train-tracking data is incomplete and therefore there may be advantages to supplementing the information gained from it with Oyster AFC data. As a new approach, this chapter develops an aggregate method to measure variability in customer journey time during disrupted periods. This metric examines journey times on all OD pairs within defined sections of the Piccadilly Line to determine variation in journey time for the *average* passenger departing at a specified time on a specified section and direction. The advantage

of this aggregate approach is that it allows far larger sample sizes—and therefore better accuracy—in analysis of short time periods than individual OD pairs do.

The chapter provides example results of this analysis for disrupted days and demonstrates the applicability of the method in identifying time periods of the most significant disruptions. Using this variability information, the chapter proposes two metrics that can be used to quantify the overall customer impact of disruptions on the system. Total disruption impact takes the integral of the journey time variation function, and allows impacts of different disruptions on different parts of the system to be compared. This measure has particular applicability for comparing actual journey time variation with the estimates derived from the NetMIS-informed Lost Customer Hour (LCH) metric currently used for retrospective disruption analysis by the London Underground. The chapter also provides a second metric that describes the relative intensity of a particular disruption.

This chapter also examines the impact of disruptions by quantifying the number of passengers using specific portions of the system. A measure of passenger accumulation shows that disrupted periods feature more active customers than normal because their travel is delayed due to slower-than-normal train service. This accumulation metric can be compared with journey time variability and NetMIS-based data to offer an overall assessment of customer impacts. Finally, Chapter 5 notes the possibility of using AFC data from disrupted periods to infer path choices and compare them to those implied by London Underground's route-assignment model. By investigating journey times under disrupted conditions for passengers taking journeys on OD pairs where there is more than one possible route, we can identify riders' likely route by examining whether or not they were delayed when one of the potential paths was disrupted and the other was not. AFC data from normal service days would not allow this type of analysis since journey times on such alternative routes are typically similar and therefore non-differentiable using Oyster data.

Chapter 6 examines the possibility of using AFC data to assess the impact of disruptions on passengers in real time. It begins with a description of the existing passenger notification system, specifically identifying the criteria used (based on NetMIS train-tracking data) to determine whether or not

passengers are informed of delays on the line. The chapter demonstrates that passengers are alerted to delays up to 10 percent of the time on parts of the Piccadilly Line, but that those delays do not always correspond directly to the incidents recorded by controllers, which raises questions about the accuracy of the current announcements.

The chapter then analyzes the possibility of using journey time variation data as proposed in Chapter 5 to monitor the performance of the rail system. This approach, however, is flawed in that it requires passengers to "tap out" of the system before their journey times can be determined. As an alternative, the chapter considers using the passenger accumulation measure first explored in Chapter 5. When compared with NetMIS data, both approaches offer useful insight into the magnitude of passenger disruptions on the line in real time, suggesting that there are useful ways for AFC data to be used in a real-time context.

7.2 LIMITATIONS FOR METHOD IMPLEMENTATION

The increasing implementation of AFC technologies on public transport systems around the world provides a major new source of information about the movement of passengers along lines. This thesis offers several insights into how AFC data might be harnessed to examine the impacts of disruptions. Yet there are several significant issues that must be recognized when using the journey time variation methods described in Chapters 5 and 6.

First, the method has limited applicability. Total journey time variation applies to the *average passenger using a section of the line in one direction*. This means that it cannot be applied ad hoc for individual passengers, though it can focus in on people who were likely affected by problems. In addition, a journey time variation in one section of the line cannot be directly compared to that of another section in terms of passenger impacts, because the number of passengers using different parts of the line varies. For instance, a large variation on the Rayners Lane branch may affect fewer riders than a small variation on the Trunk section. This could be resolved to some degree if the disruption impact figure were multiplied by the flow of passengers using each of the sections. LU's estimates of flow on the Piccadilly Line are not yet fully developed, however, though Ravichandran (2012) has explored a

potential method by which to determine passenger flow on the line for a period of time.

Second, the method relies on the sectioning of the Piccadilly Line into a series of usable parts. An effort was made to identify sections that corresponded to train reversal points based on interviews with train controllers, but the journey time variation method could be applied to a multitude of feasible sections along the line, potentially up to and including an analysis of the whole line. As such, the conclusions reached by this thesis about problems on the line may not address the concerns of most passengers. The correct sectioning of the line should be subject to further discussion to ensure that the study sectioning is appropriate to the question being asked.

There are also serious questions that remain unresolved about the value of AFC data in a real-time context. Discussion with Transport for London staff indicate that though certain Oyster data is available from several stations (almost) in real time, the majority of stations do not yet provide a direct data feed, though further analysis is necessary to define this more precisely. Though the situation is likely to improve over the coming years with the installation of new Oyster readers and fiber-optic cables, the degree to which these technologies will provide real-time information must be examined thoroughly before any solution can be implemented.

In addition, it should be cautioned that AFC data should not be used to indicate that an *incident* has occurred or been resolved. Rather, controller communications with line staff should continue to be the primary mechanism for identifying incidents. As an example, AFC data may indicate significant passenger accumulation, which may be due to the end of a football game, rather than an incident on the line. AFC data use should be confined to tools for better understanding impacts on riders. It may also be particularly useful for identifying when the disruption has ended in terms of cessation of passenger impacts.

Finally, the passenger accumulation metric, while useful for determining the build-up of riders using the line during a disruption, should not be confused with an indication of passenger flow. Because it is not direction-specific (we do not know which way passengers who board but have not exited are traveling, of course), because it does not include the Oyster boardings from interchange stations,

and because it does not account for people who transfer into and out of the Piccadilly Line it is not representative of overall ridership in any sense of the word. It provides a quick way to monitor general conditions along the line, but it should not be extended to any other analysis pertaining to rider flows.

7.3 RECOMMENDATIONS TO TRANSPORT FOR LONDON AND THE LONDON UNDERGROUND

The analyses presented in this thesis were designed to be useful to planners and analysts in London, and there are some clear ways in which the research results can be applicable to improving our understanding of the London public transport network. The following recommendations capture the lessons learned from discussions with TfL and LU staff, as well as the thesis findings.

- **Share Oyster and NetMIS data across divisions.** The thesis has demonstrated that there are concrete benefits that result from comparing AFC and train-tracking data. While each offers similar insights into the service environment, neither provides a full description of the conditions on the line. Together, however, they can effectively complement one another, one source filling in the gaps in the other source data. But currently, easy data access for the two sources is still quite siloed. TfL fares and ticketing staff are experts with Oyster, while LU service planning and analysis staff are expert with NetMIS.

   While there is clear movement towards the increased sharing of these data, the task is not yet complete. In order to expand upon the conclusions of this research and ensure that it can be as useful as possible to staff throughout TfL, it is critical to continue that movement so that data accessibility is ubiquitous.

- **Improve retrospective analysis of disruptions by including AFC data.** This thesis has argued strongly that Oyster AFC data opens an important avenue for potential evaluation of the impact of disruptions on passengers. While there is currently limited work at TfL that investigates changes in journey times for riders making specific journeys, there is no comprehensive metric by which Oyster data can be analyzed to assess the magnitude of a disruption. This thesis offers such a metric.

Though the thesis only provides the groundwork for a more comprehensive process of analysis of Oyster data for disruption evaluation, it offers a way to compare disruptions and identify which group of customers may have been most impacted by an incident.

The journey time variation metric may provide a way to evaluate the accuracy of the lost customer hour (LCH) metric currently used by LU personnel to assess disruption impacts. The LCH estimates incident impacts based on projections of rider flow and disruption duration, while the method proposed here is based on actual journey times. The LCH estimates are revised every five years or so; it would be reasonable to use the journey time variation metric or an extension of it to study the actual impacts of disruptions and then alter the LCH estimates to ensure that they accurately reflect what occurs on the system during and after an incident.

- **Investigate the applicability of AFC disrupted-day data to improve the understanding of route choice.** A large percentage of travelers on the London Underground use interchange stations to move between lines during the course of their journeys. While these interchanges may improve their commute, they cause significant difficulties for data analysis, since they are not recorded in the Oyster data and thus the way customers get from one station to another often remains unobserved. In order to correct this gap in information, TfL conducts the RODS survey to assign percentages of travelers on each OD to specific routes. Yet this thesis has raised the question of whether the results of those surveys—which are often years old—are always reliable.

This thesis shows that disrupted periods provide potentially valuable insights into traveler behavior on certain OD pairs, particularly those in central London. By examining where journey times increase significantly during disruptions but before notifications have been made to passengers, empirical evidence on route choice for certain ODs can be obtained. Using this data in addition to the current RODS results could produce improved flow analyses on the Underground system.

- **Exploring the possibility of using AFC data in real time.** The thesis documents two potential methods by which Oyster data might be used in real time: First, by using the journey time variation method, and second, by using the passenger accumulation method. Each likely requires signifi-

cant improvements to the existing faregate and communications system to be fully implemented. Yet AFC data may be able to provide important insights into disruptions, in particular their severity. Though train-tracking data may identify when a problem has occurred, it may not be specific enough to indicate how severely passengers are affected; Oyster data, if used right, can. Just as important, the Oyster data, more easily than the train-tracking data, can clarify when a disruption has ended.

That said, Oyster data should be used with caution. Too small a sample may result in misunderstanding the disruption's severity. The problem is particularly true for the journey time variation, which relies on customers to "tap out" before it can be produced, thereby creating a lag in information provision. In addition, customers should not be given too much information without considering the long-term effects of over-informing customers, which might include riders slowly abandoning their trust in the accuracy of the notifications system in representing the current conditions.

- **Reexamining the use of headways as a metric for determining whether to issue customer notifications.** The NOC uses a measure of headways as one of its primary criteria for determining whether to make an announcement to customers about the status of service on the Piccadilly Line. Headways are a useful statistic, as waiting time consumes a large percentage of many customers' journeys. Therefore, increasing headways may indicate more serious problems down the line. Discussions with LU staff suggest that NOC notifications are made with less-strict adherence to the official guidelines than might have been assumed, but headways still play a role.

Yet the evidence suggests that there are other metrics that more clearly indicate the status of service than headways at individual stations. The Oyster journey time variation and passenger accumulation metrics, once real-time Oyster data are available, may be quite reflective of actual customer experience of disruptions. Other NetMIS information, notably train travel time between individual stations, might be even more useful in making some decisions. If one TfL goal is to ensure that the information given to customers is targeted and is as accurate as possible, a broadening of the sources used to identify delays is desirable.

7.4 FUTURE RESEARCH DIRECTIONS

This thesis provides a framework for analyzing Oyster data to assess the impact of disruptions on pas-
sengers. The metrics it develops are applicable to the Piccadilly Line, but they could be modified to
assess service on other London Underground lines. Three potential studies, described below, suggest
how the metrics described could be further refined to better aid LU in understanding the travel patterns
of customers using the service.

- **Analyze the most efficient way to segment lines for aggregate analysis.** The analysis present-
  ed in the thesis divides the Piccadilly Line into five segments for analysis based on train reversal
  points (of which one section, from Rayners Lane to Uxbridge, is not used because it overlaps with
  the Metropolitan Line). This segmentation, however, could be altered in another round of journey
  time variation analysis. A valuable research direction would explore the value of using different
  sections on the same dataset and exploring the value of information provided in varying circum-
  stances.

- **Evaluate the effects of controller strategies on disruption impacts.** The V metric defined in
  Chapter 5 assesses the severity of a disruption from the perspective of customers, but it does not
  differentiate between direct impacts of the incident, "knock-on" effects, and controller actions. Fu-
  ture research should compare the customer impacts of varying strategies, given similar incidents.
  This could aid controllers in determining how best to respond to incidents while minimizing rider
  inconvience.

- **Assess whether journey time variation "tracks" through a line in a predictable manner.** The
  thesis shows certain circumstances in which journey time variation, as expected, "moves" down
  the line from one section to another, as trains move in that direction. A future research study could
  evaluate whether that progression can be reliably predicted, depending on the type of incident that
  occurred. If so, it may be possible to estimate future travel time variation on the line and inform
  customers in advance that there are likely to be problems on their journeys.

- **Determine the appropriate role of journey time variation in triggering passenger announcements.** Chapter 6 suggests potential passenger notifications based on the journey time variations recorded, suggesting that a 20% variation be considered a "minor delay" and a 40% variation a "severe delay." These correlations, however, should be further analyzed to determine whether these announcements are reflective of how the average customer *feels* about his or her journey. For instance, does a customer with a trip that takes 30% longer than normal think that he or she has had a minor or severe delay? Or is some other notification more appropriate for the situation? A survey of customer attitudes about notifications, which has yet to be conducted by TfL, might add significantly to our understanding of customer response. In addition, further investigation must be conducted to determine how journey time variation information is used in concert with train-tracking data.

- **Further analysis of the accumulation metric.** Evaluating passenger boarding data offers significant promise for determining when delays are occurring on the line, particularly in real time. This thesis offers an initial review of how to use such data, but the method should be more fully investigated to evaluate its potential applicability.

Bibliography

Andreassen, T.W. (1995) *(Dis)satisfaction with public services: the case of public transportation.* Journal of Services Marketing, Vol. 9, No. 5, pp. 30-41.

Bates, J., Polak, J., Jones, P., Cook, A. (2001). *The valuation of reliability for personal travel.* Transportation Research, Part E: Logistics and Transportation Review, Vol. 37 No. 2-3, pp. 191-229.

Bernal, M., Welch, E., Sriraj, P.S. (2009) *Understanding the Impacts of Delays on Blue Line Customers.* Paper presented at Transport Chicago. Available online at <http://www.transportchicago. org/uploads/5/7/2/0/5720074/3-understanding_the_impact_of_delays_on_blue_line_customers. pdf >.

Bonsall, P. (2004) *Traveller Behaviour: Decision-making in an Unpredictable World.* Journal of Intelligent Transportation Systems: Technology, Planning, and Operations, Vol. 8, No. 1, pp. 45-60.

Bryan, H., Blythe, P. (2007) *Understanding behavior through smartcard data analysis.* Proceedings of the Institution of Civil Engineers Transport, Vol. 160, No. TR4, pp. 173-177.

Carrel, A. (2009) *Diagnosis and Assessment of Operations Control Interventions: Framework and Applications to a High Frequency Metro Line.* Thesis, Master of Science in Transportation, MIT.

Chan, J. (2007) *Rail OD Estimation and Journey Time Reliability Metrics Using Automated Fare Data.* Thesis, Master of Science in Transportation, MIT.

Fonseca, F., Pinto, S., Brito, C. (2010) *Service Quality and Customer Satisfaction in Public Transports.* International Journal for Quality Research, Vol. 4, No. 2, pp. 125.

Frumin, M. (2010) *Automatic Data for Applied Railways Management: Passenger Demand, Service Quality Measurement, and Tactical Planning on the London Overground Network.* Thesis, Master of Science in Transportation, MIT.

Gordon, J. (2012) *Intermodal Passenger Flows on London's Public Transport Network: Automated Inference of Full Passenger Journeys Using Fare-Transaction and Vehicle-Loading Data.* Thesis, Master of Science in Transportation, MIT.

Harper, C. (2012). Communication by email with LU Head of Insight, October 19, 2012.

Hickey, S. (2011) *Improving the estimation of platform wait time at the London Underground.* Thesis, Master of Science in Transportation, MIT.

Jang, W. (2010) *Travel Time and Transfer Analysis Using Transit Smart Card Data*. Transportation Research Record: Journal of the Transportation Research Board, No. 2144, pp. 142-149.

Jesperson-Groth, J., Potthoff, D., Clausen, J., Huisman, D., Kroon, L., Maróti, G., Nyhave Nielsen, M. (2007) *Disruption Management in Passenger Railway Transportation*. Econometric Institute Report, No. EI 2007-05.

Jiang, Z., Li, F., Xu, R., Gao, P. (2012) *A simulation model for estimating train and passenger delays in large-scale rail transit networks*. J. Cent. South Univ, No. 19, pp. 3603-3613.

Kahneman, D.,Tversky, A. (1979). *Prospect Theory: an analysis of decision under risk*. Econometrica, Vol. 47, No. 2, pp. 263-291.

London First. (2012). *Crossrail 2: Supporting London's Growth. Interim report of London First's working group*. Accessed December 26, 2012. Available online at <http://londonfirst.co.uk/wp-content/uploads/2012/08/120515_Crossrail_2_-_Supporting_Londons_growth_report.pdf >.

Mayor of London. (2010) *Mayor's Transport Strategy*. Accessed December 27, 2012. Available online at <http://www.tfl.gov.uk/corporate/11610.aspx>.

Nassier, N., Khani, A., Lee, S.G., Noh, H., Hickman, M. (2011) *Transit Stop-level O-D Estimation Using Transit Schedule and Automated Data Collection System*. Submitted for Presentation and Publication, 90th TRB Annual Meeting, January 23-27, 2011.

Paul, E. (2010) *Estimating Train Passenger Load from Automated Data Systems: Application to London Underground*. Thesis, Master of Science in Transportation, MIT.

Pender, B., Currie, G., Delbosc, A., Shiwakoti, N. (2012) *Planning for the Unplanned: An International Review of Current Approaches to Service Disruption Management of Railways*. Australasian Transport Research Forum 2012 Proceedings, 26-28 September 2012, Perth Australia.

Ravichandran, H. (2012) *Estimating passenger load on the Piccadilly Line*. Unpublished work at MIT.

Schil, M. (2012) *Measuring journey time reliability in London using automated data collection systems*. Thesis, Master of Science in Transportation, MIT.

Transport for London (TfL) Heartbeat. *Heartbeat*. No specific date. Internally accessible system providing information about the status of present and historical service on the London Underground.

Transport for London (TfL) Network Operations Centre (NOC). *Guidelines for passenger notifications*. No specific date. Internally accessible system.

Transport for London (TfL). (2011a) *Developing a reliability metric of LU customers*. Customer research conducted by 2CV for TfL.

Transport for London (TfL). (2011b) *NACHS 2014 Explanatory Notes Part A: Using the CuPID Calculator*.

Transport for London (TfL). (2012a) *Budget 2012/2013*. Accessed January 2, 2013. Available online at < http://www.tfl.gov.uk/assets/downloads/2012-13-budget-tfl-march2012.pdf>.

Transport for London (TfL). (2012b) *London Rail and Tube Services Map*. Accessed January 1, 2013. Available online at <http://www.tfl.gov.uk/assets/downloads/london-rail-and-tube-services-map. pdf>.

Transport for London (TfL). (2012c) *Travel in London Report 5*. Accessed December 28, 2012. Available online at <http://www.tfl.gov.uk/assets/downloads/corporate/travel-in-london-report-5. pdf>.

Transport for London (TfL). (2012d) *Business Plan: Transport for London's plans for the next decade*. Accessed December 26, 2012. Available online at <http://www.tfl.gov.uk/assets/downloads/ corporate/tfl-business_plan-2012.pdf>.

Transport for London (TfL). (2012e) *Annual Report and Statement of Accounts*. Accessed December 29, 2012. Available online at <http://www.tfl.gov.uk/assets/downloads/corporate/tfl-annualreport-2012.pdf>.

Transport for London (TfL). (2013) *Company information*. Website, with many sources about the organization. Accessed January 12, 2013. Available online at <http://www.tfl.gov.uk/corporate/ about-tfl/4510.aspx>.

Trépanier, M., Morency, C., Agard, B. (2009) *Calculation of Transit Performance Measures Using Smartcard Data*. Journal of Public Transportation, Vol. 12, No. 1, pp. 79.

Uniman, D. (2009) *Service Reliability Measurement Framework using Smart Card Data: Application to the London Underground*. Thesis, Master of Science in Transportation, MIT.

Valdés-Diaz, D., Martínez, F., Bartolomei Suárez, S., Colucci, B. (2005) *Analysis of Disruption Causes and Effects in a Heavy Rail System*. Third LACCEI International Latin American and Caribbean Conference for Engineering and Technology, "Advances in Engineering and Technology: A Global Perspective," 8-10 June 2005, Cartagena de Indias, Colombia.

Van der Hurk, E., Kroon, L., Maróti, G., Vervest, P. (2012). *Dynamic Forecast Model of Time Dependent Passenger Flows for Disruption Management*. In Voss, S., Munoz, J.C. (Eds.). *CASPT*. Santiago, Chili: CASPT.

Von Neumann, J., Morgenstern, O. (1944). *The Theory of Games and Economic Behaviour*. Princeton, NJ: Princeton University Press.