

**From the holdings of the Department of Special Collections , Princeton University Library,
Princeton, NJ**

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyright law. We do not charge any permission or usage fees for the publication of images of material in our collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff. For more information regarding copyright, credit, and citation please review the information on the Copyright, Credit and Citation Guidelines page of our website.

<https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines>. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. <https://library.princeton.edu/special-collections/ask-us>

Master's Theses, 2022-2026

Operations Research and Financial Engineering

2023

MOVES-Style Autonomous Transit Network Design, Economic, and Operational Analysis

Dragomir, Christopher Marius

Princeton, NJ : Princeton University

<http://arks.princeton.edu/ark:/88435/dsp0112579w56x>

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. <https://www.copyright.gov/title17/title17.pdf>

MOVES-Style Autonomous Transit Network Design, Economic, and Operational
Analysis

Christopher Marius Dragomir

May 2023

Professor Alain L. Kornhauser

Submitted in partial fulfillment
of the requirements for the degree of
Masters of Science in Engineering
Department of Operations Research and Financial Engineering
Princeton University

I hereby declare that I am the sole author of this thesis.

I authorize Princeton University to lend this thesis to other institutions or individuals for the purpose of scholarly research.

Christopher Marius Dragomir

I further authorize Princeton University to reproduce this thesis by photocopying or by other means, in total or in part, at the request of other institutions or individuals for the purpose of scholarly research

Christopher Marius Dragomir

Abstract

Transportation is the realization of the demand to travel from point A to B, such that one gains utility from the destination, in exchange for expending resources to get there. We all engage in transport throughout each day, and in doing so travel to work, school, grocery stores, recreation centers, family, friends, and a host of other very specific, yet all important destinations. However, the lack of transportation in the United States, perpetuated by a lack of public transit and the large number of households that own zero or one vehicles, makes access to destinations much more challenging. People living in urban settings can often be restricted in both their job opportunities and confined to food deserts due to their lack of access to reliable and connected transportation solutions. Self-driving cars have the potential to revolutionize transportation by simultaneously reducing the time penalty of traveling somewhere, as now one can do other tasks while traveling, and lowering the financial burden that automobiles can incur through ride-sharing cost savings.

Car manufacturers have begun to implement self-driving technology in personal vehicles and heavy-duty trucks, experiencing growing pains in both the software and human elements of the process. However, it seems that the primary focus has been on personal ownership of autonomous vehicles. I argue that leveraging this technology into fleets of autonomous taxis (aTaxi) that can serve both intra-region transit and longer trips would provide an increase in ride-sharing, decrease in congestion and emissions, and provide equitable transportation access to a wide group of people who have often been left to the wayside as technologies advance. The fleets would serve to create a Mobility & Opportunity: Vehicles Equity System (MOVES) in order to best help those who need transportation the most, and facilitate mobility. This thesis documents efforts to synthesize existing work as well as new creations, in assembling economic metrics, data visualization tools, and trip simulations to allow the reader to conceptualize aTaxi networks within the urban centers of Trenton and Perth Amboy.

Acknowledgements

First and foremost, I would like to sincerely thank my advisor, Professor Kornhauser, for his guidance, not only through this thesis, but through my whole Princeton experience. Without those first days in his class, and many, many more years of direction and support, I would not be the scholar I am today. Autonomous transit was a subject I never knew about before entering Princeton, and I am beyond grateful that I am leaving with my eyes open to an incredible field, mainly because of Professor Kornhauser.

To my mother, you pulled the best out of me, knowing that success was always possible if I was just brave enough to grab it. Without your endless love and encouragement, I would never have been able to complete my graduate studies. To my late uncle, Ilie, your jovial spirit and wisdom helped me navigate through the labyrinth of life. I hope to find you one day amongst the stars. To my uncle, Bob, I am thankful for you dropping everything at a moment's notice to help me. I aspire to be a man of such character one day. I would also like to thank my family: Renata, Ilie Jr., Isabelle, and Jasmine, for being there for me during my struggles and triumphs.

Likewise, I would also like to thank Rachel, for her unwavering support throughout my journey. Princeton would not have been the same without you.

Last but not least, I would like to acknowledge Jerry, Gabe, and Bryce for their work towards an autonomous future at Princeton and with CARTS. Their discoveries and insights shaped much of what I have been able to convey in this paper.

Thank you all.

Table of Contents

Abstract	3
Acknowledgements	4
Table of Contents	5
Chapter 1: Introduction	6
Chapter 2: personTrip Synthesis	12
Chapter 3: Economic Analysis	14
Chapter 4: Mobility Disadvantaged Measures	24
Chapter 5: Comparison of Overall Urban Demand and Key Locations	27
Chapter 6: aTaxi Simulation and Results	45
Chapter 7: Conclusion	52
Bibliography	55

Chapter 1: Introduction

Through observing the U.S. energy-related carbon dioxide emissions by sector, it is evident that the top producer of CO₂ for many years, electric power generation, has fallen with respect to the constant increase of the transportation sector [1]. This trend has been facilitated by an increase in renewable and nuclear energy which do not emit carbon dioxide, while coal production has begun to wind down. This has caused a decrease in the carbon intensity, or tons of carbon created per megawatt hour of electricity production, falling “from 619 metric tons per megawatt hour (mt/MWh) in 2005 to 408 mt/MWh in 2019”, which is about a 33% reduction [1]. Transportation has had less opportunities to engage in a similar shift due to a reliance on petroleum, and as a consequence, has overtaken the production section.




Figure: CO₂ Emissions in the US by Sector

Further breaking down the emissions of the transportation sector, the figure below illustrates that passenger cars and light-duty trucks combined emit 57% of the sector’s greenhouse gasses, with heavy-duty trucks contributing a majority of the leftover. The former fulfill the criteria of being “personal vehicles”, as passenger cars are naturally owned by their drivers, and light-duty trucks are often used for personal transport rather than commercial use. This majority is logical, as the U.S.’s less dense suburban structure, in contrast to a higher amount of metropolitan areas, demands a personal vehicle in order

to access jobs, food, and services. It is clear that to reduce emissions within the transport sector, road vehicles are of primary importance, especially personal ones. Thus, in order for the United States to be able to hit ambitious climate goals in time, technological advances must be made in the transportation sectors along with stricter policy decisions such that both producers and consumers transition into a greener future.




Figure: Distribution of GHG emissions in transportation [2]

Much of the drive to decarbonize transportation has come through modification of the internal combustion engine to accept alternate, less polluting fuels. The most popular one as of recently is the electrification of vehicles, such that the car itself is powered solely by a battery. Electrification does eliminate combustion emissions, increase fuel efficiency, and provide some level of energy security over petroleum [3]. However, there are downsides that seem to be somewhat overlooked in the current discussion which often views electrification as a panacea to all of our transportation emissions. It does not inherently eliminate energy production emissions, as the energy previously generated by petroleum must now be fulfilled by the electricity grid. In regions where fossil fuels are dominant, a high saturation of electric vehicles can actually increase greenhouse gas emissions [3]. Electrification likewise does not make cars last longer, as the battery life in electric vehicles is less than a traditional engine's life, and the replacement of that battery creates waste. Finally, since electrification has the potential to reduce the cost of transport

both from a fuel perspective and from an utility perspective, it likely will not reduce the number of vehicles on the road, leading to congestion and possibly more energy use as a whole. Thus, it is imperative to couple electric vehicles with other wide-scale changes to ensure they are actually contributing towards carbon emission reduction.

Although fuel innovations have great potential, some of the mitigation of CO₂ emissions must come from changes in the transportation services offered. Public transportation currently has high potential to reduce emissions over private auto use, but the disconnected nature between the transit network and where people want to go, low costs of driving, and lower quality of public offerings compared to other countries severely impacts utilization rates [4]. This coincides with the figure that, “almost 70% of public transportation trips in the United States are concentrated in just six large metropolitan areas...public transportation agencies headquartered in the New York City area alone accounted for roughly 4 in 10 (41%) of all unlinked passenger trips”, which means public transport in other areas in the US is significantly underused [5]. The lower quality and lack of connectivity becomes even more pervasive when studying income differences between types of transport, as the “annual median household incomes for rail passengers were three times greater than for bus riders: \$72,000 vs. \$22,000” in the US, while European countries do not report a statistically significant difference [5]. This leads to a logical conclusion that buses are taken out of necessity due to lack of personal transportation, rather than as a form of transit accessed by the greater public. In order to facilitate shared transportation, which in turn would increase the use of both bus and rail, planners must focus on addressing the first and last mile trips. This is how users get to the services, and in order to increase ridership, the network must be robust enough such that the distance traveled to get on the public transit itself is not prohibitive to all but those who have no other choice.

A compelling way to address first and last mile trips, and generally shared transportation as a whole, is to create a system that is flexible enough to make access simple, intricate enough to take people to or very close to where they intend to go, and cheap enough for users to prefer it over a personal vehicle. This is where autonomous

vehicles can be leveraged as part of a fleet of autonomous taxis(aTaxi) to serve a geofenced region and provide on-demand transportation. These taxis will have the flexibility of personal vehicles that is not afforded by larger-sized public transit, but at a price not dissimilar to that of public transportation. This service would have some similarities to traditional ride-hailing services such as Uber and Lyft, but leverage autonomy to increase ridesharing and minimize costs. Private transportation network companies have difficulty in profiting off of their ride-hailing services because of the high costs associated with labor, which is the main cost factor that autonomy hopes to solve. However, because this concept of autonomous taxis is not made with a profit motive in mind, the reduced costs of the service will allow passengers to ride for a cheaper price than personal ownership, or would incentivize the government to foot the bill for certain areas and groups.

To imagine how such a system would operate, let us discuss how each of the steps of a typical trip would take place in a aTaxi. Pick-ups of passengers can be discretized to a pre-planned set of locations, named kiosks, which can take many forms, ranging from an improved bus station with opportunities for community engagement and vendoring to a small footprint sign that simply indicates the location as a pickup spot. The size and grandeur of the kiosk will be directly correlated to the traffic heading through that specific location, and the importance of that spot in gaining public interest and acceptance. These kiosks will be interspersed across the geofenced area such that they are a short walk for the majority of potential trips in the area. The aTaxi can wait for a specified time at a given kiosk, perhaps somewhere along the order of five minutes to wait for additional passengers to rideshare the trip with. Whether or not another trip will be added to the existing aTaxi will be represented through the idea of circuitry, or how much extra time the additional trip adds to the trips of others already assigned to the vehicle. Should it be too inconvenient for the current passengers to bear, another aTaxi will route towards that stop to accommodate the next trip, and the process will begin again. As for dropping off passengers, the system can be more flexible with drop offs and either make them on-route to another kiosk location, or engage in simpler decision making and operate on a solely kiosk to kiosk basis. Regardless of the operation choice

implemented, passengers will be picked up and dropped off at locations very close to their needs. Once emptied of passengers, the aTaxi will attempt to find trip requests at or close to its current kiosk location, or if demand is locally sparse, engage in dynamic empty repositioning, where it travels sans occupations to a location where it is aware that demand may soon be peaking due to both current and historical trip data. This entire process can be summarized visually through the following initial concept provided through the collaboration between Princeton University and CARTS.



Figure: Service Concept of aTaxi system, with Perth Amboy as ODD

Despite the potential strengths of an autonomous taxi system, the concept suffers from the change it would require in peoples' behaviors both in ridesharing and in transit choice, the initial legislative hurdles it may encounter, and the complexity of explaining it to an everyday audience. The behavior science behind the acceptance of driverless is complex, and requires time and outreach from the tech sector to educate the general public on operation, and prove to the masses that the technology is inherently safe and able to be operated on local streets. Any autonomous taxi project is at risk to the discontent of those who live in the system, and forcing oneself into a community will likely end disastrously. As for the legislative aspect, states must begin to address their older rules around cars, and redefine them to allow autonomous cars to operate as well, with new regulation that is appropriate and not inherently restrictive. This thesis will not deeply address these first two issues, but rather attempt to provide access to an easier and

visual way to understand how an autonomous network would operate in a city, how many cars would be needed for it to run smoothly, the cost of those cars, and other facts about operations such as number of daily trips, milage of trips, clusters of movement etc. Giving the general public and industry professionals this opportunity would hope to inspire both acceptance of the technology by a broader group of people, and generate financial interest in both the public and private sectors.

Chapter 2: personTrip Synthesis

The data that this thesis holds as its backbone could not exist without the continuous developments of several theses including “A National Disaggregate Transportation Demand Model for the Analysis of Autonomous Taxi Systems” - P. WyrOUGH and “Synthesis of Spatially & Temporally Disaggregate Person Trip Demand” - T. Mufti. These techniques and learnings were applied in “A National Hybrid Activity/Agent-Based Demand Model to Characterize the Mobility of the United States” - K. Marocchini to create a synthetic, but still incredibly representative model of every trip that happens within the US within a single day. Without this data, we would not be able to run visualizations of autonomous taxi systems in any region or perform any analysis for that matter, because we would have no knowledge of the trips individuals take in a day. Thus, I emphasize the importance of having this data, and its existence in itself opens millions of possibilities for transportation analysis.

To understand how the nation's trips are synthetically generated, I will summarize the process employed by Marocchini and his predecessors. First, census data, specifically the 2010 Census, is utilized to find distributions of demographic characteristics (age, gender, household size, etc.) and assign worker types to each person [6]. Employer census data for a region is used through a random draw process to assign workplaces to each of the people in the resident file that are of working age, through Census Journey-to-Work data that allows the program to know what country each worker commutes to their job in [6]. This same process is applied to students, using Department of Education data and then applying a probability distribution that sends the majority of students to public schools, with a smaller percent going to options of private, special, and homeschool [6]. Now that each person in the resident file has a home latitude and longitude, along with corresponding school or workplace, the process can move into populating the trips of the day. To do so, it is assumed that each person engages in a “trip tour”, or a chain ranging from zero to seven nodes, which are each a location visited by that person in a day. There exists a dictionary of weighted probabilities for each type of individual (grade school student, middle school, college on campus, worker, etc.) to randomly select what trip type they perform [6]. Trip cycles range on the shorter end from home to work and back, to


much more circuitous paths that stop at “other” labeled destinations before returning home. These other locations are selected as an attractiveness-weighted random draw, which is a ratio of their patronage (popularity among the county) to the distance that the current individual is from the location. Then, using all this information, a PersonTrip file is generated containing identifying characteristics about each trip including: latitude and longitude of origin and destination, person ID of who is taking the trip, and at what time the trip is occurring. These individual trips represent the entire market of transit demand of which an autonomous taxi service could offer service to a subset.

Even though the process from which the trips are generated is robust and well-thought out, it is important to note several limitations of the given data, such that we can recognize how that impacts further analysis of the results. First off, at a microscopic level, trips are not entirely representative of a particular person’s movement throughout one day, as some paths could be rather erratic or just travel between locations that do not inherently make logical sense. However, when we begin to view thousands and hundreds of thousands of trips concurrently, their combined behavior begins to shape into a more accurate picture of a region’s daily movement. Secondly, some kinds of trips are not compatible with the trip chains that have been assumed by the model, for instance, there are no trip chains that go from home, to school, and then to work, even though this may be a common trip for some adults that drop their children off at their place of study before heading to their own personal jobs. The absence of these is simply part of a more general simplification that is necessary to have the simulation run in a feasible amount of time across the entire population of the US. Finally, most of the data utilized in the generation is based on the 2010 Census, and has not been rerun with 2020 census data, as the coronavirus pandemic has delayed the release of much of the data of this newly created report. Much could have changed demographically and density-wise with the US population in this time, and new information is sorely needed to update the trips to something more relevant. Regardless of these issues, this data is incredibly valuable and still provides a picture of what occurs each day on our roads, trains, buses, and much more.

Chapter 3: Economic Analysis

In order to advocate and illustrate a MOVES ridesharing system to both the public body, and any local or state governments that would invest and support such a project, it is necessary to have an economic representation of what an autonomous fleet's costs would be. This allows the agency to compare it to existing forms of transportation, and how much of the market they serve at a given price point. Along with that, this provides important information for government bodies to know when applying for federal infrastructure or transportation grants to create an aTaxi system, as they need to submit appropriate cost-benefit analysis in their proposal. Autonomous ridesharing technology has the ability to greatly outperform traditional transit on purely a cost per rider basis, along with other benefits.

The initial design of this economic model was inspired by the work of ARK Investments, who claim that per their model, “a ride in an autonomous taxi will cost consumers 26 cents per mile, less than half the cost of a personal car and less than one tenth the cost of a manually driven taxi.” [7] They have utilized several key factors to generate a financial model which prices travel per mile accordingly, along with the following sensitivity analysis.



Source: ARK Investment Management LLC. 2019 | ark-invest.com

Figure: Sensitivity of Autonomous Price per Mile [7]

After examining their work, it is clear that the financial performance of autonomous ridesharing drops dramatically when utilization rate is lowered, which is logically sound, as cost per mile should naturally be most affected by usage, as that is the one of the key reasons bus transit does poorly from a financial perspective. A high utilization rate allows the price of the infrastructure and flat management costs to be amortized across a larger number of passengers, as well as increasing the likelihood of ridesharing from location to location. Likewise, I would argue that both their prices of electric vehicles and autonomous technology are irrationally optimistic, given the current state of the market, and thus, must be scaled higher to compensate. The model also does not include platform operator fees, which should be calculated and implemented appropriately. A final point of contention is that the ARK model does not account for any externalities, both positive and negative from the creation of an autonomous transit network, which could be a region for further expansion when examining the robustness of the estimation. For example, one such positive externality is that the network could allow better access to existing public transport, which in turn would reduce localized

congestion and provide an increase in alternative travel methods. However, autonomous taxis also serve to threaten existing taxi services, although this is less of a consideration since the likes of Uber and Lyft have already driven out much of traditional taxi systems through their more convenient methods.

The economic model can be split into two categories as follows: up-front capital investments(cars, kiosks, and infrastructure) and continuous operational costs(labor, maintenance, insurance, and energy). However, because some of the costs of an aTaxi network are variable in nature due to geographic constraints, such as: the cost of building kiosks, the land on which to build these structures, and the space required to park and charge such a vehicle set, the first analysis I created focused on a more limited set of inputs to focus in on the accuracy of the calculations. Likewise, in order to make the economic analysis more readily usable for later implementation in the data visualization tool, it was necessary to approach it from a price per vehicle model, rather than a price per mile traveled model. This may result in a higher cost estimation, but I believe it will create a more logical picture for the data visualization, as price per mile can be heavily affected by utilization rate, speeds of local roads, and trip lengths. An upfront cost for the entire fleet is less mutable by changes in the usage of each vehicle, as many of the constants are the same regardless of the conditions present in the particular geographic area. In software that will be able to select any region of the US to build an autonomous taxi network, it is logical to present economical estimates that are less variable in nature. Another way this model differs from ARK's implementation was to operate it over a five-year period and assess costs in that manner. This was chosen because it allows for two years of supervised "trial" operation of the system and then a remaining three years of full operation. Likewise, this seems like a relatively accurate lifetime cycle for the autonomous vehicles, as if they are used heavily each year, it is likely that they will need to be replaced after five years.

The system is modifiable on the back-end through the use of several alterable constants that can be varied either as part of a sensitivity analysis, or simply due to

updated research allowing for more precision. The following figure shows an example of the given constants.

Constants	(per car)
Cost of Autonomous Tech	100,000
Cost of Car	100,000
Labor per Year	100,000
Parking	3,700
Tires for a year	1,500
Remote network operator	1,000
Insurance	4,000
Energy	4,000
Maintenance	1,000
Number of Vehicles	100
Annual Revenue	0

Figure: Input Variables for Autonomous Vehicle Economic Analysis

The selected values are rationalized as follows:

- **Autonomous Tech:** This is a particularly challenging metric to determine, as the only real data point one could reference in a personal vehicle is Tesla's full self-driving(FSD) add-on which costs an additional \$15,000 to add when purchasing the base vehicle [8]. However, it is assumed that the cost of the initial car offsets the research cost of the autonomous technology, and it is quite likely that the software should be valued much higher than that value. As a result, we cautiously assume a price of \$100,000 per vehicle.
- **Car:** Electric vehicles created by Tesla range in price from a lower bound of \$50,000 to as high as \$120,000 [8]. However, it is important to note that a ridesharing AV will likely have a different interior design, considering vehicle occupancy capabilities are of primary priority. An equal price to the autonomous technology seems reasonable, and thus, we assume the cost of the car will be around \$100,000.
- **Labor costs:** These pay two attendants \$50,000 per year per vehicle to be behind the wheel of the vehicle for the first two years of the system only. This is mainly to bolster public confidence in the system, but these attendants also serve as key workers who can override initial errors or manage ridesharing interactions. The idea of having an attendant behind the wheel has been utilized in services such as

Waymo's Arizona deployment and the now defunct Uber program mainly as a liability measure.

- Parking: When vehicles are not in service, require charging, or cleaning and maintenance, they must have a location assigned to them. We can determine a price for this by determining the cost for a private parking spot, and then scaling it as "parking costs for fleet operators are assumed to be 133% higher than for private drivers" due to factors involving a large requirement for space and accessible facilities for management [9]. Yearly parking rates are around \$1,600 averaged across the state of New Jersey, thus, when scaling this for a fleet, we estimate parking to be about \$3,700 per vehicle [10].
- Tires: A set of tires for a Tesla are equal to this value, thus it is a reasonable initial assumption [8].
- Remote network operator: This value is part constant, part formula. We assume that two operators, working at the same attendant salary of \$50,000 per year are able to attend to 100 autonomous vehicles. Thus, this value is simply $1000 * (\text{number of vehicles}/100)$.
- Insurance: One way to determine this value is to compare it to the price of taxi insurance, which currently has an upper bound of around \$10,000 per year [11]. Autonomous vehicles will naturally experience less accidents than a human-driven vehicle. The IIHS(Insurance Institute for Highway Safety) writes that about 23% of crashes nationwide are caused by sensing and perceiving errors, along with 10% being attributed to incapacitation, which an autonomous vehicle could realistically eliminate entirely [12]. This leaves about 67% of crash scenarios remaining, and a conservative low-bound estimate would be that autonomous technology could remove about one-half of those crashes. Thus, we can conservatively estimate that an autonomous taxi's insurance rate should be 60% less than a conventional taxi's rate, which gives a figure of \$4,000 yearly.
- Energy: The long-range Tesla Model 3 requires 82kWh for 350 miles of travel [8]. The average electricity rate in New Jersey is \$0.16/kWh [13]. Thus, the cost of energy is \$0.04/mile, with an assumed yearly travel distance of 100,000 miles, equals \$4,000.

- **Maintenance:** According to AAA, the annual cost of maintenance of an electric vehicle is around the order of \$1,000 [14]. This price could increase if maintenance is done in-house, which would require additional salaried employees.
- **Number of Vehicles:** This will become an input value in the data visualization tool, as the operator will choose the quantity of vehicles needed for their own system, or the tool itself will recommend a value.
- **Annual Revenue:** This is solely dependent on whether the system is a paid one, or sponsored/subsidized by a government agency. Setting this value to zero allows for clearer cost analysis.

Likewise, to compute the net present value of one vehicle below, we use the currently established federal corporate tax rate of 21% and a more aggressive discount rate of 10%. We also assume the rate of depreciation of physical assets to be constant over the five year period.

MODEL

Autonomous Taxi Price Per Car - (5 year investment and full operation)

(units in USD)

Year	0	1	2	3	4	5
Cost of Autonomous Tech	100,000					
Cost of Car	100,000					
Labor		-100,000	-100,000	0	0	0
Insurance		-4,000	-4,000	-4,000	-4,000	-4,000
Energy		-4,000	-4,000	-4,000	-4,000	-4,000
Maintenance		-1,000	-1,000	-1,000	-1,000	-1,000
Tires		-1,500	-1,500	-1,500	-1,500	-1,500
Remote Operator Telecom Expense		-1,000	-1,000	-1,000	-1,000	-1,000
Parking		-3,700	-3,700	-3,700	-3,700	-3,700
Depreciation		-40,000	-40,000	-40,000	-40,000	-40,000
Total Costs	-200,000	-155,200	-155,200	-55,200	-55,200	-55,200
Revenue		0	0	0	0	0
Net	-200,000	-155,200	-155,200	-55,200	-55,200	-55,200
Corporate Tax Rate	21%					
Taxes						
Cash Flow	-200,000	-115,200	-115,200	-15,200	-15,200	-15,200
Cumulative Cashflow	-200,000	-315,200	-430,400	-445,600	-460,800	-476,000
Discount Rate	10%					
NPV of one vehicle	(431,174)					
Final cost of fleet for 5 year	\$43,117,368					

Figure: Financial Model of Autonomous Taxi Price per Car over 5 Years

In the model, the differentiation between the upfront costs of the vehicle and technology(year 0), along with the yearly continuous costs associated with each vehicle(years 1-5) is represented. The initial cost of the vehicle and technology are a bit smaller than the entire cost of its operation and depreciation, but the initial upfront costs of purchasing a fleet are definitely the largest single economic factor to starting the aTaxi network. Overall, with the given constants, the cost to purchase and operate one autonomous taxi is \$430,000 over a five-year period, with a 100 car fleet costing on the order of \$40 million to operate. A fleet of that size would be able to address a considerable chunk of trips in a medium population urban city such as Trenton, somewhere on the scale of 16,000 individual trips per day.

The simplicity and effectiveness of an autonomous taxi system is backed up by the fact that the infrastructure to operate this transit system, roads, already exist and are upkeep through state and federal funds. That being said, we would like to suggest that investment into some additional non-road infrastructure, the kiosks discussed in the introduction, would improve the experience of consumers and allow for better efficiency throughout. Kiosks can range in size and price from an indicator sign on the side of the road, to a small building with an information center, goods for sale, and other features.



Figure: Small, low footprint and cost kiosk versus larger multipurpose infrastructure[15]

The next logical question is, when is it the right decision to construct a well equipped kiosk, and when is it ok to use a simple design, or nothing at all. During the pilot stage of the system, it makes sense to have temporarily built kiosks to help with

onboarding, but that does not need to last forever. Permanent large kiosks would serve best in helping those who are unbanked to access the system, those who have trouble with technology, and generally with promoting the presence of the aTaxi to the general public. This last effect is seen in bus stops as well as “boardings at stops with new bike hoops were 203.5% higher; boardings at stops with new RTIS[Real-Time Information Systems] were 199.1% higher; and boardings at stops with new or improved shelters were 81.4% higher [than the control stops]” [16]. Thus, it is important to invest early in the optional infrastructure to increase public knowledge and accessibility.

In a system that would eventually expand to fifty kiosks total, we will pick the following distribution for our subsequent financial analysis. First, select the top 10 most populous or community impactful kiosks to build large kiosks at, then delegate 30 remaining kiosks to medium sized bus-shelter type and 10 simple small bus-stop kiosks to the least utilized pixels. To obtain rough estimations for each of these infrastructure solutions, I consulted currently existing projects for reference in pricing. For the idea of the large kiosk, I reference the price of NYC newsstand costs because the physical infrastructure of a newsstand would be a similar size to the kiosks one could implement. They are standardized in price, as they can be bought from a single retailer, and each cost \$35,681 as of 2022 [17]. Taking this cost for the physical structure of a larger kiosk and rounding to account for the cost of technology within the kiosk, gives us a rough estimate of \$50,000 per. Likewise, we will have to staff these larger kiosks with workers to answer questions and serve as spokespersons for the system. This will require at least two full-time-equivalent salaries per year, and we will take the same estimation as the cost for attendants for our autonomous vehicles of \$50,000 per year. In reference to smaller kiosk solutions, NJ Transit recently installed solar panel powered bus shelters in southern New Jersey at a cost of \$12,000 per unit [18], which I will then scale to \$15,000 to account for the cost of having a small electronic ordering screen installed to the station as well. For maintenance costs for both large and medium kiosk stops, we will work off an estimate given by the MTA of \$7,000 per stop [19], and scale that to \$10,000 to account for the maintenance of the technology as well. For the “small” kiosk, I reference the general cost

of standard traffic signs which is about \$30 for the sign itself, \$100-150 for the post, and \$150-250 for the foundation, rounding out to a conservative total estimate of \$500 [20].

	Large(Multipurpose and enclosed)	Medium(Bus Shelter with modifications)	Small(Just signage)
Installation(\$)	50,000	15,000	500
Labor (per year)	100,000	0	0
Maintenance(per year)	10,000	10,000	0
Total Cost over 5 years(\$)	600,000	65,500	500
Quantity	10	30	10
Net Cost(\$)	7,970,000		

Adding the sum of all of our kiosk infrastructure with the aTaxi fleet purchasing and operation costs, results in a 50 million dollar total for a five-year service. It is prudent to then compare the costs of our hypothetical system with the existing transit costs of available systems. Logically, we should compare the autonomous transit service to bus transit, as it most closely resembles the general operation of our system. Comparing the two, we observe that the approximate lifetime cost(combination of capital and operations costs, same as aTaxi model) of a electric bus is around \$1.2 million dollars [21], and it is not far-fetched to claim that a well-optimized aTaxi could reach a similar order of magnitude of completed personTrips per day as a bus, especially in regions where bus utilization continues to fall. Likewise, infrastructure costs for an aTaxi system will be generally similar to bus transit, except for the few kiosk depots that are planned to be more well equipped than a regular bus shelter. Next, in comparison to light rail, we can see that our estimated costs outperform it by an order of magnitude, mainly due to the incredibly expensive infrastructure associated with rail. For instance, the proposed Brooklyn-Queens Connector, which would consist of 16 miles of light rail through the waterfront would cost \$2.5 billion to construct, have yearly operating costs of \$26 million, and serve about 45,000-50,000 daily riders [22]. Though this system has triple the capacity of our proposed 100 aTaxi system, the scaling in costs is 50x larger, which demonstrates the dollar per passenger efficiency of an aTaxi system. Finally, it would be

useful to compare to paratransit costs, as the on-demand nature of aTaxis allows them to be more accessible to people with disabilities than rail or bus options, which in turn would compete with paratransit. It has proved more challenging to find data on a singular paratransit system's cost and individual ridership statistics, but it is estimated that these services can push upwards of \$60 per trip, which is incredibly high [23]. To compare this to an aTaxi system, we can perform some back of the envelope math: a 48 million dollar system will serve 16,000 trips a day, 300 days a year (conservative operating schedule), for 5 years. This results in a cost of approximately two dollars a trip, vastly outperforming paratransit costs. However, in order to truly serve as a replacement for ADA services, aTaxis must consider the particular needs of the mobility disadvantaged who will be utilizing the service, which I will discuss in the next chapter.

Chapter 4: Mobility Disadvantaged Measures

The current retail available self-driving vehicles, along with initial deployments of autonomous taxi services have been utilized in such a way that has made them inaccessible to all but the financially well-off, or those living in suburban areas such as Chandler, AZ(Waymo deployment), who already have ample personal transportation. Because of this, it is imperative that we consider mobility disadvantaged people when planning our systems, such that the technology is geared to help those who have ample demand for transportation, but do not have the resources to match that. We can generally divide the mobility disadvantaged into four categories that an aTaxi service would need to be tailored to serve, along with other passengers [15].

- Population with disabilities, based on US Census, including difficulties with hearing, vision, cognition, independent living, ambulation, and self-care
- Senior Citizens: ages 65 and above
- Enrolled students: grades 5-12, ages 10-18 typically
- Households with limited mobility: have access to 0 or 1 vehicles for the entire household

First, let us discuss how aTaxis can interface with people with disabilities and consider the barriers that are traditionally present in public and Mobility as a Service(MaaS) transportation systems [24]. Some inherent roadblocks in MaaS systems include:

- lack of effective programs (no policy initiative implementation)
- obstacles to special transport service and accessible taxis (limited availability and times)
- physical and informational barriers (information locked behind technology barriers, spots of transport are not accessible)
- lack of continuity in the travel chain (certain parts of the travel chain ie. Last/First mile are not accessible by all)
- lack of pedestrian access (hazards, crowds, construction, etc.)

- lack of staff awareness (operators are not prepared to assist passengers with disabilities)

Autonomous taxis by definition are accessible to pedestrians, as they can pull up to any empty curb, and can provide continuity in the travel chain by practically eliminating first and last mile trips. Likewise, there are less concerns with limited availability and times compared to traditional transport for people with disabilities, as it is clear that aTaxi will operate much more frequently than most other transit currently on the market. However, there are some barriers listed that need addressing through design and operational changes. For the pilot of such a system, it would be necessary to train vehicle attendants on how to assist individuals with disabilities to onboard and offload from vehicles, which should be built to be ADA accessible. On the software end with the development of an app to hail rides from the system, accessible measures would have to be taken to improve ease of access for those with disabilities. To help assist with both of these concerns, it could be worthwhile to station workers at certain larger “mega-kiosks” to assist with app use, guide vehicle boarding, answer questions, and help summon rides.

Next, how can aTaxi be compelling for senior citizens? According to the Committee on Disability in America, older people have become more dependent on the private car as they age, as their rates of taking any transit minimize except for the use of personal vehicles and paratransit services [25]. However, as one ages, the risk of getting into an accident increases. Although senior citizens have a lower accident rate than many age groups, especially 18-25 year olds, this is because they drive much less than other age groups, choose to drive solely in good weather conditions, and avoid longer distance trips [25]. If one was to normalize miles driven to the number of accidents, senior citizens have higher per mile crash rates than middle aged drivers. As a result, senior citizens generally engage in transportation only in a limited set of scenarios and do not interface with public transportation because of the non-connectivity of the service requiring walking or other modes of transport. When considering this in a aTaxi fleet, kiosk locations could be smartly positioned around spots that have high senior density such as living communities, and there could be options within the system to request a direct

dropoff to mitigate any apprehension about having to walk to one's destination at the end of the trip.

As for enrolled students, there are a plethora of regions in the US that do not use a traditional "yellow bus" transportation system for taking students to school, but rather make students rely on public buses or personal transportation. In these cases, and even in more supporting environments, students will face absences or lateness to school because of lack of transportation. A study on risk factors for school absences showed that living far from their school had an r mean effect of 0.476 on absenteeism, which is surprisingly higher than other potentially more impactful factors such as depression, drug abuse, and bullying [26]. Thus, it is clear that should autonomous taxis be made available, leveraging them to provide transit to certain groups of students to improve attendance would be valuable. This could come through a partnership with the local school, or as a broader municipal government initiative, and since school's start and end times are often aligned with those of work commutes(peak demand), it would be of no issue to fulfill these trips in a way that promotes high ridesharing with work trips.

Finally, we must consider providing trips to households with limited mobility, in that they own no personal vehicles, or the few that they do own cannot serve the needs of the household concurrently. Because individuals in these households lack access to transportation, they have fewer opportunities to travel for higher-paying jobs, have no access to quality groceries in urban food deserts, and generally are limited in their day-to-day tasks. This group stands to benefit the most from aTaxi, and thus, a system must interface with this group so that people are well-informed on its existence and use it. This can be facilitated through outreach events, display of the vehicles, and other experiences that all serve to cultivate a "welcoming environment" where people are not afraid or distrustful of the technology. An aTaxi system is only as effective as the utility gained by those who utilize it, and it is clear that households with limited mobility gain much more from it than others who already own and operate their personal vehicles.

,

Chapter 5: Comparison of Overall Urban Demand and Key Locations

As a transportation planner, or as an organization who would like to explore potential demand hotspots of a aTaxi network, it is critical to be able to temporally visualize trips. However, due to the nature of demand surges and valleys, it can be hard to understand an entire set of forecasted travel at a certain location, as the peak times are at such high levels, that they inherently minimize the useful information from the rest of the day. This is even more of an issue when we consider that an aTaxi service would only initially attempt to provide a mode split of 5-10% of the overall network's trips, such that periods of high demand will inherently have a smaller percentage of trips covered than in lower demand times when the active fleet can more likely cover demand adequately. Thus, it is important to have a way to visualize the minute-by-minute trip demand in a region in such a way that allows for consistent, lower trip demand to become apparent.

For this visualization exploration, I chose the urban cities of Trenton and Perth Amboy as places where an aTaxi network could be established. This is due to several factors including percentage of mobility disadvantaged citizens, high trip density, location along the critical northeastern corridor of the US, general walkability, and low personal vehicle ownership. I will justify this selection as follows. When looking at income data generated from the trip demand model, inherently census data, we find that in reference to national poverty thresholds by household size, 15% of Perth Amboy households and 30% of Trenton households match the criteria, both exceeding the nationwide average of 11.6% [27]. Another interesting phenomenon can be seen when observing all trips heading in and out of Trenton and Perth Amboy, is once they are merged with census information, there is a distinct correlation between trip distance and household income. A trend seems to hold for both, that the density of longer trips taken by a household is inversely related to the household income. This suggests that since those of lower income households have transit needs that are larger in length, but since they are less likely to have cars to complete them, and these trips are much harder on public transportation, people in these cities would benefit greatly from on-demand transportation that would allow them to access rail and long-distance bus easier.




Figure: Relationship between Trenton's Trip Distances and Household Income [27]



Figure: Relationship between Perth Amboy's Trip Distances and Household Income [27]

Adding to that, about 70% and 60% of households in Trenton and Perth Amboy respectively own either zero or one vehicles, a figure that exceeds the New Jersey average of around 45% [15]. When combining low car ownership to high household sizes as summarized in the following table, it is clear that even if one person in the family has mobility, the other members are limited and must base their travel off of the primary user's schedule. About 40-50% of households in these cities have three or more inhabitants, which indicates that when that overlaps with limited mobility options, many in the household are lacking transportation solutions.

Household Size	Trenton (percentile)	Perth Amboy (percentile)
1	26%	19%
2	44%	36%
3	58%	53%
4	68%	66%
5	80%	80%
6	90%	91%

Table: Household Size in Each City in a Percentile Basis

Finally, I examined the National Walkability Index, which is calculated as a weighted sum of intersection density, proximity to transit stops, and employment mix, to generally approximate how easily the selected census block is to navigate on foot and access desirable locations [28]. A visual analysis shows that both Trenton and Perth Amboy are almost entirely covered by Above Average and Most Walkable blocks [28]. This is an important metric for assessing what area to place an aTaxi network in, because planners would like higher walkability in the ODD region. This seems initially counterintuitive, since if the region is inherently walkable, why do individuals require micro-transit? However, I will claim that walkability allows for flexibility in aTaxi pickups, and allows an operator to create discrete kiosk locations to centralize demand, a concept which I will explain next. Likewise, regardless of walkability in a region, the distance of a trip by foot can still be a large deterrent to travel, especially if one is carrying heavy goods or cannot walk as easily as others.




Figure: National Walkability Index for Trenton and Perth Amboy

To begin city level analysis, I start with the city of Trenton, and define an Operational Design Domain (ODD) which is a geographic boundary in which an aTaxi network would operate. Taxis would circulate within these boundaries, loading and unloading passengers at designated kiosk stops that cover important locations or high density spots. Each kiosk is represented by its selected pixel, which is a discretization of latitude/longitude coordinates into a covering grid of 0.5 mile sided squares. This technique allows for better grouping of trip origins and destinations as it is unlikely to find any two trips with exactly the same origin, but zooming out slightly into a walkable range serves to aggregate trips easily for ridesharing potential. Within this ODD, assigning trips to closest available kiosks for both arrival and departure, the following distribution of trip lengths are found, where total mileage is represented by great-circle distance between the two points. Overall trip distance in a vehicle will be longer due to the geographic impossibility of driving directly between two points, which can be quantified by the following assumption that “in nearly all locations within the US, the straight-line distance is an adequate proxy for travel distance, after applying a detour index of about 1.4” [29].




Figure: Cumulative distribution of trip lengths within Trenton's ODD

Regardless, it is clear that trips between kiosks in Trenton can be completed quickly by vehicle due to their short length, and thus, if we assume a relatively conservative average aTaxi speed of 25 mph, the tail end of trips are about 20 minutes long from point-to-point. Short-length, clustered trips in an ODD is a good indicator for ride-sharing capability, thus, this initial peek at the serviceable trips has a positive outlook.

Diving deeper for the next analysis, three types of granularities were selected to cap the amount of personTrips/minute as follows: 1000 personTrips/minute with excess demand being set to a constant 1200(Large Cap), 100 personTrips/minute with excess demand being set to a constant 120(Medium Cap), and 10 personTrips/minute with excess demand being set to a constant 12(Small Cap). These caps were chosen to illustrate demand on different scales, and if we make the assumption of maximum occupancy of an aTaxi being five individuals, the cap of ten personTrips could be easily served by only two vehicles should trips be aggregated, while hundred personTrips would take a factor of ten higher vehicles. The ratio of cars needed to service a certain number of trips(1 car : 4 passengers) can only be held true on a pixel by pixel basis, as 10 trips in sparse locations in an ODD would require a significantly higher ratio to serve entirely. Because of this, in later analysis, we will observe behavior at selected pixels that demonstrate high or interesting demand patterns. The largest cap was only included as a relevant metric when observing entire ODDs through this visualization, as it is unlikely

for a single pixel to produce that much demand per minute. Each bar represents a singular minute of the day, and the following three charts summarize the behavior of all 138,377 trips that happen between the designated pixels that encompass Trenton.




Figure: Large Cap Total Addressable Market Trips in Trenton




Figure: Medium Cap Total Addressable Market Trips in Trenton




Figure: Small Cap Total Addressable Market Trips

The large cap chart paints an aggregate picture of Trenton's daily demand, which has local maxima around the hours of 8:00, 14:00, and 20:00, along with a much larger global maximum at around 17:00 which reaches the cap of 1000 personTrips/minute, and is thus, set to the predetermined value of 1200. However, other than these maximas, not as much important information can be gleaned when it comes to assessing the viability of transportation services. Thus, we zoom in on the medium cap chart and find consistent capped demand between the hours of 6:00-9:00 and 16:00-21:00, with some sparse peaks in between. This confirms the somewhat known assumption that the work commutes are full of trips that could be targeted by a novel transit service, but also shows that the demand extends past commutes, into dinner time and recreational activities. The length of this period of demand makes a strong case to the market available in evening hours and indicates that those living in Trenton have internally concentrated travel demand. Finally, we can consult the small cap chart to see how a baseline demand looks at all times of day. At almost all times of the day, we exceed the cap, which demonstrates that trips demand exists continuously. Even though constant demand is not a precursor of viability, since initial aTaxi systems would likely operate between scheduled hours, this is indicative of future potential for 24/7 service.

Continuing onto a pixel by pixel analysis, I have selected three key kiosk locations that would be prime service locations for pilot aTaxi programs, due to their

nature: the Trenton Transit Center, Trenton Central High School, and Mayor Donnelly Homes. The transit center is a major Amtrak and regional rail connection that serves as a junction between the PA and NJ borders. Likewise, it serves as commuter rail to northern New Jersey and New York City itself, and is located directly off of US-1, a major freeway that connects the northeastern region of states. Thus, integrating this location into the ODD allows us to service first and last mile trips going to the station, increasing accessibility to public transit and working in tandem to increase ridership. Then, the high school was selected as it is the primary secondary education location within the city, and as discussed previously, providing increased mobility to students increases their likelihood of graduating. Finally, the Mayor Donnelly Homes were selected as a pixel of interest due to its status as an affordable housing center operated by the Trenton Housing Authority, where residents have limited mobility options. It is geographically a closed-off community where buses, Ubers, and Lyft only stop and pick up around the periphery, such that a kiosk location could be created right on the outside of the property, and provide much needed on-demand service. Thus, through these specific selections we are covering analysis for integration with public transit, and several components of the mobility disadvantage population we previously defined. In the following graph analysis, I utilized only small and medium caps, because testing showed that expanding the cap any larger has no effect, as there are not enough minute by minute trips on a pixel basis to do so. Initial information about the type trips within these pixels is summarized in the following table. We generally see behavior correlated with expectations, the transit center and high school are dominated by work and school trips respectively. Likewise, the Donnelly homes experience many ‘Other’ trips as these are trips to groceries, recreation, and shopping, which is natural for inhabitants of a residential center.

	Number of Trips	School	Work	Other
Trenton Transit Center	39,340	11.2%	41%	47.8%
Trenton Central High School	9,189	45.4%	13.4%	41.1%
Mayor Donnelly Homes	3,503	9.8%	28.6%	61.6%




Figure: Medium and Small Cap Kiosk Trips to/from Trenton Transit Pixel




Figure: Medium and Small Cap Kiosk Trips to/from Trenton Central High Pixel




Figure: Medium and Small Cap Kiosk Trips to/from Donnelly Homes Pixel

In these capped charts, we see travel behavior at a more microscopic level throughout the day. Trenton Transit Center provides near constant small cap demand from hours 06:00 to 24:00, and reaches the medium cap of trips per minute at 11:00,

17:00, and 20:00. As a result, it is imperative that in designing a kiosk stop near the train station, this potential demand is heavily considered. When inspecting the roadway next to the transit stop, a two-lane road with a small dropoff location directly in front of the center, it is clear that it is primarily used by cars looking to drop off passengers along with a bus-stop. A kiosk should be built such that it does not interfere and create tension with the existing infrastructure, but not far enough such that potential users find the walk to be negative. This kiosk must also cover enough area to allow for multiple cars dropping and Signage inside the transit center, along with an info booth that could work in tandem to capture public transit riders and expose them to the system's capabilities.

Next, within the Central High pixel, we see strong demand that aligns with school start and closing times, peaking at over 80 trips per minute. These trips are dominated by student to school trips, and thus, a kiosk should be created at the high school that is neither in the way of school buses or parents dropping their students off. Since the aTaxi would be transporting students, however, we must be careful not to drop off students too far from the school, as this could create incidents outside of the school's grounds that are outside of the school's direct supervision. Interfacing and building relationships with the school administration itself would be imperative to have a specific reserved area for aTaxi pickup and dropoff and to educate students on the use of the technology. However, we also see some unexpected behavior in Central High pixel, with demand generated later in the day that one would usually expect students to have left. This can be justified by sports going on later in the day, janitorial services, and simply other trips finishing in residential locations that surround the high school. We know that overall Trenton demand is high around 20:00-21:00, thus demand spilling into this pixel is a logical consequence.

Finally, at the Mayor Donnelly homes, we identify much less demand overall, with typical rush hour peaks. Demand between those peaks is sparse, but still present, which would allow some aTaxis to fully fill, while others may look to incorporate demand from nearby pixels and pick up additional riders along the way to a final shared destination to bolster average vehicle occupancy. A kiosk here does not need much of a physical presence as the number of cars at the kiosk, and can likely be a sign or a small

shelter. However, this lack of consistent demand should not undermine the necessary outreach to the community through informative education. Programs should be developed in conjunction with the housing authority to engage local residents in the creation of the network, and seek to learn from their needs, rather than impose select behavior. This may justify the presence of a full blown kiosk with an attendant and other services at the Donnelly Homes, such that residents can interact more freely with the system. As a result, a transportation operator should be prepared to eat the cost or seek federal subsidy in order to increase the public good of their product, rather than just sheer profitability driven metrics. It is important to dive deeper into the area, rather than just relying only on a pixel analysis in order to address issues of accessibility and equity that have impacted previous transit solutions.

In Perth Amboy, I identified similar kiosk stops of interest using the same exploratory methodology as Trenton: Perth Amboy Station, Perth Amboy High School, and the William Dunlap Homes. To place these individual kiosk locations in context, I will generate summary figures for all the trips of Perth Amboy first.




Figure: Cumulative distribution of trip lengths within Perth Amboy's ODD




Figure: Large Cap Total Addressable Market Trips in Perth Amboy




Figure: Medium Cap Total Addressable Market Trips in Perth Amboy




Figure: Small Cap Total Addressable Market Trips in Perth Amboy

The first factor to consider with Perth Amboy's data is the total quantity of trips made daily: 51,632 trips versus Trenton's 138,377 total. This discrepancy is partly due to Perth Amboy's smaller geographical size, so when creating a containing ODD there naturally will be less trips. along with Perth Amboy having about half the population of Trenton. Likewise, as we observe in the CDF of trip distances, the average trip length in Perth Amboy is less than that of Trenton due to the smaller potential region one could travel to still be contained within the city boundaries. We must still note that as before, these trip distances will be magnified individually by a factor of 1.4 or higher due to routing, as they are direct great-circle distances, and that distance can increase even more if we consider ridesharing with different drop-off points, such that passengers are dropped off before others. All that in mind, the smaller size of Perth Amboy does not mean on-demand transportation options are superfluous, rather, we hope to leverage this density for increased ridesharing potential. Turning towards our capped total addressable market analysis(TAM) we immediately identify the smaller impact of the large capped chart in relevant data analysis, as demand peaks are not as high as in the former city. It does reveal some periods of greater than 200 trips/minute around 8:00 and 16:00 which are good to know for general planning references. Moving onto the medium cap chart, we

see those demand peaks in a clearer format, and find that demand following the peaks holds at about 40-60% of the peak level for several hours afterwards. This is promising for continual service throughout the day, rather than interrupted peaks that need to be served with little in between. Finally, from the small cap chart, we see a bit of a deviation from the behavior of Trenton that had capped trips basically for every minute. There are longer periods through the night and early morning that dip below ten trips/minute which provides an indicator that it would be potentially prudent to pull all vehicles off service late at night for refueling and maintenance, or operate a skeleton crew of cars only. Autonomous service is often erroneously assumed to be needed to operate at all times to be considered successful, but when the demand is not high enough to justify the cost and usefulness, downtime incorporation is the natural choice. This ensures that during the hours one does operate, service can be as good as possible.

In selecting specific pixels for analysis, I considered similar factors to the locations in Trenton that make these stops critical “jump-off” points for a pilot service. The Perth Amboy Station serves the same role in rail transit as the Trenton Transit Station, but with less overall traffic, as it is solely connected to the North Jersey Coast line, rather than Trenton’s connections to both NY and PA, along with Amtrak lines. Regardless of its smaller offering, Perth Amboy’s stop runs trains to New York City which itself is a large vector of trips from the area along with other sizable towns in the northern New Jersey region. This behavior results in a combination of other and work trips that up 80% of the travel to and from that pixel, as seen in the summary table below where. Likewise, Perth Amboy High School is the only high school for the public system in the city, mirroring Trenton Central’s role, and 62% of trips made in that pixel are school related. To make my final pixel choice, I found that within the boundaries of Perth Amboy operate two major subsidized housing authority complexes, the Hansen Apartments and the William Dunlap Homes [30]. The former works to assist senior citizens and disabled individuals, while the latter also includes low income families along with the prior two categories [30]. Although senior citizens and disabled individuals are two key mobility disadvantaged groups that are key demographics for autonomous service, I chose to study the William Dunlap homes due to their higher diversity of transit

needs along with the similarity to Trenton's Donnelly Homes that also provide diverse housing. Along with that, senior citizen dominated housing often already has some transit services created to assist the residents, and while that transport is far from ideal, it is more than the absence of transport for low income families. The trip classifications from the William Dunlap homes pixel are split as a majority between school and other. This begs the question of if the transit needs of students is by school provided transportation, and what part could be supplanted by additional services? Likewise, since other trips often consist of travel to places that are important for day-to-day life(groceries, recreation, and outside of work responsibilities) having an on-demand service to cover these trips in a way bus transit cannot is critical.

	Number of Trips	School	Work	Other
Perth Amboy Station	14,960	11.2%	41%	47.8%
Perth Amboy High School	8,062	61.7%	5.1%	33.2%
William Dunlap Homes	2,434	42.9%	6.9%	50.2%




Figure: Medium and Small Cap Kiosk Trips to/from Perth Amboy Station Pixel




Figure: Medium and Small Cap Kiosk Trips to/from Perth Amboy High Pixel




Figure: Medium and Small Cap Kiosk Trips to/from William Dunlap Homes Pixel

Starting with Perth Amboy Station, we see surges on the medium cap chart at the already predictable 8:00 and 16:00, along with some influx at 20:00. This is reflected in the small cap chart through near constant demand from 08:00 - 22:00, with a cooling period between 14:00 - 16:00. This indicates some easement of demand which would allow for better repositioning of vehicles to account for rush hour traffic and the need to station additional vehicles in preparation for trains. The train situation presents an interesting issue, especially at Perth Amboy station where arrivals are relatively more limited to one line. This means that there will be waves of people needing transport and then lulls in between, which is different from a larger station like Trenton presenting more constant demand as seen in its charts. Observing the location of the station through satellite and Google Street view, we observe a small u-shaped drop-off facing one direction of travel and a larger parking lot on the opposite side for commuters and day commuters to park their vehicles. It is likely that the drop-off location will not be able to

support the demand of aTaxi, and direct curbside integration on the parking lot side may cause friction with commuters leaving the lot with personal vehicles.

Next, Perth Amboy High School's medium cap charts show peak demand surrounding the hours of 08:00 and 16:00 which is consistent with expected behavior from a school. However, it was a bit unexpected for this demand to reach higher peaks than the peak demand of the station, around 30-40 trips/minute larger, while only having about half the total trips of the station. We can assume that some of this demand will be filled by the school's existing bus transit system, but this may not be a viable option for all students, which we can assist with our aTaxi system. Through street view, we can determine that the streets directly surrounding the school offer little opportunity for aTaxi pickups, as they are narrow, 2-lane roadways with street parking on both sides. Buses have no parking zones for pickup, so being able to share those lanes by staggering dismissals of students could be an option. Such an idea would spread out demand peaks, allowing for a smoother situation, but requires very close collaboration with the school, as it demands an operational change from day to day procedure. An easier option would be to create a kiosk location several blocks away from the school to add some breathing room and help decongest the immediate surroundings from an influx of trips, but adds required walking distance.

Finally, we examine the William Dunlap Homes pixel, with its smaller total number of trips, about a quarter of the total school's trips and a seventh of the station's trips respectively. Regardless of the number of trips, we emphasize the same mantra as with the Donnelly Homes, quality of service must be upheld, if not exceeded for the mobility disadvantaged in order to provide real societal good with the technology. We see no real distinguishable patterns from the medium cap chart, which is common in less active pixels, and thus, turn to the small cap graph to see behavior that greatly resembles Perth Amboy's entire behavior as seen in the medium cap chart for the whole city. This is indicative that the individuals in this pixel have temporal behavior similar to the overall population, and thus, deploying vehicles must be careful to route in the homes with the same level of consideration as the more populous pixels. If demand is the same

temporally across the city in comparison to this pixel, but population is much less, we risk under-serving this community.

In summary, the process of taking a region's data and deriving general diagnostic insights, and then exploring more granularly is modular and can accept as input any given region along with pixels of interest to be analyzed by transportation planners. This can be used to confirm suspicions about certain geographic locations or discover new behaviors that are important considerations. In theory, an initial deployment of an aTaxi system will only utilize perhaps five key stops to expose the public and collect initial real-life data about performance metrics. Then, planners can continue with the simple assumption of analyzing the top ten, twenty, or fifty largest trip-generating pixels in a region to serve as an initial general ODD, and then refine based on how many they believe are needed to serve a general region. In doing so, a service can iterate through phases of service, with full connectivity of the region being the final goal.

Chapter 6: aTaxi Simulation and Results

Now that we have an initial idea of overall demand for our selected regions, as well as a location based analysis of certain ‘focus’ locations, it is relevant to simulate and study the results of full operation of an aTaxi network within our selected ODDs. In doing so, we wish to take trip data and have some initial benchmarks on how trips are completed by aTaxis. We understand that any pilot program and its operator will seek to begin service on a smaller scale, expanding to cover more and more kiosks as it proves its efficacy. However, for academic purposes this is a forward looking approach that would like to assess the entire ODD as an aspirational goal. This is an initial look at how aTaxis would perform, their ride sharing capabilities, and the distances that will be traveled by the created routes.

A step-by-step breakdown of the process of the code to accomplish this simulation is as follows.

I. Data Preparation and Constraint Identification

- An ODD is created based on the physical boundaries of the city, creating a polygon shape through utilization of Lat, Lon coordinates as vertices.
- Trips are selected out of this polygon by checking the larger New Jersey trip set file and identifying whether origins and destinations are contained within the boundaries of the polygon.
- Several constants are predefined that can be modified for sensitivity analysis, but are currently set to **five** minutes of waiting time at each kiosk, **four** person maximum in each vehicle, and discretizing pickups and dropoffs to a **0.5 mile**(pixel basis)

II. Central Assignment Process (Main assignment hot loop modified from Gabriel Laniewski’s summer work and Junior Paper)

- aTaxi objects are created that store a index, passenger count, initial start time, location of the kiosk pickup, and locations of the subsequent drop offs
- The first aTaxi is initialized at the first chronological trip of the day, and then the code iterates forward chronologically to the next trip.

- If the following trip fulfills location(pickup and dropoff) and waiting time criteria, it will be added to the generated taxi, if not, another aTaxi will be generated at the new location.
- Should any aTaxi hit passenger maximum, it will deploy and no longer be able to be assigned additional passengers, and if not it will leave at the end of the waiting time period.
- Deployed aTaxi objects are added to a list such that we have a traceable log of each departure throughout the day.

III. Visualization

- Each filled aTaxi contains all the relevant information about the trips, such that we iterate through the list of taxis to create summary graphs and statistics.
- Occupancies and departure times require no data changes, while route distances require summing the individual trip distances within the aTaxi itself

After successfully creating and writing this code schema, the first result of interest was the distribution of aTaxi departures, which is valuable even though we already graphed the trip departure time distribution in the previous chapter, as this gives some level of indication of the relevant fleet sizes that an operator would need to staff in order to satisfy total demand. Initial expectations from these figures are that the waiting time integrated into each vehicle should result in tighter temporal clustering and the y-axis scale will be smaller due to each taxi having up to four riders.




Figure: Trenton's aTaxi departure times throughout the day by minute

Having applied aTaxi clustering, we see that this visualization resolves the problem that was apparent in the original trip departure analysis, where Trenton's large afternoon rush hour peak made the rest of the data not very readable. The staggering of the vehicles along with peak times producing maximum occupancies in many pixels, allows for the scale to not be exponentially higher for key times. We see a peak of around 140 aTaxi departures per minute at 17:00, with surrounding demand being quite high, and surging again around 20:00. Likewise, it is interesting to see that the morning's demand is elevated over a longer period of time than what the trip departure data initially showed. Overall, this chart illustrates that afternoon and evening demand is the most critical, and that operational services should be prepared to handle such an influx of trips during those hours. Next, let us observe the same chart for Perth Amboy's trips.




Figure: Perth Amboy's aTaxi departure times throughout the day by minute

It is clear that from a short glance that these two departure representations are unique. Between the hours of midnight and 7:00, it takes less than 10 aTaxis at any given point, at times dropping to only one aTaxi needing to be deployed per minute to match demand. This is not a particularly financial or logical time to operate in, but, since we know that other services also do not operate frequently during these times, offering some level of aTaxi operation could improve the optics of the program. Then, comparing these departures with Trenton's chart, we see a relatively higher impact of rush hour on aTaxi deployments, far exceeding any other demand times during the day. In Trenton, we saw that the morning rush hour was subsequently followed with somewhat equal departures. This discrepancy can be partly attributed to being a consequence of the departure charts we viewed in Chapter 5, as Perth Amboy had quite large peaks in demand during these times compared to the rest of the day. Likewise, we could possibly attribute the increased number of aTaxis potentially to disaggregate demand, such that ride sharing is more limited due to different pickup and dropoff locations. To dig deeper into this claim, and further insights, I have assembled the following figures and table to summarize the performances of each system.




Figure: Trenton and Perth Amboy distributions of aTaxi departure times by hour

Clustering departures by hour provides a smoother picture of Perth Amboy's departures in that the minute-by-minute peaks have a less drastic effect on the scale of the y-axis. This behavior is mirrored for Trenton, such that the hour buckets show a better representation of continuous departures. Likewise, for Trenton we see basically symmetric behavior to Perth Amboy during the evening hours, with a bit of a more constant aTaxi utilization present throughout the off-peak travel times, which is justified by Trenton's local maxima being more similar to the surrounding data than Perth Amboy's maxima. This departure chart also reveals the possibly concerning drastic falloff between the hours of 9:00 and 10:00, which is likely somewhat of a consequence of how the synthetic data is constructed, rather than a purely natural phenomenon. Because work and school trips are queued first in the morning, synced to school start times and assumptions on work start times, that general travel ends shortly before 9:00. Trips with "other" destinations are likely the only trips that are being made during that time period, and much less of them can be taking place, because a majority of the individuals in the data file have just completed a trip to a location they will be at for the next few hours. In day to day real world operation, although we are aware that morning traffic on roads eases up after the initial rush, we would expect similar departures to the midday/noon ones during this time period.




Figure: Trenton and Perth Amboy distributions of aTaxi Occupancies




Figure: Trenton and Perth Amboy distributions of aTaxi Route Lengths

	Trenton	Perth Amboy
Total Customers	138,377	51,632
Dispatched aTaxis	61,875	16,202
Total Driven Miles(mi)	140,822	29,138
Average Occupancy(persons/aTaxi)	2.24	3.19
Average Trip Length(miles)	2.28	1.79

Real differences begin to present themselves when we check on each aTaxi's occupancy numbers. Perth Amboy has an incredible amount of ridesharing potential, as

illustrated by the large majority of aTaxis leaving fully loaded. Trenton's figures are still good, but the relatively larger size of the city, combined with a more constant demand throughout the day, even during off peak hours, likely pushes a higher amount of cars being deployed not entirely full, or single-occupancy. When we look at total miles traveled, we find evidence to support that argument, as there are about 2.3x more passengers in Trenton than Perth Amboy, but 4.8x of the miles driven in the ODD. The average length of the trips completed also corroborates that line of thinking, as Trenton passengers' trips average a great circle distance of 2.28 miles, while Perth Amboy passengers average a 1.79 mile length. Regardless, occupancy figures for both cities are quite good and provide a great indicator for ridesharing, with 2.24 passengers/aTaxi in Trenton and 3.19 passengers/aTaxi in Perth Amboy.

Overall, the goal of this chapter was to take the same data input from the previous chapter, and instead of exploring the general and minute effects of locational demand, take the entire ODD and simulate it as a connected autonomous taxi network. In doing so, we assumed that all trips would be accepted, the operator does not have a finite set of vehicles, and that service would not experience any downtime. In reality, we only want to take a fraction of the trips in the city, have a fleet of cars somewhere on the magnitude of 100, and offer service during times of best use, accounting for maintenance and downtime. But, to convince both government officials, venture capitalists, and the aTaxi operators themselves of a location's "worthiness", I believe it can be very powerful to show the above data as an image of what a fully operational network could serve, not what it will end up serving. Showing high potential for ride sharing, demand, and revenue, based on these two visualizations, is a convincing economic argument for the existence of such a system in one's town.

Chapter 7: Conclusion

Summary of Work

In this paper, we walked through many of the steps relevant to the creation of an aTaxi network in an urban area, and provided justifications to the efficacy and need of such a system. First, we utilized industry figures for autonomous vehicles, infrastructure, and labor, along with assumptions on technological costs, to find ballpark estimations of fleet costs, kiosk costs, and overall spending for an aTaxi transportation system. In doing so, we found that an 100 vehicle fleet, that could perform a range of 10,000 to 20,000 person trips a day, would cost on the order of 50 million dollars for a five year period. Assuming an operation time of about 300 of 365 calendar days, this gives us a price per trip figure of around \$1.70 - \$3.30, depending on total daily trips. This figure competes heavily with bus fares, light rail costs, personal vehicle ownership, and any other conventional mode of transportation. However, these prices can only be accessed with good levels of adoption, smart system design, and thoughtful innovation in the autonomous vehicle market. In implementing such a system, we also discuss the importance of bringing this change in transit to the mobility disadvantaged, as they can derive the most social good from an increase in transportation availability.

Following these discussions, we implemented some algorithmic methods to learn more about trip demand and the aTaxi that could serve that demand. To perform data analysis on regions and the pixels that they contain, we relied heavily on synthetic data generated from Census data: including household demographics, journey to work trends, and socio-economic data. Because of this, the discrete accuracy of our results is limited to the assumptions utilized to create the baseline of data itself, which may lead to some localized inaccuracies in some trips. Regardless, this data allows us to have a great representation of a day's entire travel demand of the United States. Using this data, we were able to develop methods to visualize demand in such a way for aTaxi system planners to use for smarter decision making in system design. We then utilized that same data operationally in our aTaxi trip modeling and assignment. Through these two chapters, we were able to observe general demand in Trenton and Perth Amboy, zoom

into their important locations(train transit, schools, and housing), and model performance of ridesharing aTaxi in both regions. This resulted in high average ridesharing numbers(people/aTaxi) of 2.24 and 3.19 in Trenton and Perth Amboy respectively, which should signal great potential for a pilot program in either of these cities.

Extensions

The next progression for this project would be attaching a graphical interface to this code, such that users would have an easier time retrieving location metrics on their own through inputting a region of interest. This could be done by allowing users to create their own ODD polygon through placement of vertex boundaries or simply by entering a city name. Upon doing that, it would be helpful for the system to automatically grab a certain preset number of pixels with the most number of trips, and generate capped charts automatically upon mouseover. Challenges with this approach include the vast amount of data that would need to be hosted, and subsequently searched to find the exact trips contained rapidly and the pixels that encompass that region. This could be partially resolved through conditional “region-loading” where the system first asks for a state or county, enters that set of data only into analysis, and then allows the user to make choices based on that. Limiting factors that approach creates is a lack of cross-state investigation(i.e. A service that investigates travel along the Camden/Philadelphia border) but since I foresee a plan like that to be bureaucratically challenging, this limitation is not as damaging.

Then, once that feature is operational, the next step would be to take the simulation code and have it execute on the selected region. This is a bit more challenging to integrate, especially if the region is relatively large, as the hot loop of this sequence is somewhat slow to compute. For reference, Trenton’s main loop took 7 minutes of computing time while Perth Amboy’s took 5 minutes of time. Generation of summary graphs is near instant in comparison so that does not need to be considered in overall computation. Thus, there has to be a sort of ‘submit’ UI where once users have decided on the parameters, they can commit to the generation, which will return results after a set period of time, rather than something that can be dynamically adjusted. In order for a

more dynamic system, I predict significant code refactoring and parallelization needs to be done. Finally, being able to attach the economic estimation derived in Chapter 3 to the regional insights derived from Chapters 5 and 6 would be useful in that those utilizing the service could get a general gauge on the overall costs for a system at generation. This in turn allows users to get an accurate picture of how much of a transit area they can serve at what cost.

Final Thoughts

As we iterate on the idea of autonomous taxis, the results are clear. The use of the nationwide trip set in a novel way helped provide some more context to the minute operation of a system. The data shows that even though it is clearly possible to set up an autonomous taxi system, we must be careful to make specialized decisions in operation based on the location. We cannot apply the same cookie cutter approach to the quantity of aTaxis, kiosk number and size, and distribution of kiosks throughout an ODD. Keeping that in mind, I believe that an aTaxi system can get off the ground and succeed with the right decision making in initial operation, the most critical time to gain public acceptance. That being said, this thesis did not attempt to tackle the legislative and regulatory issues that must be navigated, alongside the operational and economic points. The former are critical to ensure any pilot program can get off the ground and eventually pivot into full service. Funding for aTaxi systems could potentially come from a union of federal or state funding with private operators, since the service could help bolster underperforming transit projects already being funded. Looking toward the future, Waymo has begun autonomous taxi operation to a limited set of individuals in San Francisco and Arizona, and they hope to expand after successful operation. However, the lack of competitors currently pursuing deploying service and the lack of access of this technology to the mobility disadvantaged strikes a level of concern. We hope that papers like this one provide a convincing argument for further pursuit of this idea to take off on a more prevalent scale.

Bibliography

- [1] “U.S. energy-related carbon dioxide emissions fell in 2019, mainly in electric generation.” <https://www.eia.gov/todayinenergy/detail.php?id=45836> (accessed Dec. 25, 2022).
- [2] “U.S. transport GHG emission shares by source 2020,” *Statista*. <https://www.statista.com/statistics/922655/us-transportation-greenhouse-gas-emissions-by-source/> (accessed Dec. 25, 2022).
- [3] M. Nour, J. P. Chaves-Ávila, G. Magdy, and Á. Sánchez-Miralles, “Review of Positive and Negative Impacts of Electric Vehicles Charging on Electric Power Systems,” *Energies*, vol. 13, no. 18, Art. no. 18, Jan. 2020, doi: 10.3390/en13184675.
- [4] “Public Transportation’s Role in Responding to Climate Change,” U.S. Department of Transportation, Jan. 2010. Accessed: Dec. 15, 2022. [Online]. Available: <https://www.transit.dot.gov/sites/fta.dot.gov/files/docs/PublicTransportationsRoleInRespondingToClimateChange2010.pdf>
- [5] R. Buehler, “Can Public Transportation Compete with Automated and Connected Cars?,” *J. Public Transp.*, vol. 21, no. 1, pp. 7–18, Jan. 2018, doi: 10.5038/2375-0901.21.1.2.
- [6] K. Marocchini, “A National Hybrid Activity/Agent-Based Demand Model to Characterize the Mobility of the United States,” Princeton University, Jan. 2017. Accessed: Jan. 01, 2023. [Online]. Available: https://www.dropbox.com/s/jrc8w6ygs5ex6bh/Marocchini_A_NationalHybridActivity-Agent-BasedDemandModel.pdf?dl=0
- [7] admin, “ARK Contributes Its Financial Autonomous Taxi Model to GitHub,” *ARK Invest*, Mar. 29, 2019. <https://ark-invest.com/articles/analyst-research/autonomous-taxi-model/> (accessed Dec. 24, 2022).
- [8] “Model 3,” *Tesla*. <https://www.tesla.com/model3> (accessed Dec. 26, 2022).
- [9] P. M. Bösch, F. Becker, H. Becker, and K. W. Axhausen, “Cost-based analysis of autonomous mobility services,” *Transp. Policy*, vol. 64, pp. 76–91, May 2018, doi: 10.1016/j.tranpol.2017.09.005.
- [10] L. Beir, G. Giosa, R. Goldsmith, R. Johnson, L. Morgan, and D. Sollohub, “Parking Matters: Designing, Operating, and Financing Structured Parking in Smart Growth Communities.” Accessed: Dec. 28, 2022. [Online]. Available: <https://nj.gov/state/planning/assets/docs/publications/180-parking-matters-070106.pdf>
- [11] “Commercial Taxi Insurance for Drivers & Companies.” <https://tivly.com/taxi-insurance> (accessed Dec. 27, 2022).
- [12] “Self-driving vehicles could struggle to eliminate most crashes,” *IIHS-HLDI crash testing and highway safety*. <https://www.iihs.org/news/detail/self-driving-vehicles-could-struggle-to-eliminate-most-crashes> (accessed Dec. 27, 2022).
- [13] “NJ Electric Rates: Compare Energy Rates in New Jersey - EnergyBot.” <https://www.energybot.com/electricity-rates/new-jersey/> (accessed Dec. 26, 2022).
- [14] “Is the cost of maintaining an EV less than a gas-powered vehicle?,” *WXYZ 7 Action News Detroit*, Apr. 04, 2022. <https://www.wxyz.com/news/is-the-cost-of-maintaining-an-ev-less-than-a-gas-powered-vehicle> (accessed Dec. 27, 2022).

- [15] A. Kornhauser and J. He, “Perth Amboy MOVES: Safe, Equitable, Affordable, Sustainable, High-Quality Mobility for All,” Jul. 08, 2022. [Online]. Available: <https://www.dropbox.com/s/ovy49l08lqij22b/220708%20perth%20amboy%20move%20-%20CARTS%20Princeton.pptx?dl=0>
- [16] X. Shi, A. V. Moudon, P. M. Hurvitz, S. J. Mooney, C. Zhou, and B. E. Saelens, “Does improving stop amenities help increase Bus Rapid Transit ridership? Findings based on a quasi-experiment,” *Transp. Res. Interdiscip. Perspect.*, vol. 10, p. 100323, Jun. 2021, doi: 10.1016/j.trip.2021.100323.
- [17] “license-checklist-newsstand.” <https://www.nyc.gov/site/dca/businesses/license-checklist-newsstand.page> (accessed Apr. 27, 2023).
- [18] “New Jersey Unveils Solar-Powered, USB-Equipped Bus Shelter,” *Governing*, Sep. 28, 2021. <https://www.governing.com/community/new-jersey-unveils-solar-powered-usb-equipped-bus-shelter> (accessed Apr. 27, 2023).
- [19] B. Mohl, “T notes: Bus shelter maintenance pricey,” *CommonWealth Magazine*, Apr. 09, 2019. <https://commonwealthmagazine.org/transportation/t-notes-bus-shelter-maintenance-pricey/> (accessed Apr. 27, 2023).
- [20] “Manual of Traffic Signs - Costs of Traffic Signs.” <http://www.trafficsign.us/signcost.html> (accessed Apr. 27, 2023).
- [21] J. Aber, “Electric Bus Analysis for New York City Transit,” Columbia University, May 2016. [Online]. Available: <http://www.columbia.edu/~ja3041/Electric%20Bus%20Analysis%20for%20NYC%20Transit%20by%20J%20Aber%20Columbia%20University%20-%20May%202016.pdf>
- [22] “Brooklyn-Queens Waterfront: Streetcar/Light Rail Rapid Assessment,” NYC DOT. [Online]. Available: <http://www.capitalnewyork.com/sites/default/files/Brooklyn%20Queens%20Connector%20Rapid%20Assessment.pdf>
- [23] “Transit agencies are paying the price for inefficient paratransit,” *Via Transportation*. <https://ridewithvia.com/resources/articles/transit-agencies-are-paying-the-price-for-inefficient-paratransit/> (accessed Apr. 27, 2023).
- [24] “World Report on Disability,” World Health Organization, 2011. [Online]. Available: https://disabilityinclusion.msf.org/assets/files/WorldReport_eng.pdf
- [25] M. J. Field, A. M. Jette, and I. of M. (US) C. on D. in America, *Transportation Patterns and Problems of People with Disabilities*. National Academies Press (US), 2007. Accessed: Jan. 01, 2023. [Online]. Available: <https://www.ncbi.nlm.nih.gov/books/NBK11420/>
- [26] J. Gubbels, C. E. van der Put, and M. Assink, “Risk Factors for School Absenteeism and Dropout: A Meta-Analytic Review,” *J. Youth Adolesc.*, vol. 48, no. 9, pp. 1637–1667, Sep. 2019, doi: 10.1007/s10964-019-01072-5.
- [27] G. Bickenbach, “Trip Demographic Analysis of Trenton and Perth Amboy MOVES,” Princeton University, Oct. 2022.
- [28] O. US EPA, “Smart Location Mapping,” Feb. 27, 2014. <https://www.epa.gov/smartgrowth/smart-location-mapping> (accessed Apr. 08, 2023).
- [29] F. P. Boscoe, K. A. Henry, and M. S. Zdeb, “A Nationwide Comparison of Driving Distance Versus Straight-Line Distance to Hospitals,” *Prof. Geogr. J. Assoc. Am.*

Geogr., vol. 64, no. 2, p. 10.1080/00330124.2011.583586, Apr. 2012, doi:
10.1080/00330124.2011.583586.

- [30] “Home : Perth Amboy Housing Authority.” <https://perthamboyha.org/> (accessed Apr. 22, 2023).