May 2019
Report No. 19-004

Charles D. Baker
Governor
Karyn E. Polito
Lieutenant Governor
Stephanie Pollack
MassDOT Secretary & CEO

Optimizing ADA Paratransit Operation with Taxis and Ride Share Programs

Principal Investigator
Dr. Eric Gonzales
University of Massachusetts Amherst

Research and Technology Transfer Section
MassDOT Office of Transportation Planning

U.S. Department of Transportation
Federal Highway Administration

[This blank, unnumbered page will be the back of your front cover]

Technical Report Document Page

1. Report No. 19-004	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle Optimizing ADA Paratransit Operation with Taxis and Ride Share Programs		5. Report Date May 2019
		6. Performing Organization Code 19-004
7. Author(s) Eric J. Gonzales, Charalampos Sipetas, and Justin Italiano		8. Performing Organization Report No.
9. Performing Organization Name and Address University of Massachusetts Amherst 130 Natural Resources Way, Amherst, MA 01003		10. Work Unit No. (TRAIS)
		11. Contract or Grant No.
12. Sponsoring Agency Name and Address Massachusetts Department of Transportation Office of Transportation Planning Ten Park Plaza, Suite 4150, Boston, MA 02116		13. Type of Report and Period Covered Final Report – May 2019 [November2017-May 2019]
		14. Sponsoring Agency Code n/a
15. Supplementary Notes		

16. Abstract

Rising ridership on Americans with Disabilities Act (ADA) paratransit services, such as MBTA's The RIDE, pose a challenge due to the high costs of operating this required service. This report presents the findings of a study to analyze the operations and ridership of The RIDE's van service and to optimize programs to serve paratransit trips with alternative service providers in order to minimize system cost. This study specifically looks at the operations and demand patterns before and after the implementation of the MBTA's On-Demand Paratransit Pilot Program, which allows eligible customers to use Uber, Lyft, or Curb for a subsidized trip in addition to conventional ADA paratransit service from The RIDE. Pilot participants typically use about 40% of their monthly allocation, and many continue to use The RIDE services depending on their type of disability. It is estimated that if all eligible customers were included in the program, the effect on The RIDE service would be a reduction in ADA demand by 42%, but an increase in total travel by 33% due to induced demand for the TNC service. This would correspond to a 26% reduction in costs to the MBTA. Alternatively, a model is proposed to estimate the marginal cost of each requested trip and optimize allocation to The RIDE vans or an alternative provider. The potential effect on total operating costs is estimated to be 40–48% reduction.

17. Key Word paratransit, transportation network company		18. Distribution Statement	
19. Security Classif. (of this report) unclassified	20. Security Classif. (of this page) unclassified	21. No. of Pages 81	22. Price n/a

Form DOT F 1700.7 (8-72) Reproduction of completed page authorized

This page left blank intentionally.

Optimizing ADA Paratransit Operations with Taxi and Ride Share Programs

Final Report

Prepared By:
Eric J. Gonzales, Ph.D.
Principal Investigator
(413) 545-0685; gonzales@umass.edu

Charalampos Sipetas
Graduate Researcher
csipetas@umass.edu

Justin Italiano
Graduate Researcher
jitalian@umass.edu

University of Massachusetts Amherst
130 Natural Resources Way, Amherst, MA 01003

Prepared For:
Massachusetts Department of Transportation
Office of Transportation Planning
Ten Park Plaza, Suite 4150
Boston, MA 02116

May 2019

This page left blank intentionally.

Acknowledgements

Prepared in cooperation with the Massachusetts Department of Transportation, Office of Transportation Planning, and the United States Department of Transportation, Federal Highway Administration.

The Project Team would like to acknowledge the efforts of the staff of the Massachusetts Bay Transportation Authority for providing the data and necessary assistance for this project. Staff from The RIDE provided the data that was necessary to develop the models and perform the analysis described in this report.

Disclaimer

This page left blank intentionally

Table of Contents

List of Tables

List of Figures

List of Acronyms

Acronym	Expansion
ADA	Americans with Disabilities Act
GLSS	Greater Lynn Senior Services, Inc.
JV	Joint Venture (between Thompson Transit, Inc. and YCN Transportation, Inc.)
MBTA	Massachusetts Bay Transportation Authority
NTD	National Transit Database
PDP	Pick-up and Delivery Problem
PVTA	Pioneer Valley Transportation Authority
TNC	Transportation Network Company
TRAC	The RIDE Access Center
VHT	Vehicle Hours Traveled
VMT	Vehicle Miles Traveled
VTS	Veterans Transportation, LLC
WAV	Wheelchair Accessible Vehicle

1 Introduction

Transit agencies in the United States are required to provided door-to-door paratransit service for customers with disabilities under the Americans with Disabilities Act (ADA) of 1990. The purpose of ADA paratransit is to provide service that complements conventional fixed-route transit for people who are unable to use buses, subways, or trolleys. Rising ridership with ADA paratransit services around the country poses a challenge due to the high costs of operations. In response, transit agencies are seeking ways to reorganize operations and form partnerships with alternative providers in order to contain costs while meeting rising needs.

This report presents modeling and analysis of the On-Demand Paratransit Pilot Program initiated by the Massachusetts Bay Transportation Authority (MBTA). The MBTA's ADA paratransit service is called "The RIDE," and Pilot Program allows eligible riders to make subsidized trips with ridesharing companies Uber, Lyft, and Curb. The purpose of this study is to understand

1. the effect that a cooperative arrangement with these ridesharing companies will have on demand for ADA paratransit,
2. the operations of the remaining van service, and
3. the overall cost to the MBTA of providing transportation service to eligible customers.

1.1 The RIDE: ADA Paratransit Service in Greater Boston

The MBTA operates public transit services throughout Greater Boston, Massachusetts, including buses, light rail, heavy rail, commuter rail, electric trolleybuses, and ferries. The RIDE is the MBTA's ADA paratransit service and it is generally available to customers with eligible disabilities in Greater Boston and between the hours of 5 AM and 1 AM daily.

In order to be eligible for federal funding, the ADA requires transit agencies to provide complementary paratransit service that meets the following conditions (Transportation Services for Individuals with Disabilities, 1991):

1. Eligibility – Customers who have a disability that prevents them from being able to navigate or access the conventional fixed route transit service are eligible for ADA paratransit service. This includes physical disabilities, such as the need for a wheelchair or mobility device, as well as other sensory or mental disabilities that require curb-to-curb transportation service.
2. Geographic Coverage – Service must be provided to origins and destinations within ¾ of a mile of conventional local fixed route transit. Commuter rail stations are not included for the required service area.

3. Days and Hours of Coverage – Service must be provided during the same days and hours of the day as the operation of the conventional transit service.

4. Fare – The ADA paratransit fare cannot exceed double the fare cost for conventional fixed route transit for the same origin-destination pair.

5. Response Time – All customer requests must be matched with a ride within an hour before or after the requested time of travel.

6. Constraints – Customers cannot be denied service for any of the following reasons: trip purpose, number of trips requested, length of trips requested, excessive telephone hold times for reservations, any other capacity constraint on the system.

The requirements are intended to ensure that customers with disabilities have equal access to public transportation service as the general public. Since the public is not limited in their freedom to take as many trips and for any purpose as they are willing to pay for, the same standard is held for ADA paratransit. Altogether, these requirement place a burden on transit agencies to ensure that an adequately sized fleet of paratransit vehicles and drivers are available to meet demands. As demographics change, increasing numbers of eligible customers with disabilities require agencies to dedicate increasing resources to comply with ADA requirements. The operating funds for paratransit services come out of the same budget as general transit operations, so increasing financial obligations for ADA paratransit service limits the available funding for all other transit operations. As a result, transit agencies need to find ways to cost-effectively provide service for eligible ADA customers in order to sustain transit service for all users.

The RIDE is typical of many ADA paratransit services across the United States in that operation of vehicles is provided by private operators under contract to the MBTA. In order for customers to use The RIDE, they must apply for eligibility through The RIDE Eligibility Center. Upon approval, customers are able to reserve trips by calling in from 1 to 7 days in advance of travel. Customers are offered a trip that is scheduled to provide service within one hour of the requested pick-up time.

Although the ADA only requires that paratransit service be made available within ¾ of a mile of MBTA bus and subway stops, the MBTA makes The RIDE available to customers throughout 58 towns and cities in Greater Boston as shown in Figure 1.1. This is common for many agencies, because the ¾ of a mile boundary can exclude many important origins and destinations in a region, limiting access for customers who may not have other options for travel. A distinction is made in the fares charged:

1. Local ADA one-way fare for trips with origin and destination within ¾ mile of an MBTA bus or subway stop is $3.15.

2. Premium non-ADA one-way fare for trips with an origin and/or destination further than ¾ mile from an MBTA bus or subway stop is $5.25.

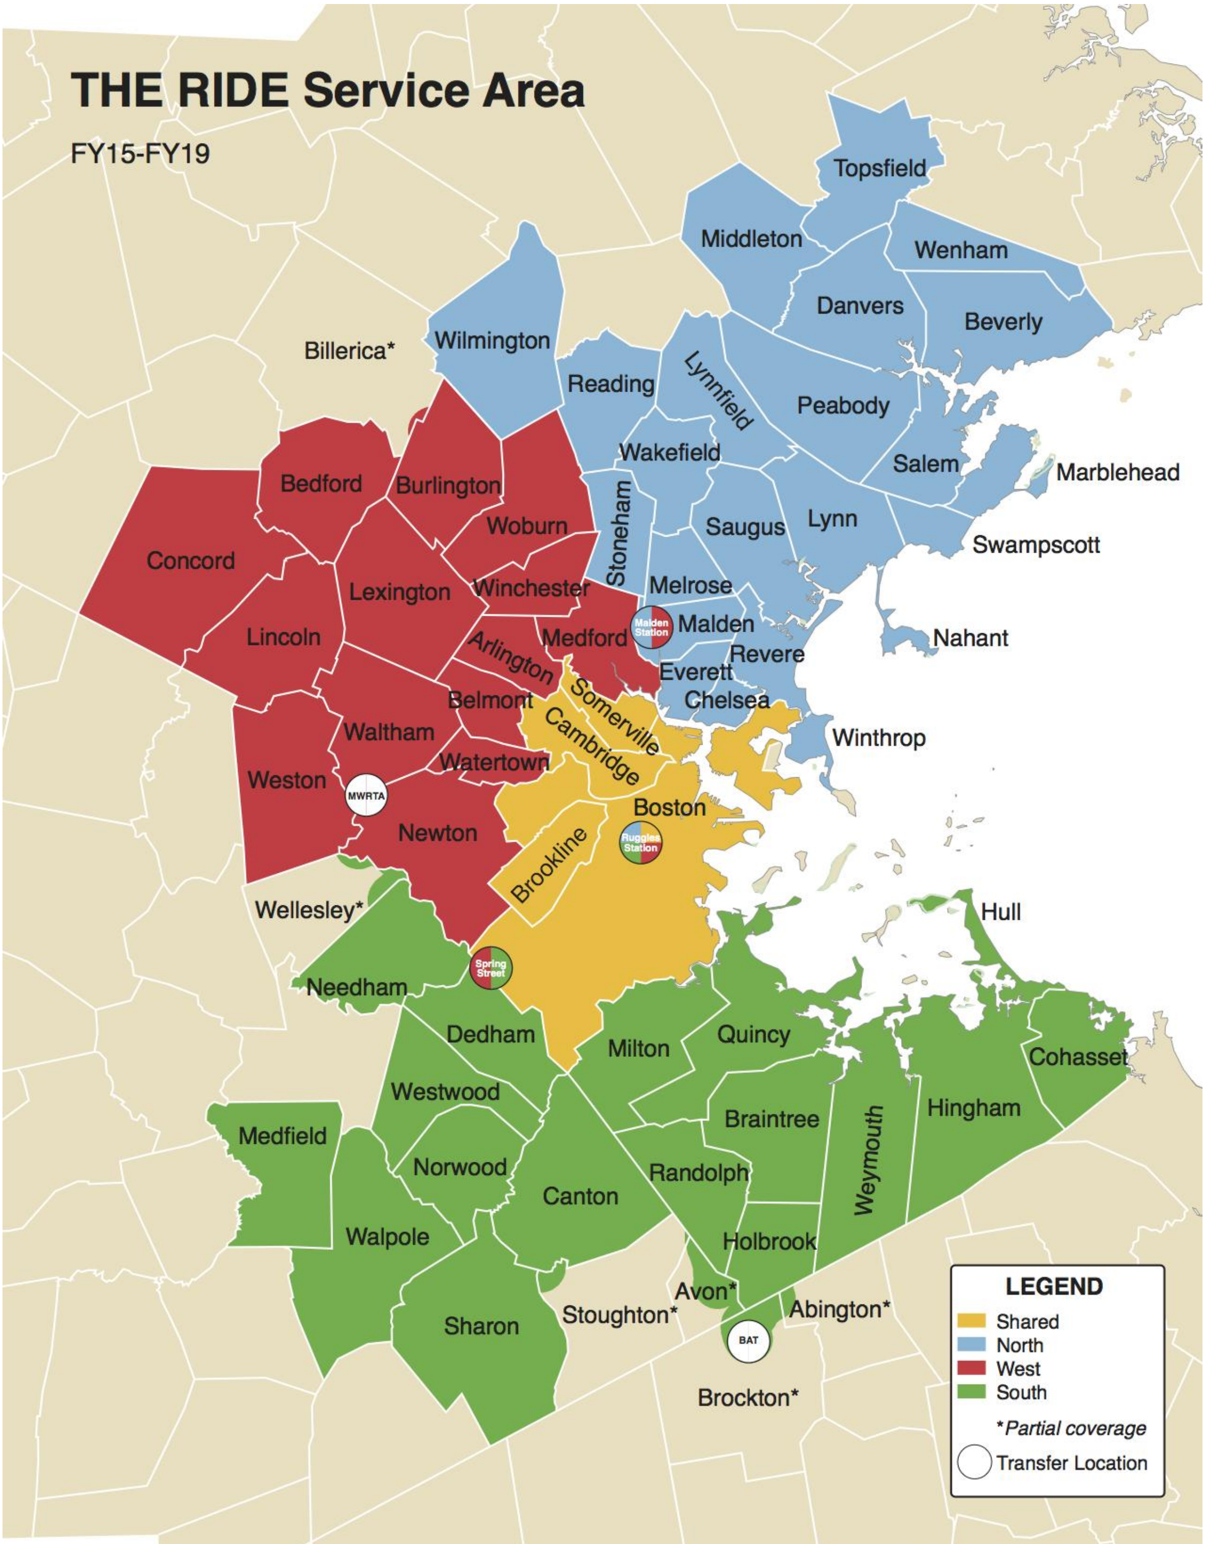

Figure 1.1 Service Area for the MBTA's ADA Paratransit, The RIDE (Source: MBTA)

Like many large public transportation systems, The RIDE's service area is divided into smaller operating regions, each served under contract by a separate provider. There are 3 regions for The RIDE: North, West, and South (see Figure 1.1). Each region also includes the shared center region shown in yellow. Trips within a single region are served by a single vehicle. Trips crossing from one region to another may require a transfer between two paratransit vehicles.

According to annual reporting in the National Transit Database from 2013 through 2017, annual ridership on The RIDE showed a gradually increasing trend (see Figure 1.2). During this time period the average operating cost per unlinked trip has fluctuated, but recently has risen to exceed $50 per trip served (see Figure 1.3). This equates to annual operating costs for The RIDE exceeding $100 million per year. This average cost is high, but within the typical range for paratransit services throughout the United States (Rodman and High, 2018).

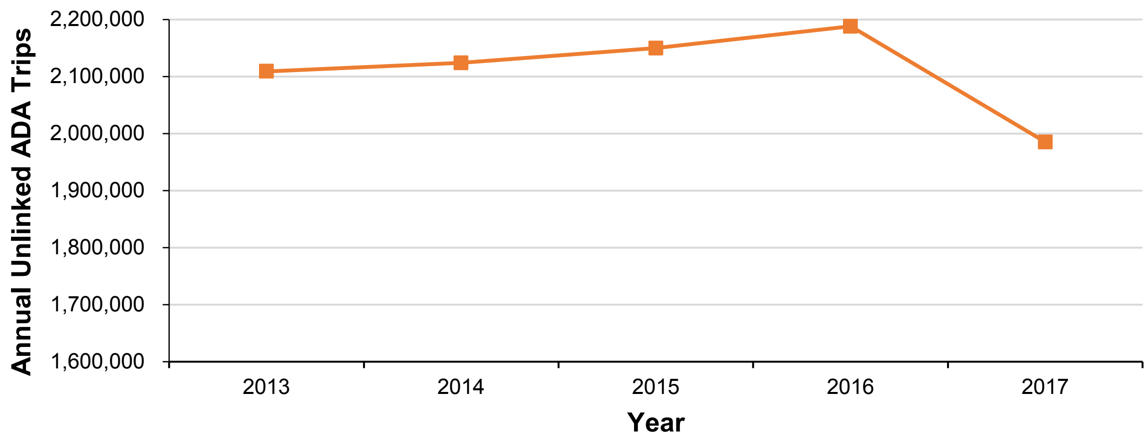

Figure 1.2 Annual Unlinked Trips on the MBTA's ADA Paratransit Service (Source: NTD, 2013-2017)

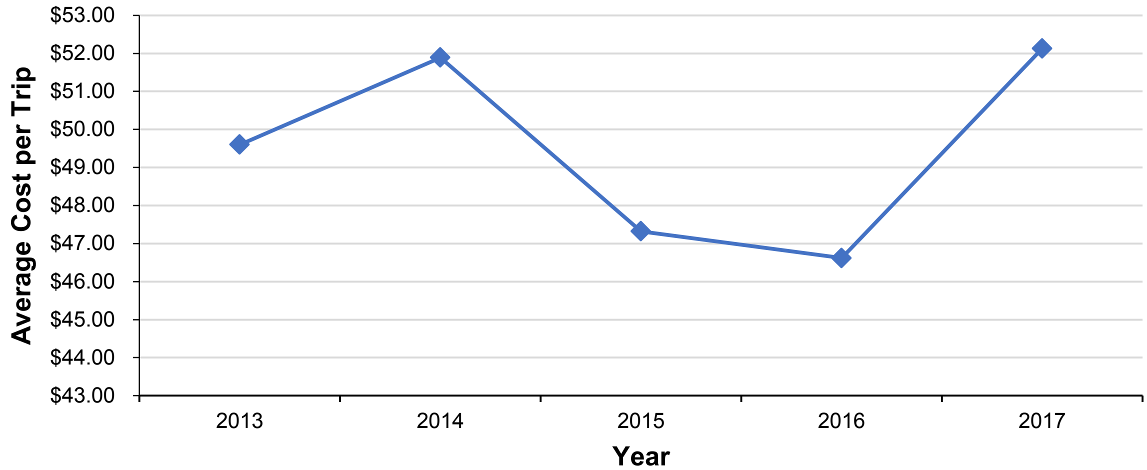

Figure 1.3 Average Operating Cost per Unlinked Paratransit Trip for MBTA (Source: NTD, 2013-2017)

1.2 The On-Demand Paratransit Pilot Program

In an effort reduce the costs of providing transportation service to eligible ADA paratransit customers, the MBTA has initiated an On-Demand Paratransit Pilot Program (referred to hereafter as the Pilot) to allow customers to use ridesharing services or Transportation Network Companies (TNCs), including Uber, Lyft, and Curb, for subsidized travel. The motivation for establishing this pilot program is two-fold:

1. Reduce Operating Costs: With the average cost exceeding $50 to serve each trip using The RIDE's conventional ADA van service, there appears to be an opportunity to reduce costs by paying the fare of a ridesharing or TNC service instead. Partnering with TNCs could therefore lead to cost savings if customers choose to make trips at lower cost than using the conventional ADA paratransit service.

2. Improve Mobility for Customers: A secondary benefit of partnering with TNCs, is that eligible ADA customers would have access to same-day on-demand service. Unlike The RIDE, which requires that customers book trips at least one day in advance and wait for the van to arrive within a pick-up window spanning 20 minutes, customers can hail TNCs using a smartphone app, and typically get picked up within a few minutes of their request.

Currently, TNCs are unable to comply with ADA paratransit requirements, because they provide a platform that links customer trip requests with a market of independent drivers. Therefore, the TNCs cannot guarantee that a vehicle will always be available to pick-up a customer request, that a vehicle will be accessible for customers needing a lift or space for wheelchair, or that drivers are trained to serve customers with special needs. As a result, The RIDE has structured the Pilot as an optional service that participating The RIDE customers can use on their own in addition to guaranteed access to conventional paratransit service.

The Pilot was initiated with limited participation in October 2016 and has been expanded to currently allow any eligible ADA paratransit customer to join and participate until July 1, 2019. The general structure of the Pilot is as follows:

1. Participants register for the Pilot in connection with one of three available TNCs: Uber, Lyft, or Curb. Uber and Lyft provide service throughout The RIDE's service area. Curb provides taxi service only within Boston, Brookline, Cambridge, and Somerville (The RIDE's shared area).

2. Participants are assigned a monthly allocation of one-way Pilot trips based on utilization of The RIDE over the previous 6 months. Allocation tiers are 2, 10, 20, 30, or 40 trips/month, which are eligible for the subsidy described below. For additional trips, customers would have to pay the full TNC fare.

3. For each trip, participants choose whether to make a reservation with The RIDE or to use a TNC through the Pilot. Trips with The RIDE are unlimited, local/premium fares are charged the same as for other ADA paratransit riders. For Pilot-based trips (using one of the TNCs), customers pay the first $2 of TNC fare ($1 for UberPOOL) and the MBTA pays the remaining fare up to a subsidy of $40. For example: a TNC trip with $20 fare would cost the customer $2 and the MBTA $18; a TNC trip with $45 fare would cost the customer $5 and the MBTA $40.

The specific structure of the Pilot has evolved over time since its inception. The timeline of changes is summarized in Table 1.1. As of March 2018, there were 1,718 Pilot participants out of approximately 40,000 total registered ADA customers. Although the Pilot program represents only a small number of all ADA customers and eligible trips, utilization provides useful insights about the potential impacts of expanding the Pilot to all users.

Table 1.1 Timeline of The RIDE On-Demand Paratransit Pilot Program

Date	Change to Pilot Program
October 2016	Initial participants allocated 20 trips/month Customer pays $2; MBTA pays the next $13
January 2017	Trip allocation assigned based on previous usage: 2, 20, 25 trips/month
March 2017	Opened to all The RIDE customers, but participants must register Customer pays $1 on UberPool
June 2017	Allocation tiers adjusted to 2, 10, 20, 30, 40 trips/month
October 2017	MBTA subsidy increased to $40 limit per trip

1.3 Study Objectives

The challenge of managing a system of demand-responsive transportation services for customers with disabilities is to accommodate the demand-side considerations related to traveler behavior with the supply-side considerations related to the system structure and costs. The research study culminating in this report has three objectives related to the operation of ADA paratransit and coordination with TNCs.

1. Supply (Operations and Cost) – Identify the structure of operations for the The RIDE's conventional ADA paratransit service. This includes estimating the effect of increasing paratransit demand on the MBTA resources required to provide the service (e.g., fleet size, vehicle-hours of operation, and vehicle miles of operation). This also includes analysis of sources of variability and the causes of delay in serving paratransit customers. The models that are developed for operations and costs are used to quantify the effect of the Pilot on the remaining operations and costs of the service.

2. Demand – Identify the effect of new TNC service options on ADA demand patterns in terms of choice of service (i.e., The RIDE versus TNCs). Specifically, the goal is to understand the extent to which Pilot participants are adopting TNCs and/or continuing to use The RIDE. A model based on observations of Pilot participants is applied to estimate the effect of expanding the program to automatically include all eligible ADA customers.

3. Optimal Service Provision – Identify efficient service strategies accounting for the effects of new service options on demand and supply. The operations models are used to identify which trips would be most cost-effectively served The RIDE versus one of the TNCs. This leads to development of a tool that could eventually be used to identify which specific trips the MBTA should either allocate or incentivize switching from The RIDE to a TNC.

The goal of this study is to provide insights about how the operation and use of the conventional system, operated as The RIDE, is changing under the Pilot and to provide insights about how a multimodal ADA program is likely to perform. Although the scope is tied closely to analysis of the MBTA system, the insights are likely to have implications for the ADA paratransit systems elsewhere in Massachusetts and around the United States.

This page left blank intentionally.

2 Research Methodology

The research approach for this study consisted of three main parts: review of existing data and technologies, collection of data, and analysis to estimate the occurrence and number of left-behind passengers due to overcrowded trains.

Section 2.1 presents a review of the literature related to crowding in transit systems and the indicators that are used across transit agencies and by the MBTA to measure the passenger experience. The relevant available data that MBTA currently collects is also summarized in this section. A review of technologies that are used for counting pedestrians is then presented in Section 2.2.

Sections 2.3 through 2.6 describe the methods for collecting data. This starts with analyzing existing data sources to identify the locations and times of day that passengers are most likely to experience being left behind. Then, the details for the manual data collection, automated video counts, and device detection are described.

Finally, Section 2.7 presents the modeling methods that are used to make the most accurate estimates of the occurrence of trains leaving behind passengers and the number of passengers that are left behind.

2.1 Literature Review

There has been extensive research over the years on demand responsive transit systems. The review of the existing literature is divided into two main parts. First, the body of research on the operation of demand-responsive transit and ADA paratransit is presented. This includes efforts to organize operations efficiently and model the demand and operations for systems of various sizes or with multiple service zones. Second, a more recent area of research is shedding light on the opportunities to coordinate ADA paratransit service with alternative transportation services. Although there is an established history of ADA paratransit providers coordinating with taxi companies, the relatively recent emergence of TNCs, presents new opportunities for collaboration and coordination.

2.1.1 Demand Responsive Transit Systems

Simulation Models

Many models have been developed to optimize the routing and zoning structures for paratransit services. The approaches used to model and study demand responsive services can be broadly classified as detailed simulation models and approximate mathematical models. Simulation has been employed to assess the effects of zoning strategies as well as time window variations on productivity measures of demand responsive transit service (Fu, 2002). Applications include evaluations of trip miles, deadhead miles, and fleet size for Los Angeles County, California (Dessouky et al., 2008). Studies of paratransit in Harris County, Texas, include comparison of centralized and decentralized operating strategies (Quadrifoglio and Shen, 2013) and the effect of zoning with transfers (Shen and Quadrifoglio, 2011; 2012). Variations in travel time and scheduling reliability can also be modeled with detailed methods (Fu, 1999; Fu and Teply, 1999). However, developing an accurate simulation is expensive because it requires detailed data inputs and a significant amount of time for model construction and calibration.

Approximate Mathematical Models

In contrast, approximate mathematical models are simpler to construct and require readily available data inputs. These types of models show the connection between data inputs and cost estimates using analytical models. A number of aggregate analytical models based on continuous approximations geometric probability have been developed to represent various operation strategies for demand responsive transit, and to predict the total distance, operating time, and fleet size required to serve a given demand. These studies date back to Daganzo (1978), which presents an aggregate analytic model to predict the average total time spent in the system including waiting and ride times using three different vehicle routing algorithms (Daganzo, 1978). The model relates total passengers' waiting times to system attributes like fleet size and vehicle capacity, but it does not allow for a provider to predict who will be experiencing delay, when they will be experiencing it, and what causes one rider to experience delay while another rider does not.

More recent studies have focused more on the operations side and have focused on improving efficiency by understanding characteristics relating to the service provider such as revenue hours and vehicle miles traveled (VMT). Rahimi et al. (2018) presents a continuum approximation model showing how agency cost is affected by demand and other service characteristics including VMT, vehicle hours traveled (VHT), and fleet size. The paper presents models for VHT, VMT, and fleet size and the agency cost as a linear function of these three variables. This model can then be calibrated—in the paper for paratransit operations in New Jersey—to reflect the operating characteristics of any specific operator and can be used to predict the likely effect of system changes.

Gonzales & Amirgholy (2015) presents an analytical model to estimate the total agency cost as a function of the capacity of the system and number of passenger requests. The capacity of the system depends on many characteristics of the system including the fleet of vehicles available to pick up passengers. The paper also presents a user cost model which can predict total system delay for passengers in terms of deviation from their desired pickup times. If the system is undersaturated, the rate at which requests for new rides are coming in is less than the rate at which the fleet can serve the rides meaning that the capacity is enough to handle requests as they come in. If the demand rate is higher than the capacity, a queueing problem arises, and the total delay to user schedules can be predicted based on the distribution of request times and the corresponding capacity of the service. The paper, however, does not go beyond assigning a pick-up time to a passenger meaning that additional delay can be experienced when the van deviates from the decided upon schedule.

Rahimi & Gonzales (2017) presents a study of how zoning effects user's travel time and travel behavior. A gravity model is used to predict travel demand between different zones and a travel time model is developed from continuous approximation models presented in Rahimi et al. (2018). The results of this paper show that trip distribution and the effect of zoning on travel demand is greatly dependent on the characteristics of the zones and their structures. The analysis showed that it is always more efficient to operate one larger zone when the two zones are small or have low demand. Additionally, when an area is split into two zones with transfers, having some overlap is optimal.

All the research mentioned up until this point have been based on aggregate models relating systemwide totals and averages, but do not provide detailed information about individual passengers or vehicles within the system. There are other means and methods that researchers have used to study paratransit operations. Demand responsive paratransit operations and specifically vehicle routing can also be considered as a Pick-up and Delivery Problem (PDP). In this type of problem, the goal is to optimize vehicle routes to handle as many requests as possible based on constraints that result in travel time or cost being minimized (Savelsbergh and Sol, 1995; Desaulniers et al., 2001; Cordeau and Laporte, 2003; Laporte, 2010).

Empirical Models

Additionally, research has been done on paratransit operations using analytic and quantitative empirical models. Daganzo (1984) and Figliozzi (2008) have used analytical models to optimize tours for a set number of vehicles to service different points. Probabilistic constraints on the number and duration of tours, the number of users, and time windows improved these models (Figliozzi, 2009). Using a quality of service constraint, Fu (2003) optimized total time and fleet size, but this was not realistic because demand is always changing. Diana (2006) and Diana et al. (2006) further developed a stochastic model to optimize fleet size based on distribution of demand.

Research has also been done on the effect of various system characteristics using real world paratransit systems. Several experiments were performed using the paratransit systems in Edmonton, Alberta to better understand how the stochasticity of travel time impacts system performance and reliability (Fu, 1999). Baily and Clark (1987), Feuerstein and Stougie (2001), and Fu and Ishkhanov (2004) investigated the effect of the number of vehicles on system performance. Other research has investigated how traffic conditions (Hauptmeier et al., 2000; Lipmann et al., 2002) and time window size (Dessouky et al., 2005) affect DRT performance.

2.1.2 Alternative Service Providers

A second relevant body of research evaluates policies to reduce the cost of ADA paratransit operations by coordinating with other service providers. Transit agencies have made agreements to partner with taxi companies to provide to provide transit services for decades. Many studies have considered the role that taxis can play to complement or supplement other public transportation services for passengers with disabilities (Chia, 2006; Burkhardt, 2010; Tuttle and Eaton, 2012; Ellis, 2016). The potential to reduce the costs of operating paratransit services by partnering with taxis has led to successful taxi voucher programs in many cities, including San Francisco Municipal Transportation Authority in San Francisco, California; Pace Transit in suburban Chicago, Illinois; Metropolitan Transit Authority of Harris County in Houston, Texas; and Washington Metropolitan Area Transportation Authority in greater Washington, D.C. (Burkhardt et al., 2010). TCRP Report 121 presents a spreadsheet toolkit for comparing the effectiveness of non-dedicated fleets (primarily taxis) for augmenting peak demand, handling long trips, or providing service in low-demand periods (Nelson/Nygaard, 2007). The toolkit provides estimates of cost savings resulting from integrating service with non-dedicated vehicles, but the tool does not explicitly consider the geographic distribution of origins and destinations within the service region. A more recent report compares contracting and service models across the United States (Rodman and High, 2018).

TNCs are rapidly changing urban mobility by using digital platforms to allow customers to request on-demand service from suppliers. A number of different business models have been developed. Some TNC services essentially mimic exclusive-ride point-to-point taxicab service, distinguished primarily by the fact the customers use a smartphone to hail or request service. Other TNC services, which are sometimes called "microtransit" seek to group riders into demand-responsive shared rides using sedans or vans that can resemble jitneys or even dial-a-ride shuttles. The aim is to achieve economical operations by simultaneously serving many passengers, although microtransit operators have struggled to stay in business as the economics of mass transit are rarely profitable. These TNC services are being added to the mix of long-existing modes which include public transportation services in the form of fixed routes, paratransit for people with disabilities (as required by the Americans with Disabilities Act of 1990, ADA), and dial-a-ride, as well as taxicabs and for-hire vehicles. Together these long-standing and emerging services provide vital urban transportation options (see Figure 2.1).

Figure 2.1 Examples of New and Existing Mobility Services

With the rapid expansion of TNCs and their adoption by the general public, there have been a number of studies evaluating their relationship to public transit and extent to which TNCs operate in competition with transit or a supplement to existing services (Taylor et al., 2015; Feigon and Murphy, 2016; Shin-Pei et al., 2016). While the potential of TNCs to substitute conventional transit trips and reduce transit efficiency is a cause for concern (Martin et al., 2010; Henao, 2017; Schaller, 2016; 2017), ADA paratransit services are so costly to provide that substitution may have the desirable effect of lowering total operating costs.

Overall, Feigon and Murphy (2016) present TNCs in a positive light, highlighting the potential for beneficial relationships with public transportation services. Other recent studies have acknowledged similar potential for TNCs to work with public transportation to improve public mobility. A study by TransitCenter also investigated the potential for public transportation and TNCs to work together, and identified similar opportunities for benefits but urged public transportation agencies to enter into agreements with a view to reinforce transit strengths and leverage agency-controlled assets such as curb space on streets (Shin-Pei, et al., 2016).

Not all research related to TNCs has been so positive. A recent study of the rapid growth of TNCs in New York City shows that they are quickly increasing vehicle-miles travelled (VMT) on the streets of the city, contributing traffic congestion and competing head to head with the tightly regulated taxicab industry (Schaller, 2017). Another recent study of TNC riders in Denver found that 22% would have taken conventional fixed route transit instead, and many others would have walked or cycled; only 19% had substituted driving alone (Henao, 2017). The Denver study also found that when accounting for the extra VMT associated with TNC trip , as compared to the modal trip replaced, the TNC increased VMT by 84%. These studies provide an indication of some of the risks associated with TNCs in terms of potential to compete directly against public transportation and exacerbate traffic congestion. The experience of taxis also provides insights related to the regulatory and policy implications of dealing with private operators. Schaller (2016) makes the case for thinking carefully about the regulatory environment for TNCs like Uber and Lyft. The current situation is that taxis are a highly regulated industry with controls on entry to the market (e.g., through medallions in New York City), fares, driver credentials and oversight, and vehicle specifications. By comparison TNCs are operating under much looser rules, and this creates some complications for cities and agencies in trying to make formal partnership agreements. This challenge is particularly pertinent to ADA paratransit services, which must be able to make contractual guarantees for service to eligible customers in order to comply with the ADA. Clearly, there are benefits and risks that must be weighed and carefully considered before establishing collaborations or partnerships between public transit operators and TNCs.

Many transit agencies have been encouraged to consider partnering with TNCs to reduce the cost of operations by study results, such as Brookings Institution report (Kane, Tomer, and Puentes, 2016), which point to the dramatically lower cost of TNC fares compared to the costs of conventional paratransit operations. Some customers have voiced strong vocal support for programs like the MBTA Pilot, because TNCs allow for trips to be scheduled minutes before they are needed, rather than requiring a reservation a day in advance. However, proposals to coordinate paratransit service with TNCs are not without critics. For example, a report by the Amalgamated Transit Union (2016) came out with sharp criticism of the Brookings Institution report, pointing out that Uber and Lyft provide low cost service with fleets of ADA inaccessible vehicles and untrained drivers. Although Uber has a platform for wheelchair accessible vehicles (WAVs), the number of available vehicles is outside of Uber's control, which relies on a platform of independent drivers who may choose to operate accessible vehicles.

With increasing interest among transit agencies throughout the United States to consider cooperative arrangements with taxis or TNCs for paratransit services, there is a pressing need for quantitative methods to assess the potential value of these cooperative arrangements. To date, the quantitative modeling of demand responsive systems has focused almost exclusively on paratransit operations as a stand-alone system. Recent studies of paratransit operations provide modeling capabilities to quantify the effect of changes in demand, such as may result for diverting some trips to taxis or TNCs, and the provide a quantitative basis for decision making (Rahimi and Gonzales, 2015; Amirgholy and Gonzales, 2016). A study of the Pioneer Valley Transit Authority (PVTA) ADA paratransit service provided an initial analysis of the potential cost savings from coordinating with taxis or TNCs in the Springfield, Massachusetts, region (Turmo et al., 2018).

It is clear from the existing literature and public debates that partnerships between transit agencies and TNCs have the potential to provide large cost savings, but there are several critical challenges that must be considered and addressed. This research study seeks to investigate the details how a partnership between the MBTA and TNCs in the Greater Boston region are likely to affect operations, demand patterns, and costs. There are several questions related to the legality and equity of partnerships with TNCs that are important but lie outside the scope of this research.

2.2 Available Data

There are three main types of data on which the models and analyses in this study are based. First, customer records from the MBTA provide demographic information about each eligible paratransit customer which can be used to associate travel patterns with personal characteristics such as age and type of disability. Second, detailed records from each ADA paratransit trip served include the locations and times of each passenger pick-up and drop-off as well as identification of the vehicle or route that served each trip. This data not only shows the temporal and spatial distribution of ADA paratransit trips, it can also be used to reconstruct the vehicle routes, which reveals the operations associated with serving the demand. Finally, some limited data about the monthly number of TNC trips made by Pilot participants shows utilization of the Pilot changes over time and how enthusiastically different customers make use of the service.

2.2.1 ADA Paratransit Customer Data

The database of eligible customers for The RIDE contains records for 40,721 individuals. Personal identifying information is not necessary for any of the analysis of this study, but the following data fields are used for the analysis:

1. Customer ID – A unique number is assigned to each eligible ADA customer. This ID allows us to track the trips that each individual makes and relate those trips to other customer characteristics.

2. Date of Birth – The customer's date of birth allows us to calculate age, which has the potential to be an explanatory factor for travel behavior.

3. Home ZIP Code – The zip code for each customer's registered home address provides an indication of where customers reside and where many of their trips are likely to start or end,

4. Disability – The qualifying disability or disabilities associated with customer are recorded, and these have the potential to be explanatory factors for travel behavior.

5. Equipment – In addition to customer disabilities, the type or types of equipment that the customer uses is listed. This includes mobility devices such as wheelchair, power chair, scooter, walker, cane, etc. This is also the field where specific vehicle requirements are listed, such as requirement of a lift or service only with a van. This field is particularly important for identifying which customers are ambulatory and which customers require a WAV.

2.2.2 The RIDE Trip Records

In addition to records about each customer, the MBTA maintains a database of all The RIDE trips. These records include a detailed accounting of where and when each customer travelled, and which vehicle or route was used to serve them. For this study, the MBTA provided the research team with all 4,012,592 trip records from January 2016 (prior to the Pilot's start in October 2016) through March 2018. Each trip record includes the following data that is used in the analysis:

1. Trip ID – Each trip is uniquely identified by an ID.

2. Customer ID – The ID for the customer requesting the trip allows each trip to be linked to the specific customer characteristics in the customer data table.

3. Trip Date – The calendar date of each requested, scheduled, and served trip.

4. Subscription – Customers that make regular trips (e.g., to and from work) are able to request their trip as a subscription rather than having to call in the same request over and over again. This data field indicates the ID of the associated subscription, if applicable.

5. Provider – Each trip is served by a private operator that works under contract with the MBTA. This field indicates which provider serves the trip. This provides an indication of the region in which the trip is assigned, because each of the three regions is initially served by a different provider. Some reorganization during the time period of observation has resulted in changing geographic coverage for each provider.

6. Pick-up Location – The address and latitude/longitude coordinates of the requested pick-up location are recorded. This is used (along with the drop-off location) to determine if the trip is within the required ADA service area or in the broader "premium" service region in which the ADA does not require service.

7. Drop-off Location – The address and latitude/longitude coordinates of the requested drop-off location. This is used along with the pick-up location to categorize trips.

8. Origin-Destination Network Distance – The estimated driving distance from the pick-up location to the drop-off location is recorded, assuming the trip can be served as a direct ride without intermediate stops. Ultimately most trips are served directly in this manner, but some vehicles are routed to share multiple rides, so the actual distance travelled by a customer may be somewhat greater.

9. Estimated Trip Time – Based on the location and time of day of the pick-up and drop-off, a travel time estimate is generated by the scheduling software for a direct trip following the network distance above. This is a travel time estimate that may be greater or less than the actual travel time for the passenger.

10. Requested Pick-up Time – This is the time that the customer initially requested to be picked up by The RIDE.

11. Promised Pick-up Time – This is the time that The RIDE offered to the customer during the booking process. Customers are expected to be prepared to board the vehicle from 5 minutes before to 15 minutes after the promised time.

12. Arrival Time at Pick-up – This is the time that the vehicle arrived at the pick-up address. As described above, the vehicle is intended to arrive between 5 minutes before to 15 minutes after the promised pick-up time. Any arrival after this time window is considered to be late.

13. Departure Time from Pick-up – This is the time that the vehicle departs the pick-up location. The difference between the departure time and the arrival time at pick-up is the time that the driver waits for the customer to get ready and to get into the vehicle.

14. Arrival Time at Drop-off – This is the time that the customer actually arrives at his or her destination. The difference between the arrival time at drop-off and the departure time from pick-up is the time that the customer spends traveling in the vehicle, including any intermediate stops. For trips that are served without intermediate stops, this elapsed time can be used with the origin-destination network distance to calculate the average speed of the vehicle in the network.

15. Vehicle ID or Route ID – Depending on the month, the data set includes a field for the vehicle ID or route ID. Within a day, all trips with a common vehicle/route ID can be grouped to identify the actual vehicle routing. By linking together trips in this way, the actual operations of all The RIDE vehicles can be deduced in terms of the total number of vehicles operating, VHT, and VMT.

The monthly demand varies slightly over the course of the year, as shown in Figure 2.2. Over the time period observed (January 2016 through March 2018) the general level of demand appears to have held steady. In fact a comparison of total trips for the first quarter of 2016 and 2018 shows a 12.5% drop over this time.

Within a day, the rate that trips are requested and served is constrained by the hours of operation of MBTA buses and subways. Generally ADA service is provided between 5 AM and 11 PM, and Figure 2.3 shows that the majority of trips are served between 6 AM and 9 PM. Although the mid-day demand is relatively flat, there are distinct peaks in the late morning and early afternoon. These peaks in demand determine the maximum number of vehicles and drivers that are needed to fulfil the required ADA service.

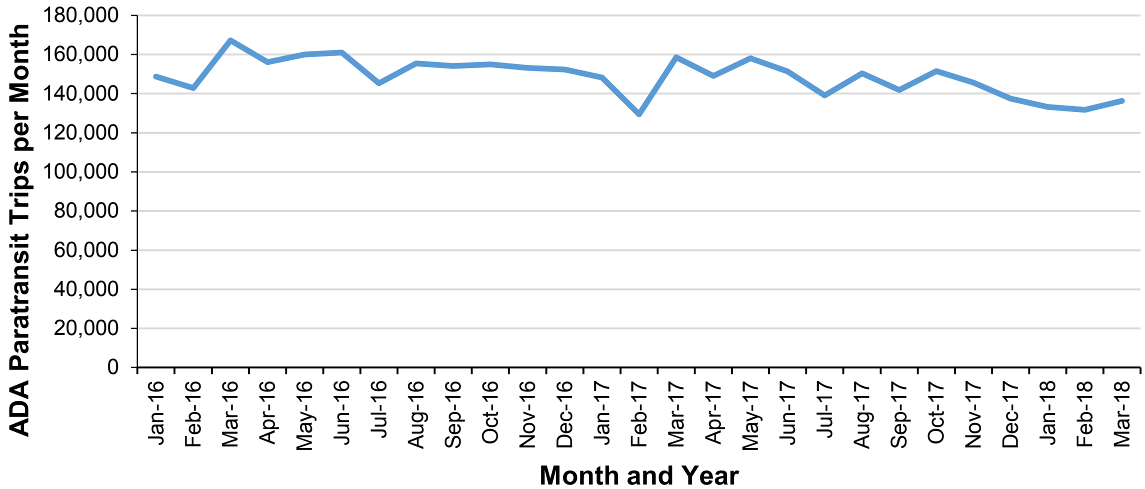

Figure 2.2 Monthly Trips on The RIDE, January 2016 – March 2018

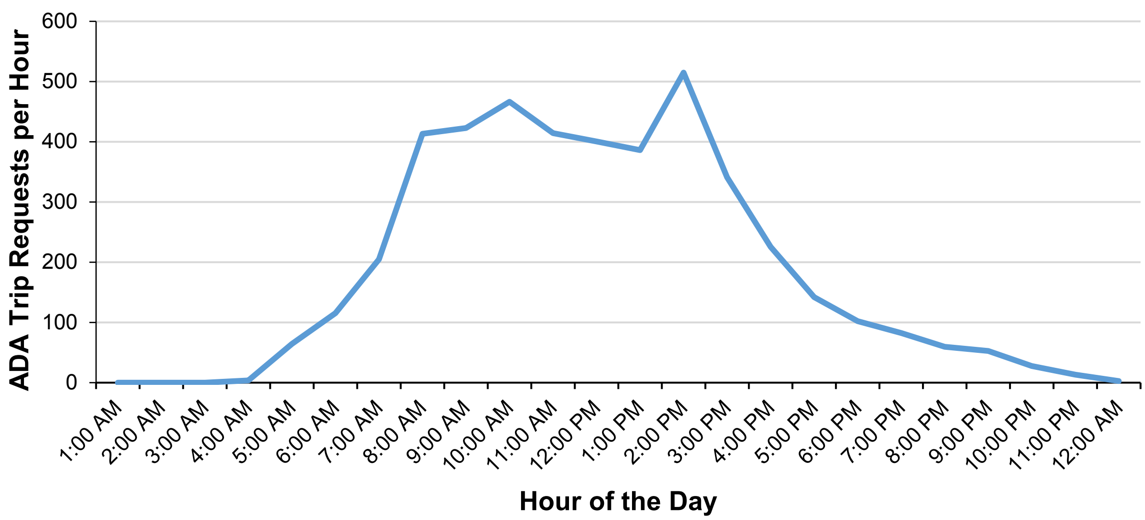

Figure 2.3 Hourly Trips Requested on The RIDE, Average of January – March 2018

2.2.3 Pilot Program Data

Data for the Pilot is the most limited, because service is provided by TNCs (specifically Uber, Lyft, and Curb) which tightly control which data are released and to whom. Although records exist of each of the trips completed by Pilot participants, only the following data for each Pilot participant was available to the research team for the analysis in this study.

1. Customer ID – The customer ID for each Pilot participant links to the table of eligible ADA customers. This allows participation in the Pilot to be linked with customer characteristics and ongoing travel booked through The RIDE.

2. Date Joined – The date that the customer joined the Pilot is recorded. Figure 2.4 shows that the program initially kicked off with a limited number of participants on October 1, 2016. Once all ADA eligible customer were able to join the program on March 1, 2017, customers have steadily joined the Pilot. This date is important for identifying how long Pilot participants have been included in the program.

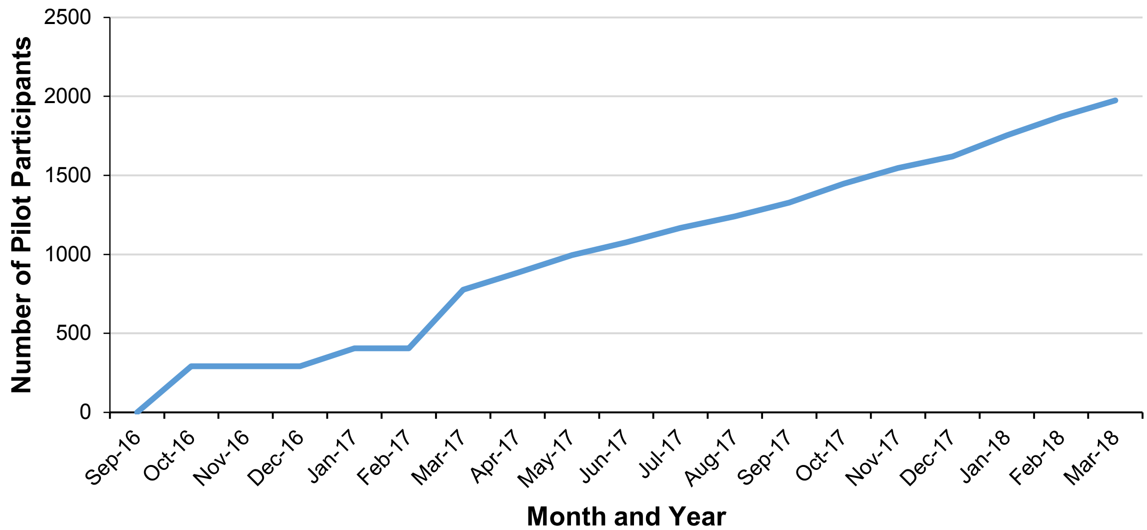

Figure 2.4 Cumulative Number of Pilot Participants

3. L6M-Pre Pilot – The average monthly number of trips completed on The RIDE in the six months prior to joining the program is recorded. The trip history is used to determine the trip allocation for subsidized TNC trips.

4. TripCap – The maximum number of monthly subsidized trips that each customers has been allocated through the Pilot program is listed. Although the tiers of trip allocations have been adjusted over time, the most recent tiers (as of June 2017) are reported. These are: 2, 10, 20, 30, or 40 TNC trips per month. For most customers, the L6M-Pre Pilot count is rounded up to the nearest tier level to determine the assigned tier for TNC trip allocation.

5. Monthly TNC Trips – For each month, the total number of subsidized TNC trips completed through the Pilot is reported for each customer. No additional details are provided about the day, time, distance, or cost of these trips. However, the monthly totals allow for analysis of general changes in travel behaviour over time.

2.3 Modeling Paratransit Operations

For the purpose of this study, there are two types of operation modeling that are utilized. The first is an aggregate operations model that relates characteristics of the service regions (including area, demand, and network traffic speeds) to estimate the number of vehicles needed for the vehicle fleet, the VHT, and the VMT required for operations. These parameters of operations are all linked to the costs of operating the paratransit service. Although the actual cost to a transit agency, like the MBTA, depends on the details of the operation, such as the number and size of garages, age and types of vehicles, salaries and benefits for workers.

The aggregate model is analytical, and it is useful for estimating the effect of changes to the system on total operating costs. Second are empirical models that are fitted to observations about operations in the data. The empirical models are useful for characterizing variability in performance that is not explicitly represented by the aggregate operations model.

2.3.1 Processing of Operations Data

Before proceeding with the analysis of the dataset, trip records were filtered to remove entries that were incomplete or did not make sense. The filters used for this procedure have been removing the following categories of data:

1. Records without a Vehicle ID were eliminated, because there is no way to associate these trips with operations.

2. Trips with negative or zero values for the actual trip time (difference between departure from pick-up and arrival at drop-off) are physically impossible. The trip record is assumed to have an error, and is eliminated.

3. Trips for which the estimated trip time is or more than 3 times larger than the actual trip time are eliminated. Since the estimated trip time is for a direct ride from the pick-up to drop-off, the actual routing cannot be more direct than this. Although it is possible for traffic conditions to be lighter than expected, a factor of 3 is used to flag trip records for which an error is suspected.

4. Trip records with loading time greater than one hour and data with negative loading times are eliminated. Although it possible for a passenger to take a long time to board a vehicle, it is assumed that reported times in excess of one hour indicate that something else was going on. Obviously, any negative values are physically impossible and indicate and error in the data.

The implementation of the aforementioned filters on each of the 12 months of 2017 led to removing the data as summarized in Table 2.1. Most the filtering criteria result in elimination of a very small number of trips. The biggest source of errors is the first filter, which eliminates records that lack a vehicle ID. This is particularly problematic in June 2017, which coincides with a reorganization of which operators serve each of The RIDE's service regions.

Table 2.1 Filtering of Raw Trip Records for Operations Analysis, 2017

Month	Raw Trips	Filter 1	Filter 2	Filter 3	Filter 4	Filtered Trips
January	148,154	18	245	206	425	147,260

Month	Raw Trips	Filter 1	Filter 2	Filter 3	Filter 4	Filtered Trips
February	129,337	224	304	206	463	128,140
March	158,499	13	393	245	508	157,340
April	149,121	2,963	395	222	455	145,086
May	157,985	22	442	244	550	156,727
June	151,343	31,425	703	169	457	118,589
July	139,131	9,219	2,179	218	483	127,032
August	150,389	0	991	285	519	148,594
September	141,821	0	1,086	242	502	139,991
October	151,446	0	1,083	201	540	149,622
November	145,571	0	811	222	542	143,996
December	137,465	0	1,484	231	567	135,183
Total *%*	*1,760,262*	*43,884* *2.49%*	*10,116* *0.57%*	*2,691* *0.15%*	*6,011* *0.34%*	*1,697,560* *96.44%*

Initially, each of the three service regions (as shown in Figure 1.1) were served by different operators: GLSS (Greater Lynn Senior Services, Inc.) in the North; VTS (Veterans Transportation, LLC) in the West; and JV (a joint venture between Thompson Transit, Inc. and YCN Transportation, Inc.) in the South. In February 2017, TRAC was introduced to replace GLSS and JV, leaving only a handful trips served by JV in March, April, and May. As of June 2017, TRAC completely replaced GLSS and JV in the North and South regions. This shift is apparent in the change in operators associated with trip pick-up locations shown in Figure 2.5.

2.3.2 Relevant Explanatory Variables

In order to calibrate the models, observed values for parameters related to the operation of the system must be measured or estimated form the available data. As some conditions change over the course of the day, the trip records broken up by operator, weekday versus weekend, and time period of the day.

For the purpose of this analysis, we consider time periods of length $t_p = 3$ hours, which results in 5 time periods per day: 6 AM – 9 AM; 9 AM – 12 PM; 12 PM – 3 PM; 3 PM – 6 PM; and 6 PM – 9 PM. Very few trips are completed outside of these hours, and they are not considered for the calibration of the aggregate operations modeling. The following subsections describe how values are estimated from the available data. Values from the analysis for each operator, day of week, time of day, and month of year are included in the Appendix. Average values by operator, day of week, and time of day are used for loading/unloading time, vehicle occupancy, and vehicle speed. The VMT, VHT, and required fleet are calculated for every day and time based on the corresponding demand and these average values.

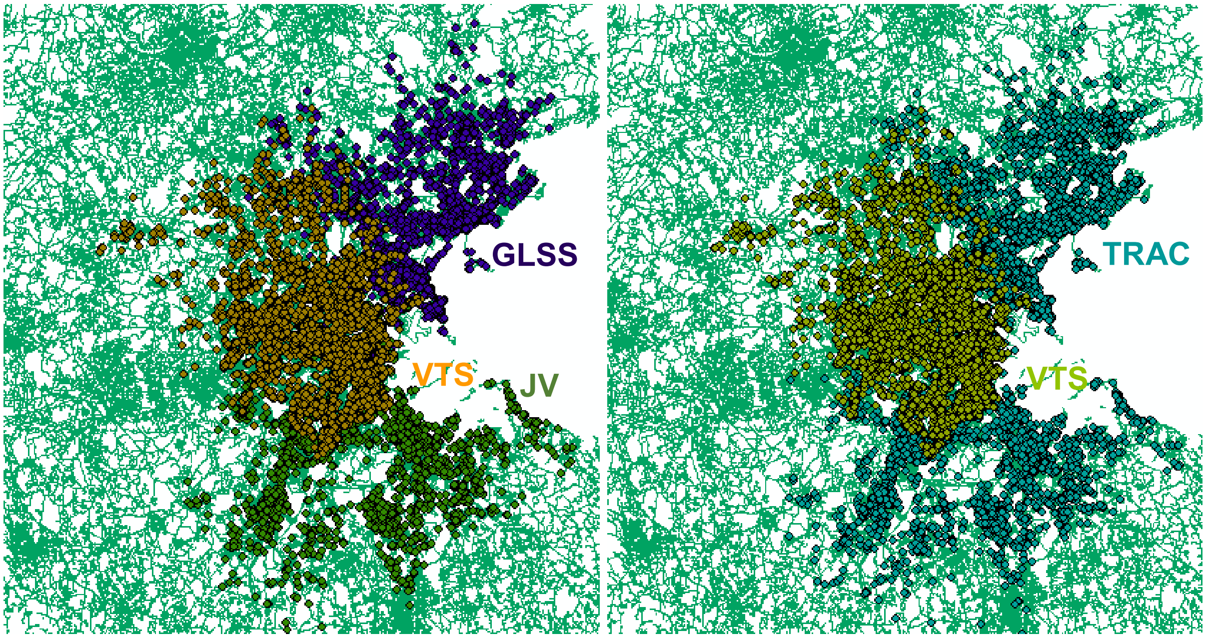

Figure 2.5 Distribution of trip origins by The RIDE operator in January 2017 (left) and June 2017 (right)

Service Region Area, *A*

The size of each service region is assumed to be constant throughout the day and is determined by summing the zip codes associated with each service region (see Table 2.2). These service areas represent the combined premium and local fare areas, because ADA paratransit service is provided throughout the full extent of the region.

Table 2.2 The RIDE Service Region Areas

Region	Area (sq. miles)
North (not including shared)	211.8
West (not including shared)	216.6
South (not including shared)	330.1
Shared	64.0

Average Loading/Unloading Time, *b*

The time for each customer to board and alight the vehicle is denoted by b, and the average is assumed to be constant over time. The trip records include the arrival time of the vehicle at the pick-up location and the departure time from the pick-up location; the difference between these values is the loading time. The unloading time is not observed, because only records of vehicle arrival time at the drop-off location are available. A working assumption is that the unloading time is one third as long as loading, because drivers do not have to wait for customers to get ready and come out to the vehicle. The average value for b is simply the average of all observed times in the trip records for the day and time period of interest. The average value of b by operator, day of week, and time of day is summarized in Table 2.3.

Table 2.3 Average Loading/Unloading Times, *b* (min/passenger)

Time Period	GLSS (North)		VTS (West)		TRAC (North/South)	
	Weekday	Weekend	Weekday	Weekend	Weekday	Weekend
6 AM – 9 AM	5.03	5.25	7.48	8.82	6.33	7.40
9 AM – 12 PM	5.37	5.65	8.13	8.91	6.95	7.22
12 PM – 3 PM	6.19	6.27	9.26	8.55	8.28	7.70
3 PM – 6 PM	6.59	6.62	9.30	9.62	8.10	8.41
6 PM – 9 PM	6.76	6.93	9.38	10.13	9.00	9.13

Average Vehicle Occupancy, n

The vehicle occupancy is estimated as the number of passengers onboard after a passenger boards. In order to make this estimation, the times of trip pick-ups and trip drop-offs for each vehicle are sorted in order to construct the actual route sequence. A cumulative count of the number of passengers onboard is tracked by increasing the count with each pick-up and decreasing the count with each drop-off. A summary of the percentage of total time by number of passengers onboard is shown in Figure 2.6. The three operators are similar, and vehicles in all regions spend most of their time without any passengers onboard at all. Although the vans are observed to carry as many as 8 passengers, the vehicles are rarely loaded with more than two passengers at a time.

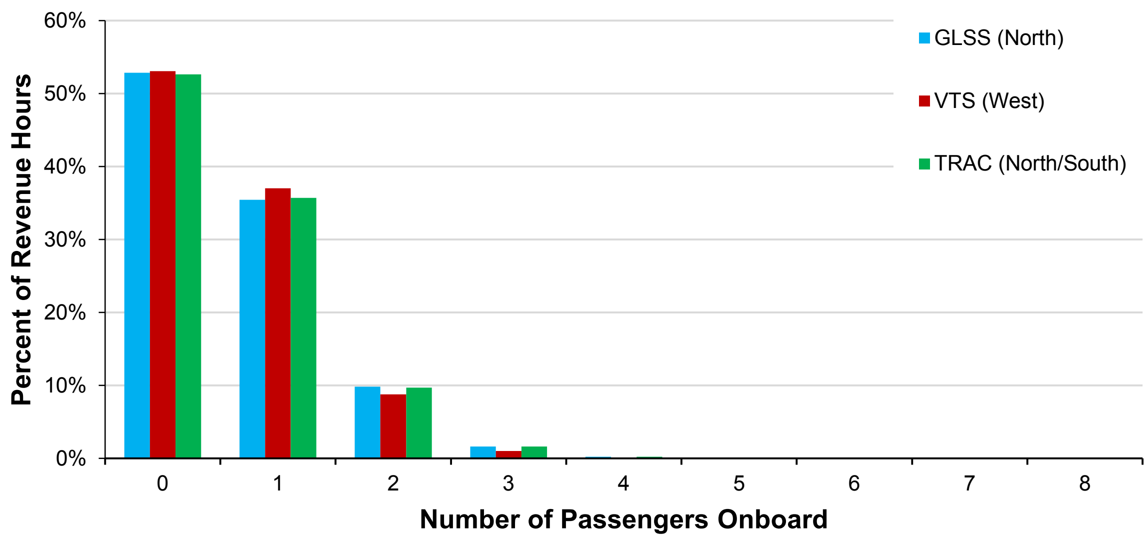

Figure 2.6 Distribution of Vehicle Occupancy by Vehicle Hours Traveled

All vehicle occupancies associated with passenger pick-ups are averaged for the days and time periods of the interest to obtain the average value for n. This represents the average number of passengers onboard after a passenger enters the vehicle, which represents the number of potential destinations for the next drop-off. If each passenger were carried individually from pick-up to drop off, then $n = 1$. The average value of n by operator, day of week, and time of day is summarized in Table 2.4.

Table 2.4 Average Vehicle Occupancy, n (passengers/vehicle)

Time Period	GLSS (North)		VTS (West)		TRAC (North/South)	
	Weekday	Weekend	Weekday	Weekend	Weekday	Weekend
6 AM – 9 AM	1.46	1.24	1.41	1.20	1.47	1.26
9 AM – 12 PM	1.35	1.33	1.27	1.22	1.31	1.28
12 PM – 3 PM	1.38	1.29	1.35	1.23	1.40	1.29
3 PM – 6 PM	1.49	1.31	1.41	1.22	1.50	1.30
6 PM – 9 PM	1.28	1.26	1.25	1.18	1.28	1.27

Average Network Speed, v

Using the reconstructed vehicle routes that were created to estimate n, the trip segments during which a vehicle carries a passenger directly from their pick-up to their destination are identified (i.e., exactly one passenger on board preceded and followed by the vehicle being empty). For these trips, the reported network distance is the actual distance traveled, whereas trips with intermediate stops are associated with additional travel distance and time. The average network traffic speed is given by dividing the network distance by the time from departure from pick-up to arrival to drop-off. The average value of v by operator, day of week, and time of day is summarized in Table 2.5.

Table 2.5 Average Network Speed, v (miles/hour)

Time Period	GLSS (North)		VTS (West)		TRAC (North/South)	
	Weekday	Weekend	Weekday	Weekend	Weekday	Weekend
6 AM – 9 AM	15.56	22.06	15.13	20.21	16.03	21.85
9 AM – 12 PM	17.83	18.69	15.57	16.50	17.60	18.48
12 PM – 3 PM	17.29	17.70	15.56	15.96	16.74	17.41
3 PM – 6 PM	14.69	18.19	13.05	16.67	14.01	17.86
6 PM – 9 PM	19.36	20.80	16.54	19.42	18.83	20.91

Demand, λ

The demand rate within a day and time period is a simple count of the number of trip records observed for the day and time of interest. This rate is expressed as a number of requested pick-ups per time. To fit the parameters of operations models for VMT and VHT, the specific demand and operational outcome from every day and time period in the data set are used. A summary of the average demand by operator, day of week, and time of day is summarized in Table 2.6.

Table 2.6 Average Demand, λ (trips/hour)

Time Period	GLSS (North)		VTS (West)		TRAC (North/South)	
	Weekday	Weekend	Weekday	Weekend	Weekday	Weekend
6 AM – 9 AM	128	43	146	41	191	71
9 AM – 12 PM	183	92	165	72	256	131
12 PM – 3 PM	181	82	188	70	262	119
3 PM – 6 PM	118	56	126	50	173	81
6 PM – 9 PM	35	27	35	23	52	41

Time Window, *T*

The window of time within which a customer pick-up is considered to be on-time is a policy variable. For The RIDE, the policy is to pick-up customers within a 20 minute window from 5 minutes before the scheduled pick-up time to 15 minutes after.

2.3.3 Observed Operational Outcomes

Network Distance Between Points

The constructed routes that were used to calculate the vehicle occupancy provide a sequential list of stops that vehicles make over the course of the day. A vehicle's travel between consecutive stops will be described as a *segment*. The straight-line distance associated with each segment is calculated based on the difference of latitude and longitude of the coordinates. Trip segments that correspond to a single customer's travel directly from pick-up to drop-off have a corresponding network distance reported in the data set. By comparing the straight-line distance and the network distance for these segments, a network circuity factor can be estimated. This when multiplied by the straight-line distance, this factor provides an estimate of the actual network distance traveled. Figure 2.7, Figure 2.8, and Figure 2.9 show the relationship between straight-line distance and network distance for route segments by GLSS, VTS, and TRAC, respectively.

Observed Vehicle Miles Traveled, *VMT*

Although the actual VMT are not reported in the data, the constructed routes based on Vehicle/Route ID are the actual vehicle routings were operated each day. The best estimate of the actual vehicle miles traveled is to sum the network distance associated with each of the identified route segments (some of which are traveled with passengers onboard and others are traveled empty). For route segments that correspond to a direct customer pick-up to drop-off, the network distance reported in the trip data is used. The network distance for all other segments is estimated by calculating the straight-line distance between points and multiplying by the network distance factor derived above. A summary of the average VMT per three-hour time period by operator, day of week, and time of day is summarized in Table 2.7.

Table 2.7 Average Vehicle Miles Traveled per Time Period, *VMT* (vehicle-miles)

Time Period	GLSS (North)		VTS (West)		TRAC (North/South)	
	Weekday	Weekend	Weekday	Weekend	Weekday	Weekend
6 AM – 9 AM	2399	857	2267	644	3713	1455
9 AM – 12 PM	3093	1581	2299	1008	4378	2262
12 PM – 3 PM	3118	1478	2814	1051	4757	2205
3 PM – 6 PM	2091	1072	1976	817	3324	1641
6 PM – 9 PM	705	618	572	442	1125	975

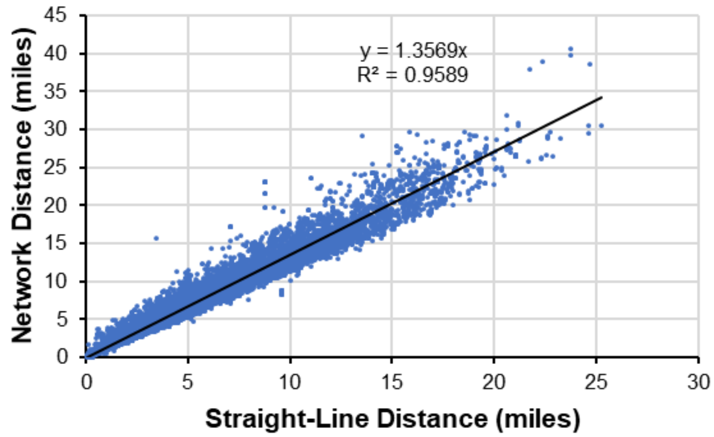

Figure 2.7 Network versus Straight-Line Distance for GLSS (North), 2017

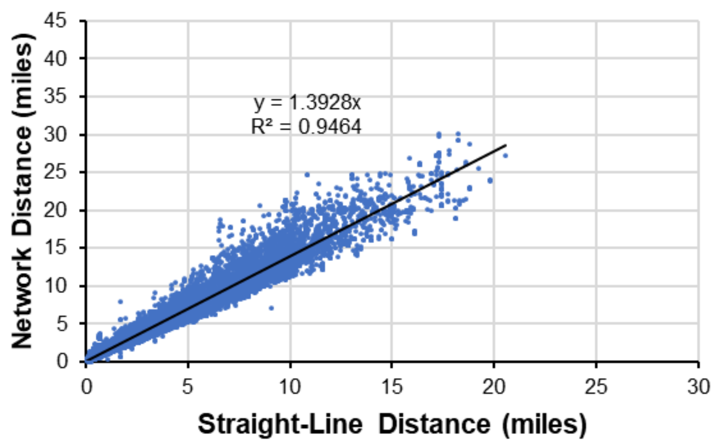

Figure 2.8 Network versus Straight-Line Distance for VTS (West), 2017

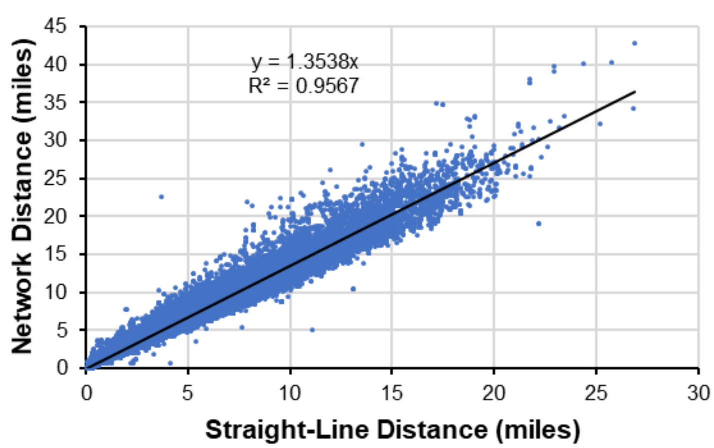

Figure 2.9 Network versus Straight-Line Distance for TRAC (North/South), 2017

Observed Vehicle Hours Traveled, VHT

The observed revenue hours of operations, or VHT, are easier to tabulate than the VMT. Each of the constructed route has a first and last stop. The difference between these times are the VHT associated with each route. When broken up across the time periods of a day, each route segment and its corresponding VHT is associated with the time period at the start of the segment. A summary of the average VHT per three-hour time period by operator, day of week, and time of day is summarized in Table 2.8.

Table 2.8 Average Vehicle Hours Traveled per Time Period, VHT (vehicle-hours)

Time Period	GLSS (North)		VTS (West)		TRAC (North/South)	
	Weekday	Weekend	Weekday	Weekend	Weekday	Weekend
6 AM – 9 AM	419.6	165.9	460.4	167.8	617.1	277.1
9 AM – 12 PM	538.6	268.8	481.2	209.4	799.7	394.6
12 PM – 3 PM	529.4	270.1	541.9	232.7	880.8	444.0
3 PM – 6 PM	405.1	201.3	469.1	194.9	646.2	324.3
6 PM – 9 PM	151.2	101.3	148.1	96.1	269.6	181.9

Observed Fleet Size, M

The observed size of the fleet is simply the number of vehicle/route IDs observed during a day or time period.

2.3.4 Aggregate Operations Model

Models of aggregated VMT, VHT, and fleet size are based on geometric probability and the resources needed to serve a density of demand over each operator's service area. The models are of the form introduced in Daganzo (1978) and Rahimi et al. (2018). These models are based on simplifying assumptions about the distribution of demand in each service regions and the operating algorithm for serving requested trips. What the model lacks in detail and realism, it makes up for in providing an analytical formula that physically relates explanatory factors to operational outcomes. This approach is valuable, because only two parameters (one for the VMT model and another for the VHT and fleet model) must be calibrated to fit the data. All of the other variables are measurable quantities.

The aggregate model builds on the basic operating assumptions for a dial-a-ride system presented in Daganzo (1978). Demand is uniformly distributed within a roughly circular region with area A, and conditions do not change significantly within an analysis time period. For this study, we break each day into time periods of length t_p, within which the demand rate, λ, and network traffic speed, v, are assumed to be constant. At any time, all of the demand within a pick-up window of duration T are potential customers to pick-up. Each vehicle is assumed to operate by first picking up the nearest waiting customers until the target vehicle occupancy, n, is reached. Then, the vehicle alternates between dropping off the on-board customer with the nearest destination and picking up the next nearest waiting customer. In this way, each the number of passengers onboard the vehicle is maintained at a near constant level, and the vehicle is approximately minimizing distance and time traveled by always proceeding to the next nearest stop.

Vehicle Miles Traveled, VMT

The total VMT operated within a time period is the sum of the distances traveled to pick-up each requested trip and then to drop-off each requested trip. From geometric probability, the average distance to the nearest of n uniformly distributed points within an area of size A is:

$$E(d|n, A) = \frac{r}{2}\sqrt{\frac{A}{n}} \tag{1}$$

where r is a unitless adjustment factor for the network that can be thought of as the ratio between the actual network distance and the straight-line distance. The value of r can also be affected by the distribution of the points within the area. If the points are not uniformly distributed, r may be bigger if they are clustered together (e.g., many points near the center of area), or r may be smaller if they are overdispersed (e.g., many points are spread around the edge of the area).

The distance traveled to pick-up a customer is associated with the nearest among λT potential customers. The drop-off is associated with the nearest among n customers on-board. Therefore, the total VMT within an analysis period of duration t_p is given by

$$VMT = r_{VMT}\frac{1}{2}\left(\frac{1}{\sqrt{\lambda T}} + \frac{1}{n}\right)\lambda t_p\sqrt{A} \tag{2}$$

where r_{VMT} is the factor that is calibrated to fit the observed data for the region. This model forms a linear relationship between the right-hand side expression and the VMT, so the value of r_{VMT} can be estimated using linear regression.

Vehicle Hours Traveled, VHT

The model for VHT is based on the VMT model in equation (2) with three important changes. First, the distance traveled is converted to travel time by dividing by the average network speed, v. Second, the time required for loading and unloading each passenger, b, is added. Finally, the calibration factor is replaced by a new parameter r_{VHT}, which allows for the relationship between travel time variables to differ from the relationship between travel distance variables.

$$VHT = \lambda t_p\left[b + r_{VHT}\frac{1}{2v}\left(\frac{1}{\sqrt{\lambda T}} + \frac{1}{n}\right)\sqrt{A}\right] \tag{3}$$

In theory, $r_{VHT} = r_{VMT}$ if there is no wasted time or slack in the system schedule. That is to say, the minimum possible VHT would correspond to a system in which vehicles are always in one of two states: moving toward the next pick-up or drop-off location, or in process of loading or unloading a customer. In reality, such efficiency is impossible to achieve, because there are inevitably gaps in the schedule during which some vehicles must wait until the next customer is ready to be picked up. Therefore, in practice we always expect $r_{VHT} > r_{VMT}$, and the degree to which they differ provides some indication of how efficiently the system is operating compared to an unachievable baseline.

The model for VHT in equation (3) is also linear. For estimation of r_{VHT}, it can be useful to rearrange the terms as follows:

$$\frac{VHT}{t_p} - \lambda b = r_{VHT}\frac{1}{2v}\left(\frac{1}{\sqrt{\lambda T}} + \frac{1}{n}\right)\lambda\sqrt{A} \tag{4}$$

where the slope relating the right-hand side expression to the left-hand side expression is the calibrated value for r_{VHT}.

Fleet Size, *M*

The number of vehicles in operation is closely related to the VHT. Within a time period, operations are assumed to be in roughly steady state conditions, meaning that there are no peaks within each interval of length $t_p = 3$ hours. In this case the fleet required during each time period is

$$M = \frac{VHT}{t_p} \tag{5}$$

because each vehicle is assumed to be fully occupied for the entire time period. For example, if 30 vehicle-hours are operated during a 3-hour time period, this would equate to 10 vehicles in operation. The required fleet size for a region is the maximum fleet size required over the course of a day, so the busiest time period determines the necessary resources.

Operating Cost

The total costs of operating a paratransit service are based on the magnitude of the operational components that are modeled. The annual operating cost of the paratransit system is based on the estimates of annual VHT, VMT, and the fleet size. These operations parameters are associated with dollar costs, which can be estimated from agency operations and cost data. The total cost model takes the following form:

$$TC = \alpha_0 + \alpha_1 VMT + \alpha_2 VHT + \alpha_3 M \tag{6}$$

where α_0 are the fixed costs associated with setting up a paratransit operation in a region, and α_1, α_2, and α_3 are the incremental cost of each vehicle-hour, vehicle-mile, and vehicle in the fleet.

The actual costs to an agency depend on the details of the operating contracts. For example, an agency may enter agreements with subcontractors to pay a rate vehicle-hour of revenue service, which would equate to associating all costs with α_2. On some level, however, the underlying costs of operating a demand responsive transportation service following a pattern as shown in equation (6). In the long run, these costs are likely to be reflected in the rates that subcontractors bid to operate services. Therefore, organizing paratransit operations to minimize the total operating costs will lead to the lowest long run costs for agencies.

2.3.5 Analysis of On-Time Performance

The aggregate operations model provides estimates of the resources required to serve demand under typical conditions. By calibrating the model to data from a specific agency, such as the MBTA, it provides unbiased estimates of total operating parameters over long periods of time. This is useful for estimating monthly or annual fleet requirements and costs, for example. On a day-to-day basis, demand patterns, network conditions, and operating outcomes can vary greatly.

One aspect of system operations and performance that the aggregate model does not quantify of account for is on-time performance. From the extensive set of available trip records, data on on-time performance can be extracted and monitored. Each customer is offered a scheduled pick-up time, and vehicles are supposed to arrive for the pick-up in a 20 minute time window from 5 minutes before to 15 minutes after. Any vehicle arriving after this time is considered to be late. It useful to understand if there are systematic causes of lateness that can related to the operations model, because this would allow us to incorporate an important aspect of the user's experience into the model of costs associated with paratransit operations.

There are two hypotheses that are analyzed as part of this study:

1. Do network traffic speeds affect on-time performance? The hypothesis is that slower traffic speeds are an indication of traffic congestion, which could delay vehicles and make them more likely to arrive late to serve a customer.

2. Does the intensity of vehicle utilization affect on-time performance? The hypothesis is that vehicles are more likely to be late when there is less empty time between trips served. When vehicles are very busy, there may be little slack in the routing schedule. Although this would likely lead to better efficiency in terms of VHT per trip, it also means that there is less down time for vehicles to recover from an disturbance and get back on schedule.

In order to conduct these analyses three values are calculated from the trip records. These are then compared in order to identify whether or not any systematic relationship exists. For the analysis of on-time performance, trip records are aggregated by day to identify if daily characteristics of traffic or demand provide an indication of on-time performance.

Percent of Trips that are Late

Each trip record includes the promised pick-up time and the actual arrival time of the vehicle at the pick-up location. Any time the vehicle arrives more than 15 minutes after the promised time the trip is considered to be late. On a daily basis, the percentage of trips that are late are tracked as a measure of on-time performance.

Average Network Speed

Although an aggregated average network speed is used for modeling the VHT associated with service, once a day has passed trip records can be used to calculate the actual average speed for that day. Specifically, each of the trip segments for which network distance and actual travel time are observed (as described in Section 2.3.2) are used to calculate the average network speed for the time period of interest.

Percent of Revenue Hours that are Occupied

The constructed routes used to calculate parameters for the operations model also reveal which route segments have at least one passengers onboard each vehicle. While an efficient system will seek to keep vehicles occupied with passengers and productively moving customers to their destinations, there are inevitably some times in the day when the vehicle is empty, either while moving from a drop-off location to the next pick-up or when waiting for the next scheduled customer pick-up. Figure 2.6 shows that vehicles are empty more than 50% of the time, on average, but the utilization does vary by day and time. The ratio of occupied revenue hours to total revenue hours provides a metric for how busy a system. This measure has the benefit of being unitless, so a very large region or a very small region can be similarly busy depending on how well the supply of vehicles matches the demand.

2.4 Demand for Paratransit and TNCs in the Pilot

The analysis of demand for TNCs through the Pilot is based on comparing the trip-making behaviors of Pilot participants before and during the Pilot with the general group of non-participants. Limited by access only to monthly trip totals on the Pilot program by each participant, there is insufficient data to develop a trip-level choice model to identify which service a customer is likely to choose in response to time of day, location, length of trip, or fare. Instead, the analysis focuses on estimating the aggregated effects of the program on monthly trip demand and the costs for the MBTA.

The first method of analysis is simply to compare aggregated patterns of the numbers of trips made before and after on The RIDE and by TNCs. The second method of analysis is to develop a modeling process in which clustering is used to identify groups of similarly-behaving customers, and logistic regression is used to identify which cluster a customer is likely to be in based on their observable characteristics. This model is used to forecast participation among customer who are eligible, but not currently registered as participants.

2.4.1 Ridership Changes Among Pilot Participants

A detailed look at Pilot participants requires that we account for a couple of key characteristics that affect their trip making behavior: the date that they joined the Pilot; and the monthly allocation of subsidized TNC trips. The process for analyzing the data is as follows:

1. The rate of monthly trip making by TNCs is plotted over time from the date that participants join the Pilot in order to identify how long it typically takes for customers to settle on a level of use. Typically, when a new transportation choice becomes available to people, it can take some time for people to try it out and decide how often they want to use the service.

2. Trip totals will be compared for the same months of the year. Since the last three available months of Pilot ridership data are January, February, and March 2018, travel patterns of participants will be compared with the same three months in 2016, before the Pilot program started. The values that are compared are the monthly trips on The RIDE in 2016, the monthly trips on The RIDE in 2018, and the monthly trips on TNCs in 2018.

3. Trip totals for customers that are not participating in the Pilot are compared for the same time periods (i.e., monthly trips in The RIDE in 2016 and 2018) as a control group. Any changes in the control group need to be corrected in the Pilot participant data set in order to determine how much demand is substituted from paratransit to TNCs and how much new demand is induced by the TNCs.

2.4.2 Modeling Ridership Behavior

As of March 2018, only a small number of the total eligible ADA customers were registered for the Pilot. In order to estimate the effect of expanding the Pilot to all eligible customers, it is necessary to predict how the current non-participating customers would behave. It is expected that the initial Pilot participants are likely to be enthusiastic early adopters, who may utilize the TNCs more than the average paratransit customer. In order to estimate the potential behaviors of all eligible customers, a model is needed to relate observable customer characteristics to utilization patterns among the Pilot participants. The following subsections describe the modeling steps to estimate likely customer behaviors.

Clustering Analysis of Pilot Participants

For each Pilot participant, three relevant observations of travel behavior are calculated as described in Section 2.4.1: monthly The RIDE trips in 2016 (before the Pilot); monthly The RIDE trips in 2018 (during the Pilot); and monthly TNC trips in 2018 (during the Pilot). A clustering analysis using K-means clustering forms clusters of participants that minimize the difference between each of these three numbers within the groups. The goal of the clustering analysis is to systematically identify types of users with common travel behaviors in terms of the numbers of trips they make and their propensity to use TNCs instead of conventional paratransit. Determining the correct number of clusters is subjective, but the decision is based on a balance between minimizing the differences within clusters and maximizing the differences between clusters (Hartigan and Wong, 1979).

For each cluster, the average of the three values is calculated and compared. Two important values are calculated for each cluster k: the percentage of initial ADA trips that remain on The RIDE, $P_{ADA,k}$; and the number of TNC trips expressed as a percentage of the initial ADA trips, $P_{TNC,k}$. Typically, clusters can be used to characterize various general demand profiles. These clusters will also be used for the next modeling step.

Logistic Regression to Determine Cluster Assignment

Each Pilot participant is a registered ADA paratransit customer with personal characteristics recorded in the customer database. Once clusters have been defined based on a common demand patterns, a logistic regression will be used to identify which personal characteristics have the strongest power to predict which cluster a person will be assigned. A logistic regression differs from other types of regression in that the model provides estimated probabilities associated with discrete possible outcomes. Mathematically this is the same as the multinomial logit model commonly used to model transportation mode choice decisions.

The structure of the logistic regression is based on fitting the parameters of a log-odds function for each of the possible outcomes. In this case, we seek to estimate the probability that customer will assigned to a cluster k out of K total clusters. For each cluster k, a linear log-odds function is specified as follows for customer c

$$V_{c,k} = \beta_0 + \sum_{i=1}^{I} \beta_i x_{c,i} \tag{7}$$

where $x_{c,i}$ is one of I independent variables for customer c, β_i is the model parameter associated with x_i, and β_0 is an alternative specific constant. The types of independent variables that would be appropriate to consider for this model would be any values for which we have observations for the general group of customers that are not participating in the Pilot. These variables include: the initial monthly number of trips on The RIDE, customer age, customer gender, and type of equipment or vehicle requirement.

For each customer, the predicted probability of being assigned to cluster k is given by the following expression:

$$p_{c,k} = \frac{e^{V_{c,k}}}{\sum_{k=1}^{K} e^{V_{c,k}}} \tag{8}$$

The model parameters are estimated using maximum likelihood estimation. This is an optimization to maximize the log-likelihood of the model given by the sum of the natural log of the estimated probabilities associated with each customer's actual assigned cluster in the data sample for estimation.

$$LL = \sum_{c=1}^{C} ln\left(p_{c,k=\text{assigned cluster}}\right) \tag{9}$$

In the context of the Pilot program, C is the number of Pilot participants composing the data sample used to estimate the parameters of the logistic regression model. The β parameters are selected to maximize the estimated probabilities associated with the clusters to which each participant was actually assigned in the K-means clustering step.

Estimation of Demand Impacts of Pilot on All Customers

On its own, the logistic regression is of limited value, because it provides estimated probabilities associating each customer with a cluster that is that has been defined only within this analysis. The value of this estimate, however, is that each cluster is characterized by a different behavioral response to the introduction of TNCs as an alternative to paratransit. Specifically, the data summary following the clustering task provides an estimate of the percentage of initial ADA paratransit trips that will remain on The RIDE and the number of TNC trips that will be made (also expressed as a percentage of the initial demand). The estimated probabilities associated with each cluster are used to calculated the expected value of the number of ADA paratransit trips, N_{ADA}, and the number of TNC trips, N_{TNC}, that will be made upon inclusion in the Pilot. These values are given by

$$N_{ADA} = \sum_j \sum_k p_{j,k} P_{ADA,k} N_j \tag{10}$$

$$N_{TNC} = \sum_j \sum_k p_{j,k} P_{TNC,k} N_j \tag{11}$$

where N_j is the number of initial monthly paratransit trips completed by customer j, and the individual estimates are summed across all customers to obtain an aggregated systemwide estimate of demand impacts.

2.5 Optimizing Allocation of Trips to Paratransit and TNCs

The aggregate operations model described in Section 2.3.4 provides unbiased estimates of the total operating parameters associated with serving a level of demand in a service area. These totals are useful for estimating the total monthly or annual costs of operations, but the model is not sensitive to specific variations in the timing and location of requested trips. By its very nature, the aggregate model treats all trips as equivalent components of the total demand λ.

A previous study of the PVTA utilized the aggregate model to identify the number of trips that should be shifted to taxis or TNCs to minimize the combined cost of the system (Turmo et al., 2018). However, that model is lacking because not all paratransit trips have the same impact on operations and operating costs. For example, a trip that happens to fall along the path of an otherwise empty vehicle can be served at very little cost to the agency. On the other extreme, a requested trip at the edge of the service area during early, late, or peak hours might all require an additional vehicle to be put into service to drive out to serve the requested trip at great cost.

Although the regulatory landscape does not yet allow transit agencies to assign ADA paratransit trip requests to TNCs, it is conceivable that this may be possible in the future.[1] In order to decide which trips to allocate to TNCs versus keep on the ADA van service, it is necessary to estimate the marginal cost of each ADA paratransit trip and the corresponding cost of service by TNC. In order to do this, the specific routing of vehicle must be known so that the incremental effect on cost of unilaterally eliminating each requested trip can be calculated. The trip records provide information about the actual vehicle routes that are operated each day, but the task of allocating all trips requires the routes can be incrementally re-optimized each time a requested trip is shifted to a TNC.

---

[1] One way that this option could be structured is to allow customers to self-select into program that allows the MBTA to assign trips to either conventional ADA service or a TNC. The details of how this might be accomplished are beyond the scope of this study.

The following subsection present a proposed approach to quickly create a routing plan for vehicles based on a set of actual trip requests in the region. Then the marginal cost of each trip is estimated for each trip as a result of this routing and compared against the estimated TNC fare of the same trip. The trip with the greatest cost benefit for switching is removed from the pool of ADA trips, and the routes are re-optimized. In this manner, trips are incrementally shifted to TNCs until no cost savings can be achieved. It is possible that all trips should ultimately be shifted to TNCs or that some subset of the total ADA demand should shift. This approach is designed in such a way that the algorithm could be run daily as part of the vehicle routing solution.

2.5.1 Algorithm to Construct Representative Routes

A fast algorithm is needed to construct the hypothetical vehicle routes, because the procedure will be run iteratively each time a trip is allocated to TNCs. The problem with using the observed routes is that these observations are retrospective of operations that have already occurred, and we need an algorithm to predict operations and costs for a set of demanded trips that has not yet been served. A fast algorithm for constructing routes is a Greedy Algorithm, which is a heuristic in which each vehicle route is constructed in sequence by choosing among available trips that result in the most efficient route. The works as follows:

1. Daily trip data within a region is sorted chronologically by requested time.

2. The first vehicle route starts from the first requested trip of the morning. Assuming the first pick-up is on-time, the arrival time at the drop-off location is estimated based on the straight-line distance factored up by the network circuity factor and divided by the average network speed. Upon drop-off, the vehicle becomes available to serve the next customer.

3. The time each other unserved pick-up request is calculated by adding together the estimated travel time (straight-line distance factored up for network circuity and divided by average speed) and then additional waiting time until the requested pick-up time. Any trips that could only be served with negative waiting time are eliminated as infeasible next pick-ups. The trip with the shortest time from drop-off to pick-up is selected as the next trip in the route.

4. Steps 2 and 3 are repeated until one of two constraints are reached: the duration of the route has reached the maximum length of a shift (if such a constraint is desired), there are no more trips at the end of the day left to serve.

5. Steps 2, 3, and 4 are repeated to construct each route until there are no unserved trip requests left.

6. The daily totals for VMT, VHT, and the required fleet size are calculated from these constructed trips in the same manner used for the actual vehicle routing plan. Although the Greedy Algorithm does not exactly match the observed operations, the model produces estimates that are proportional (see Figure 2.9).

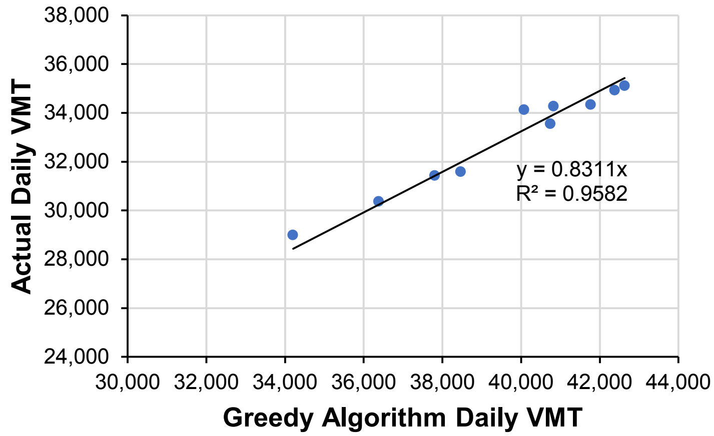

Figure 2.10 Comparison between Greedy Algorithm and Observed Operations, GLSS (North) January 23, 2017

2.5.2 Estimation of Marginal Cost of Each Trip

Marginal Cost of Each Paratransit Trip

The marginal cost of each trip is estimated by considering the effect of unilaterally removing the trip on the remaining costs of operations. Each trip falls into one of three cases, each having different degrees of impact on operations and cost.

1. Type 1 – Trips are in a route that contains only that trip are the costliest. Eliminating the trip reduces the required fleet size by 1 vehicle; eliminates the VMT associated with going to pick-up, drop-off, and loaded travel in-between; and eliminates the VHT associated with the route. These trips have a very high marginal cost, because reducing the number vehicle in the fleet saves a lot of money. Type 1 trips are associated with the peak demand times during which all other vehicles are occupied and an additional vehicle must be brought into service to serve a single requested trip.

2. Type 2 – Trips that are at the beginning or end of a route have a moderate cost. Eliminating a Type 2 trip does not affect the fleet size, but it does eliminate the VMT associated with serving the trip and reduces the VHT by allowing the vehicle to start operating later or stop operating sooner. Once all of the Type 1 trips have been eliminated, Type 2 trips are the most likely to have high marginal cost.

3. Type 3 – Trips that are served in the middle of a route typically have the lowest cost, because eliminating the trip only affects VMT. The fleet size and VHT is unchanged because the vehicle must still be out in operation to serve the preceding and following trip. The effect of removing a Type 3 trip is only the change in VMT associated with deviating the vehicle's route for the pick-up, to carry the passenger, and after drop-off. This saving is offset by the distance that would have traveled anyway from the previous drop-off to the next pick-up.

Cost of TNC Trip

The cost of serving a trip by TNC varies depending on the specific service provider, time of day, and length of trip. It is not possible to know exactly what the trip will cost, because prices can fluctuate in real-time in response to the relative supply and demand. The basic TNC fare is relatively consistent. For example, Uber publishes that standard fares for various Uber services in Boston. For UberX, the basic Uber service, the estimated cost of a trip is

$$F = \max\{6.85, 3.95 + 0.36t + 0.88d\} \tag{12}$$

where t is the trip duration in minutes and d is the network trip distance in miles. Uber's fares are structured so that a minimum of \$6.85 is charged no matter how short or fast the trip is.

The estimated fare translates directly to an estimated cost to the MBTA, because the fare policy is to charge the first \$2 to the Pilot participant and pay the next \$40 of fare. For the majority of trips, this amounts to $F - 2$. Since the travel time and network distance have been calculated for every requested trip and reported in the trip database, estimation of the subsidy for each trip is a straightforward calculation using equation (12) and subtracting \$2 for each trip.

For a sample of trips from GLSS (North) on January 23, 2017, the average Uber fare was \$18.02, which translated to an average cost to MBTA of \$16.02 if all trips were shifted to TNCs. This is well below the average cost of the ADA paratransit operation

2.5.3 Procedure to Allocate Trips to Paratransit or TNC

Equipped with a method to estimate the marginal cost of each trip on the ADA paratransit service and the cost of the subsidy to serve it with a TNC, the trips with the greatest benefit of shifting to TNCs can be identified. The procedure for optimally allocating trips is as follows.

1. Group all of the requested ADA paratransit trips in a region into routes using the algorithm described in Section 2.5.1.

2. Identify the trip with the greatest estimated cost saving associated with a switch to service with a TNC using the cost calculations presented in Section 2.5.2.

3. Eliminate the trip from the pool of requested ADA paratransit trips and repeat steps 1 and 2. Each time updating the total cost estimate for the ADA paratransit operations and adding the cumulative cost of all of the trips shifted to TNCs. This process can be repeated until there are no trips remaining on the ADA paratransit service.

In practice, the total cost to the agency is minimized when it is no longer possible to save money by transferring trips from ADA paratransit to the TNC. Although it may appear at the first iteration that there are many ADA trips with very low marginal cost, this incremental approach shows how this cost increases other trips are removed. As Type 2 trips are removed from a route, formerly Type 3 trips become new Type 2 trips. Eventually, when one trip is left in the route, it becomes a costly Type 1 trip. This means that the marginal cost of each trip depends on all of the other demand that is served. Trips that appear to be very cost efficient with one set of demand may become very costly as the trips around are shifted to TNCs.

The final challenge is that it may not be possible to shift all trips to TNCs either because the vehicles are not accessible to some customers or because some customers are reluctant to use an alternative service provider. In this case, the same procedure is implemented with the only difference being that only feasible trips are actually eliminated from the pool of requested ADA paratransit trips and shifted to TNCs. The process of shifting trips must then stop when no feasible trips remain.

3 Results

3.1 Operations

The models for The RIDE's paratransit operations were developed as described in Section 2.3. Using the data available from the trip records, aggregate operations models were estimated for each of the operators during 2017. An additional analysis of on-time performance revealed a relationship between the level of utilization of the vehicles in the system and the percent of trips that were served on-time.

3.1.1 Aggregate Model of The RIDE Paratransit Operations

Modeled Vehicle Miles Traveled

The VMT model expressed in equation (2) was estimated by calculating the right-hand side expression for each time period in each day using the average parameter values listed in the Appendix. This right-hand side expression is denoted x_{VMT}, so the regression is used to estimate r_{VMT} by regression of the following form

$$VMT = r_{VMT} x_{VMT} \tag{13}$$

where the VMT for each data point is estimated from the specific vehicle routes operated.

The aggregated results across all months of 2017 are plotted for each operator in Figure 3.1, Figure 3.2, and Figure 3.3. These results show an overall good fit of the model with VMT being well approximated by a linear relationship with r_{VMT}. Some scatter, especially for VTS (West) in Figure 3.2, raises some questions about whether or not there may be systematic variation in r_{VMT} by time of day or day or day of the week. A systematic comparison of model fit by month of year (Table 3.1) shows little variation in the value of r_{VMT} by month.

Table 3.1 Modeled Value of r_{VMT} by Month, 2017

Month	GLSS (North)		VTS (West)		TRAC (North/South)	
	r_{VMT}	R^2	r_{VMT}	R^2	r_{VMT}	R^2
January	1.11	0.99	0.92	0.97		
February	1.10	0.99	0.92	0.98	0.94	0.97
March	1.11	0.99	0.91	0.97	0.93	0.97
April	1.11	0.99	0.92	0.97	0.99	0.94
May	1.13	0.98	0.92	0.97	0.97	0.98
June			0.92	0.97	0.98	0.98
July			0.92	0.97	0.98	0.98
August			0.93	0.96	0.98	0.98
September			0.93	0.96	0.98	0.98
October			0.94	0.97	0.97	0.98
November			0.94	0.97	0.97	0.98
December			0.91	0.97	0.96	0.98
All Months	*1.11*	*0.99*	*0.92*	*0.97*	*0.97*	*0.98*

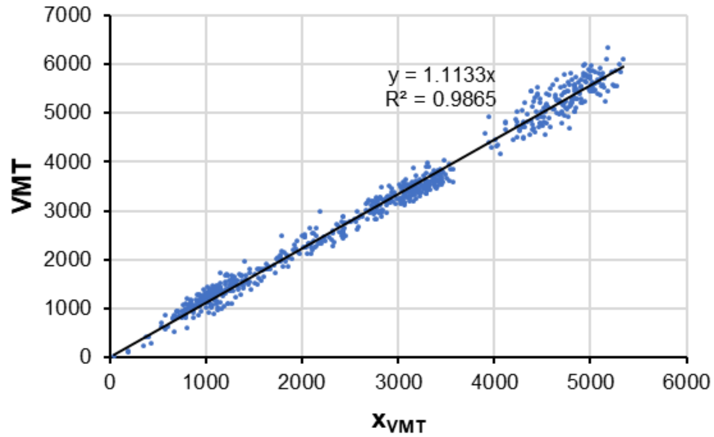

Figure 3.1 VMT Model for GLSS (North), 2017

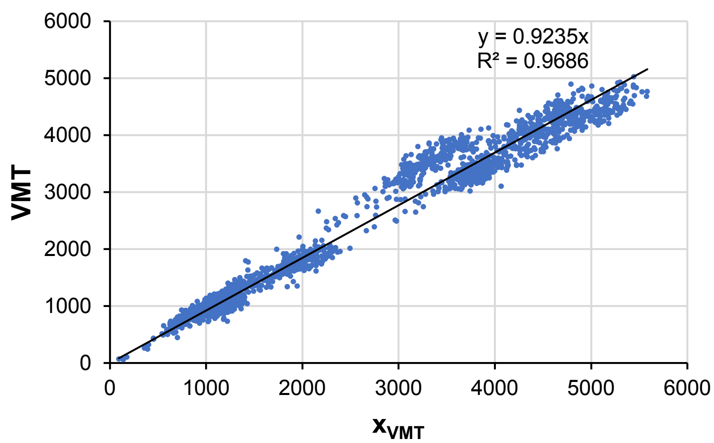

Figure 3.2 VMT Model for VTS (West), 2017

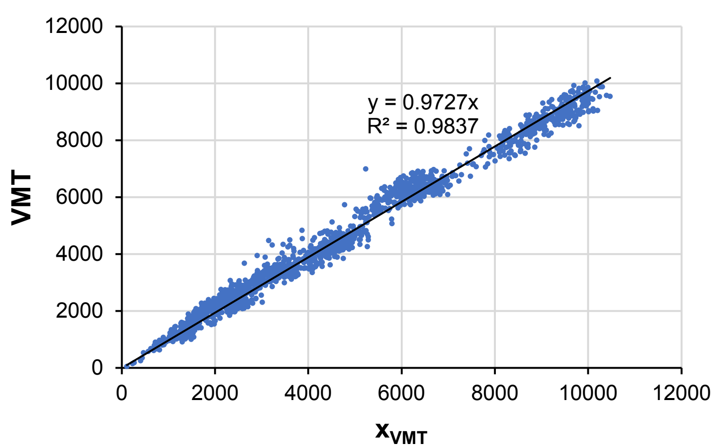

Figure 3.3 VMT Model for TRAC (North/South), 2017

Table 3.2 shows that there is not very much difference between weekdays and weekends. When comparing the value of r_{VMT} by time of day, there is more variability, with relatively greater values in the afternoon and evening hours. Table 3.3 shows that this pattern is consistent across all operators and most pronounced for VTS.

Table 3.2 Modeled Value of r_{VMT} by Day of the Week, 2017

Day	GLSS (North)		VTS (West)		TRAC (North/South)	
	r_{VMT}	R^2	r_{VMT}	R^2	r_{VMT}	R^2
Weekdays	1.11	0.98	0.93	0.95	0.97	0.98
Weekends	1.10	0.96	0.90	0.92	0.97	0.96
All Days	1.11	0.99	0.92	0.97	0.97	0.98

Table 3.3 Modeled Value of r_{VMT} by Time of Day, 2017

Month	GLSS (North)		VTS (West)		TRAC (North/South)	
	r_{VMT}	R^2	r_{VMT}	R^2	r_{VMT}	R^2
6 AM – 9 AM	1.08	0.98	0.88	0.99	0.96	0.99
9 AM – 12 PM	1.14	0.98	0.94	0.98	0.97	0.99
12 PM – 3 PM	1.09	0.98	0.87	0.99	0.93	0.99
3 PM – 6 PM	1.11	0.97	1.06	0.98	1.05	0.98
6 PM – 9 PM	1.24	0.88	0.98	0.87	1.12	0.95
All Times	1.11	0.99	0.92	0.97	0.97	0.98

Although the model fits the data generally quite well, it is possible to improve the model's fit and predictive power by estimating distinct values of r_{VMT} for different time periods of the day. Fitting five different models is excessive for most applications, but separating the afternoon/evening time periods after 3 PM from the rest of the day results in a better looking fit for the model with less estimation error.

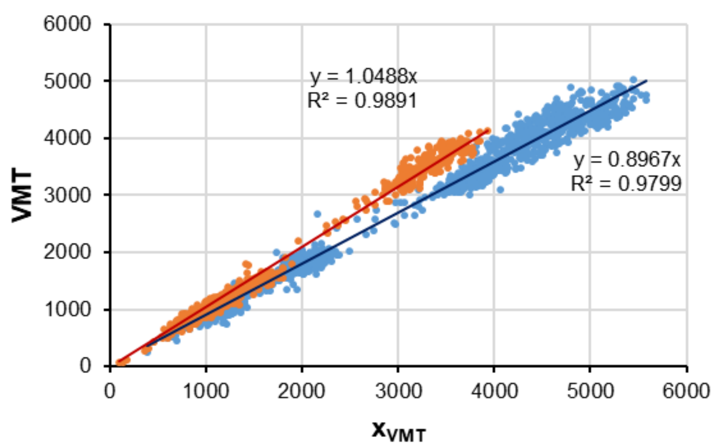

Figure 3.4 VMT Model Separated by Time Period for VTS (West), 2017

Modeled Vehicle Hours Traveled

The VHT model expressed in equation (4) was estimated by calculating the right-hand side and left-hand side expressions for each time period in each day using the average parameter values listed in the Appendix. This right-hand side expression is denoted x_{VHT} and the lefthand side expression is denoted y_{VHT}. The regression is used to estimate r_{VHT} by regression of the following form

$$y_{VHT} = r_{VHT} x_{VHT} \tag{14}$$

Just as for the VMT model, the aggregated results for VHT across all months of 2017 are plotted for each operator in Figure 3.5, Figure 3.6, and Figure 3.7. These results show an overall good fit of the model with VHT being approximated by a linear relationship with r_{VHT} with fits that are similar to the VMT model. Compared to r_{VMT}, which is depends on the spatial distribution of demand and the circuity of the road network, there is more variation in the estimated values of r_{VHT}. Some of this variability is a reflection of extra slack time in the route schedules when demand is lower. Although the value varies little by month (see Table 3.4), the value of r_{VHT} is noticeably greater on weekends (Table 3.5) and evening hours (Table 3.6) when demand rates are lower and vehicles are not as busy.

Table 3.4 Modeled Value of r_{VHT} by Month, 2017

Month	GLSS (North)		VTS (West)		TRAC (North/South)	
	r_{VHT}	R^2	r_{VHT}	R^2	r_{VHT}	R^2
January	1.96	0.95	1.60	0.96		
February	1.93	0.94	1.60	0.93	1.80	0.91
March	1.93	0.96	1.61	0.96	1.78	0.92
April	1.87	0.96	1.57	0.95	1.80	0.92
May	1.85	0.96	1.55	0.95	1.70	0.96
June			1.57	0.95	1.64	0.95
July			1.59	0.94	1.74	0.95
August			1.61	0.94	1.73	0.95
September			1.59	0.95	1.66	0.94
October			1.60	0.95	1.63	0.93
November			1.62	0.95	1.70	0.94
December			1.65	0.94	1.72	0.93
All Months	1.91	0.95	1.59	0.95	1.69	0.95

Table 3.5 Modeled Value of r_{VHT} by Day of the Week, 2017

Day	GLSS (North)		VTS (West)		TRAC (North/South)	
	r_{VHT}	R^2	r_{VHT}	R^2	r_{VHT}	R^2
Weekdays	1.90	0.94	1.59	0.93	1.68	0.95
Weekends	2.03	0.81	1.72	0.47	1.85	0.83
All Days	1.91	0.95	1.59	0.95	1.69	0.95

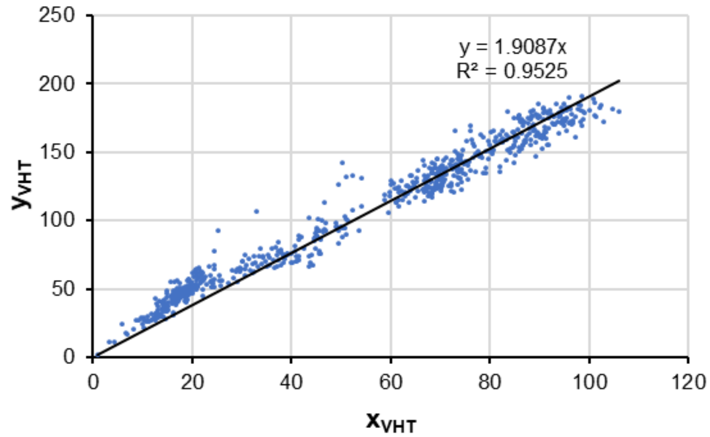

Figure 3.5 VHT Model for GLSS (North), 2017

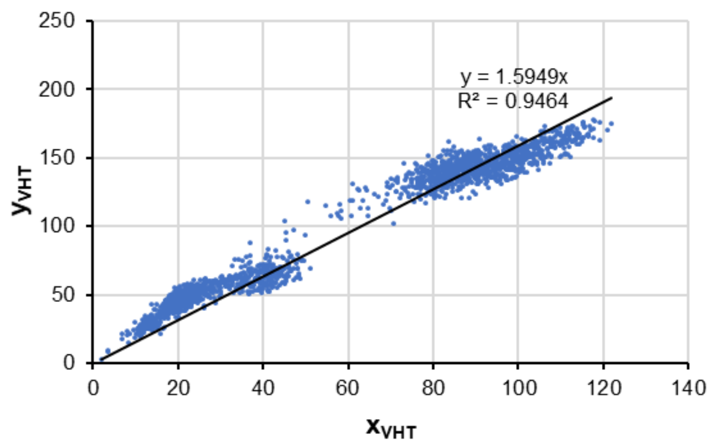

Figure 3.6 VHT Model for VTS (West), 2017

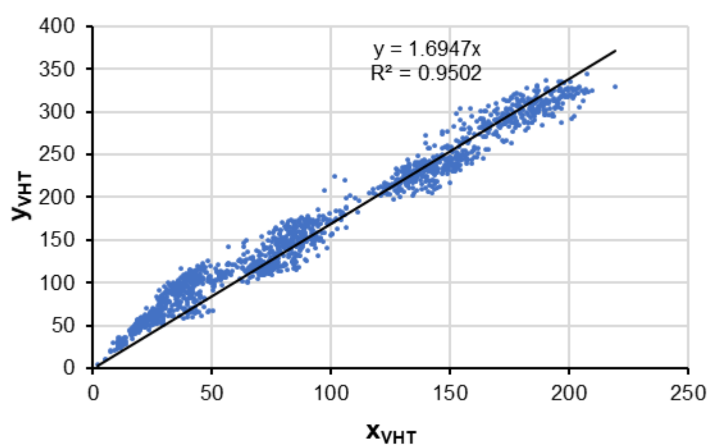

Figure 3.7 VHT Model for TRAC (North/South), 2017

Table 3.6 Modeled Value of r_{VHT} by Time of Day, 2017

Month	GLSS (North)		VTS (West)		TRAC (North/South)	
	r_{VHT}	R^2	r_{VHT}	R^2	r_{VHT}	R^2
6 AM – 9 AM	1.95	0.92	1.71	0.94	1.67	0.94
9 AM – 12 PM	1.92	0.93	1.50	0.96	1.65	0.97
12 PM – 3 PM	1.85	0.94	1.52	0.96	1.72	0.96
3 PM – 6 PM	1.89	0.94	1.69	0.95	1.66	0.96
6 PM – 9 PM	2.60	0.85	2.11	0.87	2.60	0.92
All Times	1.91	0.95	1.59	0.95	1.69	0.95

Note that in all cases the value of $r_{VHT} > r_{VMT}$, and the difference is greater during time periods when demand is lower. This is especially true at the end of the day when there may be many vehicles operating from the busier peak in afternoon demand. In the later hours many of these vehicles are not fully utilized, and the wasted time is reflected in a greater VHT factor. Also of note is that the model fit as represented by the R^2 values is lower for r_{VHT} than for r_{VMT}, especially at times with lower demand. This is an indication of greater variability in the data, which leads to greater uncertainty in model estimates.

The fleet size is directly related to the VHT, so the same model outcomes are used to estimate the number of vehicles using equation (5). Weekday afternoons between 12 PM and 3 PM consistently have the greatest levels of demand, so the required fleet size is determined by the VHT during that time period.

3.1.2 Aggregate Model of The RIDE Paratransit Costs

Without detailed cost information from the MBTA's operators, it is necessary to make cost estimates based on data from other operators. To estimate costs, we know that the total annual operating cost for the MBTA's demand responsive services in 2017 was reported to be $103,493,764 (NTD, 2017). For the purpose of illustration, we use cost factors estimated from the PVTA in Springfield, Massachusetts (Turmo et al., 2018):

- Cost per Vehicle Mile of Operation, $\alpha_1 = 0.518$ $/veh-mile;
- Cost per Vehicle Hour of Operation, $\alpha_2 = 19.89$ $/veh-hour;
- Cost per Vehicle, $\alpha_3 = 150.81$ $/veh-day or $55,046 $/veh-year (fleet size cost)

During 2017, there were 1,465,092 trips records per month in the dataset.[2] Applying the operations model by operator, time of day, and day of week (as presented above), the modeled operations requirements for 2017 were:

- 15,120,876 vehicle-miles traveled (VMT) in 2017

---

[2] The National Transit Database reports 1,985,115 unlinked trips, but the operations model for The RIDE accounts for linked trips, not including the TNC Pilot.

- 1,832,642 vehicle-hours traveled (VHT) in 2017
- 530 vehicles, minimum required fleet size

Based on these operations values for 2017 and the cost coefficients from PVTA, listed above, the total operating cost associated with these physical components, the remaining costs are the fixed operating cost that is independent of vehicle operations.

- Fixed Annual Cost, $\alpha_0 = 30{,}017{,}847$ \$/year

If anything, unit cost parameters for the MBTA are likely to be higher because Greater Boston is a more costly place to employ labor and maintain a fleet of vehicles than Springfield. Therefore, any marginal cost estimates based on these values are likely to be lower bounds, and the cost savings of shifting trips to TNCs could be higher.

3.1.3 On-Time Performance

The analysis of on-time performance data revealed a lot of variation in the percent of trips that are late from day to day. The first hypothesis is that network traffic speeds affect on-time performance, because traffic congestion may cause paratransit vehicles to fall behind schedule. Figure 3.8 shows that there is wide variation in the daily on-time performance from around 2% late trips to more than 20% with no discernable relationship to network speeds.

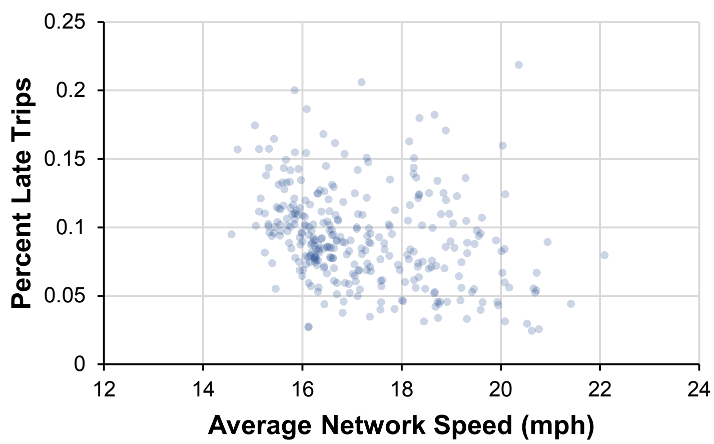

Figure 3.8 On-Time Performance and Network Speeds for TRAC (North/South), 2017

Figure 3.9 shows the relationship between the percent of revenue hours in which vehicles are occupied with customers and the on-time performance. The range of explanatory values (between 0.4 and 0.6) is not very wide, which make it difficult to fit a model to the data. Nevertheless, there does appear to be an increasing relationship between the variables, albeit a noisy one. The slight curvature of the data appears to suggest that 0.5 occupancy is a sweet spot, below which performance hovers around 5% late trips, and above which delays quickly increase. There is too much variation in the sample data make a statistically significant model to estimate delays. Observed patterns were similar for the other providers.

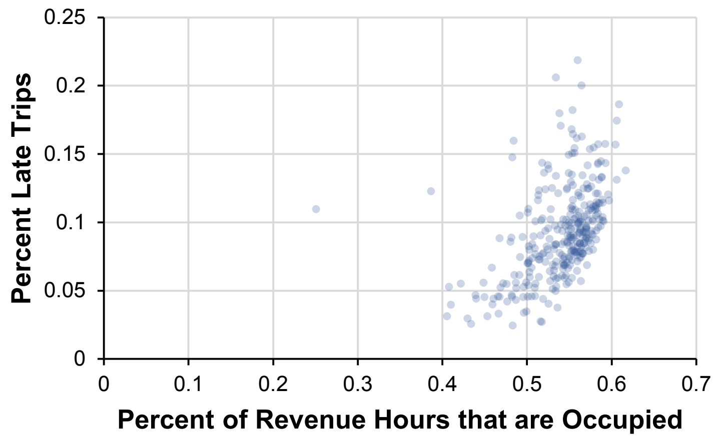

Figure 3.9 On-Time Performance and Ratio of Occupied Revenue Hours for TRAC (North/West), 2017

3.2 Effect of the Pilot on Demand

Overall, Pilot participants made up a relatively small percentage of the total eligible ADA customers registered with the MBTA. As of March 2018, there were 1,975 Pilot participants out of 40,721 total registered customers, which amounts to a participation rate of about 4.85%. Over time some customers stop using the paratransit system, and by the very fact that they have elected to participate in the program, Pilot participants are likely to be more active users. Figure 3.10 shows that even with stead growth since the start of the Pilot in October 2016, TNCs made up 7.44% of total travel supported by the MBTA in March 2018. This is well above the percentage of customers participating in the Pilot, so these participants are traveling more than the average registered ADA customer.

3.2.1 Pilot Participants

First, the rate of Pilot trips is compared over time from the date each customer joined the program. The purpose of this comparison is to determine how long it takes for Pilot participants to settle on a steady number of trips per month using TNCs. This is important, because the comparison of travel demand before and after participation in the study should avoid including transition months during which customers' behaviors are still evolving.

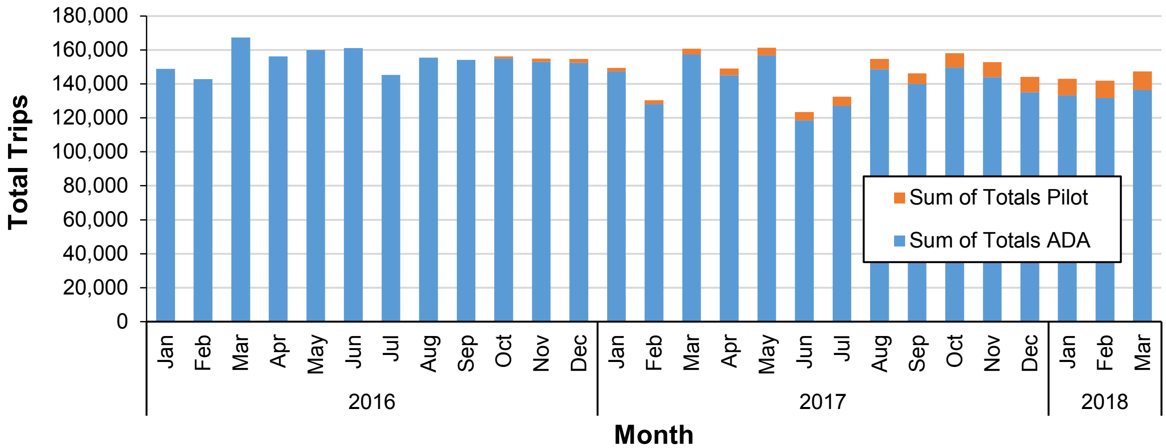

Figure 3.10 Monthly Trips on The RIDE and the TNC Pilot

Figure 3.11 shows the monthly trip counts as a percentage of the maximum TNC trip allocation for customers grouped by their starting month. Although there is a lot of variability in the data, it generally appears to take about two months for the trip count to rise to a more-or-less stable average. Therefore, in order to compare travel patterns among Pilot participants before and during the Pilot program, we consider only participants that have joined at least two months prior to the months of compared.

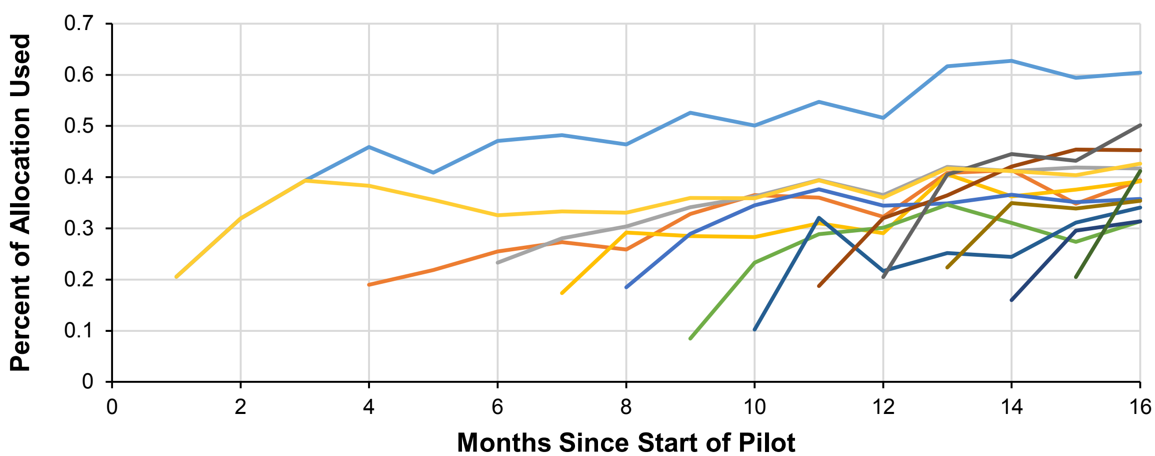

Figure 3.11 Monthly Trips on the TNC Pilot as a Percentage of Allocation for Customer Grouped by Starting Month

Although some of the early participants tend to use more of their trip allocation (as shown by the light blue line across the top), most Pilot participants use well under their maximum allocation. In fact, the average customer uses about 40% of their allocated TNC trips each month. The first cohort of Pilot program participants utilize approximately 60% of their allocation, confirming that the first Pilot participants are more enthusiastic users than average.

Comparing the travel patterns among from January through March 2016 with January through March 2018, it is possible to quantify the effect of the Pilot program on aggregate travel behavior. In order to make a meaningful comparison, trips are only included for customers who are active users in both time periods. Any customers without any trip records in 2016 are assumed to have joined later and any customers without any trip records in 2018 are assumed to have discontinued use of ADA paratransit services altogether. Furthermore, among Pilot participants, only customers that joined the Pilot more than 2 months before January 2018 are included.

A comparison of the total number of trips completed by each group of customers in January through March of 2016 and 2018 is shown in Table 3.7. Among the control group of non-Pilot customers, demand dropped 11.3% over the 2-year time period. It is not clear what the reason for this change is, although it is possible that some customers are already utilizing TNCs for some of their travel outside of the Pilot program. The Pilot customers demonstrated continued utilization of the ADA service, but with fewer trip per month. Meanwhile, the Pilot participants took more monthly trips by TNCs than they did on ADA paratransit and their total monthly number of trips across both services increased.

Table 3.7 Comparison of Monthly (Jan–Mar) Trips by ADA Paratransit and TNCs

Group	ADA Trips 2016	ADA Trips 2018	TNC Trips 2018	Total Trips 2018
Non-Pilot Customers (Control)	137,290	121,734		121,734
% of Initial		88.7%		88.7%
Pilot Participants	12,723	8,229	9,080	17,309
% of Initial		64.7%	71.4%	136.0%
% Adjusted for Control		72.9%	80.5%	153.4%

In order to characterize the extent to which the availability of TNCs through the Pilot resulted in substitution of trips away from ADA paratransit or induction of new trips that were not served before, it is necessary to account for the 11.3% decrease in the control group. The number of trips made by each mode in 2018 is expressed as a percentage of the initial 2016 ADA trip demand adjusted down by 11.3% to an effective 11,281 trips that would have been made in the absence of the Pilot. Compared to this number, the effects on demand were the following:

- 72.9% of ADA paratransit trips remained on The RIDE's ADA paratransit service.
- 27.1% of ADA paratransit trips were substituted by TNCs.
- 53.4% additional demand was induced by the availability of the TNCs.

In general, more trips represent increased mobility, so the MBTA is supporting more travel for customers with disabilities. An important question is what the effect on demand would look like if the Pilot program were expanded to include all customers that are registered with The RIDE.

3.2.2 Modeled Effect on All Customers

Cluster Analysis

Using the K-means clustering approach described in Section 2.4.2, the travel demand of each individual included in the tabulation of Pilot participants in Table 3.7 is used to identify clusters of similar travel behaviors. Selection of the number of clusters to use is based on comparing the sum of squares differences within the clusters. Increasing the number of clusters allows this value to be reduced. The goal is to identify the smallest number of clusters for which the change in sum of squares is not too big; i.e., pick a point nearest to the "elbow" of the curve. This relationship for the Pilot program clusters is shown in Figure 3.12. Although there is not a distinct "elbow" point, four clusters appear to be a good compromise between the number of clusters and the error within clusters.

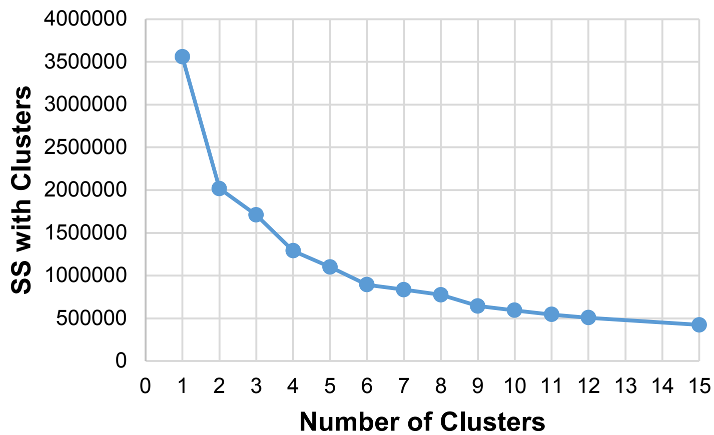

Figure 3.12 Sum of Squares within Clusters for K-Means Cluster Analysis

A three dimensional plot in Figure 3.13 shows each pilot participants travel behavior as a point representing monthly trips by ADA paratransit before the pilot, by ADA paratransit after the pilot, and by TNC after the pilot. The two-dimensional plots are presented to show the three orthogonal views. The assigned clusters are indicated by color. The broad scatter of points indicates that there are a wide range of travel behaviors represented among the pilot participants. The clustering shows how these participants are grouped. Most participants are in Cluster 2, which are the users that made the fewest ADA paratransit trips and have the smallest monthly allocation.

The characteristics of the four clusters are quantified in Table 3.8 in terms of the average number of trips per month. Some characteristics stand out. For example, Cluster 3 represents users that drastically reduce travel on ADA paratransit and more compensate through increased travel on TNCs. In contrast, Cluster 4 represents customers that appear reluctant to make the switch by maintaining similar numbers of trips on ADA paratransit, taking relatively few trips by TNCs. Induced demand is shown by the increase in total trips in 2018. This is most notable for Clusters 2 and 3, which are the customers that enthusiastically adopt TNCs.

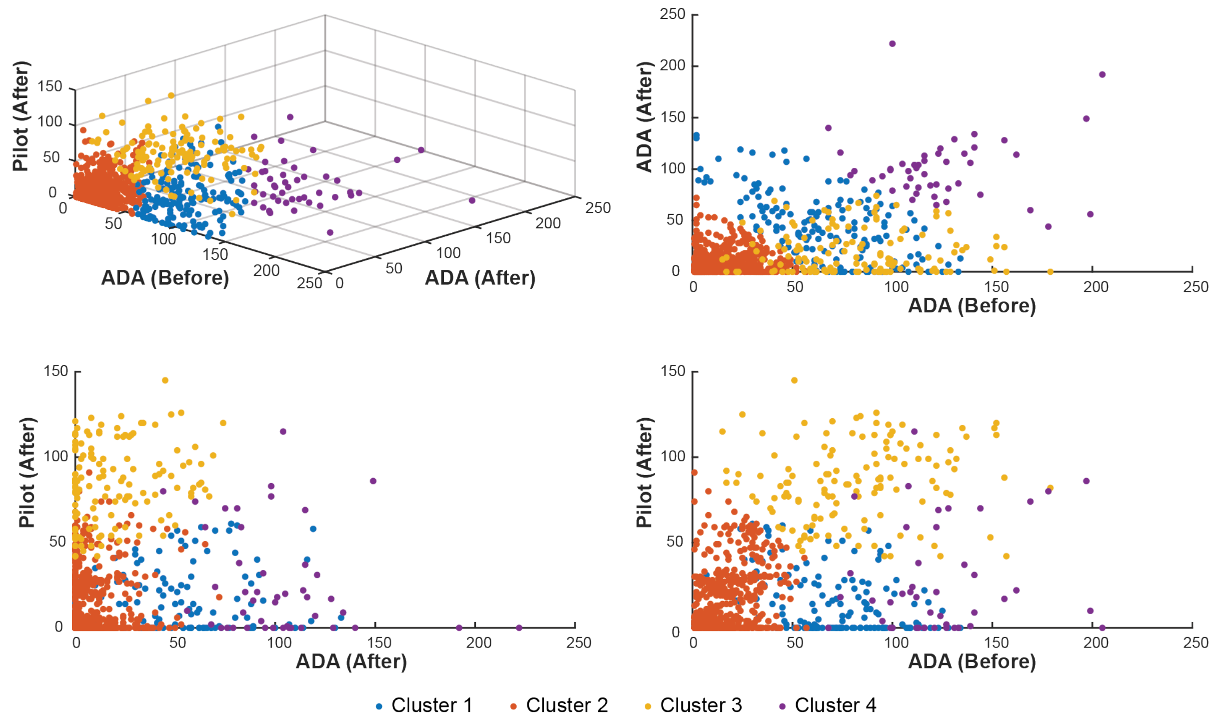

Figure 3.13 Plot of Clusters Based on Monthly Trips by ADA and TNC Pilot from 2016 (Before) and 2018 (After).

Table 3.8 Average Monthly Trips of Pilot Participants by Cluster

Group	Number of Customers	ADA Trips 2016	Adjusted 2016 Trips	ADA Trips 2018	TNC Trips 2018	Total Trips 2018
Cluster 1	168	67	59	43 71.9%	13 21.8%	56 93.7%
Cluster 2	643	15	13	8 57.7%	13 97.0%	21 154.7%
Cluster 3	137	81	72	19 26.7%	85 117.7%	104 144.5%
Cluster 4	47	130	115	106 92.5%	26 22.5%	132 115.0%

Logistic Regression

A logistic regression is conducted to identify the customer characteristics that have the strongest ability to determine which cluster a customer is assigned to. The structure of the model is as described in Section 2.4.2. The parameters of the logistic regression are presented in Table 3.9. Cluster 1 is treated as the base case, so all parameters express the effect of a characteristic in assigning a customer to one of the other clusters relative to Cluster 1. Positive parameter values indicate increased likelihood and negative parameter values indicate decreased likelihood.

Table 3.9 Logistic Regression Model for Clusters

Parameter (Cluster)	β Value	t-statistic	p-value
Alternative Specific Constant (2)	5.09	9.88	0.00
Alternative Specific Constant (4)	-6.94	-8.32	0.00
Monthly ADA Trips in 2016 (2)	-0.124	-12.98	0.00
Monthly ADA Trips in 2016 (3)	0.0139	4.32	0.00
Monthly ADA Trips in 2016 (4)	0.0574	7.67	0.00
Customer Age (2)	0.00924	1.24	0.22
Customer Age (3)	-0.0215	-4.47	0.00
Customer uses Wheelchair or Lift (3)	-2.00	-3.17	0.00
Customer uses Wheelchair or Lift (4)	1.17	2.21	0.03
Model Log-Likelihood	-470		
Null Log-Likelihood	-1379		
ρ^2	0.659		

The omitted parameters in Table 3.9 are associated with values that lack statistical significance. One exception is left in the model: customer age for Cluster 2, because value improves the model fit despite its lack of statistical significance (as indicated by the p-value exceeding 0.05).

A few parameter values are worth drawing attention. Note from Table 3.8 that Cluster 3 represents a group of eager TNC adopters who largely leave behind the ADA service and make an increased number of trips by TNC. The negative parameter for wheelchair or lift use indicates that these customers are much less likely than average to use such a device. On the flip side, Cluster 4 is associated with customers that make many ADA paratransit trips (which would grant them high allocations for TNC trips in the Pilot) but continue to use ADA paratransit for most of their travel. The positive parameter on wheelchair or lift use for Cluster 4 indicates that these customers are much more likely to use such a device. Together, this suggests that ambulatory customers are more likely to adopt TNCs, but there appear to be barriers for customer using wheelchairs, power chairs, scooters, or other devices that require a WAV with a lift. As of early 2018, it appears that the TNCs are not providing a service that is equally appropriate or appealing to all eligible ADA paratransit customers.

Estimating the Effect of Expanding the Pilot

Since most of The RIDE's eligible customers are not registered with the Pilot program, it is expected that expanding the program to all customers would cause a more widespread change in travel totals. The results of the logistic regression are used to estimate the number of ADA paratransit trips and TNC trips that would be made through an extension of the program to all users. Using data from 2017 travel, without the Pilot program, an estimated 118,613 trips per month would be served by The RIDE.

Figure 3.14 shows that the number of trips served by ADA paratransit (The RIDE) would likely drop by 42% if the Pilot were expanded to all customers. These rides would be more than made up for by TNC travel. In fact, it is expected that induced demand would result in a 33% increase in total trips served per month. Despite this increase, the effect on operating costs is expected to be a net reduction of approximately 26% (see Figure 3.15), because the cost of TNC trips is only $16.02 on average.

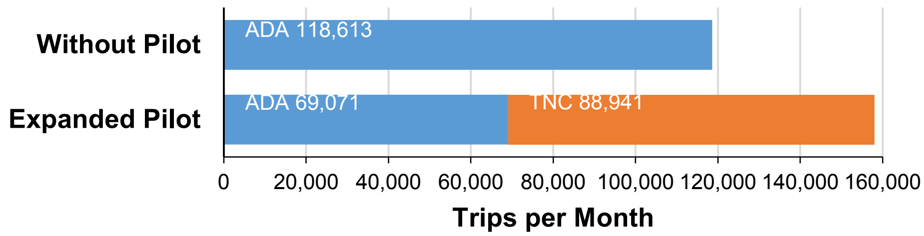

Figure 3.14 Comparison of Monthly Travel Without and With an Expanded TNC Pilot

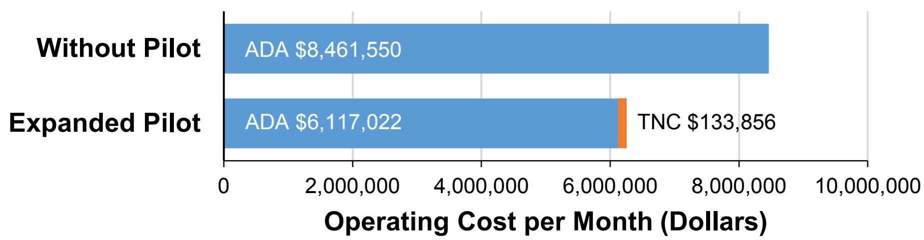

Figure 3.15 Comparison of Monthly Cost Without and With an Expanded TNC Pilot

If the estimated unit costs of VMT, VHT, and fleet size are greater than the values from the PVTA, the effect on costs may be an even greater reduction. High operating costs associated with ADA paratransit are easy to beat with the average cost a trip at $16.02 (based on Uber's reported fare structure). This is especially true for any trip reduction that allows that paratransit operator to reduce the size of the fleet.

3.3 Optimized Allocation of Trips to Paratransit and TNCs

In order to implement that algorithm proposed in Section 2.5, cost estimates are needed for each of the ADA paratransit operating parameters. Without access to detailed cost records of the operators, these coefficients can be difficult to estimate. The MBTA contract is structured to pay a rate per VHT with other requirements on the conditions of service. Although VHT-based costs may be applicable in the short-run, it is more meaningful to estimate the various types of costs that are related to fleet size, VHT, and VMT, because these distinguish the different types of impacts that eliminating trip can have on the cost of operating the system.

3.3.1 Marginal Cost of Paratransit Trips

The proposed algorithm is first implemented on the entire set of requested ADA trips for January 23, 2017, in the North region operated by GLSS. By sequentially creating hypothetical vehicle routes, the marginal cost of each trip is estimated and compared with the fare that would be charged if the trip were served by a TNC (based on Uber's fare structure). Figure 3.16 shows the distribution of the net marginal cost of each trip based on the estimated cost savings from shifting the trip away from The RIDE, offset by the estimated cost of the TNC fare.

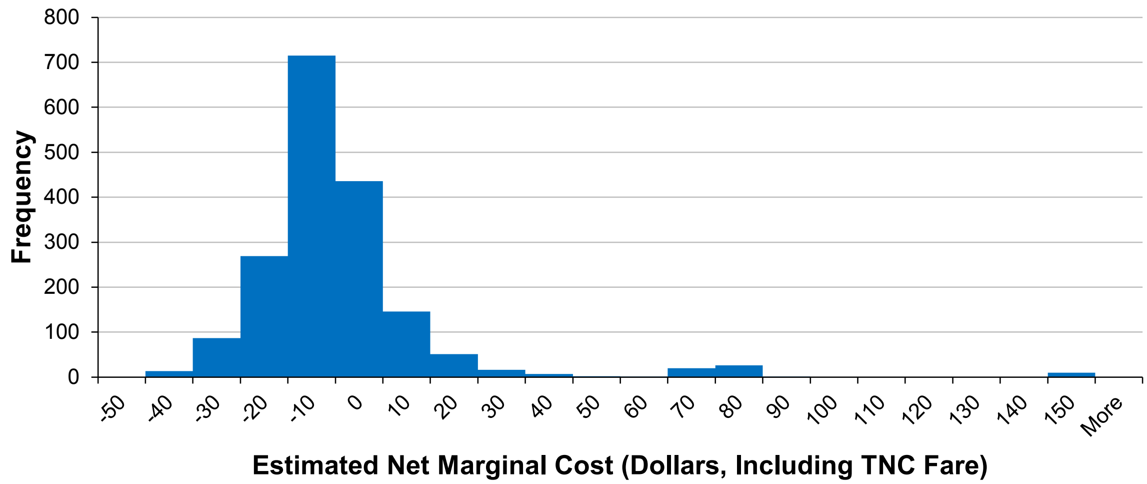

Figure 3.16 Distribution of Estimated Marginal Costs for All Trips, GLSS (North) January 23, 2017

A positive value in Figure 3.16 indicates that the marginal cost of paratransit operations exceeds the expected TNC fare, and shifting the trip would save money. A negative value indicates that the expected TNC fare would exceed the marginal operating cost. The greatest values are for the small number of Type 1 trips (for which an extra vehicle is needed to serve a single trip. Many trips with negative net marginal costs are the Type 3 trips within a route, which can be served at relatively low cost by the ADA fleet, because the vehicles are already out on the road.

By the proposed algorithm, the costliest trip should be shifted to a TNC, and then the routing process must be recalculated to estimate the new marginal costs. Therefore, trips with low (or negative) net marginal cost at the first iteration may become more beneficial contenders for shifting to TNCs as the routes change.

3.3.2 Effect of Incrementally Shifting Trips to TNCs

As the costliest trips are shifted from The RIDE to TNCs, the effect on total agency cost is calculated. The sequence of cost changes is shown in Figure 3.17 (blue curve), where the horizontal access indicates the cumulative number of trips shifted to TNCs and the vertical access is the total agency cost, including subsidies paid for TNC trips. The costs drop most dramatically for the first few trips as inefficient routes serving peak demand are eliminated. For the particular date selected, and the cost parameters used, the agency costs continue to decline until all demand has been shifted to TNCs. In this case, the lowest cost is achieved by shifting all trips from ADA paratransit to TNCs.

It is not always possible (or desirable) to shift all trips to TNCs. For example, the analysis of Pilot participants in Section 3.2 shows that some customers are reluctant to choose TNCs even when the option is available to them. Part of this is may be due to general attitudes or preferences regarding the modes, but the evidence suggests that customers with wheelchairs, power chairs, scooters, or other devices requiring a WAV with a lift are unable or uncomfortable using a TNC. Applying the same procedure for optimally allocating trips to TNCs while leaving all wheelchair and lift customers on ADA paratransit, the red curve in Figure 3.17 shows the sequence of changing costs. Overall the pattern is similar, with a steep initial decline in agency costs associated with eliminating the most inefficient routes during peak demand. Then the cost savings accrue more slowly and the effect of shifting trips levels off before all of the feasible trips have been shifted. The prevailing pattern is still that costs are minimized when as many trips as possible are shifted to TNCs (although this may not necessarily happen in all regions with all demand patterns). In this case, because the are some customers that must always be served by the ADA paratransit van fleet, there are a small number of general trips that can be more efficiently served by the vans in combination with the other trips than shifting to TNCs.

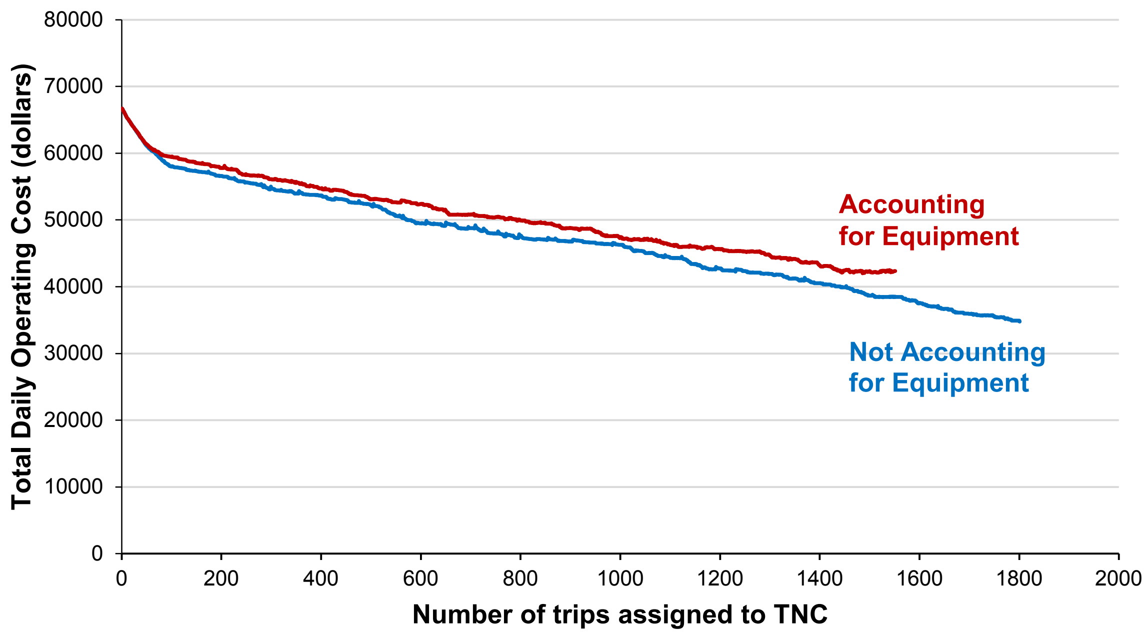

Figure 3.17 Change in Cost by Incrementally Shifting Trips with Greatest Net Marginal Cost to TNC, GLSS (North), January 23, 2017

Based on the cost parameters used for this example, it appears that agency costs could be reduced by approximately 48% if all trips could be shifted to TNCs. If equipment limits the shift, the potential cost savings are lower; 40% reduction for the example day shown. Overall, this is a large reduction in operating cost. The challenge moving forward is understanding how demand patterns may change in response to fundamentally changing the character of ADA paratransit service.

The effect of shifting trips strategically to TNCs (i.e., by incrementally shifting the trips with the greatest net marginal cost) can be compared with alternative patterns. Figure 3.18 shows the same blue curve for the total costs associated with the optimized reallocation of trips to TNCs. The orange curve shows the resulting cost if the trip with the lowest net marginal cost were shifted at each iterations (i.e., the opposite of the optimized strategy). In addition, the effect on total cost of randomly shifting trips was calculated for 10 realizations. The mean and 95% confidence interval based on these realizations is shown in gray. A few points regarding the change in total cost are noteworthy:

1. All cases have a general downward trend in cost (even the series corresponding to shifting the least costly trips). Therefore, total costs are expected to decline with increasing utilization of TNCs, at least with the estimated cost parameters.

2. The trajectory of total operating costs for the "worst" case, in which the least costly trip is always shifted, is comparable to randomly shifting trips for all but the last trips. Another way to view this is that allowing customers to self-select to TNCs (which would effectively be random selection of trips) is as bad for total costs as selecting the least costly trip to shift away from The RIDE at each iteration. This "worst" case only performs demonstrably worse for the very first trips and at the end, when there are few trips remaining on The RIDE.

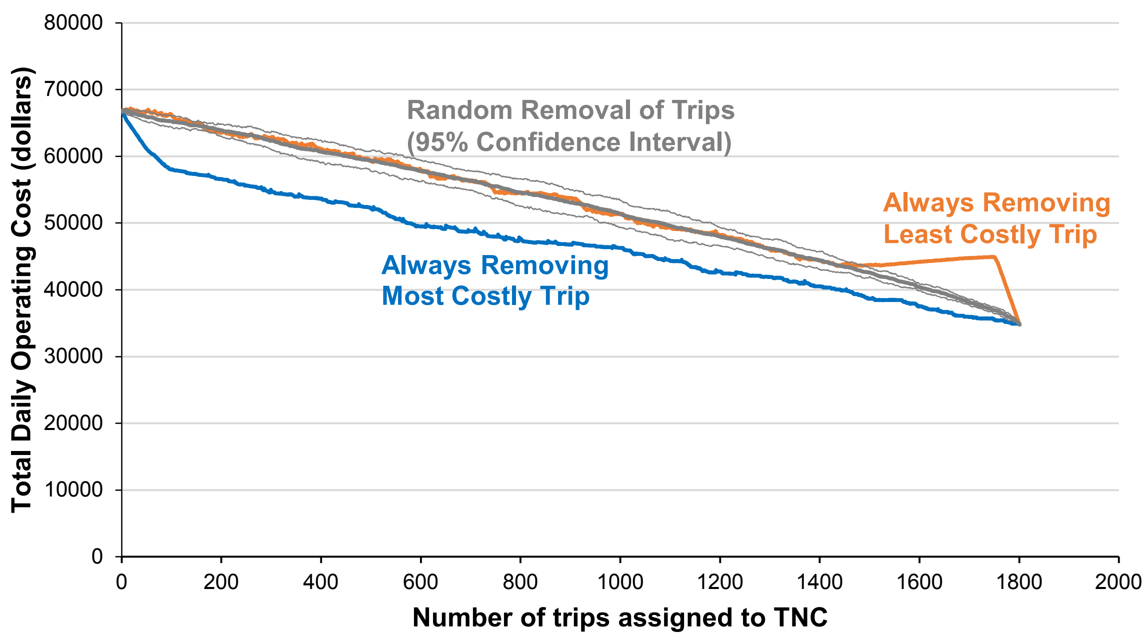

Figure 3.18 Comparison of Costs by Incrementally Shifting Trips to TNC in Random Order, GLSS (North), January 23, 2017

The plots of total cost in Figure 3.17 and Figure 3.18 show that impact on total costs of selectively shifting trips to TNCs. It is also useful to look at the characteristics of the trips that are shifted. For example, the distribution of shifted trips by requested pick-up time is shown in Figure 3.19. Each curve in the figure shows the distribution of trip start times after a number of trips have been shifted to TNCs in the optimized order. The curve for all trips represents that existing case that all demand is served by The RIDE, and this curve exhibits two distinct peaks: a late morning peak at 11 AM and an afternoon peak at 3 PM.

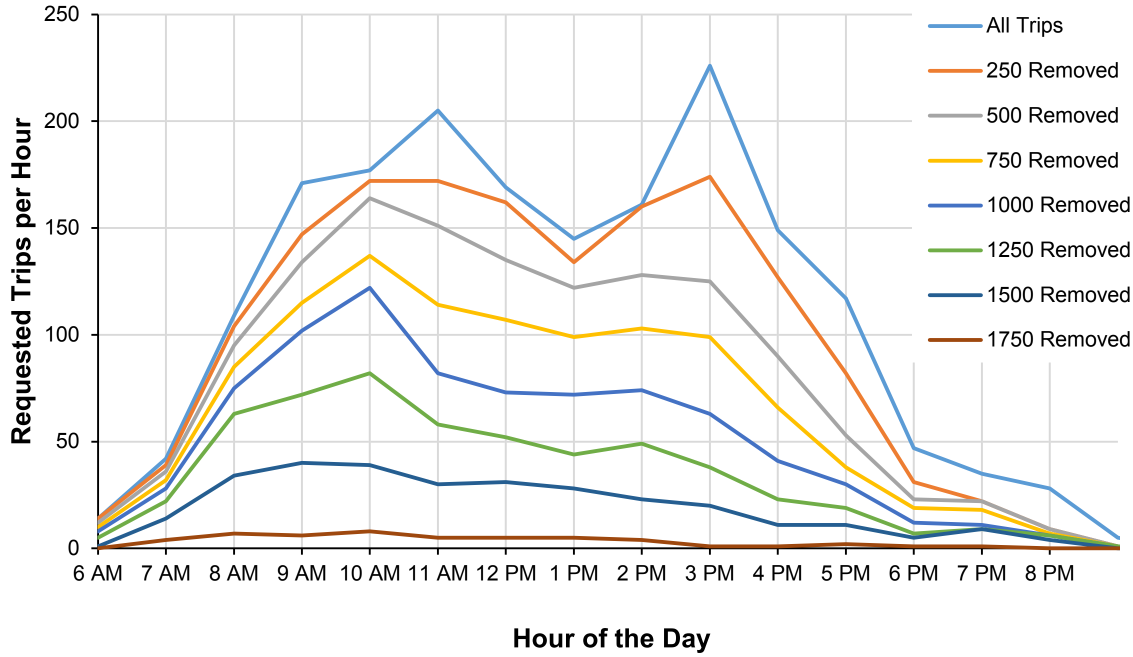

Figure 3.19 Distribution of Remaining ADA Paratransit Trips by Time of Day, GLSS (North), January 23, 2017

The first 250 trips to be shifted from The RIDE to TNCs are mostly Type 1 trips from the peak of the peak and Type 2 trips from the end of the day. The effect of removing these trips is to flatten the peaks and drop demand faster at the end of the day (as shown by the curve labeled "250 Removed"). As trips are sequentially removed, the resulting demand pattern for The RIDE is a more uniform distribution, which allows vehicles to be used more consistently throughout the day.

A second analysis of the distribution of the shifted trips is to look at the geographic locations of shifted trips within the region. Figure 3.20 shows a series of maps of the North and Shared regions that were served by GLSS in January 2017. Each map shows the locations of requested trip pickups, and the colored points indicate the trips that are selected to shift to TNCs. The color indicates the time period of the day when the trip was requested, so the busy middle periods of the day are represented by orange, yellow, and green, while early morning and evening trips are represented by red and blue, respectively. Black points are trips that remain on The Ride (ADA paratransit) and gray points are the trips that have already been removed in earlier iterations.

There is not an obvious geographic pattern to the trips being removed, because trips in the suburbs and in the city center are selected for removal at each stage. In general, there seems to be a trend to eliminate suburban trips sooner than city center trips, because the remaining points appear to be increasingly clustered near Boston's city center at higher iterations. This is expected, because the requested trips in the suburbs tend to be longer in distance and more spread apart, which makes them costlier to serve with the ADA vans.

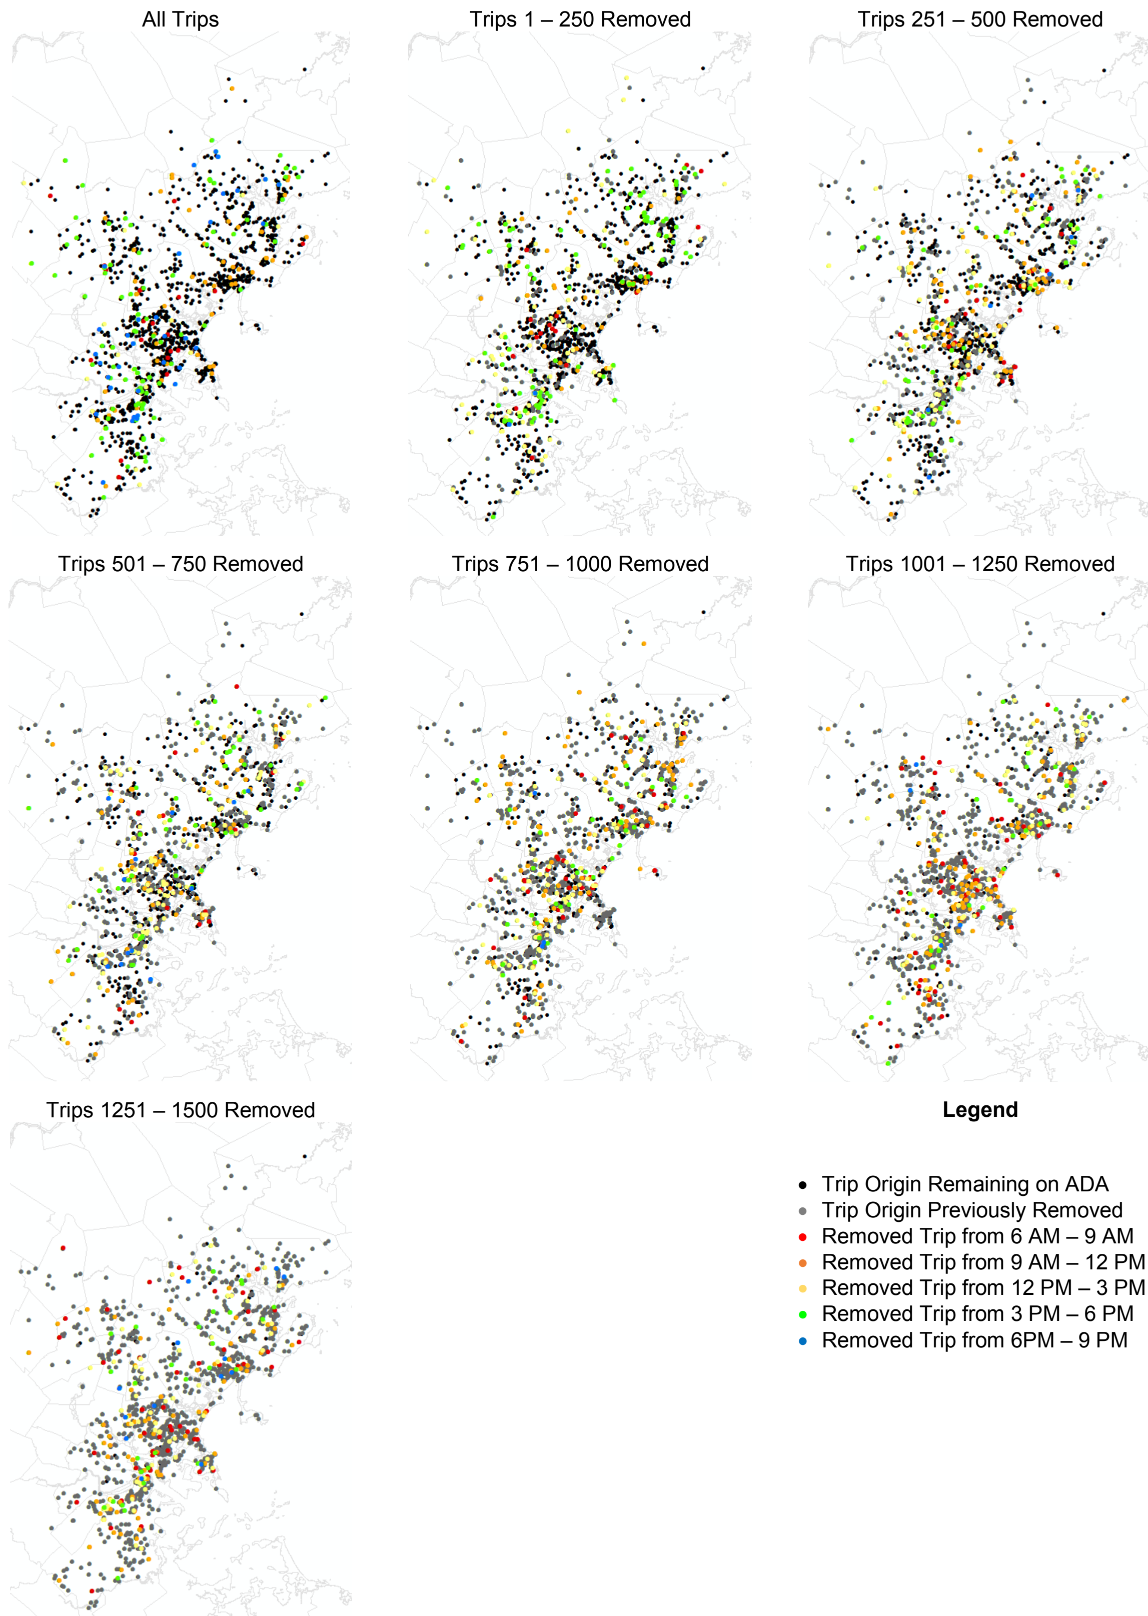

Figure 3.20 Distribution of Removed ADA Paratransit Trips by Location, GLSS (North), January 23, 2017

4 Conclusions

The MBTA's Pilot program to allow eligible ADA customer to make subsidized trips using Uber, Lyft, or Curb is intended to provide customers with a more flexible range of mobility choices while reducing costs for the agency. The Pilot program is structured to enable participants to make limited number of subsidized TNCs rides per month based on their previous ADA paratransit travel history.

Although participation in the Pilot is still relatively small compared to the total number of registered ADA paratransit customers, the data show that the program is popular among most users. In addition to substituting TNCs for many existing ADA trips, the improved flexibility of the TNCs induces customers to travel more. This is a double-edges sword. On the one hand, more trips being served to customers with qualifying disabilities represents an improvement in mobility for a population that is often disadvantages in this regard. On the other hand, increasing travel represents increased costs for the MBTA, which is seeking to the contain growing costs associated with ADA paratransit operations.

The modeling developed in this study culminates in two important lines of analysis. On the demand side, the travel behavior of Pilot participants was analyzed in order to identify the relationships between customer characteristics and response to the TNC Pilot. Applying these models to the general population of ADA customers, it appears that opening up the Pilot to all customers could lead to a drop in conventional ADA trip requests by 42%. This reduction in conventional ADA trips is more than made up for in new TNC trips, and it is expected that under the existing trip allocation structure approximately 33% additional trips would be completed in total. This would result in more trips being served by TNCs than the conventional van service. Despite the increase in travel, the low cost of TNCs relative to ADA paratransit would result in an estimated 26% reduction in net expenditures by the MBTA.

An important insight from the demand analysis is that not all customers are equally enthusiastic about engaging TNC services. Although younger, ambulatory customer are generally enthusiastic about the TNCs and are induced to make roughly 50% more trips, customers with wheelchairs, power chairs, scooters, or otherwise in need of a lift are more likely to continue using the conventional ADA paratransit service. Until the TNCs can provide enough WAVs and trained drivers to offer a comparable level of accessible service, this difference is likely to remain a challenge.

A second line of analysis was on the operations side, developing an algorithm to optimally allocate trips between conventional ADA paratransit service and TNCs. Although The RIDE is not currently structured in a way to assign riders to TNCs, this could be a potential future operating strategy. An algorithm is developed to estimate the marginal cost of each paratransit trip in the context of the vehicle routings so trips can be incrementally reassigned to TNCs when the costs make it advantageous to do so. For the example day shown, the result was reallocation of all trips to TNCs with an expected cost savings of approximately 48%. If some trips cannot be assigned to TNCs, because customers are unwilling or unable to use Uber, Lyft, or Curb, the potential cost savings are lower. When limited to ambulatory customers, the example case resulted in cost savings of approximately 40%.

Overall, this study shows that there are great opportunities for the MBTA to continue coordinating with TNCs to provide service. The key challenge moving forward is determining how limits, if any, should be placed on utilization of the TNCs for subsidized trips. Since the ADA does not allow constraints to be placed on customers for travel, the current Pilot is strictly elective and all customers continue to be guaranteed service on conventional vans. To understand how changing the specific details of the Pilot would affect travel demand would require analysis of disaggregated customer choices.

5 References

Amalgamated Transit Union (2016). Abandoning Paratransit Service to Save It? How Partnering with Lyft and Uber Undermines the Mission of Transit Agencies. Available from https://www.atu.org/atu-pdfs/fieldmobilization/ParatransitandTNCs.pdf

Amirgholy, M., Gonzales, E.J. (2016). Demand responsive transit systems with time dependent demand: User equilibrium, system optimum, and management strategy. *Transportation Research Part B*, 92:234–252.

Bailey Jr., W. A., & Clark Jr., T. D. (1987). A Simulation Analysis of Demand and Fleet Size Effects on Taxicab Service Rates. In: Proceeding of the 19th Conference on Winter Simulation, ACM Press, 838-844.

Burkhardt, J.E. (2010). Potential cost savings from taxi paratransit programs. Social Research in Transport, Institute of Transport Studies, Monash University.

Burkhardt, J.E., J. Doherty, J.M. Rubino, and J. Kim (2008). Accessible Community Transportation in Our Nation: A Survey on the Use of Taxis in Paratransit Programs. National Aging and Disability Transportation Center, Washington, D.C.

Chia, D. (2006). Policies and Practices for Effectively and Efficiently Meeting ADA Paratransit Demand: A Synthesis of Transit Practice. TCRP Synthesis 74. Transportation Research Board, Washington, D.C.

Chia, D. (2008). Policies and Practices for Effectively and Efficiently Meeting ADA Paratransit Demand. Washington, D.C.: Transportation Research Board.

Committee for Review of Innovative Urban Mobility Services; TRB; National Academies of Sciences, Engineering, and Medicine. (2015). Between Public and Private Mobility: Examining the Rise of Technology-Enabled Transportation Services. Washington, D.C.: National Academies of Sciences, Engineering, and Medicine.

Cordeau, J., & Laporte, G. (2003). A Tabu Search Heuristic for the Static Multi-Vehicle Dial-a-Ride Problem. Transportation Research Part B 37 (6), 579-594.

Daganzo, C. F. (1978). An Approximate Analytic Model of Many-to-Many Demand Responsive Transportation Systems. Transportation Research.

Daganzo, C. F. (1984). The Distance Traveled to Visit n Points with a Maximum of c Stops Per Vehicle: An Analytic Model and an Application. Transp. Sci. 18 (4), 331-350.

Desaulniers, G., Desrosiers, J., Erdmann, A., Solomon, M., & Soumis, F. (2001). VRP With Pickup and Delivery. Society for Industrial and Applied Mathematics, 225-242.

Dessouky, M., L. Quadrifoglio, and F. Ordez (2008). A simulation study of demand responsive transit system design. *Transportation Research Part A*, 42(4):718–737.

Dessouky, M., Ordonez, F., & Quadrifoglio, L. (2005). Productivity and Cost-Effectiveness of Demand Responsive Transit Systems. Tech Rep., California PATH Research Report, UCB-ITS_PRR-29.

Diana, M. (2006). The Importance of Information Flows Temporal Attributes for the Efficient Scheduling of Dynamic Demand Responsive Transport Services. J. Adv. Transp. 40 (1), 23-46.

Diana, M., Dessouky, M. M., & Xia, N. (2006). A Model for the Fleet Sizing of Demand Responsive Transportation Services With Time Windows. Trans. Res. Part B 40 (8), 23-46.

Ellis, E. (2016). *Use of Taxis in Public Transportation for People with Disabilities and Older Adults.* TCRP Synthesis 119. Transportation Research Board, Washington, D.C.

Feigon, S., and Murphy, C. (2016). *Shared Mobility and the Transformation of Public Transit*. TCRP Research Report 188. Transportation Research Board, Washington, D.C.

Feuerstein, E., & Stougie, L. (2001). On-Line Single-Server Dial-a-Ride Problems. Theoret. Comput. Sci. 268, 91-105.

Figliozzi, M. A. (2008). Planning Approximations to the Average Length of Vehicle Routing Problems with Varying Customer Demand and Routing Constraints. Trans. Res. Rec. 2089, 1-8.

Figliozzi, M. A. (2009). Planning Approximations to the Average Length of Vehicle Routing Problems with Time Window Constraints. Trans. Res. Part B 43 (4), 438-447.

Fu, L. (1999). Improving paratransit scheduling by accounting for dynamic and stochastic variations in travel time. *Transportation Research Record*, 1666:74–81.

Fu, L. (2002). A simulation model for evaluating advanced dial a ride paratransit systems. *Transportation Research Part A*, 36(4):291–307.

Fu, L. (2003). Analytical Model for Paratransit Capacity and Quality-of-Service Analysis. Trans. Res. Rec. 1666, 81-89.

Fu, L., & Ishkhanov, G. (2004). Fleet Size and Mix Optimization for Paratransit Services. Transp. Res. Rec. 1884, 39-46.

Fu, L., and S. Teply (1999). On-line and off-line routing and scheduling of dial-a-ride paratransit vehicles. *Computer-Aided Civil and Infrastructure Engineering*, 14(5):309–319.

Gonzales, E. J., & Amirgholy, M. (2015). Demand Responsive Transit Systems With Time-Dependent Demand: User Equilibrium, system optimum, and Management Strategy. *Transportation Research Part B*, 65, 45-67.

Gonzales, E. J., Turmo, V., Rahimi, M., & Armstrong, P. (2018). Evaluating Potential Demand and Operational Effects of Coordinated Americans with Disabilities Act Paratransit and Taxi Service. Washington, D.C.: Transportation Research Board.

Hauptmeier, D., Krumke, S. O., & Rambau, J. (2000). The Online Dial-a-Ride Problem Under Reasonable Load. In: Lecture Notes in Computer Science, Vol. 1767, 125-133.

Hartigan, J.A., Wong, M.A. (1979). Algorithm AS 136: A k-means clustering algorithm. *Journal of the Royal Statistics Society, Series C*, 28(1):100-108.

Henao, A. (2017). Impacts of Ridesourcing – Lyft and Uber – on Transportation Including VMT, Mode Replacement, Parking, and Travel Behavior. PhD Dissertation. Civil Engineering Program, University of Colorado at Denver.

Henao, A. (2017). *Impacts of Ridesourcing – Lyft and Uber – on Transportation Including VMT, Mode Replacement, Parking, and Travel Behavior*. PhD Dissertation. Civil Engineering Program, University of Colorado at Denver.

Kane, J.A., Tomer, A., and Puentes, R. (2016). How Lyft and Uber Can Improve Transit Agency Budgets. Report. Brookings Institution. Available from https://www.brookings.edu/research/how-lyft-and-uber-can-improve-transit-agency-budgets/

Koffman, D., Lewis, D., Chia, D., Burkhardt, J., & Bradley, M. (2007). Improving ADA Complementary Paratransit Demand Estimation. Washington, D.C.: Transportation Research Board.

Laporte, G. (2010). A Concise Guide to the Traveling Salesman Problem. J. Oper. Res Soc 61 (1), 35-40.

Lipmann, M., Lu, X., de Paepe, W. E., & Sitters, R. A. (2002). On-Line Dial-a-Ride Problems Under a Restricted Information Model. In: Lecture Notes in Computer Sciences, Vol. 2461, 674-685.

Martin, E., S. Shaheen, and J. Lidicker, J. (2010). Impact of carsharing on household vehicle holdings: Results from North American shared-use vehicle survey. *Transportation Research Record*, 2143:150–158.

Merzbach, S. (2018, July). PVTA Rates Increase for First Time In Over a Decade. Retrieved from The Recorder.

NADTC. (2019). ADA & Paratransit. Retrieved from National Aging and Disability Transportation Center (NADTC): https://www.nadtc.org/about/transportation-aging-disability/ada-and-paratransit/

Nelson/Nygaard Consulting Associates (2007). *Toolkit for Integrating Non-Dedicated Vehicles in Paratransit Service*. TCRP Report 121. Transportation Research Board, Washington, D.C.

PVTA. (2017). 2017 ADA Paratransit Survey Results. Springfield, MA: Pioneer Valley Transit Authority.

Quadrifoglio, L., and C. Shen (2013). Evaluating centralized versus decentralized zoning strategies for metropolitan ADA paratransit services. *Journal of Transportation Engineering*, 139(5):524–532.

Rahimi, M., & Gonzales, E. J. (2017). Modeling Trip Distribution for Demand Responsive Transit 2 Systems: Evaluating Zoning with Transfers. Transportation Research Board Annual Meeting.

Rahimi, M., Amirgholy, M., & Gonzales, E. J. (2018). System modeling of demand responsive transportation services: Evaluating cost efficiency of service and coordinated taxi usage. Transportation Research Part E, 112:66-83.

Rahimi, M., Gonzales, E.J. (2015). Systematic evaluation of zoning strategies for demand responsive transit. 15-4023. *Transportation Research Board 94th Annual Meeting*, 11–15 January, Washington, D.C.

Rodman, W., High, W., (2018). *ADA Paratransit Service Models: A Synthesis of Practice*. TCRP Synthesis 135. Transportation Research Board, Washington, D.C., 2007.

Savelsbergh, W., and Sol, M. (1995). The General Pickup and Delivery Problem. Transportation Science, 29(1):17-29.

Schaller, B. (2016). *Unfinished Business: A Blueprint for Uber, Lyft and Taxi Regulation*. Schaller Consulting, New York.

Schaller, B. (2017). *Unsustainable? The Growth of App-Based Ride Services and Traffic, Travel and the Future of New York City*. Schaller Consulting, New York.

Shen, C., and L. Quadrifoglio (2011). The coordinated decentralized paratransit system: formulation and comparison with alternative strategies. Paper Number 11-2825. *Transportation Research Board 90th Annual Meeting*, Washington, D.C.

Shen, C., and L. Quadrifoglio (2012). Evaluation of zoning design with transfers for paratransit services. *Transportation Research Record*, 2277:82–89.

Shin-Pei, T., Accardi, Z., Schaller, B., and Hovenkotter, K. (2016). *Private Mobility, Public Interest: How Public Agencies can Work with Emerging Mobility Providers*. TransitCenter: New York.

Taylor, B., et al. (2015). *Between Public and Private Mobility: Examining the Rise of Technology-Enabled Transportation*. Special Report 319. Transportation Research Board, Washington, D.C.

Transportation Services for Individuals with Disabilities (ADA), 49 C.F.R. § 37.131 (1991)

Tuttle, D., Eaton, K. (2012). *Optimizing Accessible Taxi Service to Augment Traditional Public Transit Services in Delaware*. DCT 222. Institute of Public Administration, University of Delaware.

USDOT. (2014). Pioneer Valley Transit Authority (PVTA). Retrieved from Federal Transit Administration: https://www.transit.dot.gov/ntd/transit-agency-profiles/pioneer-valley-transit-authority-pvta

6 Appendix: Operations Parameters

Table 6.1 Average Boarding/Alighting Time, *b* (min/pass.) for GLSS(North), 2017

Month	6 AM – 9 AM		9 AM – 12 PM		12 PM – 3 PM		3 PM – 6 PM		6 PM – 9 PM	
	w'day	w'end	w'day	w'end	w'day	w'end	w'day	w'end	w'day	w'end
January	5.07	5.33	5.49	5.75	6.48	6.21	6.71	6.72	6.63	7.96
February	5.24	5.19	5.49	5.61	6.29	6.20	6.76	6.17	6.96	6.88
March	5.07	5.23	5.48	5.81	6.33	6.32	6.69	6.77	6.84	6.67
April	4.84	5.20	5.11	5.57	5.92	6.27	6.41	6.35	6.72	6.41
May	4.96	5.31	5.25	5.48	5.93	6.37	6.40	7.11	6.67	6.73
Average	*5.03*	*5.35*	*5.37*	*5.65*	*6.19*	*6.27*	*6.59*	*6.62*	*6.67*	*6.93*

Table 6.2 Average Boarding/Alighting Time, *b* (min/pass.) for VTS(West), 2017

Month	6 AM – 9 AM		9 AM – 12 PM		12 PM – 3 PM		3 PM – 6 PM		6 PM – 9 PM	
	w'day	w'end	w'day	w'end	w'day	w'end	w'day	w'end	w'day	w'end
January	7.63	9.44	8.35	9.63	9.55	9.21	9.41	9.80	9.55	10.13
February	7.68	9.19	8.39	9.20	9.36	8.64	9.52	9.96	9.71	10.83
March	7.53	9.12	8.15	9.39	9.13	9.25	9.43	9.83	9.03	10.61
April	7.48	9.37	7.79	8.63	9.11	8.32	9.15	9.31	8.91	10.57
May	7.44	8.71	7.99	8.60	9.04	8.25	8.87	9.37	9.31	9.64
June	7.48	9.03	7.92	9.11	8.89	8.39	8.97	9.64	8.97	10.19
July	7.31	8.57	8.19	8.55	9.24	8.05	9.28	9.20	9.64	10.77
August	7.75	8.72	8.31	8.89	9.33	8.37	9.15	8.92	9.85	9.36
September	7.25	8.29	7.77	8.37	9.07	7.84	9.05	9.36	8.77	8.88
October	7.25	8.48	7.91	8.57	9.27	8.65	9.53	9.88	8.97	9.33
November	7.41	8.31	8.27	8.85	9.49	8.71	9.64	9.81	9.72	10.55
December	7.51	8.63	8.51	9.11	9.60	8.88	9.64	10.35	10.16	10.64
Average	*7.48*	*8.82*	*8.13*	*8.91*	*9.26*	*8.55*	*9.30*	*9.62*	*9.38*	*10.13*

Table 6.3 Avg. Boarding/Alighting Time, *b* (min/pass.) for TRAC(North/South), 2017

Month	6 AM – 9 AM		9 AM – 12 PM		12 PM – 3 PM		3 PM – 6 PM		6 PM – 9 PM	
	w'day	w'end	w'day	w'end	w'day	w'end	w'day	w'end	w'day	w'end
February	6.68	7.59	7.67	7.15	9.31	8.69	9.04	8.84	9.35	10.19
March	6.77	8.80	7.89	8.12	9.13	8.52	8.77	8.80	9.85	10.28
April	6.81	8.97	7.64	8.27	8.92	8.71	8.68	9.13	9.56	10.49
May	6.40	7.39	6.91	7.31	8.17	7.53	7.77	8.37	9.08	9.39
June	6.05	6.93	6.47	6.44	7.77	7.12	7.67	7.85	7.75	8.03
July	6.39	6.92	6.64	6.81	7.85	6.93	7.71	7.79	9.04	8.08
August	6.44	6.77	6.75	6.69	7.83	7.12	7.79	7.96	8.59	8.59
September	5.87	6.69	6.23	6.65	7.83	7.37	7.56	8.13	8.47	7.85
October	5.85	7.29	6.28	6.77	7.63	7.11	7.60	8.45	8.95	8.61
November	6.12	6.89	6.89	7.41	8.27	7.63	8.16	8.20	8.92	9.43
December	6.29	7.09	7.12	7.75	8.43	7.93	8.31	8.99	9.43	9.49
Average	*6.33*	*7.40*	*6.95*	*7.22*	*8.28*	*7.70*	*8.10*	*8.41*	*9.00*	*9.13*

Table 6.4 Average Vehicle Occupancy, n (pass./vehicle) for GLSS(North), 2017

Month	6 AM – 9 AM		9 AM – 12 PM		12 PM – 3 PM		3 PM – 6 PM		6 PM – 9 PM	
	w'day	w'end	w'day	w'end	w'day	w'end	w'day	w'end	w'day	w'end
January	1.47	1.23	1.34	1.31	1.38	1.29	1.48	1.31	1.26	1.24
February	1.45	1.22	1.34	1.33	1.36	1.26	1.47	1.30	1.31	1.23
March	1.46	1.23	1.35	1.35	1.40	1.27	1.49	1.30	1.30	1.27
April	1.47	1.24	1.37	1.33	1.39	1.33	1.51	1.33	1.27	1.32
May	1.47	1.26	1.34	1.35	1.37	1.3	1.49	1.32	1.26	1.25
Average	*1.46*	*1.24*	*1.35*	*1.33*	*1.38*	*1.29*	*1.49*	*1.31*	*1.28*	*1.26*

Table 6.5 Average Vehicle Occupancy, n (pass./vehicle) for VTS(West), 2017

Month	6 AM – 9 AM		9 AM – 12 PM		12 PM – 3 PM		3 PM – 6 PM		6 PM – 9 PM	
	w'day	w'end	w'day	w'end	w'day	w'end	w'day	w'end	w'day	w'end
January	1.39	1.19	1.26	1.21	1.34	1.22	1.41	1.22	1.22	1.16
February	1.38	1.21	1.27	1.21	1.33	1.23	1.40	1.21	1.24	1.15
March	1.40	1.21	1.27	1.22	1.35	1.23	1.39	1.23	1.25	1.16
April	1.41	1.22	1.28	1.24	1.34	1.24	1.42	1.23	1.25	1.18
May	1.42	1.20	1.28	1.23	1.35	1.26	1.41	1.21	1.25	1.15
June	1.40	1.19	1.28	1.22	1.35	1.24	1.40	1.23	1.23	1.20
July	1.41	1.19	1.27	1.21	1.35	1.22	1.40	1.22	1.20	1.16
August	1.40	1.20	1.26	1.24	1.35	1.23	1.39	1.22	1.22	1.22
September	1.43	1.21	1.27	1.23	1.35	1.25	1.39	1.21	1.26	1.18
October	1.42	1.21	1.26	1.21	1.36	1.24	1.41	1.21	1.28	1.20
November	1.42	1.21	1.27	1.23	1.36	1.22	1.42	1.23	1.30	1.16
December	1.42	1.21	1.26	1.20	1.36	1.21	1.42	1.21	1.26	1.19
Average	*1.41*	*1.20*	*1.27*	*1.22*	*1.35*	*1.23*	*1.41*	*1.22*	*1.25*	*1.18*

Table 6.6 Average Vehicle Occupancy, n (pass./vehicle) for TRAC(North/South), 2017

Month	6 AM – 9 AM		9 AM – 12 PM		12 PM – 3 PM		3 PM – 6 PM		6 PM – 9 PM	
	w'day	w'end	w'day	w'end	w'day	w'end	w'day	w'end	w'day	w'end
February	1.54	1.26	1.35	1.32	1.41	1.30	1.50	1.29	1.29	1.32
March	1.57	1.22	1.35	1.27	1.42	1.31	1.55	1.29	1.27	1.24
April	1.53	1.26	1.34	1.28	1.41	1.29	1.53	1.27	1.28	1.22
May	1.45	1.24	1.30	1.26	1.38	1.25	1.46	1.29	1.28	1.29
June	1.46	1.29	1.30	1.29	1.40	1.29	1.49	1.31	1.32	1.31
July	1.44	1.25	1.30	1.28	1.40	1.26	1.48	1.27	1.25	1.27
August	1.46	1.27	1.31	1.28	1.39	1.29	1.50	1.31	1.27	1.29
September	1.46	1.28	1.31	1.31	1.41	1.29	1.51	1.33	1.29	1.26
October	1.48	1.29	1.32	1.27	1.40	1.30	1.52	1.34	1.28	1.26
November	1.44	1.25	1.29	1.25	1.38	1.28	1.49	1.33	1.29	1.24
December	1.39	1.23	1.29	1.24	1.37	1.28	1.47	1.30	1.25	1.24
Average	*1.47*	*1.26*	*1.31*	*1.28*	*1.40*	*1.29*	*1.50*	*1.30*	*1.28*	*1.27*

Table 6.7 Average Network Speed, v (miles/hour) for GLSS(North), 2017

Month	6 AM – 9 AM		9 AM – 12 PM		12 PM – 3 PM		3 PM – 6 PM		6 PM – 9 PM	
	w'day	w'end	w'day	w'end	w'day	w'end	w'day	w'end	w'day	w'end
January	16.02	22.27	18.75	18.92	17.85	17.94	15.17	18.32	19.73	21.75
February	15.88	21.76	17.62	18.34	17.09	17.08	14.94	17.79	19.53	19.51
March	15.56	21.80	18.00	19.15	17.66	18.10	15.05	18.18	19.12	20.72
April	15.29	22.96	17.39	18.58	17.36	17.59	14.43	19.09	19.01	21.31
May	15.06	21.49	17.40	18.44	16.50	17.78	13.86	17.55	19.39	20.69
Average	*15.56*	*22.06*	*17.83*	*18.69*	*17.29*	*17.70*	*14.69*	*18.19*	*19.36*	*20.80*

Table 6.8 Average Network Speed, v (miles/hour) for VTS(West), 2017

Month	6 AM – 9 AM		9 AM – 12 PM		12 PM – 3 PM		3 PM – 6 PM		6 PM – 9 PM	
	w'day	w'end	w'day	w'end	w'day	w'end	w'day	w'end	w'day	w'end
January	15.15	20.10	16.35	17.06	16.53	15.98	13.80	16.84	16.99	19.32
February	15.29	20.49	15.58	16.37	15.95	15.51	13.70	16.29	16.30	18.77
March	15.19	20.60	15.82	16.92	16.17	16.45	13.52	17.15	16.59	19.49
April	14.92	19.99	15.64	16.54	15.74	16.11	13.00	16.95	16.55	19.90
May	15.00	20.31	15.57	16.15	15.23	15.52	12.33	16.90	16.22	19.79
June	15.03	20.22	15.09	15.72	14.92	16.01	12.23	16.73	16.50	20.10
July	15.62	19.89	15.34	17.08	15.20	16.57	13.19	16.81	16.97	19.57
August	15.85	20.67	15.55	16.52	15.24	16.35	13.33	17.25	16.83	19.85
September	14.44	20.39	15.23	16.13	15.33	15.98	12.92	16.77	16.17	19.16
October	14.47	19.71	15.40	16.98	15.20	15.80	12.84	15.90	16.29	19.18
November	14.75	20.09	15.74	16.26	15.71	15.24	12.78	16.13	16.54	19.63
December	15.81	20.06	15.47	16.29	15.45	15.95	12.97	16.35	16.54	18.25
Average	*15.13*	*20.21*	*15.57*	*16.50*	*15.56*	*15.96*	*13.05*	*16.67*	*16.54*	*19.42*

Table 6.9 Average Network Speed, v (miles/hour) for TRAC(North/South), 2017

Month	6 AM – 9 AM		9 AM – 12 PM		12 PM – 3 PM		3 PM – 6 PM		6 PM – 9 PM	
	w'day	w'end	w'day	w'end	w'day	w'end	w'day	w'end	w'day	w'end
February	16.83	20.92	18.16	18.43	17.96	17.68	14.77	17.82	18.88	21.08
March	16.80	22.97	18.29	19.57	17.67	18.46	14.90	18.89	19.40	22.65
April	16.54	22.63	18.02	18.65	17.40	18.00	14.28	18.93	19.27	22.40
May	15.75	22.30	17.49	18.51	16.36	17.24	13.82	17.96	19.43	21.76
June	15.45	21.86	17.06	18.37	15.91	17.21	13.09	17.84	18.58	20.60
July	16.64	21.72	17.39	18.38	16.62	17.43	14.24	17.50	19.03	20.24
August	16.75	21.35	17.58	18.67	16.61	17.65	14.25	18.23	18.74	20.64
September	14.92	21.65	17.20	17.93	16.38	17.34	13.74	18.01	18.50	20.24
October	14.71	20.87	16.95	17.75	16.44	16.22	13.52	17.14	18.17	19.88
November	15.48	22.37	17.69	18.81	16.42	16.96	13.69	17.00	18.33	20.68
December	16.43	21.72	17.82	18.19	16.36	17.37	13.84	17.09	18.85	19.83
Average	*16.03*	*21.85*	*17.60*	*18.48*	*16.74*	*17.41*	*14.01*	*17.86*	*18.83*	*20.91*