Planning for a Shared Automated Transportation Future

by

James Fishelson

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Urban and Regional Planning)
in the University of Michigan
2018

Doctoral Committee:

Professor Jonathan Levine, Chair
Professor Joe Grengs
Assistant Research Professor Robert Hampshire
Professor Henry Liu

James Fishelson

jamesfis@umich.edu

ORCID iD: 0000-0002-9630-6500

Acknowledgements

First and foremost, I would like to thank my adviser, Jonathan Levine. He provided invaluable support and feedback for this dissertation work, helping to shape an unwieldy project into something manageable and hopefully useful to the broader community. He patiently listened to my sometimes rambling digressions, always asked the right questions (including ones I never would have considered), and made me a better writer, presenter, and scholar. I would also like to extend thanks to the rest of my committee: Joe Grengs, Robert Hampshire, and Henry Liu. They were all very generous with their time and support. Their suggestions, especially their various pieces of advice on how to make the work more accessible and policy-relevant, made my dissertation far, far better than it would have been otherwise.

I want to extend my sincere gratitude to Louis Merlin. We worked together while he was a post-doc here at Michigan; he contributed extensively to both the literature review and the preliminary versions of model. He was a wonderful person with whom to discuss various modelling options and approaches, and more importantly, he has been a very good friend. I could not have completed this dissertation without him. I would also like to thank my various professors at the University of Michigan, especially within the Urban Planning department and the Complex Systems program. My friends here at the University likewise provided a wealth of support and comfort in matters both professional and personal. I would like to single out my two roommates that I had in my house in Michigan, John Scrudato and Michael Koscielniak. Both were with me through good times and bad, and listened to me talk about my research for more hours than I care to admit.

Furthermore, I would not be here were it not for my family: my father Michael Fishelson, my mother Karen Herrmann-Fishelson, my sister Elise Kaufman, and my brother-in-law Scott Kaufman. They supported my dreams when they sometimes seemed crazy, and helped keep me going when I was down, always convincing me that I could succeed and that I have the ability to do great things. I want to thank my mother for two things in particular. First, for being my official proofreader. She caught all of my various writing errors, and moreover, she let me know when I became unclear and unprecise. I could not have asked for a better editor. Second, she returned my beloved beagle Cassie, who is also hers to a large extent. Mom knew that I needed Cassie's companionship, snorting, occasionally naughtiness, and desire for frequent walks to help me finish.

Lastly, I want to thank my girlfriend, Sabrina Peczonczyk. Though we have only been together for a short while, I will be forever grateful to you. Most of my PhD friends told me that the final push of modelling, writing, revisions, and formatting was the most stressful time of their lives. Because of you Sabrina, these past four months have been the most deliriously happy times that I can remember.

Table of Contents

Acknowledgements ii

List of Tables vi

List of Figures vii

Abstract ix

Chapter 1 — Introduction 1

Research Overview 2

Research Questions and Approach 6

Hypotheses 10

Chapter 2 — Literature Review 16

Different Modeling Approaches 18

The Invisible Influence of Modeling Choices 21

Service Characteristics 23

Context Characteristics 28

Market Characteristics 32

Performance Metrics 37

Summary of Modeling Work to Date 43

Recommendations for Future Modelling Efforts 48

Relationship between Previous SAV Models and This Dissertation 54

Conclusion 55

Literature Review Appendix 58

Chapter 3 — Model Overview 70

Model Construction 71

Model Usage 76

Overview of Model Performance 82

Benefits of a Simplified Model 86

Chapter 4 — Effects of Varying Fleet Size 93

Importance of Modelling Fleet Size ... 93
Fleet Size Model Runs ... 95
Results of Varying Fleet Size ... 96
Results of Varying the Density of Unassigned Vehicles ... 105
Conclusion ... 109
Chapter 5 — Effects of Varying Urban Density ... 112
Understanding and Modelling Urban Density ... 112
Urban Density Model Runs ... 117
Results of Varying Urban Density ... 118
Conclusion ... 130
Chapter 6 — Effects of Varying Urban Form ... 133
Importance of Varying Urban Form ... 134
Urban Form Model Runs ... 136
Results of Varying Urban Form ... 138
Conclusion ... 143
Chapter 7 — Effects of Considering Mode Choice ... 146
Importance of Modelling Mode Choice ... 147
Implementation of Mode Choice ... 149
Calibration Runs ... 154
Varying Fleet Size Runs ... 155
Varying Density without Multimodality ... 159
Varying Density with Multimodality ... 164
Conclusion ... 170
Chapter 8 — Conclusion ... 172
Core Research Findings ... 173
The Future of SAVs ... 178
Future Work ... 186
Planning for SAVs: Failures and Possibilities ... 187
Bibliography ... 196

List of Tables

Table 1: Modelling Synthesis 44

Table 2: Service Parameters for SAV Papers 58

Table 3: Context Parameters for SAV Papers.. 60

Table 4: Market Parameters for SAV Papers... 64

Table 5: Performance Metrics for SAV Papers 67

Table 6: Passenger and Vehicle Model Behavior 76

Table 7: Baseline Holistic Metric Calculation Assumptions... 81

Table 8: Model Inputs, Assumptions, and Outputs.. 83

Table 9: Comparing Previous SAV Models with Simplified Model Developed Here..... 91

Table 10: Varying Fleet Size Parameters... 95

Table 11: Density of Selected Real World Cities 115

Table 12: Varying Urban Density Simulation Runs 118

Table 13: Varying Holistic Metric Assumptions 124

Table 14: Break-Even Fleet Sizes and Densities 127

Table 15: Comparing Different Urban Density Types 138

List of Figures

Figure 1: Example of Initial Random Taxi Distribution 72

Figure 2: Daily Travel Demand Pattern 73

Figure 3: Snapshot of a Simulation Run 75

Figure 4: Distribution of Taxis in a Centralized City 78

Figure 5: Average Wait Time vs. Fleet Size 97

Figure 6: Queueing over Simulation Runs 99

Figure 7: Separating Hailing and Relocation Wait Time 101

Figure 8: Relocation Percentage vs. Fleet Size 102

Figure 9: Holistic Performance vs. Fleet Size 104

Figure 10: Relocation Percentage vs. Density of Unassigned Vehicles for High Fleet Size Runs 106

Figure 11: Relocation Percentage vs. Density of Unassigned Vehicles for Primary Runs 109

Figure 12: Fleet Size vs. Density for Hail Time Intervals 119

Figure 13: Fleet Size vs. Density for Total Wait Time Intervals 121

Figure 14: Relocation Percentage vs. Density for Hail Time Intervals 122

Figure 15: Relocation Percentage vs. Density for Total Wait Time Intervals 123

Figure 16: Holistic Performance vs. Density for the Base Case 125

Figure 17: Optimal Fleet Sizes for Different Densities 126

Figure 18: Holistic Performance vs. Density with Different Costs per Relocation Kilometer ... 129

Figure 19: Holistic Performance vs. Density with Different Profits per Revenue Kilometer 130

Figure 20: Vehicle Distribution for Standard Deviations of 10, 18, and 26 km 137

Figure 21: Base, Compact, and Centralized Cities 137

Figure 22: Standard Deviation vs. Relocation Percentage.. 139

Figure 23: SAV Performance for Base, Centralized, and Compact Cities as a Function of Fleet Size.. 141

Figure 24: Tipping Point Behavior for Base, Centralized, and Compact Cities........................ 143

Figure 25: Mode Share vs. Fleet Size for Ridesourcing 156

Figure 26: Mode Share vs. Fleet Size for SAVs.. 158

Figure 27: Mode Share vs. Density for Ridesourcing... 160

Figure 28: Mode Share vs. Density for SAVs 161

Figure 29: Effects of Ridesourcing and SAVs on Transit and PAV Mode Share 163

Figure 30: Transit vs. PAV Mode Share.. 166

Figure 31: PAV Mode Share with Ridesourcing/SAV Only 167

Figure 32: Ridesourcing, Transit, and PAV Mode Share with Multimodality.......................... 168

Figure 33: SAV, Transit, and PAV Mode Share with Multimodality 169

Abstract

Vehicle automation represents the greatest revolution in transportation since the automobile itself. However, the greater the potential impact of a new technology, the harder the future is to predict; the nature of revolutionary advances is that they are not incremental, and they will be used in far different ways than what came before them. For automated vehicles, this could include wholesale shift towards shared automated vehicles (SAVs), similar to self-driving taxis. While SAVs are not a panacea for problems of the transportation system, this dissertation is based on the premise that an SAV future is far more desirable than one dominated by privately owned and operated AVs (PAVs). Automation in concert with private ownership would encourage more sprawling development, more vehicle kilometers traveled, more congestion, more emissions, and more inequality.

This dissertation seeks to understand and model the conditions under which SAVs are more likely to succeed. It constructs a simplified and flexible agent-based model to test system performance under a wide variety of situations, including varying fleet size, urban density, and urban form. By performing sensitivity analyses on these independent variables, this model is able to identify tipping points and other non-linear performance trends. For example, when fleet sizes increase, overall wait times and relocation percentages (the time/distance the vehicle must travel while empty) decrease, albeit at sharply reducing rates. As long as fleet sizes are sufficiently large to avoid queuing, which occurs when there are more trip requests than available vehicles, further fleet size increases does not substantially improve performance. When urban density

increases, overall SAV system performance also increases at decreasing rates. SAV systems do not appear to be viable for densities much lower than approximately 500 people per km^2. However, there is not a huge increase in performance when going from medium density (e.g. Ann Arbor) to high density (e.g. Manhattan); SAV systems could work in both kinds of places. Varying urban form has a somewhat more complex relationship with SAV performance. The most important points are that is that trips to the less dense outskirts of a city are more difficult and expensive to serve than those in the city center, and that once trip distance is controlled for, the benefits from serving the denser central areas are more than outweighed by the costs of having to serve the outskirts when compared to a base city with constant density throughout.

The final modelling runs examine mode choice, comparing SAVs with transit and PAVs across different urban densities. These results suggest that SAVs obtain their greatest mode share over medium densities between approximately 500 and 4,000 people per km^2, but never more than 13% of all trips. While on their own, SAVs are not sufficient to break the dominance of private vehicles, they can act as a supplement to transit, helping to provide rides that would be otherwise too difficult or expensive for transit to serve. A combination of transit and SAVs can encourage people to go car-free, providing superior service to either of the two alone. To encourage greater SAV usage, the gold standard is public-private partnerships combining transit and SAVs, but other enabling policies include requiring data sharing from ridesourcing companies, regulations favoring sharing over private ownership (e.g. pricing private use and parking), and establishing guiding principles that include resilience, sustainability, equity, and accessibility.

Chapter 1 — Introduction

Automated vehicles are coming. A cherished dream of automakers and the general public for at least 75 years, the idea has made massive strides in the past decade (Beiker, 2014). Driverless "Google Cars" have been operating in mixed traffic since 2013, and every major auto manufacturer has an automated vehicle (AV) work stream. There are wide-spread efforts coming from the tech sector: Google, Uber, Apple, Tesla, and too many startups to name. The interest of cities and potential consumers is surging as well; the huge number of AV-focused proposals for the Smart Cities Challenge, awarded to Columbus in June 2016, the exploding number and size of AV-conferences, and the near-constant coverage in the media are only some examples. Promised benefits from automation include, but are not limited to: improved vehicle safety (Lutin & Kornhauser, 2014), increased roadway capacity (Carbaugh, Godbole, & Sengupta, 1998), higher roadway speeds (Shladover et al., 1991), and drivers' freedom to multitask (Wachs, 2010). Yet the forms automated vehicles will take, and how they will operate, is far from certain.

The owner-operator model of private vehicles has dominated the past century, and is one option for the future. Owning one's vehicle is certainly a reasonable approach with the current technology, since a traditional car does someone very little good if it is not parked nearby. Almost all trips have a walking component, even if it is just a few steps from garage to the front door. However, this walking step can be eliminated by fully automated vehicles, those that can drive themselves even without an occupant (i.e. Society of Automotive Engineers Level 5 automation). Instead of you going to your car, the can come to you. The car does not need to be in your garage, and it does not even have to be yours. For a car to come to you currently requires another driver, typically paid, as with a taxi or Lyft. By providing door-to-door service without a driver, shared automated vehicles (SAVs) represent a decoupling of the ownership of cars and the transportation they provide. There will always be a mix of private and shared vehicles, but this work argues that a greater proportion of shared is desirable. More sharing would help increase accessibility and reduce emissions, vehicle-kilometers traveled (VKT), and total transportation costs. Sharing has the potential to counteract some of the negative effects of private AV domination, such as increased sprawl and endlessly circulating empty vehicles waiting for their wards.

Much of the existing AV work focuses on the potential benefits, but this dissertation does not assume prima facie that an automated world will be better than our current one. How automated vehicles are used will matter far more than when they arrive, since as Kranzberg (1986) wrote, "technology is neither good nor bad; nor is it neutral." As automated vehicles creep closer to the roadway, planners have a responsibility to understand and identify how SAVs will fit into the new transportation environment, moving from the question of "if" SAVs can work to "how" they will work. This dissertation explores the conditions under which SAVs are most likely to succeed, seeking to provide practical and actionable knowledge to planners, governments, and private companies on what types of SAV systems can effectively be deployed and where. Ideally, this work will help encourage a future that tilts towards more and better sharing of our transportation system.

RESEARCH OVERVIEW

> To every man is given the *key* to the gates of *heaven*. The *same key* opens the gates of *hell*. And so it is with science. In a way it is a *key* to the gates of

> *heaven*, and the *same key* opens the gates of *hell*, and we do not have any instructions as to which is which gate (Feynman, 1998).

Vehicle automation is a fundamentally disruptive technology, promising the biggest transformation in transportation since the automobile itself. However, a promise is a not a guarantee. Predicting the deployment of such a revolutionary technology is a difficult task, since the bigger the potential impacts of a new technology, the greater the uncertainty of those impacts (Rogers, 2003). A meta-analysis by the National Renewable Energy Lab shows this uncertainty, showing a range of studies that estimate automation could have anywhere from an approximately 50% decrease to a 200% increase in energy consumption; sharing tends to lower these impacts (Brown, Repac, & Gonder, 2013). The ultimate goal of this dissertation is to address and narrow this uncertainty, helping to guide the deployment of automated technologies in ways that will be more beneficial for society. In light of the uncertainty, it takes a broad view of automation, positing that higher degrees of sharing will be more beneficial and seeking to identify the conditions and types of cities that are most amenable to SAVs. Though she did not explicitly reference Feynman, Robin Chase (2014) of ZipCar termed the two extremes of shared and private ownership of automated vehicles as the "heaven or hell" scenario. If automation is a key to both heaven and hell, this work hopes to contribute to the search for the doors to the good place.

This dissertation selects a few key issues of analysis, which have both been relatively unstudied to date and are key to understanding the necessary conditions for successful SAV adoption. In particular, it examines fleet size, urban density, urban form, and mode choice and the relationships between them: e.g. how SAV performance varies as a function of fleet size of different urban densities. The fleet size analyses help provide an in-depth view of SAV performance, including the effects of variations in travel demand throughout the day and the tension in determining optimal fleet sizes, where larger fleet sizes can increase some aspects of performance (e.g. passenger wait times) but decrease others (e.g. the utilization of the vehicles). The investigations into urban density and urban form are linked, showing how SAV performance varies across different types of cities, with respect to both their absolute density and how people are distributed throughout the city. The mode choice work borrows from the evaluations of these three issues and the work estimates how many trips will be taken by SAVs as opposed to other modes of travel across a limited number of types of cities and assumptions of AV costs.

The core of this dissertation's methodology is a flexible agent-based model (ABM) coded in NetLogo, which simulates the operation of a single-occupancy SAV system (i.e. automated taxis), providing the equivalent of door-to-door taxi service anywhere in the service area. It is an intentionally simplified and flexible model to allow for more rapid simulation and thus detailed sensitivity analyses; it thus sacrifices complexity and detailed specificity of a given city environment for broader applicability, examining SAV performance across a wide spectrum of different conditions and types of cities. Past work has tended to examine a single city environment, such as Austin (D. J. Fagnant & Kockelman, 2014), Zurich (Boesch, Ciari, & Axhausen, 2016), Lisbon (Martinez & Viegas, 2016), Singapore (Spieser et al., 2014), Seattle (Childress, Nichols, Charlton, & Coe, 2015), and all of New Jersey (Brownell, 2013). These papers typically ask some version of the question, "How would the performance of an SAV system compare with current transportation options?" By comparison, this dissertation takes a more urban planning centered approach, asking "Under what conditions will SAVs be successful compared to both current and future transportation models?"

This dissertation is structured as follows. Chapter 1 includes the research questions and hypotheses. Chapter 2 is the literature review, which organizes, structures, and summarizes previous modelling work on SAVs and related topics. These previous efforts provide the basis and justification for the research presented in the following chapters, especially the use of a broad, generalizable model that can perform sensitivity analyses on the effects of varying fleet size, urban density, and urban form. Chapter 3 describes the construction of the agent-based model, including independent variables, fixed inputs, assumptions, outputs, and performance metrics. It also compares the performance of this model to that of the more detailed and specific models described in Chapter 2 as a form of proto-calibration, showing that the simpler model used here provides equivalent results. Chapter 4 provides and analyzes the model results on the effects of varying fleet size, including an explanation of the importance of avoiding queueing, the situation that arises when there are more passenger requests than available vehicles. Chapter 5 provides and analyzes the model results on the effects of varying a wide range of urban densities, from effectively rural areas to city centers. It shows that denser places do allow for better performance – lower wait times, smaller relative fleet sizes, and fewer empty relocation distances driven by the vehicles – but with sharply reducing marginal returns. Chapter 6 provides and analyzes the model results on the effects of varying urban form, examining how the relative degree of urban centrality affects system performance. It show how SAV performance differs in a city with a true dense center and declining densities from that point, as with the Burgess (2008) model of cites of concentric zones vs. a more center-less city with relatively consistent density throughout. Chapter 7 modifies the model developed in Chapter 3 to include a mode choice component, including the unique contribution of comparing shared automation with privately owned and operated automated vehicles, as opposed to SAVs against the status quo as a few other studies have done. It also presents these mode choice results, showing that SAVs struggle as an attractive complete replacement of private vehicles, but that they become more attractive to a wide group of users and cities if people use a combination of SAVs and public transit for their transportation needs. Chapter 8 is the conclusion; it provides a summary of the findings and makes recommendations for both policy makers and future research.

RESEARCH QUESTIONS AND APPROACH

The overall goal of this dissertation is to identify the conditions under which SAV adoption will be more or less likely, including which city forms are more amenable to SAVs, how SAVs compare with privately owned and operated AVs, and how SAVs can viably function in early stages when adoption rates are quite low. Simply showing that the system is desirable with near-full adoption of a new technology is not enough; there must be a reasonable deployment pathway, including a means to attract early adopters. This is shown by past work on the general diffusion of innovations (Christensen, 2013; Rogers, 2003) and for automated and electric vehicles in particular (Fishelson, Freckleton, & Heaslip, 2013; Zoepf & Heywood, 2012). This section describes the overall research approach taken by this dissertation. It first details four research questions that seek to shed light on viable deployment paths for shared automation. It then details the importance of identifying tipping points and performing sensitivity analyses, the value of using a simple and flexible ABM to perform such analyses, the SAV performance metrics used in this dissertation and how they relate to each other, and the planning implications of the research findings. The following section details the hypotheses that correspond to each research question.

Research Questions

This dissertation poses the following four research questions each with a corresponding hypothesis and chapter. By evaluating each question via the developed ABM, this dissertation seeks to show the various conditions under which SAV adoption is more or less likely.

1) ***What are the effects of varying fleet size on the performance of SAV systems (Chapter 4)?***
2) ***What are the effects of varying urban density on the performance of SAV systems (Chapter 5)?***
3) ***What are the effects of varying urban form on the performance of SAV systems (Chapter 6)?***
4) ***How does the likelihood of people choosing SAVs over other modes of travel vary as a function of urban form and density (Chapter 7)?***

Tipping Points and Sensitivity Analyses

When a small change in input values produces a large shift in output values over a given range of input values, a "tipping point" exists over this range. The narrower the range and the larger the shift in output values, the more extreme the tipping point behavior. For example, if high fleet sizes all yield fairly low wait times, but wait times start to skyrocket once fleet size falls below a given value – that value is the tipping point. Treating fleet size as an independent variable is the focus of the first research question and its corresponding hypothesis.

Chapter 3 describes this tipping point phenomenon in more detail. However, two major points deserve mention here. First, when evaluating new and transformative technologies like SAVs, identifying performance trends in general (i.e. how performance varies as a function of different input values) and tipping points in particular, is very useful. Precise point predictions on costs, minimum required fleet size, and more may be impossible, but identifying general trends can help guide the development and deployment of new technologies in positive directions. For example, with respect to fleet size, if a tipping point exists so that below a given number of vehicles passenger wait times skyrocket, then any operator must be careful to ensure a fleet size greater than this point. By comparison, if the work shows that a tipping point exists for density, but at a relatively low level (i.e. less than 500 people per km^2), then an operator should not be overly concerned about putting an SAV system in a medium density area (e.g. Ann Arbor at 1,500 people per km^2). Such a medium density place might have slightly worse performance than a denser city (e.g. San Francisco at 6,200 people per km^2), but since the tipping point occurs at a lower density level, the difference in SAV performance in these two cities would not be that great. The second major point is that these tipping points can only be effectively identified via a sensitivity analysis that utilizes a parameter sweep, such as testing performance across a wide range of narrowly spaced different densities. Chapter 2 details other SAV models to date, but none of them have undertaken such an analysis as performed here.

Simplified Agent-Based Model (ABM)

Sensitivity analyses are needed to identify tipping points, but they require large number simulation runs to perform the parameter sweeps. For example, one of the goals of Chapter 5 is to identify the fleet size that provides an equivalent level of SAV performance for different urban densities. This requires varying both fleet size and density, over 10,000 simulation runs in total. A simplified model is necessary to run these simulations in a reasonable time frame. The more complicated the model, such as by moving from single-occupancy trips to ridesharing, the higher the computational load and the harder it becomes to perform parameter sweeps. Even with the simplified model used in this dissertation, the above simulation runs took over 3 days of computational time. Additionally developing a simplified ABM allows for the testing of a wider variety of city types, varying both the magnitude of density and the distribution of people throughout the city (i.e. urban form). Chapter 3 provides greater detail on the construction of this ABM and the justifications for making it simple and flexible.

Performance Metrics

The first three research questions imply performing sensitivity analyses on three independent variables — fleet size, urban density, and urban form — and how varying them affects SAV performance. The fourth question extends the core ABM to consider mode choice and see how that affects SAV performance as a function of density. This work considers three primary performance metrics: wait times, relocation percentage (how often the vehicle is traveling empty), and a holistic performance metric denominated in dollar terms. Wait times refer to how long the passenger is waiting for a vehicle to pick them up, and relocation percentage to the additional distance a vehicle must travel while empty and driving to pick up the passenger (i.e. a relocation percentage of 10% means that for every 100 km a vehicle travels while carrying a passenger, it travels 10 km while empty). These first two performance metrics are direct outputs of the ABM. Notably, they affect the users of the system (i.e. passengers) and the system operator differently. Passengers obviously care about how long they have to wait, and the system operator would prefer a lower relocation percentage, as they would not be able to generate revenue during those times.

Note that higher fleet sizes should correspond to lower wait time and relocation percentages. There are more empty vehicles, so it is more likely that a passenger will have to wait less and the vehicle travel while empty less before a pick-up can occur. Therefore, fleet size can also be treated as a form of a performance metric, such as the aforementioned fleet size necessary to maintain a given performance level (e.g. average wait times less than one minute) for a given density. The third, holistic performance metric — of which there are two forms, called "profit-maximizing" and "societally optimal" — provides a mechanism to consider fleet size by giving a value in dollar terms per vehicle in the fleet. The profit-maximizing version considers the value only from the system operator's perspective, including the daily fixed costs of vehicle operation, the costs of empty vehicle relocation, and revenue generated from ferrying passengers. The societally optimal version is the same as the profit-maximizing, with the addition of considering a passenger's wait time costs. These metrics are not intended to incorporate all costs and benefits of an SAV system, but by using two versions, they can help show how different decisions could be made when considering only the system operator's perspective (i.e. with a privately owned and operated system) and when the user's needs are directly considered (i.e. with heavier public sector involvement).

HYPOTHESES

The following four hypotheses correspond with the above research questions, and each is tested by the agent-based-model. The hypothesis testing pays attention to any larger trends that emerge and what could explain them.

Hypothesis #1: Fleet Size

Hypothesis #1 states that "***Performance increases at a decreasing rate with fleet size.***" This hypothesis is tested in Chapter 4.

Past work has typically taken a satisficing approach to fleet size, such as determining the necessary number of vehicles to avoid any passenger waiting more than a set maximum time, but a number of efforts also tested fleet sizes both higher and lower than the satisficed values (Boesch et al., 2016; D. Fagnant, 2014; Zhang, Spieser, Frazzoli, & Pavone, 2015). These models find that increasing the fleet size from a given base level showed only slight improvements in performance, such as marginally reduced averaged wait times. However, reductions in fleet size from the base caused very sizable drops in performance, such as average waits that exceeded 30 minutes. This dissertation explores these patterns in a more rigorous fashion at different starting trip densities using the three main performance metrics.

This dissertation hypothesizes that this non-linear pattern of performance changes as a function of fleet size, as found by past work, is robust and a fundamental feature of SAV systems. Furthermore, it proposes that this pattern is explained by queueing and its potential to create very large wait times. Queues form when there are more passenger requests than available vehicles, so that people must wait to be assigned to a vehicle until they reach the head of the queue. If queueing exists, even small increases in fleet size cause large improvements in performance by reducing the extent of queuing. However, once the fleet size is sufficiently large to prevent queuing, then further increases in fleet size will only bring slight improvements by reducing the average relocation distance and time: how far the empty vehicle has to travel to the waiting passenger.

Hypothesis #2: Urban Density

Hypothesis #2 states that, "***Performance increases at a decreasing rate with urban density.***" This hypothesis is tested in Chapter 5.

That a taxi-like system performs better in denser than in less-dense areas seems self-evident. A number of SAV models have borne this out, albeit across a fairly narrow spread of densities, such as both double and half of the base density level (Boesch et al., 2016; Burns, Jordan, & Scarborough, 2013; Chen, 2015; Levin & Boyles, 2015). The patterns of this improved performance at higher densities remains unclear, and this dissertation hypothesizes that higher densities bring improved performance but with decreasing marginal returns. If this hypothesis holds, then starting from a low density, even a small increase in density will yield a large performance increase. Yet if starting from a relatively high density, the gains will be substantially less. This pattern may be asymptotic or nearly so (i.e. there is a given “maximum” performance reached at an infinite density level), but there is not enough information to claim as such in this hypothesis.

Assuming this hypothesis is supported, this work also investigates if tipping points exist, which would be indicated by more extreme decreases in marginal returns for SAV performance as a function of density. This work expects that a tipping point does exist, more specifically that performance is relatively binary with respect to density – very low performance at low densities, but relatively constant performance with only minor increases once above a given minimum viable density. If a tipping point does not exist, then performance will increase in a non-trivial manner for all densities considered here. The more extreme the decreasing marginal returns, the more likely a tipping point is to emerge.

Hypothesis #3: Urban Form

Hypothesis #3 states that, “***More compact cities lead to higher SAV performance, and more centralized cities lead to lower SAV performance, when compared to cities with constant density throughout.***” This hypothesis is tested in Chapter 6.

While urban density indicates how many people live in a given area, urban form as used here looks at how people are distributed throughout the city. This work treats degree of urban centrality as the independent variable that measures differences in urban form. Higher degrees of centrality mean that more people are located in the center of the city than at the outskirts, and lower degrees of centrality mean that people are more randomly distributed throughout the city. This work looks at two broad types of cities with higher degrees of centrality: compact cities and centralized cities. These are compared with cities with no centrality, so people are evenly distributed throughout; such constant density cities are addressed by the previous two hypotheses. Compact cities have higher degrees of urban centrality and the same city size and number of people, with the result that they have the same density but lower average trip lengths as compared to a base city with constant density throughout. Centralized cities have higher degrees of urban centrality and are larger, so that their average density is lower but their average trip length is the same as compared to the base city with constant density throughout.

This work hypothesizes that more compact cities will have better SAV performance, most notably smaller fleet sizes to provide equivalent levels of service to the users. This is because the of the shorter average trip lengths of more compact cities, where more people start and end their trips closer to the city center. However, it hypothesizes that if these differences in trip length are controlled for, as with centralized cities, then SAV performance will be lower. This is because the benefits gained from providing the shorter trips in denser areas are overwhelmed from the costs of providing longer trips in the lower density city outskirts.

Hypothesis #4: Mode Choice

The fourth and final hypothesis deals with issues around mode choice, and is split into two constituent parts: Hypothesis #4a and #4b, both of with are tested in Chapter 7. These sub-hypotheses compare the performance of SAVs against transit, traditional private vehicles, and privately owned AVs (PAVs) under different sets of conditions. Comparing SAVs with PAVs is especially important; if SAVs are to successfully reach a non-trivial share of all trips, they must compete not just with the existing modes of travel, but those likely to exist in the future. If self-driving technologies continue to advance and eventually reach the roadway, both shared and privately owned AVs will likely coexist. SAVs could be a superior option to all existing modes, but if they pale in comparison to PAVs, then the supremacy of the private ownership model is likely to predominate. This is a unique extension; most previous SAV modelling efforts avoid mode choice, and those that do consider it have avoided comparisons with PAVs. As a result, they have been fairly optimistic on the potential of SAVs to carry a sizable number of total trips, e.g. over 40% (Childress et al., 2015; Gucwa, 2014; Horl, Erath, & Axhausen, 2016; International Transport Forum, 2016).

Hypothesis #4a states that ***"SAVs shows the greatest improvement over other travel modes in middle density areas, from 2,000 to 8,000 people/km^2.***" This hypothesis looks at the mode share of SAVs compared to these other modes under different densities, claiming that SAVs show the greatest improvements at medium density levels, from 2,000 to 8,000 people per km^2. This work expects that at lower densities, the poor absolute performance of SAVs means that private vehicle usage will predominate. At the other end, for higher densities, Hypothesis #2 holds that there is not a massive improvement for SAVs when moving from medium to high densities. If this is the case, then traditional public transit might predominate in the highest density areas, where heavy overlapping of trips supports high-occupancy transit like BRT and metros. Importantly, this hypothesis makes a relative claim with respect to density; it does not argue the conditions under which SAVs predominate over private-vehicles, which is the focus of the Hypothesis #4b.

Hypothesis #4b states that, "***If SAV passengers rely exclusively on that mode of travel, privately owned AVs are likely to predominate. However, if people use SAVs to supplement transit trips, then SAVs will become more popular, lowering the feasible density for people to go car-free.***" This hypothesis looks at the conditions under which SAVs could encourage people to go car-free, examining if a SAV-transit combination can be more attractive to users than either of the two on their own. An important aspect of mode choice is the different payment structures for the different modes of travel. Private vehicle ownership has relatively high fixed costs but low variable costs; a person must purchase a vehicle, but once they own it, it is relatively inexpensive to use. By comparison, both SAVs and transit have low/no fixed costs but high variable costs; a person does not need to purchase a vehicle, but they typically have to pay per-use of the transit system. This hypothesis argues that the high variable costs of SAVs are a deterrent for someone relying on them for all trips. As such in a competition between only SAVs and only PAVs, the PAVs are likely to win. However, as argued in this hypothesis, SAVs in combination with public transit can compete with PAVs. SAVs effectively increase the number of people that can feasibly go car-free; they may use transit for most trips, and take SAVs for the trips where traditional transit is not feasible.

Chapter 2 — Literature Review

The past five years have seen an explosion of interest in transportation modes that combine shared mobility and vehicle automation. Neither are strictly new; dreams of automation stretch back to the 1939 World's Fair (Beiker, 2014) and for shared mobility, taxis have existed almost since the automobile (Hodges, 2009), jitneys serve passengers in developing countries with semi-fixed routes (Clayborne, 2012), and carsharing has made huge strides over the past two decades (S. Shaheen & Cohen, 2007). However, adding automation could make such services far more affordable and convenient. Research suggests that automated taxis will have lower costs per mile compared with human-driven taxis, largely enabled by avoiding labor costs, and some authors have even found lower costs than privately owned vehicles (Burns, Jordan, & Scarborough, 2013; Hars, 2015). In congested urban conditions, using an automated vehicle could also obviate the need to park one's vehicle (W. Zhang, Guhathakurta, & Ross, 2017). Ultimately, the services combining shared mobility and vehicle automation, could allow people to avoid the hassle and cost of private vehicle ownership while maintaining the timeliness, freedom, and flexibility of point-to-point travel on demand.

These new modes occupy a spectrum between traditional public transit and privately owned vehicles. Like transit, they would be collectively used by the travelling population over the course of a day. Yet they would also be assigned to meet passenger demand flexibly in time and space, as opposed to running along fully fixed routes and schedules as with traditional transit. Various terms have been used to describe the different services – robot taxis, automated mobility on demand, automated shuttles, etc. This dissertation uses the most common one, shared automated vehicles (SAVs), as an umbrella to cover the wide range of potential services, categorizing them in three groups, ranging from most private-vehicle-like to most transit-like: automated taxis, automated shared taxis, and automated shuttles.

A small but rapidly expanding literature has emerged that seeks to understand the implications of shared automated vehicle systems. In the absence of empirical data on how such systems perform and how much market share they will attain, the best way to understand these hypothetical new modes is via modeling and simulation. By necessity, such models must extrapolate and hypothesize ways these new modes will interact with the existing transportation system. As the field has grown, a variety of research methods and questions have been introduced, and different research efforts are limited by their particular context, modeling methods, and its travel behavior assumptions. These differences make it difficult to generalize about SAV systems; notably, many models report starkly different findings. Our goal here is to understand and synthesize this literature as a whole by creating a common framework. This framework is used to place and compare the various studies, identifying where possible general and broadly applicable trends and findings.

This literature review chapter helps examine the implications and operations of shared automated vehicle systems, including the current state of the research, the need for standardization in reporting, and the current gaps in the research. It reviews 34 different papers, with a publication date cutoff of March 2017. This includes multiple related but functionally different papers from the same group (e.g. UT Austin); such papers are grouped together in the appendix summary tables. The review is organized as follows. First, it describes three general modeling approaches. Next follows a discussion of modeling choices or assumptions, including how these choices can interact with one another. It then reviews the different papers' modelling approaches and findings, splitting them into service, context, market, and performance categories. The penultimate section discussed previous modelling efforts and findings specifically with respect to the research questions addressed by this dissertation, and the conclusion entails a broader overview of the ramifications of the various modeling efforts, their strengths and weaknesses, and areas of need for future work. The Appendix includes summary tables for each of the major categories and we encourage the reader to use them as a reference.

DIFFERENT MODELING APPROACHES

This section examines the major groups of models that have been applied to shared automated vehicle systems, proceeding from most to least detailed.

Agent-Based Models (ABMs)

Agent-based models describe the behavior of individual agents in a simulation environment, offering the greatest level of descriptive detail for potential new modes (Railsback & Grimm, 2011). All SAV models include at least two agent "classes," vehicles and travelers, operating according to a given set of rules with the context/environment that the modeler has created. For example, the model might be seeded with travelers with set origins and destinations. Vehicles might follow a rule to travel to the nearest traveler to pick-up, follow the fastest route to the destination, and then repeat. Time is explicitly modeled via steps. ABMs allow for a great deal of flexibility in setting agent behaviors, which is a large part of their attraction. This flexibility is useful for testing a range of hypothetical situations, such as various different service types (International Transport Forum, 2016), service areas (D. J. Fagnant & Kockelman, 2014), or vehicle relocation and staging strategies (Chen, 2015).

ABMs are especially well-suited to modeling detailed interactions between the traveler and the vehicle, such as simulating pick-up, drop-off, and ridesharing behavior. The modeler can set the agent rules, and see how that will change overall system behavior and performance. Note though, that this inclusion of two distinct agent types is somewhat novel for transportation modeling; traditionally only the traveler is modeled and vehicles do not move autonomously from travelers. Additionally, because behaviors are described and decided at the level of individual agents, optimization with ABMs is difficult if not impossible. For example, ABMs can be effective in comparing the different ways vehicles might relocate when not serving a passenger, but methodologically they cannot identify the *best* such approach.

The ABMs reviewed here typically assume simple, rule-based travel behaviors. For example, travel mode is typically assigned exogenously, and it is generally assumed that all travelers have the same behavior. This is a reasonable choice, but not an inherent limitation of the ABM model type; more complex traveler behaviors could be modeled within this framework. Though abstract ABMs are possible, and are in fact common in the larger complex systems literature (Railsback & Grimm, 2011), the ABMs reviewed here are fairly specialized. Each examines a single city and includes a high level of spatial detail, generally trying to mimic the details of the existing transportation system.

Network Assignment Models

Network assignment models have at their core a transportation network made of links and nodes; this can be an actual street network, but they are typically at a lower resolution, such as each node being a transportation analysis zone representing a few thousand people. The primary goal of these models is to assign vehicles across the network to serve passenger demand, thus estimating the flow of vehicles and people across the transportation network. In traditional transportation engineering, both macroscopic network models like CUBE and activity-based models like MatSim fall under this category; the difference is that activity-based models examine and generate travel demand at much greater detail, down to the individual household level. As with ABMs, network assignment models for SAVs must consider both the vehicle trip to the passenger and the trips while taking the passenger. Traditional taxi models take this sort of network assignment approach, with the heaviest focus on the taxi traveling to the passenger (Yang & Wong, 1998).

Network assignment models use algorithmic approaches to reach their solutions, as opposed to the time-step simulation approach of ABMs. They are very strong at optimization, but generally have less flexibility in modeling a variety of service types and are less exploratory than ABMs. For example, network assignment models can be used to optimize a given aspect of system performance, such as the optimal way to relocate vehicles when not in use (Spieser et al., 2014). Many researchers (e.g. Brownell, 2013; Martinez, Correia, & Viegas, 2015; Spieser et al., 2014) adopt a hybrid modeling approach, utilizing a network assignment model to first generate demand and then assign it throughout the network, which yields a detailed, spatial origin-destination matrix and roadway speeds. In the next step, they employ another model, usually an ABM, to estimate vehicle movements and passenger interactions.

Aggregate Models

Aggregate models use combinations of raw data, assumptions, and equations or assumed relationships, to deterministically estimate system-level performance metrics: costs, time, and more. They include so-called “Excel models,” as they can be implemented in spreadsheets. For example, Greenblatt & Saxena (2015) use existing taxi data as a starting point, adding assumptions about the performance of electric SAVs to estimate the national effects on greenhouse gas emissions. Princeton researchers (Zachariah, Gao, Kornhauser, & Mufti, 2014) take a more detailed and hybrid approach, using a network assignment model to develop a time-based trip schedule for all passengers, and then using this schedule as an input for an aggregate model that estimates how many trips could be shared under different assumptions. Overall, aggregate approaches allow for studying larger areas (i.e. entire states or counties) and easier comparisons between different cities, but they are highly dependent upon initial assumptions, which may limit the ability to model new and unexpected system behaviors.

THE INVISIBLE INFLUENCE OF MODELING CHOICES

Models can provide useful pictures of hypothetical transportation system performance, but the ultimate credibility and generalizability of any model depends on the reliability of its assumptions. This is true for transportation models in general, but the SAV models present special challenges. SAV services are effectively nonexistent, so there is no real-world passenger or trip data for the models to use. A limited amount of taxi data is available, notably for New York City, but Uber, Lyft, and other privately operated shared mobility services fiercely protect the privacy of their data (Henao, 2017). A wide number of modeling assumptions and choices must be made, defining not just the type of service and its operating parameters, but the market served by the system and the context in which it operates, often with little to no empirical basis. These choices can and do have wide-reaching implications for the models' findings.

Very detailed simulations can be developed, but the higher the detail, the more input data and assumptions that is required. The easiest path therefore is to use current transportation system as a base, such as how Fagnant (2014) and the International Transport Forum (2016) assume that SAVs will operate in an environment with the same land use patterns, roadway networks, travel speeds, and alternative competing modes as those today. However, in an automated vehicle future, these characteristics can and likely will change in significant ways. Additionally, the lack of real-world data makes calibration of model parameters exceedingly difficult, especially with respect to sharing-specific operations or demand estimates: e.g. how long does boarding and alighting take in a shared taxi, and how far will people walk to a pickup point? The absence of calibration is not a fatal flaw for the viability of these models, but it does mean that real-world performance could vary significantly from the results put forth in simulation models. Ideally, models should address some of this uncertainty through a sensitivity analysis of key assumptions (Sterman, 1991).

To help frame the different modeling assumptions around SAVs, this chapter divides them into three general modeling choice categories — service, context, and market — and a fourth category for performance metrics. The first three deal with true choices (i.e. inputs to the models and descriptions of how it will perform), and performance category deals with what kinds of model outputs are monitored and reported. The definitions for the four categories, each of which has its own summary table in the appendix, are as follows:

- **Service:** The set of the vehicles – and their attendant behaviors – that comprise the system. This includes assumptions like vehicle size/number, their response to passenger demand, and how/whether rides are shared (ridesharing).

- **Context:** The environment in which the system will operate, including the geographic area of service, the population being served, the infrastructure (i.e. streets) over which the service is provided, the presence of other modes of travel, and any policy context that may influence operations.

- **Market:** The set of travelers – and their attendant behaviors – that are served by the system. This includes assumptions such as total number of requested trips, spatial and temporal travel demand patterns, and traveler mode choice behavior – for example, whether travelers would chose to take an automated taxi over other options).

- **Performance:** The different metrics describing the operation of the system, including vehicle miles traveled, level of sharing, environmental effects, system costs, and more.

SERVICE CHARACTERISTICS

This section provides overviews of the different service characteristics; the corresponding service table in the appendix more extensively characterizes the modelling choices made by each of the papers reviewed here. The following sections on SAV context, market, and performance take the same approach. These choices and metrics for each of these categories are defined broadly in order to ease comparison and generalization across studies and to help provide external validity. Many of the papers also include more specific choices to answer specific research questions, with the performance metrics showing the broadest range.

Service Type:

There are three basic service types for SAVs: automated taxis, automated shared taxis, and automated shuttles. SAVs, like any mobility service, must perform a few basic tasks; drive to the passenger (or have the passenger come to it), pick them up, transport them, and drop them off, but each of these service types can provide them in somewhat different ways. In general, the higher the degree of sharing, the more transit-like the service, and these types are most easily distinguished by passenger occupancy. Automated taxis only serve one request at a time, though there may be multiple passengers per request, and where specified have passenger occupancies of 1-4. Automated shared taxis are the same, with the exception that they can service multiple trip requests at once, typically with different origins and/or destinations. Automated shuttles are larger vehicles, with passenger occupancies greater than 4. Some studies allow for combinations of different services, such as International Transport Forum (2016), which has both shuttles and shared taxis operating at the same time. In general, the higher the degree of sharing, the more transit-like the service; e.g. shuttles are more likely to have some form of fixed routing and are more likely to have a person walk to a pickup point.

Most models to date assume taxi or shared-taxi services with dynamic, real-time booking; as a rule these mimic traditional taxis, Uber, Lyft, and/or their shared versions (e.g. UberPool). For the most part these are door-to-door "go-anywhere" services within the defined service area. A notable exception is the Princeton work, which explicitly assumes a taxi stand network, all of which are a set distance apart (i.e. 0.8 km), and which require passengers to walk to them (Brownell, 2013; Ford, 2012; Zachariah et al., 2014). The MIT group models shared taxis in a sparse network with stations spaced relatively far apart, but how passengers arrive at the nodes is not an explicit focus of their work (Spieser et al., 2014; R. Zhang, Spieser, Frazzoli, & Pavone, 2015).

Efficient door-to-door service is harder to execute the more sharing the systems hopes to achieve, higher sharing is more likely to correspond with some form of fixed routing and/or requirements for person walk to a pickup point. One method used by the International Transport Forum's (ITF, 2016) automated shuttles is for the model to require passengers to walk to/from assigned pick-up/drop-off point along the shuttle's route; these routes are determined on an ad-hoc, on-demand basis. Moving closer to traditional transit, three models consider low speed operation, which is often cited as a reasonable first use case for automated vehicles. All three have fixed routes, with the vehicles operating both in mixed traffic on city streets (Levine, Zellner, Shiftan, Arquero de Alarcon, & Diffenderfer, 2013; Winter, Cats, Correia, & van Arem, 2016) and with some degree of grade separation (Shepherd & Muir, 2011). Notably, boarding and alighting times play a large role in determining the effectiveness of shared systems and increases in overall travel times (Walker, 2011), but few models explicitly include them outside of a few exceptions (e.g. Winter, Cats, Correia, & van Arem, 2016).

Two studies have fairly unique service types that bear mentioning. Levin & Boyles (2015) assume privately owned AVs that either return to the user's home after completing a trip or stay at the destination and pay a parking fee. A following paper builds on this model with more standard automated shared taxis (Levin, Li, Boyles, & Kockelman, 2016). Maheo, Kilby, & Van Hentenryck (2016) develop a non-automated hub-and-spoke concept, combining high-frequency hub-to-hub buses with feeder taxis to provide superior service than either of the two could provide alone. Roundabout coverage transit routes across areas of low demand are eliminated, and the taxis serve last mile trips and trips not near the hubs. Contra Costa Transportation Authority started a pilot program in 2014 to similarly replace lightly used bus routes with dynamic shuttles, but this does not include modelling or optimization work (Lazarus et al., 2018).

Vehicle Allocation:

All models require some logic for vehicle allocation: the mechanism under which vehicles are assigned to passenger requests. In most cases, vehicles are allocated in response to dynamic passenger requests, though two studies allow for the vehicles to be booked in advance (Bagg & El-Geniedy, 2016; Liang, Correia, & van Arem, 2016). For taxis, the most common approach is for the nearest available (empty) vehicle to a passenger to pick them up. This heuristic has the advantage of being simple and quick to compute. Some models develop more complex algorithms to reach more optimal allocations, such as mixed integer programming (Maheo et al., 2016), local minimization of passenger travel time via constraint programming (Martinez et al., 2015), and linear programming (R. Zhang et al., 2015).

The logic for forming shared rides is also part of vehicle allocation. Constraint-based rules represent the most common approach, so that shared rides are made when neither the existing nor prospective passengers are heavily inconvenienced. D. Fagnant (2014) allows sharing when it does not cause existing passengers to exceed a given additional trip time of 20%. W. Zhang, Guhathakurta, Fang, & Zhang (2015c) tweaks this approach by assuming that that individuals will only share rides when it improves their individual generalized time plus money costs to do so. There are a wide variety of ways to consider and implement these constraints, and to date, few of the modelers have taken an in-depth view into the effects of varying these constraints. As seen in the performance section, there is a great variation in achieved occupancies in the different models; stricter constraints should logically lead to less sharing, but might be appropriate where there is stiffer modal competition. Additionally, different sharing algorithms can substantially add to the computational load as opposed to private rides only, and ensuring real-time assignment is likely a prerequisite for real-world dynamic assignment (Linares, Montero, Barceló, & Carmona, 2016). Lastly, while dealing with erratic passenger behaviors will necessarily be part of any real-world system, no models consider how the vehicles would respond to lateness of the passengers or other inaccurate requests.

Behavior When Not Serving Passengers:

The models must also describe the vehicle behavior when not directly serving passengers — after they have dropped off their passenger(s) and have no immediate trip requests to which they must respond. “Park in place” while waiting for another request is a common approach, but a sizable number of papers also consider some form of empty vehicle relocation in order to rebalance the fleet across the service area and better serve anticipated demand. There is a large literature on the need and approach to rebalancing from car- and bike-sharing literature (e.g.

Chemla, Meunier, & Calvo, 2013; Dell'Amico, Hadjicostantinou, Iori, & Novellani, 2014; Febbraro, Sacco, & Saeednia, 2012). Rebalancing is especially helpful when travel demand patterns change throughout the day, such as more trips ending in the city center in the morning peak but originating there in the afternoon peak.

D. Fagnant (2014) tests four rebalancing strategies, and picks the best performing for the primary simulations. Spieser et al. (2014) explicitly optimizes vehicle rebalancing in order to minimize the total rebalancing distance necessary to serve the passenger demand. Maheo et al. (2016) use integer programming to optimize vehicle allocation in order to minimize total system wait time. Chen (2015) examines electric automated taxis, so accounts for the vehicles' need to relocate to charging stations; this includes when vehicles decide to recharge, how long they spend recharging, and where charging stations should be located. Beyond these efforts, most models ignore vehicle behavior after drop-off, implying that the vehicles stay in place. This presents some real-world challenges with respect to the availability of parking, though only W. Zhang et al. (2015c) accounts for the demand for parking spaces needed throughout the day to accommodate the system. In a different related study, the same authors are also unique in considering keeping vehicles in motion cruising for passengers for a fixed period of time (W. Zhang, Guhathakurta, Fang, & Zhang, 2015b). None of the reviewed papers model the possibility of off-site staging (i.e. outside of the dense downtown) where parking is more available.

Connections with Transit:

The most common approach of studies is to have the SAV systems be self-contained, rather than integrated or coordinated with existing public transit. However, a number of models do allow for connections with transit. Both Levine et al. (2013) and Maheo et al. (2016) have their vehicles act as feeders for the public transit system, and the latter is unique in imagining a revamping of the public transit system itself, calling for high-frequency hub-to-hub bus routes in order to reduce transit wait times and lower costs. Shepherd & Muir (2011) also imagines a first-mile, last-mile access system to extend the reach of existing public transit; the entire program is called "CityMobil." One last exception is International Transport Forum (2016), which assumes that those close to metro stations will continue to use transit, while those living elsewhere will prefer SAVs; they also discuss the possibility of using the SAVs as feeders.

General Comments

In defining service behavior, there is an inherent tension between the travelers, the operators, and the system as a whole. For example, is the modeler assuming individuals seek their own goals without consideration of system-wide impacts? In general, modelers have focused on the travelers, both at an individual and a system level, such as by setting maximum wait times as constraints and by measuring total average passenger wait time. Some though, have focused on the system operator, such as with Liang et al. (2016) trying to maximize profits. In general, the ABMs are better for investigating individual-level behaviors and their implications, while the network assignment models succeed in optimizing for the entire system.

CONTEXT CHARACTERISTICS

This section describes the various context characteristics that SAV models have addressed.

City Population and Density

Each SAV system operates in a given service area, and the characteristics of this area can have a large implication for how effectively such transit services operate. A city's population and density influence the total travel demand for SAV systems how that demand is spatially distributed. All models to date assume a [mostly]-fixed geographic, policy, and infrastructure context set in the present day built environment, limited to developed world cities in the US, Europe, or Singapore. Additionally, most models look at a single location, the exceptions being Burns et al.'s (2013) comparisons between Babcock Ranch, Ann Arbor, and New York and City Mobil's four European cities (Shepherd & Muir, 2011). R. Zhang et al. (2015) looks at both Singapore and Manhattan, but these have very different market penetrations of all vehicle trips vs. existing taxi trips respectively, and are not compared directly. There is a wide range of geographic service areas from the neighborhood scale (Levine et al., 2013) to an entire state (Brownell, 2013), but most models focus on cities or metro areas. Densities covered include a corresponding wide range from about 600 to 28,000 people per km^2.

City population and density directly pertains to this dissertations research questions and hypothesis, along with all of the characteristics in the following market section — market size and density, alternative market size, fleet size, and mode choice. Together, all of these help describe issues around fleet size (the focus of Hypothesis #1), urban density (the focus of Hypothesis #2), urban form (the focus of Hypothesis #3), and mode choice (the focus of Hypothesis #4). For city population and density, the most important difference between previous efforts and the approach taken in this dissertation is that the previous efforts have a specific geography that they study. They look at a single city, often with a detailed street network as described below, and typically have a baseline market size with some variations, as detailed in the following market section. None treat the city itself as an independent variable, with a sensitivity analyses across different densities as done in Chapter 5. Furthermore, though the variety of studies look at a wide range of different cities, including very-high density places (e.g. Singapore and New York City), none look at very low densities, such as 100 people per km^2.

Street Networks and Vehicle Speeds

The various models describe how the vehicles move through the transportation system in different ways. Street networks can be directly modeled at a high level of detail (i.e. microscopic models), a higher-level network used instead (i.e. macroscopic models), or the vehicles can even move in straight lines in open space with no network at all. The more detailed the transportation network, the higher the computational burden. Modelers must also decide what factors influence or determine vehicle travel speeds, including any possible congestion effects. Almost all models treat speed as exogenous, either assumed from the outset or estimated in a separate network assignment model. In both cases, the speeds used are either present-day prevailing street speeds or at some proxy that represents them. If straight-line travel distances are used, as opposed to a street network, some calibration factor should be necessary, as done by W. Zhang et al. (2015b) but not by Burns et al. (2013). A few studies do endogenous model speeds, where the operation of the automated transit system itself is assumed to have an effect on roadway congestion. Levin & Boyles (2015) assume that AV's will enable higher roadway capacity, thus ameliorating potential congestion, while Gucwa (2014) and Childress, Nichols, Charlton, & Coe (2015) assume AVs have identical congestion impacts as current vehicles.

Costs of Vehicle Automation

Papers report the cost(s) of automating the vehicles in a number of different ways. The most common approach is to assume a fixed marginal cost to convert an existing conventional vehicle into an automated one, typically around $5,000-$10,000, though there are a few outliers. Chen (2015) has a baseline cost of $10,000 per vehicle, but performs a sensitivity analysis for costs up to $100,000. A few other studies also varied costs, and as a general rule found that the effects of the cost of automation are dwarfed by the variable operational costs, such as per-distance usage. Two studies actually assume automated vehicle costs by distance; Childress et al. (2015) uses $1.03 per km for automated taxis, using this value for their mode choice model, and D. Fagnant (2014) varies the per-distance cost as part of the economic analysis of system viability. However, no model performs an in-depth examination of the potential operational costs for automated vehicle transit fleets, such as via consideration of regulatory costs, labor costs for system managers, dedicated automation infrastructure, etc.

Average Trip Length

The average trip length is how far the average person travels from their origin to destination. Most models have average trip lengths of less than 10 km, but there was substantial variation across the papers reviewed. These distances have some correlation with the size of the service area; larger ones tend to have longer trips. Longer trip lengths may increase empty vehicle repositioning, making automated transit systems less efficient. Many of the models did not report this distance or make it easily calculable from the data they provided.

Model Type

As defined earlier, this review separates all models as being agent-based, network assignment, aggregate, or some hybrid combination. There is a reasonable mix among the three, though aggregate models are in the minority. The appendix provides further definition of the model type where appropriate, such as describing the specific type of network model.

General Comments

The potential effects of policy choices are generally avoided in these papers, though the hope of guiding policy is an oft-stated goal. Regulations may prove especially important, such as those surrounding mode-switching, ridesharing, and allowable operating parameters. There is a long history of transit agencies trying to protect their most profitable routes (Walker, 2011). A strict regulatory regime effectively bankrupted the jitneys of the 1920s, as they were "stealing" rides from the streetcars (Chambliss, 2008). Many cities regulate the supply of taxis, presumably on grounds of increased efficiency, and regulations for ridesourcing firms like Uber and Lyft are steadily increasing.

MARKET CHARACTERISTICS

This section describes the various context characteristics that SAV models have addressed. As all of these characteristics pertain to the research questions and hypothesis of this dissertation, the general comments subsection synthesizes the general approaches taken by previous models and briefly discusses how they pertain to this dissertation's modelling efforts. The penultimate section of this chapter provides a more holistic review of how all the SAV models detailed in this chapter relate to this dissertation.

Daily Baseline Market Size and Density

As with the service area, the models have a wide range of then baseline number of total daily trips served by the SAV system, from 110 (Kang, Feinberg, & Papalambros, 2016) to over 32 million trips (Brownell, 2013). How the market size is reached also varies; some studies estimate total market size by replacing a percentage of existing car trips; others replace existing transit trips. Some models, most notably Fagnant's and related UT Austin papers, and Zhang's and related Georgia Tech papers, take approximately 2% of total car trips as an approximation of the percentage of total trips in cities currently served by taxis, i.e. taxi mode share. At the other end, some models assume a near-total market share, such as 100% capture of all private vehicle trips (e.g. Martinez et al., 2015; Spieser et al., 2014). Such high levels of adoption are not reasonable in any foreseeable time horizon without dramatic regulatory action (Litman, 2015).

In tandem with service area, daily baseline market size also defines the trip density, defined here as the daily number of trips per km^2. There is a wide range of trip densities, from 0.35 trips/km^2/day (Kang et al., 2016) to 9,965 trips/km^2/day (Martinez et al., 2015); this is related to the wide variety of both market sizes and service areas. In general, the models show that as trip densities increase, empty vehicle repositioning decreases and rideshare opportunities increase; vehicles are closer to the passenger requests. Average trip length also affects the ease of serving trips; all else equal, a given fleet of vehicles can more easily serve 2 km trips than 20 km trips.

Alternative Market Size

The most common approach in the modeling literature is to assume a fixed market size, such as a percentage of current trips. However, some models do also test alternate market sizes from their baseline, such as D. Fagnant (2014) doubling the initial 1.8% of total vehicle trips, Bagg & El-Geniedy (2016) progressively adding more private vehicle trips to the baseline of all current transit trips, and Boesch, Ciari, & Axhausen (2016) varying trip demand from 10% to 100% of the regional total for private vehicles.

Fleet Sizes and their Determination

The majority of modeling efforts obtain a baseline fleet size by a trial and error method, adjusting fleet sizes in order to obtain a given performance parameter or constraint. Of such constraints, maximum allowable wait time is the most common; e.g. fleet size is set so that no passenger must wait for more than 5 minutes. Though not always reported, a common method is to add additional vehicles until the constraint is met (D. Fagnant, 2014; Martinez et al., 2015). The Princeton work does not actually consider real vehicles at all, but simply assumes that a vehicle is present anytime a trip is made (Brownell, 2013; Ford, 2012; Zachariah et al., 2014). Kang et al., (2016) is unique in that the fleet size is part of the solution to an optimization problem that maximizes profit for the system operator. A sizable number of papers do not report the logic they used to determine fleet size. Some models do consider varying the fleet size from the baseline to varying degrees, and the effects of varying fleet sizes are discussed in the following "performance" section.

Mode Choice

The logic and conditions under which people choose to take SAVs as opposed to other modes of travel is definite by mode choice. Since the dynamics between new SAV systems and market demand are uncertain, most of the modeling approaches have avoided in-depth considerations of mode choice, instead assuming from the outset that SAVs would serve a given set of trips. Some include low levels of mode choice, such as with a few ABMs where passengers will "leave" the system under some predefined rules, such as a set maximum waiting time (D. Fagnant, 2014) or proximity to a transit station (Martinez et al., 2015). Brownell (2013) uses an aggregate model, and assumes that all trips to and from Manhattan from New Jersey, occur via public transit. The network models of Chen (2015), Childress et al. (2015), and Gucwa (2014) model mode choice more extensively as part of the traditional four-step transportation forecasting process, using logit models considering monetary costs, wait times, and in-vehicle travel times. For Childress et al. (2015) and Gucwa (2014), this is at the expense of not directly modeling the SAV service; it is instead given a set monetary and time cost. Chen (2015) uses a hybrid modeling approach, taking the outcome for the 4-step network model as inputs for an ABM, and therefore there is an explicit mode choice component. Childress et al. (2015) is the only research to consider both private AVs and SAVs, though as separate model runs, so that people never choose between the two automated modes. Levine et al. (2013) presents an ABM looking at automated shuttles providing first-mile service to the metro system in Chicago, directly modeling both mode choice and SAV service, by integrating a multinomial logit model within their ABM. Horl, Erath, & Axhausen (2016) also endogenously model mode choice as part of their activity model, with detailed and variable utility functions, showing the potential for automated taxis to serve up to 60% of all trips in their hypothetical Sioux City network.

Mode choice can have a large effect in determining model outcomes and parameters, such as market size and trip density. SAV trips could include people shifting from other modes – i.e. taking an SAV instead of the bus or instead of a private car – as well as new "induced" trips – i.e. taking an SAV instead of staying at home. Rayle, Dai, Chan, Cervero, & Shaheen (2016) perform a survey on the effects of Uber and Lyft in the Bay Area, finding that each of the above are occurring, but could not conclude on their relative extent. More recent non-academic studies have echoed this uncertainty; for example, one survey found that at least some people have given up their private vehicles due to Uber and Lyft, but the extent on this on vehicle-kilometers traveled (VKT) and mode share remains unknown (Henderson, 2017).

General Comments

In addition to greatly differing urban contexts, the vastly different assumptions about market share further contributes to difficulties comparing between models. Mode choice concerns are vital here; under what conditions would people choose to take the new automated system as opposed to private vehicles, traditional transit, walking, etc.? Additionally, low demand levels are useful for understanding how new SAV systems might work under early stage deployment. While some shuttle work does consider low demand levels, automated [shared] taxi models typically use high baseline demand. The smallest market shares modeled are Boesch et al. (2016) at 1% of total trips and Liang et al. (2016) at an unstated market share of total trips that yields 2,061 trips/day.

Each of this dissertations four research questions and related independent variables have some relationship to market characteristics — fleet size, urban density, urban form, and mode choice, the results of which are shown in Chapters 4 through 7 respectively. Very few of the previous efforts treat any of these as independent variables, though the most effort has been spent on varying fleet sizes. SAV papers that examine fleet size typically pick a few alternate sizes from the baseline, such as 100 and 200 more and less vehicles than the baseline. However, the lack of a detailed sensitivity analysis makes it difficult to identify tipping points, where small changes in the independent variable lead to large change in outputs. Less papers examine variations in density, the ones that do are addressed in the above alternative market size subsection. These efforts have the same issue as with fleet size; they typically pick a few alternative values, but do not perform a detailed sensitivity analyses as with the variations of density in Chapter 5. Additionally, all of the automate taxi and shared taxi efforts have relative high minimum market sizes: none are much higher than 600 people per km^2 for density and 26 daily trips per km^2 for market size / trip density. None of the SAV papers reviewed here examine variation in urban form, as done here in Chapter 6. Additionally, where these SAV papers include mode choice, they do it for a single context, as opposed to a host of different densities as done in Chapter 7. Also unlike Chapter 7, none of the mode choice models explicitly consider the possibility of private autonomous vehicles (PAVs).

PERFORMANCE METRICS

This section summarizes the various the various performance metrics that that SAV models have used and measured.

Average Passenger Wait Time

Passenger wait time is the time between when they make a request and when they are picked up by the vehicle. The average wait time for all passengers is by far the most commonly used wait time metric. Some papers define peak wait times as well, though none delve into the relation between peak and average times. Most studies have average wait times under five minutes. Oftentimes, this is due to the model design requiring a given level of performance, such as ABM's setting maximum allowable wait time as a constraint. A few models do examine how wait times change due to shifts from the baseline assumptions. Some look at the effects of varying fleet sizes, and find a highly non-linear relationship with wait times (D. Fagnant, 2014; R. Zhang et al., 2015; W. Zhang et al., 2015c). Namely, if there are sufficient vehicles to avoid queuing during the peak demand, then additional vehicles only produce slight decreases in average waiting time. However, if queues do start to develop, when there are more trip requests than vehicles, then wait times can skyrocket, to more than 2.5 hours in R. Zhang et al.'s (2015) Manhattan case.

One way to modulate peak demand is by allowing for ridesharing, i.e. going from automated taxis to automated shared taxis. D. Fagnant (2014) shows a reduction from 1,977 to 1,715 vehicles as well as slightly reduced average wait times when ridesharing is introduced. Brownell (2013) show larger drops in required fleet size, up to over 50%, by more rigorously requiring sharing through their taxi stand model, though this model does not actually capture individual vehicle movements. None of the papers consider much variations in wait time, such as reporting how they vary throughout the region or allowing for traveler's to pay more to have shorter waits. Lastly, very few consider how passenger wait times or travel distances/times change with ridesharing

Cost per Passenger-Kilometer

Many of the models either ignore the cost component of SAVs or treat it as an assumption, such as by fixing the price from the outset at $0.25 per km (W. Zhang et al., 2015c). Roughly half though directly estimate per-distance costs of operation for SAV services, which include both variable costs (e.g. fuel consumption) and amortized fixed costs (e.g. total vehicle cost). Many of these studies find costs in the $0.25-30/km range, though Spieser et al. (2014) is higher at up to $0.66/km, and Liang et al. (2016) the highest at $1.14/km; a potential explanation is Liang's unusually low trip density compared to other models. In general, the vehicle cost, even at relatively high automation expenses, are a relatively small part of total operational expenses; fuel and other operating expenses predominate. No studies consider the potential effects of varying system cost on other aspects of performance. For example, as the system becomes more costly, this would affect people's mode choice whether or not to take an SAV. D. Fagnant (2014) does show that the system he modeled could be operated profitably even at a higher per-mile cost, assumed to be $0.47/km.

Daily Vehicle Utilization

There are two basic types of vehicle utilization given by the models, though many do not report one or either. First is hourly: how many hours each 24 hour day a vehicle is in motion, either relocating or transporting passenger(s). Second is by vehicle: how many passenger trips does each vehicle serve in a given 24 hour day. A few models also report the daily distance traveled, often around 300 km per day (D. J. Fagnant, 2014; Hars, 2015), but not enough did so to make any broader claims.

For hourly utilization, models tend to report 7-8 hrs/day, which would equate to percentage utilization of about 30% over the 24 hour day. Temporal variations in demand appear to be the reason that the vehicles cannot operate for larger shares of the day; a fleet large enough to reasonably serve demand during morning rush hour will have a large percentage of its vehicles sitting empty at 2 AM. Requiring relocation to depots seems to increase utilization; both Martinez et al. (2015) and Chen (2015) show utilizations up to 17 hrs/day, though note that latter actually only looks at a 3.5 hr peak period, so this utilization value is scaled. More ridesharing also corresponds with higher hourly vehicle utilization rates, apparently because the vehicles may be filling up during peak demand times, while in off-peak times they are not utilizing all the seats but are still taking a passenger or two. Two models show approximately 15% higher utilization rates when moving from automated taxis to automated shared taxis (D. J. Fagnant & Kockelman, 2014; Martinez et al., 2015). W. Zhang et al. (2015b) actually calculates lower utilization for shared automated taxis, though this is likely a result of the regular automated taxis having much higher waiting times than their shared brethren. For per-vehicle utilization, automated taxis and automated shared taxi models mostly report in the vicinity of 10-40 trips/veh, with values around 30 trips/veh being the most common. Due to higher occupancies, automated shuttles have more trips per vehicle, estimated at up to 360 trips/veh/day (Levine et al., 2013). The biggest takeaway from vehicle utilization is that it is difficult to run SAVs throughout the day filled with passengers.

Vehicle to Passenger Distance Ratio

Many studies also report the vehicle to passenger distance ratio, which is the ratio of the distance traveled by the vehicle fleet to that that traveled by the passengers. Nominally, private vehicle trips should have a ratio of one, since the vehicle and driver always travel together. Empty vehicle relocation between pickups and drop-offs makes the vehicle/passenger ratio higher, and ridesharing makes it lower. Ratios greater than one imply that the SAV system increases the total vehicle distance traveled on the roadways as compared to private vehicle trips, which creates both congestion and environmental concerns. For automated taxis, the ratio is always greater than one, due to empty vehicle relocation between pickups and drop-offs. For most models, this relocation requirement increases distances traveled from about 7-15% (ratios of 1.07-1.15). Two factors tend to increase this vehicle/passenger distance ratio. First, lower trip densities tend to require longer relocation distances; this is seen in the low density operation of Bagg & El-Geniedy's (2016) automated taxis, whose ratios range from 1.45 to 1.61. Second is more aggressive relocation strategies; requiring relocation to depots as opposed to allowing parking in places appears to be the reason for Martinez et al.'s (2015) high 1.44 ratio.

Ridesharing lowers the ratios by having a single vehicle at least occasionally is transporting more than one person at once. The magnitude of reduction depends on the extent of this sharing and by extension the average achieved occupancies. In general, sharing seems to have a greater effect on fleet size than on vehicle/passenger distance ratio. Sharing in D. Fagnant (2014) only reduces this ratio by 3%, from 1.11 to 1.08, while the fleet size drops by a much larger 13%. Heavy sharing scenarios though, dropped the ratio as low as 0.40 (Levin & Boyles, 2015) and 0.34 (Zachariah et al., 2014), but these are uncommon except for shuttles. Overall, most shared taxis have ratios greater than one; this is because their average occupancy is quite low, and the vehicles are deadheading more often than they are taking more than one person at a time. The few papers that report average occupancy for shared taxis have them around 1.2.

Vehicle Replacement Ratio

Many of the models estimate how many traditional privately owned vehicles an automated taxi or shuttle could replace. 10:1 was the most common vehicle replacement ratio, but the models showed wide variation , as low as 2.5:1 for automated taxis in Singapore (R. Zhang et al., 2015) and as high as 31:1 with shared taxis (Levin & Boyles, 2015) and 20-30:1 with a combination of shuttles and shared taxis (International Transport Forum, 2016). Vehicle replacement ratios appear higher when SAVs claim a larger share of the market, when market demand is spatially concentrated through higher trip densities, and of course when current private car ownership rates are higher. Private vehicle ownership rates are notably lower in Singapore than in the US, for instance, so the vehicle replacement ratio is logically lower as well. Vehicle replacement ratios for carsharing services like ZipCar, which can provide a substitute for car ownership, are estimated to be a maximum of about 11:1 (S. A. Shaheen, Chan, & Micheaux, 2015). Many of the studies claim, though do not explicitly model, that replacing private vehicles is an environmental good. However, reducing the total number of vehicles on the roadway is not equivalent to reducing congestion, VKT, or air pollution caused by vehicular traffic. It is in fact possible for the vehicle fleet to be smaller while VKT increases due to much higher vehicle utilization.

Economies of Scale

SAV systems display economies of scale when their average performance improves as a function of increasing market size. This can be manifested as lower costs per passenger mile, shorter wait times, and/or smaller fleets relative to the number of passengers served, though it is difficult to generalize findings across studies due to very different urban contexts. For example, D. Fagnant (2014) finds only slight economies, such as a 5% reduction in required fleet size due to a doubling of the number of trips, and a 9% reduction when number of trips went from 10% to 100% of the total. Burns et al. (2013) find that system costs rise by 25% in Babcock Ranch when system daily trips are reduced from 10,000 to 2,300. Only D. Fagnant (2014) and (W. Zhang et al., 2017) considered the related issue of economies of geographic scope, where the geography distribution of trips could affect performance, finding in a general way that SAV systems work better in the higher density parts of the city. One interesting point by W. Zhang et al. (2017) is that assuming paid parking changes the spatial distribution of parking lots, such as encouraging more lots in low-income areas. More work is needed to address general issues of spatial impacts and potential equity concerns for SAV systems, especially where policy choices must be made. Economies could also occur with implementation costs, such as if per-vehicle purchase costs go down when bought in bulk, though no studies have considered this yet.

Environmental Effects

The models report a wide range of environmental effects, largely because some assume electrified vehicles while others do not. At one end, (Greenblatt & Saxena, 2015) assume that if all trips would occur via automated *electric* shared taxis, greenhouse gas emissions could decline by 70-90%; Burns et al. (2013) has similar numbers with the addendum of "right-sized" vehicles smaller than they are to and thus more energy efficient. Studies without electrification show more moderate reductions in emissions, generally concentrated at under 10%, though International Transport Forum (2016) estimates a 34% reduction for a combined automated shared taxi and shuttle system. Where reported, higher sharing understandably leads to lower emissions, with the extent of reduction depending on the extent and type of sharing.

Most papers also find environmental improvements due to higher vehicle utilization rates than privately owned vehicles. Higher utilization would presumably result in a newer and more efficient vehicle fleet with fewer environmental impacts. However, if a system does not include ridesharing, then the need for vehicle relocation between trips increases total distance traveled, which could have potentially negative environmental impacts. As discussed above in the market section, few papers consider induced demand (i.e. more total trips) or likely mode choice impacts, such as diverting riders from conventional public transit, and the environmental and congestion effects these modal shifts could have. An exception here is Horl et al. (2016) which estimates that by "stealing" rides from both transit and private vehicles, automated taxis could gain a 60% mode share and increase vehicle distance driven by 60% as well. Lastly, most papers assume no impact of AV fleets on roadway congestion, one exception being Gucwa (2014), and none investigate the possibility of risk-averse AVs increasing congestion.

SUMMARY OF MODELING WORK TO DATE

Shared automation would profoundly impact transportation systems, but there is a great deal of uncertainty about both the type and the magnitude of its impacts. This section seeks to provide clarity to policy makers and modelers on work to date and potential paths forward for both research and deployment. It synthesizes the major assumption, modelling choices, and findings of past modelling work, and offers some key takeaways, such as on the effects of fleet size on performance and the benefits of ridesharing. The following and final section is geared specifically towards modelers, providing guidance and recommendations for future modeling and research efforts.

Synthesizing Results

Shared Automated Vehicles (SAVs) encompass automated taxis, automated shared taxis, and automated shuttles. SAV modelling is a new and emerging field; correspondingly, there is very little standardization around the system design and performance reporting. The research community cannot even agree on the names of such services. This chapter provides an organizational framework for shared automated vehicles by describing the models via service, context, and market characteristics and performance metrics. These characteristics and the corresponding modelling treatments/findings are summarized in Table 1. The reader should refer to the appendix for the specific treatments/findings of each individual modelling study. Lastly, note that Chapter 3 includes a table comparing this modelling synthesis of previous SAV efforts with the model developed in this dissertation.

Table 1: Modelling Synthesis

SERVICE	
Service Type	Automated Taxi, Automated Shared Taxi, and Automated Shuttle.
Vehicle Allocation	Most use nearest available vehicle and satisficing rules for sharing. Some optimization of assignment with respect to different objectives: wait times, profit, relocation distance, etc.
Behavior When Not Serving Customer	Most have parking in place, with relocation to serve expected demand where appropriate. Some relocation to dedication stations.
Connections with Transit	Mostly none, though some shuttles connect with transit as a first-mile solution. One study explicitly allows for one mode switch for AV to/from transit, and another, non-automated study revamps the transit network to high-frequency bus routes supported by taxis.

CONTEXT	
City Population and Density	~80,000 to 8 million people. Most used entire city, metro area, or equivalent. ~100 to 27,000 ppl/km^2, most in 1000 to 6000 range.
Street Network and Speeds	Most use either abstract or real street networks, with speeds either fixed or given by a travel demand model, ranging from ~20-60 kph, typically assuming present-day roadway capacities.
Marginal Cost of Vehicle Automation	Either not reported or a one-time cost of ~$2,000-15,000, with some allowing for higher figures.
Average Trip Length	Either not reported or from ~4-10 km, with a min of <2.4 km and a max of 34 km.
Model Type	Aggregate, Network, Agent-Based, or Hybrid.

MARKET	
Daily Baseline Market Size and Density	Wide ranging, from ~20,000 to 11 million trips, ~1% to 100% of all trips in the region. Densities also wide ranging, from ~ 100 to 10,000 trips/km^2.
Alternate Market Sizes	Most use a fixed market size, but some do perform sensitivity analyses using selected different market sizes (i.e. double and half of baseline).
Fleet Size and Determination	From ~50 to 300,000. Typically determined by a satisficing approach. A few studies vary fleet size from the baseline.
Mode Choice	Generally not considered. Where mode choice is included, it is typically as a rule-based approach as opposed to a logit model or similar method to estimate user behavior.

PERFORMANCE	
Average Passenger Wait Time	Typically under 3 minutes, often under 1 minute. Some evidence of substantially higher peak wait times (i.e. over 15 minutes).
Cost per Passenger KM	Where reported, from approximately $0.20 to $0.65/km, with some higher outliers.
Daily Vehicle Utilization	Taxis and shared taxis generally in motion from ~6-12 hrs/day, most frequently from ~7-8 hrs/day, and serve from ~10-40 trips/vehicle, mostly centering around 30 trips/vehicle. More sharing leads to more utilization to varying degrees; shuttles can serve up to 360 trips/vehicle.
Vehicle/Passenger Distance Ratio	For automated taxis, generally ~1.05-1.2, with some higher outliers if required relocation to depots. For shared taxis, typically slight reductions from sharing (i.e. from 1.11 to 1.07). Substantially lower values for automated shuttles.
Vehicle Replacement Ratio	Mostly 1 AV replacing ~6-12 traditional vehicles, but numerous higher and lower outliers.
Economies of Scale	Only a minority of studies report, showing a wide range from miniscule to moderate economies of scale, including smaller relative fleet sizes as market size increases.
Environmental Effects	Most report only small reductions in emissions of 0-15%, though a few outliers report massive improvements, especially if SAVs allow for switching to an EV fleet.

Key Takeaways

These six key takeaways provide an overview of SAV model findings to date, coming from both the reporting of the models results and close analysis of the details of the models' functioning.

1) **System Viability/Profitability:** Every single study indicates that the SAV system they evaluate will be an improvement on the status quo; in other words, they claim that SAV systems are not just viable, but desirable. All studies that include cost estimates conclude that these systems could substantially outperform taxis, and some show that they could be

less expensive than current private-vehicle ownership costs. This mostly appears due to the elimination of driver labor cost. No studies take an explicitly critical approach to SAVs.

2) **Vehicle Relocation:** For single occupancy trips, vehicles typically must travel an additional 5-20% over what the passengers travel by themselves, due to empty "relocation" travel (aka deadheading), driving to where the passenger is waiting. Requirements for vehicles to return to a depot when not in use can substantially increase this number, but there is not enough research in this area to make general claims. Active relocation strategies, moving vehicle in response to expected demand levels as opposed to only in response to direct trip requests, can be important in improving passenger wait times, especially where there are directional variations in travel demand throughout the day: i.e. into the city center in the morning, out of the center in the afternoon. Such active relocation does not appear to add to total distance traveled compared to if it were not used.

3) **Sharing:** The biggest benefit of automated shared taxis is their ability to effectively service peak hour demands with smaller fleet sizes. There is substantial variation in travel demand over the course of the day in most cities, so shared-ride taxis may operate as single-occupancy taxis for most of the day, when there are typically more available vehicles than there are trip requests. However, during the peak times of high demand, they may serve two, three, or four people at a time. As such, sharing acts as a reservoir of additional system capacity. Correspondingly, ridesharing does not actually increase vehicle occupancies by much, as they transport only a single party the majority of the time. The amount of deadheading (i.e. travelling while empty) is greater than the amount

of sharing, so that even automated shared taxis travel longer distances than would private vehicles for equivalent trips. These SAV model findings are supported by Henao's (2017) real-world research, which estimates that only 5-10% of UberPool and LyftLine vehicles carrying passengers actually have more than one party in them at the same time.

4) **Fleet Size:** When fleet size falls substantially below a "satisficing" level necessary to provide service to meet peak demand, queues of waiting passengers form. This results in massive performance degradation, including excessive average wait times. Fleet sizes above the satisficing level do not seem to improve performance by much.

5) **Trip Density:** Studies have a wide variation in travel demand density, both due to different underlying urban population densities and different assumptions about the size of the market served (i.e. ranging from less than 2% of private vehicle trips to 100% of private vehicle trips). Increasing trip density seems to bring relatively small improvements in performance for automated taxi performance (i.e. doubling the density allows for a 10% reduction in relative fleet size), but it is difficult to say anything more conclusive based on work to date. More work is also needed on the relationship between ridesharing and trip density, though they seem to be positively correlated. This is likely because as densities increase, it becomes far more likely that two or more trips will be overlapping in both time and space.

6) **Vehicle Utilization:** Most studies have the vehicles in operation from about 6 to 12 hours a day, with the remainder of the time spent parked. Having vehicles in operation for a larger proportion of the day seems difficult due to variations in daily travel demand. Overall though, in-depth examinations of vehicle utilization across different assumptions are sparse.

RECOMMENDATIONS FOR FUTURE MODELLING EFFORTS

Overall, researchers to date have shown that shared automated vehicles *COULD* improve urban transportation. The greatest need for modelling efforts in the future is on *HOW* these systems will be implemented, including their secondary impacts on both transportation and land-use patterns. As with any new technology, SAVs will follow an adoption curve, starting with the early adopters; modelling 100% market penetration is useful from an academic perspective, but does not explain how SAVs will get to that point from their current status of 0% adoption. SAVs will shape and be shaped by the current and future transportation systems and built environment. Total travel demand will likely due to new, more flexible transportation modes at lower costs. Trip patterns themselves may change. On one end, privately owned and operated AVs could reduce the perceived cost of travel by effectively giving people a virtual chauffer and allowing for in-vehicle travel time to be free time, thus encouraging more travel overall: e.g. longer acceptable commutes. By comparison, shared automation could effectively expand the market size of transit, offering low cost taxi-like service that makes it feasible for people to eschew private vehicles. This could encourage more but shorter trips, since people would switch from a high fixed-cost but low variable-cost model of private ownership to a mobility on-demand system.

This section recommends two mechanisms to help answer the "how" question. First is a greater adherence to benchmarking model parameters and results, which can allow for easier comparisons across modelling efforts and more detailed meta-analyses, such as attempted by this chapter. Second is varying modelling choices under service/market/context framework, especially in the interest of examining potential deployment pathways and the conditions under which SAVs are more or less likely to succeed.

Benchmarking Automated Flexible Transit Services

Model results are highly dependent on the context of the analysis, behavioral assumptions behind the models, and the sometimes hidden details of how the models work. Therefore, as modeling research in this area proliferates, we would encourage the virtues of transparency and adherence to common reporting standards. Such standardization is important for modelers and the broader public, a vital tool to allow for comparisons across research efforts. By splitting system characteristics into service, context, market, and performance categories, as detailed above in Table 1, we hope to provide a useful conceptual framework. Beyond a general call for openness and transparency, we recommend three areas where more detailed disclosure is merited:

Specifying Model Inputs vs. Outputs: It can occasionally be difficult to see what inputs are assumed from the beginning and how these relate to the performance metrics reported. For example, if both demand profile and fleet size are fixed from the outset, then trips per vehicle is a modelling choice that is similarly fixed, though it might at first glance appear to be a model output. Detailed descriptions of model assumptions are also useful, especially with respect to any restraints in the model, such as minimum acceptable wait time or number of rejected trips.

Time and Distance Values: Easier benchmarking could be achieved if researchers reported more descriptive statistics for travel times and distances (i.e. min/max, mean, and standard deviations). A system with an average wait time of two minutes may seem desirable, but it becomes far less attractive if it also has a peak wait time of an hour.

Temporal and Geographical Demand Variations: Many of the models describe, or at least allow for the calculation of, the daily trip demand and average trip densities. However, few models provided sufficient detail to describing how trip demand varies over space and time across the service region.

Varying Context, Market, and Service Profiles

To understand the potential and limitations of shared automation, we need to examine the range of contexts, potential markets served, and services provided as independent variables of analysis, asking questions like, "What types of systems will work better in what types of cities?" Only a handful of papers have analyzed different services in the same city (e.g. automated taxis vs. automated shared taxis), and fewer still have analyzed how the same service might perform in different urban contexts. One broad challenge is that the models presented here are typically quite detailed and data/computationally intensive. This tends to make them specific to a given service, context, and market, with only a few possible variations from the baseline case. It is difficult to perform a sensitivity analysis on the effects of fleet size, one of the easiest parameters to vary, if each distinct model run takes an hour, let alone more complicated variables like geographic scale or service type. Simpler models could potentially allow for more explicit comparisons across contexts, markets, and services. The remainder of this section details the most important of these comparisons, proposing specific modelling recommendations for each of the service, context, and market categories.

Service

Early Stage Services Types: Automated vehicles that can travel all over the city in mixed traffic carrying thousands of people a day present large technological hurdles, and is not a likely early deployment scenario. More work is needed into early stage SAV services, including analyses of low-speed operation, separation from traditional vehicles, and geographically limited areas. Start-ups have been promoting low-speed automated shuttles and implementing real-world trials: e.g. Olli by Local Motors, NuTonomy, an MIT spinoff that partnered with Lyft in June 2017, NAVYA, a French firm that is starting a trial in Ann Arbor, and EasyMile, a spinoff from the CityMobil project. While not extensively modeled, three papers do consider low-speed operation (Levine et al., 2013; Shepherd & Muir, 2011; Winter et al., 2016).

Interactions with Transit: The concept of using SAVs as "last-mile" connections with transit is often discussed by the studies presented here and the broader AV literature, but interactions with transit are less frequently modelled. More detailed analyses are merited, such as examining how the how transit service itself might be reshaped, especially as integrated systems where the SAVs and transit work together. Though not an automated system, Maheo et al. (2016) provide a good example of this; they seek to optimize bus routing towards high-frequency hub-to-hub travel in order to provide the best possible traditional transit system to work in tandem with a taxi service to provide transit to the greater Canberra region in Australia.

Context

Varying Urban Contexts: More work is needed to examine how performance can vary both across different cities — only Burns et al. (2013) explicitly does this – and across different parts of the city. D. Fagnant (2014) and W. Zhang et al. (2017) include some work on variations across the city, but neither consider varying the service area, e.g., dense downtown vs. whole metropolitan region. Simpler models would likely be more transferable across different contexts.

Effects of Automated Flexible Transit on Traffic Flow: Similar to the "last-mile" many papers discuss the potential for SAVs to reduce congestion, but few model it outright. One large area of uncertainty is the effects automated vehicles will have on roadway capacities, especially with respect to the service types: e.g. mixed-traffic vs. grade-separated or regular vs.

low-speed operation? Only one study assumes that automation could increase roadway capacity (Levin & Boyles, 2015). More detailed considerations and sensitivity analysis are need, including the possibility that automation could actually decrease capacity in mixed traffic; risk-averse and rule-bound AVs could have lower speeds and higher headways (Townsend, 2014).

Policy Implications: This is a catch-all category to consider how government regulation and policy might either promote or hinder new transportation services like SAVs. Of special interest are the possibilities for variable pricing — such as via distance or congestion-based taxes — and prioritized access, such as via HOV lanes. Equity concerns are another major issue; e.g. how will the SAV system affect accessibility to underserved areas, and what can the government do to promote this? Other issues include but are not limited to ADA requirements for people with disabilities, curbside access, existing taxis/transit/TNC rules and regulations, and relationships between the government and the system operator — e.g. the system could be fully public, fully private, or a highly regulated private monopoly with transit companies in Hong Kong and Singapore.

Market

Range and Size of Trip Demand Profiles: Some of the models do consider a range of demand levels, most notably Boesch et al. (2016), which varies from 10% to 100% of total trips in the region by 10% steps. However, the generalizability of their findings is limited by their assumption that people who wait more than 10 minutes leave the system and their attempt to serve a very large geographic market. More work is merited where market size is an independent variable, especially with respect to modelling early stage deployment and how SAV systems might evolve over time.

"Peakiness" of Demand: In addition to the absolute levels of demand, another important aspect is how demand varies both spatially and temporally. Neither has been considered as an independent variable. There are more trips in the city center than in the outskirts, more trips during 8 am than at 2 am, and the spatial and temporal patterns are themselves linked: more trips into the city center in the morning, more trips out in the afternoon, and more recreation-focused trips at night. Spatial peakiness is closely linked with varying contexts, and two studies have done some preliminary work on spatial variation of demand (Boesch et al., 2016; W. Zhang et al., 2017). Temporal peakiness has been studied even less, though International Transport Forum (2016) does show how system performance varies throughout the day. Note that demand patterns may also change as a function of trip type; commuting trips have different peaks and spatial orientation than taxi trips.

Mode Choice: Moving beyond exogenously fixed demand to endogenously modeling which people will opt into using SAVs is a big opportunity for future work. While mode choice is placed under the context section here, it also has sizable impacts in the service and context categories; e.g. how private rides vs. ridesharing vs. dynamic shuttles or living in the city center vs. the outskirts affects people's desires to take SAVs. Of special interest is how SAVs would compare not just with current modes of travel but also with future ones, especially privately owned AVs To address these issues, we propose three broad lines of inquiry in ascending orders of complexity. First is the application of existing mode choice models to SAVs, such as via generalized nested logit models, allowing for the fact that not every traveler has the same utility function. Second is the potential for SAVs to induce travel, rather than just replacing existing trips; if SAVs lower the cost of transportation as promised, economic principles dictate that the total amount of travel should increase. The third and most complex task is to integrate modeling of mode choice with land use change. Past transportation innovations, such as the automobile itself, have brought about significant changes in land use patterns, and SAVs will likely be no different.

RELATIONSHIP BETWEEN PREVIOUS SAV MODELS AND THIS DISSERTATION

This chapter is intended as a general literature review of SAV models to date. It can act as a stand-alone document to help summarize the state of SAV modelling research for other modelers, policy makers, and general academic audiences. As an emerging field, SAV research is rapidly expanding and changing, and there are numerous opportunities for novel efforts that to help fill critical research gaps —far more gaps and opportunities than can be addressed by this dissertation alone. However, the four research questions posed in the previous chapter are necessary additions to the field. This section summarized what previous work has been done on these questions, what remains to be done, and what methodological innovations are needed to support this new work.

These four research questions laid out in Chapter 1 are reproduced here:

1) ***What are the effects of varying fleet size on the performance of SAV systems?***
2) ***What are the effects of varying urban density on the performance of SAV systems?***
3) ***What are the effects of varying urban form on the performance of SAV systems?***
4) ***How does the likelihood of people choosing SAVs over other modes of travel vary as a function of urban form and density (Chapter 7)?***

The first three questions include three distinct independent variables: fleet size, urban density, and urban form. The fourth question deals with mode choice, with is a model extension rather than a true independent variable. Instead of having SAV travel demand as a fixed input, including mode choice makes it a model output. Total travel demand remains fixed, but the percentage of trips served by the different modes (i.e. the mode share) can vary. As discussed earlier, in the real world, total travel demand could also vary, such as due to induced demand, but it is not considered by this dissertation. What is considered in the fourth question is how SAV's mode share could change due to variations in the other independent variables, namely density and urban form. Therefore, all four questions require sensitivity analyses of three independent variables, across a wide range of different values with small separations between them.

None of the above papers perform this sort of sensitivity analyses. They are typically detailed and complex models, with very little or no variations in the independent variables. For example, some might pick a baseline market size with three or four alternate sizes. Between the three independent variables, there is some consideration of different fleet sizes, limited considerations of different densities, and no consideration of different urban form. Relatively few papers look at mode choice either, and none with either considerations of private autonomous vehicles (PAVs) or of sensitivity analyses of the independent variables (e.g. how SAVs mode share changes as a function of density). The modes detailed above are all fairly complex and detailed, and are thus too computationally intensive and cumbersome for sensitivity analyses. Better is a comparatively simplified model that allows for quicker run times, as proposed in Chapter 3. For example, the simple and flexible ABM developed there allows for over 10,000 simulation runs, for densities ranging from 100 to 12,500 people per km^2 to take approximately 60 hours of computation time. Chapter 3 also provides a broader theoretical justification for the use of simpler ABMs, borrowing heavily from Complex Systems literature.

CONCLUSION

This chapter reviews modelling efforts to date on shared automated vehicles: automated taxis, automated shared taxis, and automated shuttles. This is a new and emerging field, with a number of challenges for researchers. Since shared automated vehicles are not currently on the roadway, modelling may be the best method to understand their potential operation, but any real-world calibration of these models is impossible. Furthermore, these services will likely look and act differently than current travel modes; "self-driving Ubers" is but one of many possibilities for SAVs. These challenges have led to a wide variety of different approaches to modelling SAVs, such as using agent-based vs. network models, having widely different mechanisms for estimating the travel demand profiles, and reporting performance metrics in different ways. This chapter creates and populates an organizational framework to help standardize SAV modelling work and guide future efforts. It splits the modelling choices into service, context, and market categories, and attempts to standardize the respective choices in each of these categories as well as fourth category on performance metrics.

These results show the potential for these services to revolutionize transportation for the better, promising drastic cost savings compared to traditional taxis or even private vehicle ownership, low wait times and reliable service, reduced environmental impacts, higher utilization of vehicles as opposed to sitting in driveways all day, and more. However, these results are highly dependent on the assumptions made by the respective modelers. This importance of assumptions, combined with the diverse modelling choices in the respective models, means that it is hard to make general statements about SAVs, such as which types of services would work the best in which places. To help guide both researchers and policymakers, this chapter strongly recommends more standardized reporting and more treatment of key assumptions as independent variables, including sensitivity analyses. Key assumptions include, but are not limited to: magnitude of travel demand (i.e. market size), spatial and temporal variation of travel demand, acceptable wait times, acceptable conditions under which a shared trip will be made, and mode choice. The latter requires a comparison of overall performance of SAVs with other modes of travel. Human-driven vehicles will likely remain common for many decades into the future as well (Litman, 2015), but it is likely that as vehicle automation proceeds, there will be growth in both the market for SAVs as well as privately owned automated vehicles. The work reviewed here has shown wide range of possibilities for SAVs and their viability as a new mode of travel, and future work can help guide the actual deployment of these services, putting the right types of SAV systems in the right types of environments.

Literature Review Appendix

The below tables give the parameters and findings of the SAV models in the four main categories: Service, Context, Market, and Performance. Throughout the tables, "not reported" refers to something that model considered but did not report and/or cannot be estimated from what they did report, while "NA" means that the choice is not applicable to or considered by the respective model. Additionally, even if the modeling papers did not provide the values in the exact form as in these tables, we provide them here as long as they could be easily calculated from the data the papers did provide.

Table 2: Service Parameters for SAV Papers

Model Name	Service Type	Vehicle Allocation	Behavior When Not Serving Passengers	Connections with Transit	Special Limitations or Considerations
(Bagg & El-Geniedy, 2016)	Automated Taxi	Advanced booking; Closest available vehicle	Park in place	Automatic transfer to regional transit	Limited service area for taxi fleet – local neighborhood trips only
(Bischoff & Maciejewski, 2016)	Automated Taxi	Closest available vehicle	Park in place	NA	Exclusive area for AVs in city center considered as an alternative option
(Burns et al., 2013); Ann Arbor	Automated Taxi	Closest available vehicle	Park in place	NA	Assumes trip origins and destinations are evenly distributed across the region
(Burns et al., 2013); (Babcock Ranch)	Automated Taxi	Closest available vehicle	Park in place	NA	Assumes trip origins and destinations are evenly distributed across the region
CityMobil: (Shepherd & Muir, 2011)	Automated Shuttle (20 passengers): Fixed Route	Not reported	Park in place	Assumes AVs provide first/last mile service into existing transit systems	NA
ETH Zurich: (Boesch et al., 2016)	Automated Taxi	Closest available vehicle	Park in place	NA	Includes all travelers who journey through Zurich, regardless of their origin or destination
ETH Zurich: (Horl et al., 2016)	Automated Taxi	Closest available vehicle	Park in Place	NA	Considers dynamic demand and mode choice
Georgia Tech: (W. Zhang, Guhathakurta, Fang, & Zhang, 2015a)	Automated Shared Taxi: 25%-100% willing to share ride (Max occupancy not reported)	Closest available vehicle, up to 10 minutes away; Sharing if passenger wait and travel time constraints satisfied	Cruise to out-of balance district; minimum cruising time	NA	Focus on impacts of parking
Georgia Tech: (W. Zhang et al., 2015b)	Automated Shared Taxi (Max occupancy not reported)	Closest available vehicle; Sharing if cost savings greater than time costs for all passengers.	Cruise to out-of-balance district; minimum cruising time	NA	Examines ridesharing based on real-time dynamic choice logic
Georgia Tech: (W.	Automated Shared Taxi	Closest available vehicle; Sharing if cost	NA	NA	Assume EVs. Focus on temporal and

Model Name	Service Type	Vehicle Allocation	Behavior When Not Serving Passengers	Connections with Transit	Special Limitations or Considerations
Zhang et al., 2017)	(Max occupancy not reported)	savings greater than time costs for all passengers.			spatial distribution of parking demand with both paid and free parking.
(Greenblatt & Saxena, 2015)	Automated Shared Tax (Assumes 10% trips with 2 passengers)	NA	NA	NA	Considers whether internal combustion, hybrid electric, or battery electric vehicles are cost competitive at different levels of annual usage
(Gucwa, 2014)	Privately Owned Automated Vehicles	NA	NA	NA	Activity-based model takes into account how travel demand changes with introduction of Level 3 automation
(Hars, 2015)	NA	NA	NA	NA	Focus on cost comparison of private and fleet-owned vehicles
(Maheo et al., 2016)	Non-Automated Taxi: Also Includes Fixed-Route Buses	Optimized via mixed integer program on a known set of trips	Not reported	Human-driven taxis and buses integrated together	Changes bus routes to high-frequency express routes, with taxis covering service gaps and first/last mile
ITF: (Martinez et al., 2015)	Automated Taxi and Automated Shared Taxi (2, 5, and 8 passengers)	Local minimization of total passenger travel time; Passenger wait and excess travel time constraints	Automated Shared Taxi: Relocates to nearest station; Automated Taxi: Parks on street	NA	Model built on highly detailed synthetic travel data, including traveler age, income, trip purpose, time of day, origin, and locations
ITF: (International Transport Forum, 2016)	Automated Shared Taxi (3 passenger) AND Automated Shuttle (8 or 16 passengers)	Shared Taxis: local minimization of VKT with minimum insertion Hamiltonian path; Taxi-Bus: Local maximization of vehicle occupancy; Passenger wait and travel time constraints for both.	Shared Taxi: Relocates to nearest station, 60 stations in the city; Taxi-Bus: Stays at station.	Transfers between shared taxi/taxi-buses and the rail system are permitted	Model built on highly detailed synthetic travel data
(Javanshour & Dia, 2016)	Automated Shared Taxi: (Assume 75% ridesharing, no max occupancy)	Not reported	Privately owned vehicles return to home; Shared vehicles relocate to designated staging areas	NA	Models local neighborhood trips only
(Kang et al., 2016)	Automated Taxi	Optimize customer wait time; Charging requirement restraints	If charge level below bound, moves to available charging station. Otherwise waits in place.	NA	Automation used only to facilitate easier use of car share services
(Levine et al., 2013)	Automated Shuttle (no max occupancy)	No allocation; fixed route travel through neighborhoods.	Fixed route	Shuttles feed into a single train station	Only serves first mile/last mile trips to train station
(Liang et al., 2016)	Automated Taxi	Optimized to maximize profits	Can wait or relocate to serve passengers.	Taxis feed into train stations	Only serves first mile/last mile trips to train station
MIT-Singapore: (Spieser et al., 2014)	Automated Taxi	Optimized rebalancing distance using lumped spatial queuing model	Rebalance every 30 minutes while minimizing distance traveled	NA	NA
MIT-Singapore: (R. Zhang et al., 2015)	Automated Taxi	Optimized rebalancing distance using lumped spatial queuing model	Constant rebalance as needed	NA	NA

Model Name	Service Type	Vehicle Allocation	Behavior When Not Serving Passengers	Connections with Transit	Special Limitations or Considerations
MIT-Singapore: (Iglesias, Rossi, Zhang, & Pavone, 2016)	Automated Taxi	Optimized rebalancing distance using lumped spatial queuing model within a capacitated road network	Constant rebalance as necessary; pickups at 50 centers in city	NA	NA
Princeton: (Brownell, 2013)	Automated Shared Taxi & Automated Shuttle (both with variable occupancy)	Assumes vehicles are always present at taxi stands; Rides shared if have same origin and destination for taxis, and if along same route for shuttles	NA	NA	NA
Princeton: (Ford, 2012)	Automated Shared Taxi	Assumes vehicle always present at taxi stands; Rides shared if have same origin and destination	NA	NA	NA
Princeton: (Zachariah et al., 2014)	Automated Shared Taxi (Varies 1 - 6 passengers)	Assumes vehicle always present at taxi stands; Rides shared if have same origin and destination	NA	NYC/ Philadelphia commuters taken to nearest NJ train station	NA
(Childress et al., 2015)	Automated Taxi	NA	NA	Not reported	Examines 4 policy scenarios: 1) AVs may increase roadway capacity, 2) reduce time costs of travel, 3) reduce parking costs, or 4) VKT fee for all vehicles
(Shen & Lopes, 2015)	Automated Taxi	Extension of closest available; use "expand and target" to identify nearest expected available taxi (may be currently occupied)	Not reported	NA	Focus on different vehicle redistribution algorithms
UT Austin: (Chen, 2015)	Automated Taxi	Closest available vehicle; Vehicle is unavailable if needs to recharge	Relocate to closest available charging station if need, else relocate to areas of expected need	NA	Considers electric SAVs in comparison to internal combustion SAVs, with variable charging times and ranges, and locations of charging stations.
UT Austin: (D. Fagnant, 2014)	Automated Shared Taxi (4 passengers)	Closest available vehicle; Sharing if passenger wait and travel time constraints are met	Relocate according to block-balancing approach	NA	NA
UT Austin: (D. Fagnant & Kockelman, 2014)	Automated Taxi	Closest available vehicle; Sharing if passenger wait and travel time constraints are met	Relocate according to block-balancing approach	NA	NA
UT Austin: (D. J. Fagnant & Kockelman, 2014)	Automated Taxi	Closest available vehicle	Test four relocation strategies. Decide on block-balancing approach	NA	Environmental impacts of shared taxis in comparison with privately owned vehicles
UT Austin: (Levin & Boyles, 2015)	Private Automated Vehicle	Not applicable, as vehicles are not shared	AVs either park and pay or return to the origin	NA	Models mode choice between transit, automated transit, and private cars
UT Austin:(Levin et al., 2016)	Automated Shared Taxi (4 passengers)	Closest available or occupied with estimated arrival less than 10 minutes; Sharing if passenger wait and travel time constraints are met	No active relocation	NA	Integrated traffic assignment and SAV routing
(Winter et al., 2016)	Automated Shuttle (from 2 to 40 passengers as independent variable)	Closest available with space and wait time constraint; Booked in advance and dynamically	No Active Relocation	NA	Low-speed campus-only shuttle, with only two predetermined pickup and drop-off locations.

Table 3: Context Parameters for SAV Papers

Model Name	City Population and Density	Street Network and Speeds	Marginal Cost of Vehicle Automation	Average Trip Length	Model Type
(Bagg & El-Geniedy, 2016)	Saint-Jean-sur-Richelieu (Part of Metro Montreal); 87,500 people; 387.1 ppl/km^2	Microscopic street network; Speeds fixed by link type	$5,000 per year	Not reported	Agent-Based
(Bischoff & Maciejewski, 2016)	Berlin; 3.502 million people 3,930 people/km^2	Macroscopic street network; Link speeds modeled exogenously	NA	Not reported	Agent-Based
(Burns et al., 2013) (Ann Arbor)	Ann Arbor, MI; 285,000 people; 845.7 people/km^2	No street network; Fixed straight-line speeds: 48 kph	$2,500	9.3 km	Hybrid: Aggregate built upon Agent-Based
(Burns et al., 2013) (Babcock Ranch)	Babcock Ranch, FL; 50,000 people; 726.7 ppl/km^2	No street network; Fixed straight-line speeds: 40 kph	$2,500	5.6 km	Aggregate
CityMobil: (Shepherd & Muir, 2011)	Madrid: 5,846,000 people; Vienna: 2,755,000 people; Gateshead: 1,451,000 people; Trondheim: 150,000 people;	Vehicles operate on dedicated guideway (network); 10-15 kph	NA	Not reported	Agent-Based
ETH Zurich: (Boesch et al., 2016)	Zurich; 1,300,000 drivers; ~14,792 drivers/km^2	Macroscopic street network; Speeds modeled exogenously for drop-offs, but straight-line fixed speeds for pick-ups: 40.6 kph	NA	Wide variation in trip length, most under 10 km.	Agent-Based
ETH Zurich: (Horl et al., 2016)	Sioux Falls (artificial); 84,110 people; 675 ppl/km^2	Macroscopic street network: Link speeds not stated, but appear exogenously obtained	Fixed at $0.528/km (cost to user)	Not reported	Hybrid: MatSim Activity Model
Georgia Tech: (W. Zhang et al., 2015a)	Hypothetical City; ~370,000 people; ~ 1,429 people/km^2	Abstracted network (0.8 km grid); Speeds fixed by time of day: 48 kph off-peak and 34 kph peak	Fixed at $0.249 per kilometer total cost	Not reported	Agent-Based
Georgia Tech: (W. Zhang et al., 2015b)	Hypothetical City; ~282,000 households ~ 1,089 hh/km^2	No street network reported; Speeds fixed by time of day: 48 kph off peak and 34 kph during peak	Fixed at $0.249 per kilometer total cost	8.7 km	Agent-Based
Georgia Tech: (W. Zhang et al., 2017)	Atlanta, GA; 447,841 people; 1,288 ppl/km^2	Macroscopic street network; Link speeds exogenously obtained	Fixed at ~$0.25 per km; $0.5 /min and $0.3/min cost to user for private/shared	Not reported	Agent-Based
(Greenblatt & Saxena, 2015)	Entire United States	No	$5,000 (in 2030)	Not reported	Aggregate
(Gucwa, 2014)	San Francisco Bay Area, 2030 (Nine+ county region) Population NA	Microscopic street network; Link speeds and congestion endogenously modeled	NA	Not reported	Network: Metropolitan Planning Commission's Travel Model One
(Hars, 2015)	NA	No street network modeled; Assumes typical speeds of < 40 kph	$2,200	< 15 km	Literature Review

Model Name	City Population and Density	Street Network and Speeds	Marginal Cost of Vehicle Automation	Average Trip Length	Model Type
(Maheo et al., 2016)	Canberra, Australia Population NA	Macroscopic street network; Link speeds calibrated to exogenous real-world data	NA	~8 km (13.8 minutes)	Network
ITF: (Martinez et al., 2015)	Lisbon; 565,000 people; 6,678 ppl/km^2	Macroscopic street network; Link speeds and congestion calibrated to exogenous real-world data	NA	~ 6.6 km	Agent-Based
ITF: (International Transport Forum, 2016)	Lisbon; 565,000 people; 6,678 ppl/km^2	Macroscopic street network; Link speeds calibrated to exogenous real-world data.	Not reported	~ 2.3 km (includes walk trips)	Agent-Based
(Javanshour & Dia, 2016)	Melbourne, single neighborhood; Population NA	Microscopic street network; Calculation of speeds not reported.	NA	4.0 km	Agent-Based
(Kang et al., 2016)	Hypothetical City; 313 km^2 (Population NA)	No street network; Straight-line constant speeds of 34 kph, chosen to approximate real-world street network speeds	$2,500	Not reported	Network (Optimization)
(Levine et al., 2013)	Chicago, 4 urban neighborhoods; Population NA	Abstracted grid street network; Fixed speed: 24.1 kph	NA	≤ 2.4 km	Agent-Based: Includes Multinomial Logit Mode Choice Model
(Liang et al., 2016)	Delft, Netherlands; 100,000 people; 4,167 ppl/km^2	Macroscopic street network; Link speeds calibrated to exogenous real-world data	Not reported	All trips < 15 minutes from station.	Optimization Model (Integer Program)
MIT-Singapore: (Spieser et al., 2014)	Singapore; 1,144,000 households 1,591 hh/km^2	Microscopic street network; Link speeds calibrated to exogenous real-world data	$15,000	6.4 - 13.3 km	Network
MIT-Singapore: (R. Zhang et al., 2015)	Singapore (Same as above); Manhattan; 1.626 million people 27,812 ppl/km^2	Macroscopic street network; Link speeds calibrated to exogenous real-world data	$15,000	Not reported	Network
MIT-Singapore: (Iglesias et al., 2016)	Manhattan; 1.626 million people 27,812 ppl/km^2	Macroscopic street network (357 nodes); Link speeds endogenously modeled	NA	Not Reported	Network
Princeton: (Brownell, 2013)	State of New Jersey ~8.8 million people ~467 ppl/km^2	Abstracted grid street network; Fixed speed: 48 kph.	$100,000 (total vehicle cost)	26 - 34 km	Hybrid: Aggregate Built upon Network
Princeton: (Ford, 2012)	Mercer Country, Princeton, NJ; ~370,000 people ~624 ppl/km^2	NA	NA	Not reported	No SAV Modeling
Princeton: (Zachariah et al., 2014)	State of New Jersey; ~8.8 million people; ~467 ppl/km^2	Abstracted grid street network; Fixed speeds: 48 kph	NA	Not reported	Hybrid: Aggregate built upon Network
(Childress et al., 2015)	Puget Sound Region (Seattle) ~ 3,733,580 people	Macroscopic street network; Link speeds endogenously modeled	Fixed at $1.03 per kilometer total cost	9.3 - 12.7 km	Network (Activity-Based)

Model Name	City Population and Density	Street Network and Speeds	Marginal Cost of Vehicle Automation	Average Trip Length	Model Type
	~230 ppl/km^2				
(Shen & Lopes, 2015)	New York City; 8,550,405 people 7,043 ppl/km^2	Microscopic street network from Open Streets; Speeds not reported	NA	Not reported	Agent-Based
UT Austin: (Chen, 2015)	Abstracted Austin Metro Region 2.9 million people 111.8 ppl/km^2	No street network reported; Speeds vary by location and time of day, 24-58 kph	Baseline: $10,000. Sensitivity analysis with costs up to $100,000.	~16 km	Hybrid: Agent-Based built upon Network
UT Austin: (D. Fagnant, 2014)	Austin, TX (service area slightly larger than city); ~931,831 people ~1,297 ppl/km^2.	Macroscopic street network; Link speeds and congestion exogenously modeled.	$10,000	Not reported	Agent-Based
UT Austin: (D. Fagnant & Kockelman, 2014)	Austin, TX (same as above)	Macroscopic street network; Link speeds and congestion exogenously modeled.	NA	7.69 km	Hybrid: Agent-Based built upon Network
UT Austin: (D. J. Fagnant & Kockelman, 2014)	Abstract Version of Austin 256 km^2 (Population NA)	Abstracted network (0.4 km grid); Speeds vary by time of day: 53 kph off-peak, 34 kph peak	NA	8.84 km	Hybrid: Agent-Based Model built on Network Model
UT Austin: (Levin & Boyles, 2015)	Austin, TX metro region	Macroscopic street network; Speeds endogenously modeled according to Greenshield's, averaging 41.2 kph. Assumes AVs increase link speeds relative to human-driven vehicles.	Not reported	Not reported	Network
UT Austin:(Levin et al., 2016)	Austin, TX metro region	Macroscopic street network; Speeds endogenously modeled according to Greenshield's, averaging 41.2 kph.	NA	3.65 km	Network: Event-Based Cell Transmission Model Intended to be Overlaid on Various Simulation-Based Traffic Models
(Winter et al., 2016)	Wageningen, Netherlands (Campus)	Tested on two node network, applicable to macroscopic street network; Speeds fixed at 30 kph	$114,000 (total vehicle cost); $228,000/year in total infrastructure cost	7 km	Optimization of fleet size using ε-constraint and simulation

Table 4: Market Parameters for SAV Papers

Model Name	Daily Baseline Market Size and Density	Alternative Market Sizes	Fleet Size(s) and Determination.	Mode Choice
(Bagg & El-Geniedy, 2016)	7,585 trips (all local transit trips); 34 trips/km^2	Progressively add private vehicle trips, up to 180,299 for all trips	586 vehicles (baseline); 11,022 vehicles (all trips); All passengers assigned at their intended departure time.	NA
(Bischoff & Maciejewski, 2016)	278,000 trips (10 % of private vehicle trips); 2,803 trips/km^2	174,000 (10 % city center vehicle trips); 323,000 (10% vehicle/transit trips); 2.5 million (all vehicle trips)	6,500-7,000 vehicles for city center; 11,000 vehicles for baseline; 12,000 vehicles for10% vehicle/transit; 100,000 vehicles for all vehicle; Meet 95th percentile wait time constraint (< 15 minutes) and minimizes fleet cost.	No, but tests different levels of AV capture of existing transit trips
(Burns et al., 2013); Ann Arbor	528,000 trips (all trips); 1,567 trips/km^2	Sensitivity analysis on trip rate, from 50 to 50,000 trips per hour	18,000 vehicles; Meet average peak-hour wait time constraint of 2 minutes, plus 10% contingency.	NA
(Burns et al., 2013); (Babcock Ranch)	115,000 trips (all trips; 1691 trips/km^2	NA	3,500 vehicles; Meet average peak-hour wait time constraint of 2 minutes with 10% contingency	NA
CityMobil: (Shepherd & Muir, 2011)	Not reported	NA	Number not reported; Meet maximum peak wait time of 5 minutes	Considers growth in transit mode share
ETH Zurich: (Boesch et al., 2016)	366,124 trips (10% of car trips); 409 trips/km^2	1-10% of current car trips	Ranges from 10% - 100% of current car fleet size; “Reasonable” service if 95% of people have 10 minute waits or less	NA
ETH Zurich: (Horl et al., 2016)	Dynamic demand; all trips potentially served by AVs	NA	Baseline 1,000 vehicles; Independent variable from 0 to 8,000 vehicles	Detailed and variable utility functions; Private car, public transit, walk, and automated taxi
Georgia Tech: (W. Zhang et al., 2015a)	28,046 trips (2% of metro car trips); 108 trips/km^2		Ranges from 500 to 800 vehicles; Meets unstated wait time constraint;	NA
Georgia Tech: (W. Zhang et al., 2015b)	29,911 trips (2% of metro car trips); 115 trips/km^2	NA	700 vehicles; Meets unstated wait time constraint.	NA
Georgia Tech: (W. Zhang et al., 2017)	32,365 trips (~3.7% of vehicle trips); 92 trips/km^2	NA	Not Reported; Meets unstated wait time constraint	NA
(Greenblatt & Saxena, 2015)	Not reported	NA	Not Reported	Exogenously considered as a cause of energy efficiency gains
(Gucwa, 2014)	Not reported	NA	Not reported	Part of traditional 4-step network model; Automated taxi, private vehicle, transit, or walk/bike
(Hars, 2015)	NA	NA	Ranges from 63 - 93 vehicles per 1000 inhabitants.	NA

Model Name	Daily Baseline Market Size and Density	Alternative Market Sizes	Fleet Size(s) and Determination.	Mode Choice
(Maheo et al., 2016)	~21,000 trips; Trip/km^2 not reported	NA	32 buses Set as fixed initial assumption.	NA
ITF: (Martinez et al., 2015)	840,000 trips (all private vehicle and bus trips); 9,965 trips/km^2	50% of baseline; Add train trips (1,104,000 trips)	21,120 vehicles (if shared taxi) or 34,082 vehicles (if taxi); Method not reported, but presumably to meet wait and travel time constraints.	Rule-based mode choice; Walk, metro, automated taxis or shuttle
ITF: (International Transport Forum, 2016)	643,487 trips, 56% shared taxi, 44% shuttle bus (all private vehicle and bus trips); 6,702 trips/km^2	No total trip variation, but does test serving all trips by shared taxis only	Number not reported; Meets wait time constraints.	Rule-based, stochastic mode choice; Walk, subway, automated shared taxi, or automated shuttle
(Javanshour & Dia, 2016)	2,136 trips; Trips/km^2 not reported	NA	1,217 vehicles; Method not reported.	NA
(Kang et al., 2016)	110 round trips 0.35 trips/km^2	36 and 73 round trips	19 vehicles; Optimized to maximize profit	Individualized mode choice model Station-based car share or automated car share
(Levine et al., 2013)	30,000 trips; 416 trips/km^2	NA	8 shuttles; Method not reported, except to maintain 3 minute headways.	Individualized mode choice calibrated to stated preference survey; Drive, shuttle bus to train, walk to train, cycle to train, or cycle
(Liang et al., 2016)	2,061 trips; 45 trips/km^2	NA	From 20 to 80 vehicles	NA
MIT-Singapore: (Spieser et al., 2014)	6,000,000 trips (all private vehicle trips); 8,343 trips/km^2	NA	From 200,000 to 300,000 vehicles	NA
MIT-Singapore: (R. Zhang et al., 2015)	Singapore: 6,000,00 trips 8,343 trips/km^2 Manhattan: 439,950 trips (all weekday taxi trips); 5,057 trips/km^2	NA	Singapore: From 200,000 to 300,000 vehicles Manhattan: From 6,000 to 8,000 vehicles	NA
MIT-Singapore: (Iglesias et al., 2016)	22,416 trips (7-8 AM) ~430 trips/km^2	NA	From 0 to 4,000 vehicles; Independent variable	NA
Princeton: (Brownell, 2013)	32,770,528 trips (all New Jersey trips); 1,416 trips/km^2	NA	From 1.6 to 4.4 million vehicles; Minimum necessary to serve peak-hour demand with there always being an available vehicle. Number varies by service type and repositioning time.	NA
Princeton: (Ford, 2012)	1,298,727 trips; (all Princeton area trips) 2,190 trips/km^2	NA	NA	NA

Model Name	Daily Baseline Market Size and Density	Alternative Market Sizes	Fleet Size(s) and Determination.	Mode Choice
Princeton: (Zachariah et al., 2014)	~30 million trips (all New Jersey trips); 1,416 trips/km^2	NA	NA: Assume available vehicle will always be at each taxi stand	Rule based mode choice; Short trips by walking and transfer to train for commuters
(Childress et al., 2015)	~11.4 million trips (all metro Seattle trips); ~ 750 trips/km^2	Trip market changes slightly based on costs; Amount is not reported	NA	Part of traditional 4-step network model; Automated taxi, privately owned AVs, private vehicle, transit, or walk/bike
(Shen & Lopes, 2015)	~340,216 trips (scaled from all annual NYC taxi trips); ~280 trips/km^2	NA	12,216 vehicles (Roughly # of vehicles in the NYC Yellow Cab Fleet)	NA
UT Austin: (Chen, 2015)	3.9 million trips (all metro Austin trips); ~26 trips/km^2	Ranges from 3.62 - 4.26 million trips, depending upon cost and value of travel time	29,939 vehicles for regular vehicles, from 31,859 to 57,279 vehicles for EVs; Fleet size based on # charging stations, wait time, and vehicle range/charging time constraints;	Part of traditional 4-step network model; Automated taxi, private cars, or transit
UT Austin: (D. Fagnant, 2014)	56,324 trips (1.3% of metro Austin total); ~ 76 trips/km^2	2x and 4x of baseline	~1,700 vehicles; Meets 10 minute max wait time constraint	NA
UT Austin: (D. Fagnant & Kockelman, 2014)	57,000 trips (1.3% of metro Austin total); ~ 76 trips/km^2	NA	1,715 vehicles (shared taxi), 1,977 vehicles (taxi), varies taxi from 1,187 to 2,217 vehicles; Meets 10 minute max wait time constraint	NA
UT Austin: (D. J. Fagnant & Kockelman, 2014)	65,350 trips ; ~ 255 trips/km^2	¼ x, ½ x, and 2x baseline	From 1,400 to 1,908 vehicles Meets 10 minute max wait time constraint	NA
UT Austin: (Levin & Boyles, 2015)	62,826 trips; ~ 45 trips/$km^{2;}$ Both for 2 hour AM peak	Private AV trips determined by mode choice model	Fleet size unknown; Models trips, not vehicles: Ownership of AVs based on wealth.	Part of traditional 4-step network model; Privately owned AVs, private vehicles, or transit, with 10 classes of wealth, vary who can afford AVs
UT Austin:(Levin et al., 2016)	62,826 trips; ~ 45 trips/$km^{2;}$ Both for 2 hour AM peak	NA	From 1,000 to 60,000 vehicles; "Reasonable" service at 17,500 private-ride, 2,000 shared ride.	NA
(Winter et al., 2016)	3,693 trips trips/km^2 NA	0.5x, 0.75x, 1.25x and 1.5x baseline	224 vehicles baseline; Varies from 95 – 409 vehicles; Determined by optimization algorithm	NA

Table 5: Performance Metrics for SAV Papers

Model Name	Average Passenger Wait Time	Cost per Passenger Kilometer	Daily Vehicle Utilization	Vehicle to Passenger Distance Ratio	Vehicle Replacement Ratio	Economies of Scale	Environmental Effects
(Bagg & El-Geniedy, 2016)	1 - 2 min; 6 min maximum wait time	Not reported	Hours not reported; 13 trips/veh	1.66 (Transit + 25% vehicles); 1.51 (Transit + 50%); 1.45 (All Trips)	NA	No economies shown: 10% run is equivalent to 100% run	Appears to be entirely from electrification – Automated EVs 15x cleaner than internal combustion engines.
(Bischoff & Maciejewski, 2016)	3 min (100% scenario); 14 min 95th percentile	NA	7.6 hours; 25 trips/veh	1.15	1:11	9% reduction in required relative fleet as go from 10% to 100% of the population	Not reported
(Burns et al., 2013); Ann Arbor	<1 min, even during the peak	\$0.24 - 0.26	8.4 hour; 29 trips/veh	1.01	1:11	Yes, but magnitude not reported	75 - 90% less energy use with small vehicle, electric fleet.
(Burns et al., 2013); (Babcock Ranch)	<1 min, even during the peak	\$0.28 - \$0.31	~7.2 hours; 33 trips/veh	1.05	Not reported	20% reduction in costs as go from 2,300 to 10,000 trips	Not reported
CityMobil: (Shepherd & Muir, 2011)	1.5 - 3 min for large city; 5 min for small city	NA	NA	Not reported	NA	Not reported	0 - 2% reduction in carbon emissions by 2035.
ETH Zurich: (Boesch et al., 2016)	3.1 min; Passenger leave if wait more than 10 min	NA	7.7 hours; trips/veh not reported	1.14	1:10 at high levels of demand	32% reduction in required relative fleet size as go from 1% to 10% of trips	None
ETH Zurich: (Horl et al., 2016)	~3.5 min, ~15 min at peak. With more vehicles. Average wait drops slightly, and peak more sharply.	Fixed at \$0.528/km (cost to user)	Utilization not reported	~1.3 (numbers unclear)	Not reported	Not strictly reported, but graphs show decreasing gains in mode share AV fleet increases (e.g. 39% at 1000 veh, 65% at 8000 veh)	Can show large increases in VKT, up to 60%, due to reductions in transit and relocation distances.
Georgia Tech: (W. Zhang et al., 2015a)	10.2 min at peak (No Sharing); 5.5 min at peak (50% sharing); 3.0 minutes at peak (100% sharing)	\$0.249 (assumed)	~12.1 hours; ~35 trips/veh (No Sharing) ~12.5 hours; ~40 trips/veh (Sharing)	1.234 (No Sharing); 1.059 (Sharing)	1:14	NA	88 - 97%, reduction in parking demand, depending upon system configuration.
Georgia Tech: (W. Zhang et al., 2015b)	2.0 min, 4.2 min at peak (No Sharing; 1.7 min, 2.6 min at peak (Sharing)	\$0.249 (assumed)	~10.8 hours; trips/veh not reported (No Sharing); ~10.3 hours; ~43 trips/veh (Sharing)	1.114 (No Sharing); 1.071 (Sharing)	NA	NA	1.0% reduction in GHG & energy use (No Sharing); 5.6% reduction in GHG & energy use (Sharing)
Georgia Tech: (W. Zhang et al., 2017)	~ 2.5 min, 6.0 min at peak (Free Park); ~1 minute longer for Paid Parking	~\$0.25 (assumed)	~8.1 hours; ~30 trips/veh; ~ 1 hour in parking lot	Not reported	NA	NA	~5% reduction in parking spaces with ~5% trips served, with heavier reductions in the downtown;

Model Name	Average Passenger Wait Time	Cost per Passenger Kilometer	Daily Vehicle Utilization	Vehicle to Passenger Distance Ratio	Vehicle Replacement Ratio	Economies of Scale	Environmental Effects
(Greenblatt & Saxena, 2015)	NA	\$0.20 - 0.27 for BEVs in 2030	NA	0.958	NA	NA	70 - 90% reduction in GHGs by 2050 if most of the fleet converts to shared battery electric vehicles
(Gucwa, 2014)	NA	NA	NA	NA	NA	NA	4-8% increase in VKT and fuel use
(Hars, 2015)	~1 minute	\$0.17 (only vehicle costs)	6.9 hours; trips/veh not reported	NA	1:4 - 1:7	NA	Smaller electric vehicles will become more common.
(Maheo et al., 2016)	10 min	\$1.74 (paid human drivers)	Not reported	Not reported	NA	Not reported	Not reported
ITF: (Martinez et al., 2015)	Average not reported; Max of 5 min	NA	15.6 hours; 25 trips/veh (No Sharing); 17.5 hours; 40 trips/veh (Sharing)	1.44 (No Sharing); 1.06 (Sharing); (Some bus travel replaced with SAVs.	1:6 to 1:10	NA	44% increase in VKT (No Sharing); 6% increase in VKT (Sharing)
ITF: (International Transport Forum, 2016)	Average not reported; Max of 5 min for shared taxi, 10 min for shuttle bus	~\$0.32 for shared taxis	12.8 hours; trip/vehicle not reported (Sharing); Utilization not reported for shuttles	Not reported.	1:20 to 1:30 (for shared taxi and shuttle combined)	NA	16% reduction in CO_2 (No Sharing); 34% reduction in CO_2 (Shared Taxis and Shuttles)
(Javanshour & Dia, 2016)	Not reported	NA	Hours not reported; 1.8 trips/veh (2 hour peak)	0.71 (combined private and shared vehicles)	1:2.3 for shared vehicles	NA	29% reduction in VKT
(Kang et al., 2016)	4.9 min	\$4.40 (car share)	Hours not reported; 5.8 trips/veh	Not reported	No	9.5% increase in profits as market size triples	Not reported
(Levine et al., 2013)	1.5 min (half shuttle headway time)	NA	Hours not reported; 360 trips/veh	Not reported	NA	NA	7 - 29 % reduction in private vehicle travel - higher decreases in more auto-oriented neighborhoods.
(Liang et al., 2016)	None; trips are booked in advance, vehicles arrive instantaneously	\$1.14	6.2 hours; 26 trips/veh (80 vehicles); Both per 16 hour day	Not reported	NA	NA	NA
MIT-Singapore: (Spieser et al., 2014)	3 min, 30 min max (250 K veh)	\$0.66	Hours not reported; 20 trips/veh	Not reported	~1:3	Yes, but magnitude not reported	Not reported
MIT-Singapore: (R. Zhang et al., 2015)	Singapore: 13 min at peak hour (300 K veh); Manhattan: 2.5 min at peak (8,000 veh)	\$0.66 vs. \$0.96 with a driver (Singapore)	Singapore: 20 trips/veh; Manhattan: 54 trips/veh; Hours not reported	Not reported	Singapore: ~1:2.5 (all vehicles) Manhattan: ~1:1.4 (taxis)	Not reported	Not reported
MIT-Singapore: (Iglesias et al., 2016)	Not Reported	NA	Varies with fleet size; utilization stays high as fleet size starts to increase, then drops swiftly	Not clearly reported; ~1.2	NA	NA	NA

Model Name	Average Passenger Wait Time	Cost per Passenger Kilometer	Daily Vehicle Utilization	Vehicle to Passenger Distance Ratio	Vehicle Replacement Ratio	Economies of Scale	Environmental Effects
Princeton: (Brownell, 2013)	NA	\$0.26 - \$0.46;	Hours not reported; 7.4 trips/veh	0.463 – 0.788 (varies with vehicle sharing assumptions)	NA	NA	NA
Princeton: (Ford, 2012)	NA	NA	NA	NA	NA	NA	NA
Princeton: (Zachariah et al., 2014)	0 - 5 min	NA	NA	0.34 - 1.0	NA	NA	NA
(Childress et al., 2015)	NA	\$1.03 (assumed)	NA	NA	NA	NA	Increases in delay and VKT with private ownership, even with optimistic assumptions; Reductions in delays, VKT, and total trips in shared vehicle case.
(Shen & Lopes, 2015)	5.62 - 8.14 min; "Expand and Target" routing algorithm reduces waiting time by ~27%.	NA	Hours not reported; 28 trips/veh	Not reported	NA	NA	NA
UT Austin: (Chen, 2015)	Not reported	Not Reported	~ 8.6 – 17.5 hours; ~37 trips/veh	Ranges from 1.07 - 1.14.	1:6.8 - 1:37	NA	7 - 14% increase in VKT
UT Austin: (D. Fagnant, 2014)	0.2 - 1.18 min; Ridesharing reduces wait times by ~20-50%.	\$0.16 - \$0.31	~7.75 hours; 33 trips/veh	~1.11 (No Sharing); ~1.08 (Sharing)	1:9 (No Sharing) 1:11 (Sharing	4% reduction in required relative fleet as go from 1.3% to 2.6% of total trips	Not reported
UT Austin: (D. Fagnant & Kockelman, 2014)	0.48 min, 0.5% people wait 10+ min (2,217 veh); 2.4 min, 0.5% wait 10+ min (1,817 veh)	NA	~7.75 hours; 29 trips/veh	~1.10	1:9	NA	Not reported
UT Austin: (D. J. Fagnant & Kockelman, 2014)	0.26 minutes; 0.4% wait 5+ minutes	NA	~7.2 hours; ~34 trips/veh	~1.09	1:10	5% reduction in required relative fleet size and wait time decrease from 0.26 to 0.14 min as go from 1.3% to 2.6% of total trips	6% reduction in CO_2 emissions
UT Austin: (Levin & Boyles, 2015)	Not reported	~\$0.26 - \$0.30	Not reported	1.83 to 1.99	NA	NA	271% increase in personal vehicle trips.
UT Austin:(Levin et al., 2016)	10.3 min (17,500 veh); 5 min (25,500 veh) (No Sharing)	Not reported	Hours not reported; 3.6 trips/veh	~3.00 (No Sharing) ~0.4 (Sharing)	1:3.6 (No Sharing); 1:31.4 (Sharing)	Not Reported	Not reported
(Winter et al., 2016)	2.2 min (Baseline) Varies from 0.72 - 5.01 min	\$0.59 per passenger (Baseline)	Hours not reported; ~16 trips/vehicles	Not reported	NA	Small reductions in required relative fleet size, but magnitude unclear.	NA

Chapter 3 — Model Overview

This chapter describes the construction of an agent-based model for SAV service. This model is intended to be flexible and simple, enabling detailed sensitivity analyses for key variables such as fleet size and allowing for the consideration of different types of cities. As discussed in the previous chapter, SAVs represent a wider range of potential mobility services, from automated taxis to automated shared taxis to automated shuttles. The model presented below represents the automated taxi subgroup of SAVs; the vehicles are ordered on-demand, provide "door-to-door" service, and only transport a single party at any one point in time (i.e. they do not include ridesharing). Table 9 at the conclusion of this chapter provides an in-depth comparison of the modelling choices and results between the models reviewed in the previous chapter and the model developed here.

This chapter first describes the construction of the model, including the various pick-up and drop-off behaviors of both the vehicles and passengers. It then details the model usage, including descriptions and quantifications of the various model inputs, outputs, and performance metrics. The chapter concludes with a discussion of the benefits of using a simplified model, in two broad areas. First is how a simplified model is better suited to answering the research questions laid out in Chapter 1, with a theoretic justification from complex systems literature. Second, it compares this simplified approach with the comparatively complex models detailed in the previous chapter, performing a proto-calibration to show how the results of this simple model can effectively duplicate those efforts.

MODEL CONSTRUCTION

This section describes the construction and parameters of the model, which simulates single passenger pick-ups and drop-offs. Passengers "hail" an available vehicle; the vehicle then becomes "assigned" to the passenger and thus unavailable, travels to the passenger, picks them up, drives them to their destination, and drops them off, at which point the vehicle becomes an available vehicle again.

Model Setup and Trip Generation

The model is coded in NetLogo, a dedicated agent-based modelling platform. A hypothetical city represents the service area, 10 x 10 km split into 100 m cells. Trip origins and destinations are independently assigned for each trip across the space of the city. There is no street network, so that vehicles move in open space; this and other model simplifications are explained and justified in the following section. Both vehicle and passenger locations are given as continuous variables. Initial vehicle distribution matches that of the trip distribution (i.e. random vs. more centralized). Figure 1 gives an example of 400 vehicles randomly distributed throughout the city.

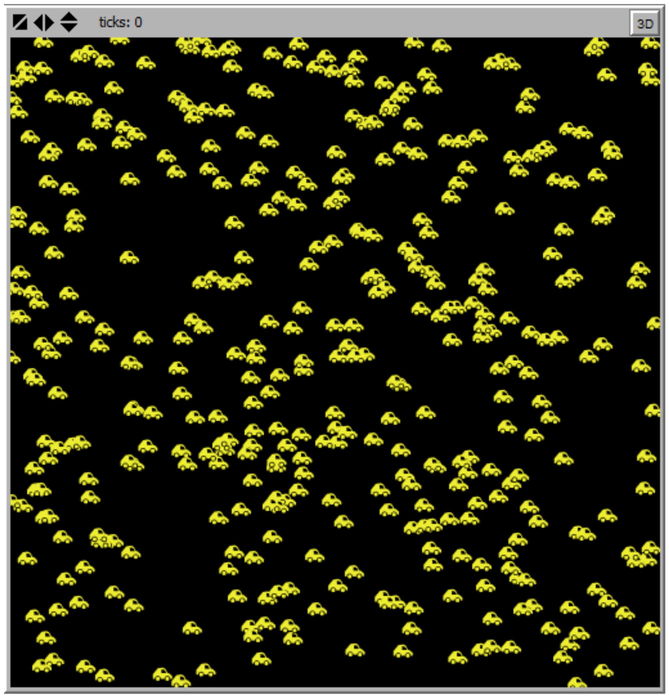

Figure 1: Example of Initial Random Taxi Distribution

Trip generation varies across the day as according to the minute-by-minute demand patterns of yellow taxis in New York City on June 1st, 2016 using a Poisson distribution. As such, the actual number of generated trips might vary slightly simulation run to run. When a trip is generated, the passenger will have both an origin (their starting location) and a destination. The peak hour occurs in the evening, and represents 6.5% of total daily trips. This daily travel demand is given in Figure 2, split into one-minute intervals.

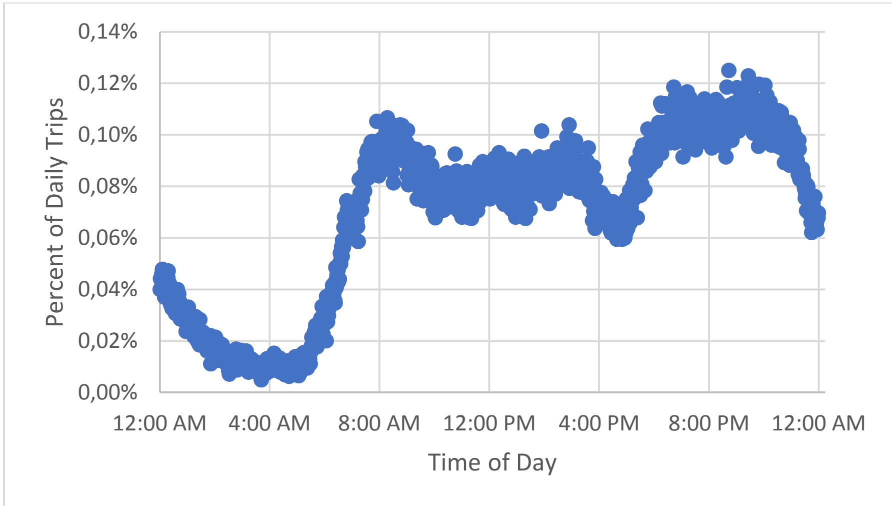

Figure 2: Daily Travel Demand Pattern

Passenger and Vehicle Behavior

This model uses two agent classes: vehicles and passengers. Their linked behavior for pick-up, travel, and drop-off and corresponding status at each step of their travel is described as follows:

Trip Assignment

For each time step, there is a list of all "available" vehicles: those that are not assigned to or otherwise serving a passenger. Once a trip is generated, the passenger's status is set as "hailing." Each hailing passenger is assigned to the nearest available taxi, with the proviso that the passenger(s) that have been waiting longest are assigned first. This occurs under conditions of queueing, defined as when there are more hailing passengers than available taxis so that some passengers have to wait to be assigned. The time that passengers wait to be assigned is given by the "hailing wait time." Once a passenger has been assigned to a vehicle, the passenger's status goes from "hailing" to "assigned" and vehicle's status goes from "available" to "assigned."

Vehicle Relocation and Pick-Up

After the vehicle and passenger have been assigned to each other, the vehicle travels to the passenger. The vehicle picks up the passenger once they are in the same 100 m radius. The pick-up takes 1 minute, during which both the passenger's and vehicle's status go from "assigned" to "pick-up."

In-Transit

After the pick-up has been completed, both the passenger's and vehicle's status becomes "in-transit," and they travel together towards the passenger's destination.

Drop-Off and Trip Completion

Once the vehicle and the passenger reach the cell that contains the passenger's destination point, a drop-off occurs; as with pick-ups, the drop-offs take 1 minute, and during this time both the passenger's and vehicle's status are "drop-off." After the drop-off is complete, the passenger's status becomes "arrived," and the vehicle's status reverts back to "available."

Behavior Overview

Vehicle speed is fixed at 28 km/hr; this is slightly lower than the expected true speeds to account for the straight-line distances used in this model. Here and elsewhere we use the metric conventions. Each simulation runs for a 24 hour day, with one-minute time steps. Figure 3 gives an example of the simulation run at a single point in time. The population density is 2,500 people per km^2 and the fleet size is 2.0 vehicles per 100 trips, which equates to 400 total vehicles, as shown in Figure 1. The vehicles are all in yellow, passengers who have arrived are in red, and passengers who have not arrived yet are in white. White lines represent the path that the vehicle will take to the passenger's destination once the vehicle has picked up the passenger. Table 6 summarizes the passenger and vehicle statuses during their various actions.

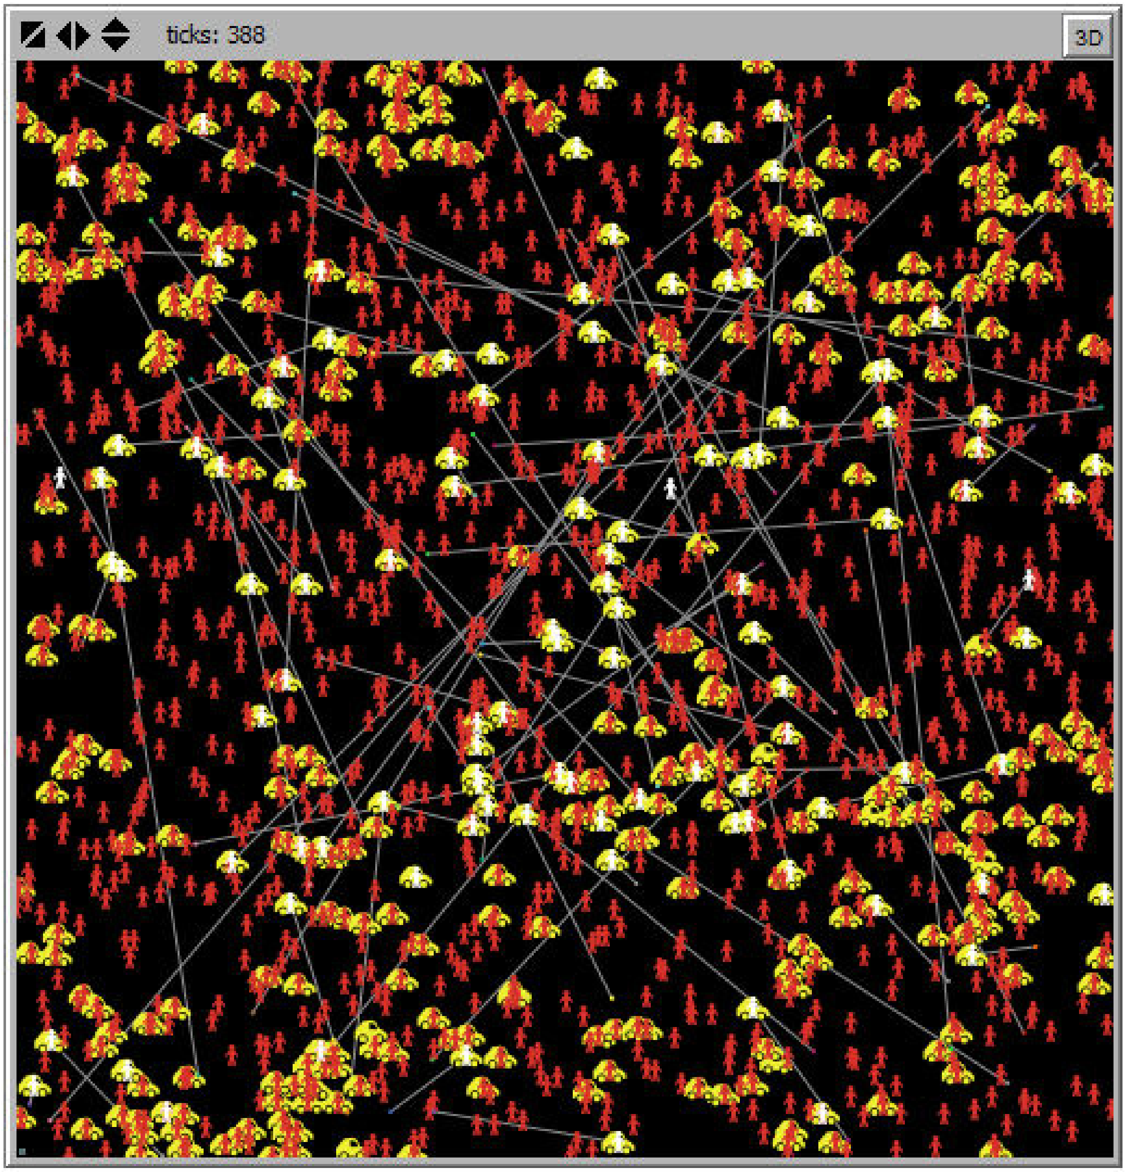

Figure 3: Snapshot of a Simulation Run

Table 6: Passenger and Vehicle Model Behavior

Passenger	Vehicle
Hailing: The passenger is requesting a ride. If delays occur here, they are *hailing wait* times, and indicate queueing (more passengers than available vehicles).	***Available:*** The vehicle is empty, staying in place, and looking for a passenger.
Assigned: A vehicle has been assigned to the passenger. If delays occur here, they are *relocation wait times.*	***Assigned:*** The vehicle has been assigned to a passenger and is travelling to them.
Pick-Up: The passenger is in the process of being picked up.	***Pick-Up:*** The passenger is in the process of being picked up.
In-Transit: The passenger is travelling in the vehicle to their destination.	***In-Transit:*** The passenger is travelling in the vehicle to their destination.
Drop-Off: The passenger is being dropped off.	***Drop-Off:*** The vehicle is dropping off the passenger. After this is done, its status reverts to "available."
Arrived: The passenger has arrived at their destination.	

Model Usage

This section describes how the model is used to test the effects of varying fleet size, urban density, and urban form, addressing Hypotheses #1, #2, and #3. It provides the model inputs, outputs, and derived performance metrics.

Model Inputs

First, to define the three main categories and how they are implemented in the model:

Fleet Size

In order to allow for comparisons across different densities, fleet size is given here as a ratio of vehicles per 100 completed trips. The absolute number of vehicles can always be calculated as need be, based on the total number of trips; the following section describes how population density relates to total trips. As an example, a population density of 1,250 people per km^2 equates to 10,000 total trips, so a fleet size of 2 vehicles per 100 trips equates to approximately 200 vehicles. The true provided fleet size value will actually be slightly different, both due to the random nature of trip generation and because some of the generated trips will not be completed by the time the 24 hour model run is finished.

Urban Density

There are two basic types of density: trip density is given as trips per km^2, and population density is given as people per km^2. Unless otherwise stated, where this dissertation uses "density," it is referring to population density. Additionally, this chapter assumes a fixed conversion between population and trip densities, assuming that for every person, there are 4 trips per day, and that SAVs will take 2% of total trips; this conversion is in line with other models that have looked at SAVs replacing current taxi trips (D. Fagnant & Kockelman, 2014; R. Zhang, Spieser, Frazzoli, & Pavone, 2015). This fixed population to trip density conversion is used in Chapters 4-6, while Chapters 7 and 8 include mode choice, so SAV travel demand is a dependent rather than independent variable there.

Urban Form

There are many different ways to consider variations in urban form. For simplicity's sake, this work only considers monocentric cities, and allows for different forms by varying the degree of urban centrality. Trip distributions occur via a normal distribution with a variable standard deviation from the city center, given in km; an infinite (∞) standard deviation corresponds to random trip distribution throughout the city. Smaller standard deviations correspond with more centralized cities, with both origins and destinations thus being more likely to occur in the city center than its outskirts. As opposed to Figure 1, which gave a random initial distribution of 400 vehicles, Figure 4 gives an example of the taxis initially distributed in a centralized city with a standard deviation of 15 km.

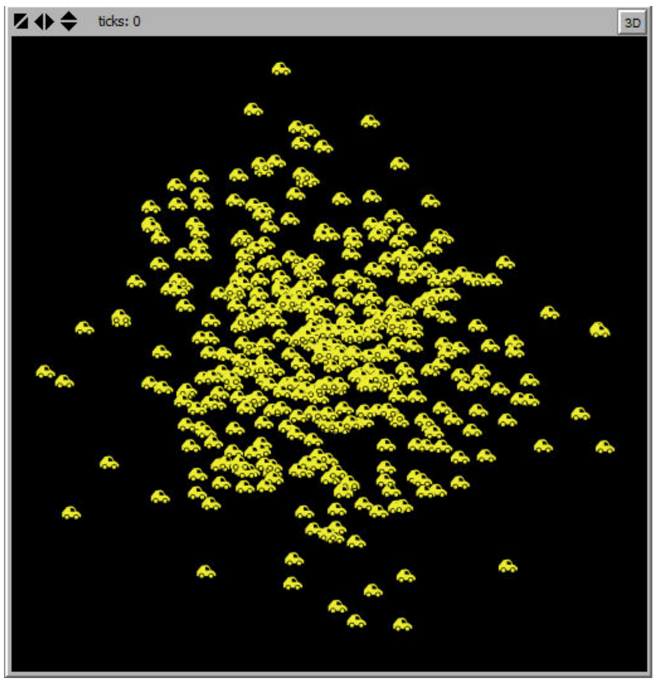

Figure 4: Distribution of Taxis in a Centralized City

Model Outputs

The model considers two basic classes of outputs: times and distances traveled. It records the times spent and distances traveled for both the vehicles and the passengers during each of their states (e.g. assigned and in-transit) as detailed in Table 1. Some of these outputs are effectively fixed by the model assumptions and inputs; pick-up and drop-off times are explicitly fixed at one minute and, and the combination of urban density centrality and city size defines the average trip distance, and by extension average travel time once the fixed speed of 28 km/hr is considered. For example, with a random trip distribution and a city size of 10 by 10 km, average trip distance is 5.23 km and average trip time is 11.3 minutes. Again, the random nature of the simulation means that there can be some slight variation from run to run.

In addition to the time and distance a vehicle travels while taking a passenger, it has an additional relocation requirement; time it takes for the vehicle to travel from its original location to the passenger, when the passenger and vehicle are in the "assigned" state. These relocation distances and times are not fixed and are thus dependent variables. Passenger wait time, defined as the time from when a passenger hails a vehicles to when it is picked up, is also a dependent variable. This work splits these wait times into two parts. First is the *hailing wait time* – the time it takes for a vehicle to be assigned to a passenger, when the passenger is in the "hailing" state. This value is typically zero unless a queue exists when there are more passenger requests than available vehicles. Second is the *relocation wait times* – the time it takes for the vehicle to travel from its original location to the passenger, when the passenger is in the "assigned" state. Additionally, the model records the *relocation percentage*—the additional time/distance that the vehicle travels by itself while empty when travelling to the passenger. For example, a relocation percentage of 30% means that for every mile the vehicle takes with a passenger, it also travels 0.3 miles while empty; fixed speeds means the same is true for additional minutes traveled. This value is given as a percentage to maintain consistency throughout different model runs and for comparisons with other modelling work that have different daily times/distances traveled. It is equivalent to the vehicle/passenger distance ratio described in the previous chapter. A relocation percentage of 30% is the same as a vehicle/passenger distance ratio of 1.3; the vehicle travels 30% more than the passenger.

All of these outputs can be considered individually and in aggregate. For wait time, this dissertation considers average hail, relocation, and total wait time, as well as the 95th percentile total wait time. For distances, the dissertation considers average trip distance and average relocation distances/percentage.

Performance Metrics

Two of the outputs are used directly as metrics: wait time and relocation percentage. Past modelling efforts have typically set wait times as constraints, such as establishing the fleet size necessary to ensure that no person has to wait more than 5 minutes (D. J. Fagnant, 2014). This model treats them as dependent variables, though as stated earlier, will often use a high 95th percentile wait time (e.g. 90 minutes) as a constraint for computational efficiency. Obviously, the passengers of any system would prefer lower wait times, and high wait times could negatively affect the attractiveness of any SAV system once mode-choice considerations are added; if people have to wait 15 minutes for an SAV, private vehicles look far more attractive. Similarly, lower relocation percentages are desirable for the passengers, system operator, and society overall. Fewer “empty” miles mean less relocation wait times for the users, fewer wasted costs to the system operators, and fewer emissions.

The third and final metric is a holistic performance metric, denominated in dollar terms per vehicle and coming in two forms. The first, a “societally optimal” metric, is the broadest measure, looking at the performance of an SAV system from the perspective of both the system operator and the travelers. It combines wait times, relocation distances, and vehicle utilization rates (how often the vehicle is serving a passenger). The second is a “profit-maximizing” metric; it excludes wait time costs, and thus only considers SAV performance from the perspective of the system operator, but is otherwise the same.

Creating a holistic metric is an attempt to consider the overall performance of a system. It is a necessary addition to the first two metrics of wait time and relocation percentage because it also considers fleet size. Taken to an extreme, of one vehicle per person, there would be very low wait times and relocation percentages, and would effectively be modelling a system of individually owned vehicles. One approach to this problem is to determine the smallest fleet size necessary to meet acceptable levels of wait time and relocation percentage, an approach that is also taken by this work to show SAV performance (e.g. the minimum fleet size necessary to maintain sub-minute wait times for each urban density level. However, by directly considering costs, including allowing for "break-even" analyses, the holistic cost metrics help show SAV performance in a more direct way. Furthermore, by considering both the societally optimal and profit-maximizing variants, this work can indicate how SAV systems would be deployed in different ways with different motives (e.g. with or without heavy public sector involvement). The baseline values for the holistic performance metric are given in Table 7.

Table 7: Baseline Holistic Metric Calculation Assumptions

Waiting Time	$10/hour
Vehicle Fixed Cost	$10/day
Profit per Revenue Kilometer	$0.15
Cost per Relocation Kilometer	$0.35

The waiting time is a common value used in transportation economics, and for relocation costs $0.35/km corresponds to the government reimbursement rate of $0.54/mi. The profit per revenue km is harder to determine, but this is in the vicinity of the 10-20% cut that ride-hailing companies take on their rides. The vehicle fixed cost, or how much cost to assign to the vehicle per day even if they are not operating, is harder to justify, and must be taken with the largest grain of salt. However, changing this costs does not actually change the shapes of the holistic metric curves, but only shifts them up or down. Overall, the point of these assumptions is not to make an accurate prediction of the expected profitability or societal value of a shared automated system, but rather to provide a means to show the general trends of how performance changes as a function of fleet size, urban density, and urban form.

OVERVIEW OF MODEL PERFORMANCE

This section provides a brief overview of key model parameters and performance, as well as giving some key conversion equations that are used to convert between different inputs and outputs, such as between urban trip densities to urban population density.

Overall Model Performance

Table 8 first provides a summary of the inputs and outputs as used in Chapters 4-6 on the effects of varying fleet size, urban density, and urban form. Chapter 7 takes the base model described here and add a mode choice component. Table 8 also gives the key assumptions made in this modelling work, including which chapters these assumptions are relaxed, as the fixed 2% SAV mode share that is relaxed in Chapter 7. The corresponding chapters details how the assumptions are relaxed and explains their implications. Overall, the assumptions made here help to simply the model used in this dissertation; the following section justifies the use of a simplified model.

Table 8: Model Inputs, Assumptions, and Outputs

Inputs		Outputs
Dependent Variables	***Derived Inputs***	**Hailing *Wait Times*** $\sim 0 - 12\ min$ $\max\ of\ \sim 60\ min$
Fleet Size: $\mathbf{1.6 - 4\frac{veh}{100\ trips}}$	**Daily Revenue Trips, Hours, and Distance** $\sim 25 - 65\ trips/day\ (per\ vehicle)$ $\sim 6 - 12\ hrs/day\ (per\ vehicle)$ $\sim 130 - 340\ km/day\ (per\ vehicle)$	**Relocation Wait Times** $\sim 0 - 8\ min$
Urban Population Density: $\mathbf{10 - 12{,}500\ ppl/km^2}$	**Urban Trip Density** $\sim 0.8 - 1{,}000\ trips/km^2$	**Total Wait Times** $\sim 0 - 20\ min$ $\max\ of\ 75\ min$
Density Distribution (Standard Deviation) $\mathbf{2 - \infty\ km}$	**Per Passenger Average Fixed Distance and Time** $\sim 5.23\ km$ $\sim 11.3\ min$	**Relocation Percentage** $\sim 5 - 25\%$ $\max of\ 80\ \%$
General Model Assumptions		
Category	***Assumption***	***Relaxed***
Trip Generation	$Origins\ and\ destinations\ randomly\ generated$	Never
Trip Distribution	$Trips\ are\ evenly\ distributed\ throughout\ city$	Chapter 6
Daily Travel Demand Pattern	$Daily\ variations\ in\ demand\ come\ from\ New\ York\ City\ taxi\ patterns$	Never
SAV Travel	$SAVs\ move\ in\ a\ straight\ line\ (no\ street\ network)$	Never
SAV Speed	$28\ km/hr$	Never
Pick-up/Drop-Off Time	$1\ min\ (for\ each)$	Never
Service Area	$10\ x\ 10\ km$	Chapter 6
SAV Mode Share	2%	Chapter 7
Daily Trips per Person	$4\ trips/person$	Never
Holistic Performance Estimations	See Table 7	Chapter 5

Note that the derived inputs column describes values that are obtained via a combination of the dependent variables and the assumptions. For example, per vehicle revenue times, or how many hours a day a vehicle is ferrying a passenger, is a function of the fleet size and the demand profile, so that the same fleet size will always yield the same revenue hours. By comparison, the outputs are true dependent variables. Where a maximum value is given, the range above it refers to the “typical” range of values observed in the results. For example, hailing wait times typically range from 0 to 12 minutes, but in a few instances are as high as 60 minutes. By comparison, for relocation wait times, the typical range is between 0 and 8 minutes, but no maximum is given since they never get much higher than this range. For the per vehicle revenue and per passenger values in the derived inputs, as well as all of the outputs, the table assumes a random distribution of trips.

Converting between Various Input and Output Values

The below equations show some of the key conversions between different input and output values.

Conversion Equations

The following eight equations provide some common conversions used throughout this paper. Note that Eq. 1, which gives the trips per resident, provides a means to convert between trip density and urban density. Since this work makes the initial assumption of 4 total trips per day per person and that SAVs will transport 2% of total trips, it thus assumes that there are effectively 0.08 trips per day for each person living in the city. This is a fixed conversion, but the other seven equations show conversions that are dependent on the model outputs (e.g. relocation percentage) and/or on the variable model inputs (e.g. fleet size).

$$Trips\ per\ Resident = \%\ of\ total\ trips * Daily\ personal\ trips \tag{1}$$

$$Trips\ per\ Resident = 2\% * 4\frac{trips}{person} = 0.08\frac{trips}{person}$$

$$Total\ Daily\ Trips = Urban\ Density * City\ Size * Trips\ per\ Resident \tag{2}$$

$$Daily\ Trip\ Density = Urban\ Density * \frac{Trips}{Resident} \tag{3}$$

$$Total\ Fleet\ Size = \frac{Fleet\ Size * Total\ Trips}{100} \tag{4}$$

$$Total\ Daily\ Revenue\ Hours = Total\ Trips * Trip\ Time \tag{5}$$

$$Revenue\ Hours\ per\ Vehicle = \frac{Total\ Revenue\ Hours}{Total\ Fleet\ Size} \tag{6}$$

$$Total\ Daily\ Operating\ Hours = Total\ Trips * [(Trip\ Time * (Relocation\ \% + 1) + Boarding\ and\ Alighting\ Times) \tag{7}$$

$$Daily\ per\ Vehicle\ Utilization = \frac{Total\ Daily\ Operating\ Hours}{Fleet\ Size} \tag{8}$$

Sample Conversions

The below calculations provide sample conversions using the above equations, using an example case of a density of 2,500 people per km^2 and a fleet size of 2.0 vehicles per 100 trips; this case is simulated in Chapter 4. Note that one of the outputs from this simulation is that this density-fleet size combination corresponds to a relocation percentage of 12%. Note also that all simulations with randomly generated origins and destinations (as done in Chapters 4 and 5) have an average trip time of 11.3 minutes.

$$Total\ Trips = 2{,}500\frac{ppl}{km^2} * 100km^2 * 0.08\frac{trips}{resident} = 20{,}000\frac{trips}{day}$$

$$Daily\ Trip\ Density = 2{,}500\frac{ppl}{km^2} * 0.08\frac{trips}{resident} = 200\frac{trips}{km^2}$$

$$Total\ Fleet\ Size = \frac{2.0\frac{veh}{100\ trips} * 20{,}000\ trips/day}{100} = 400\ veh$$

$$Total\ Revenue\ Hours = 20{,}000\frac{trips}{day} * 11.3\,\text{min} = 3{,}767\ hours$$

$$Revenue\ Hours\ per\ Vehicle = \frac{3{,}767\ hours}{400\ veh} = 9.4\ hours/day$$

$$Total\ Daily\ Operating\ Hours$$

$$= 20{,}000\frac{trips}{day} * [11.3\ \text{min} * (1.12) + 1\ \text{min} + 1\ min)]$$

$$= 4{,}885\ hours$$

$$Daily\ per\ Vehicle\ Utilization = \frac{4{,}885\ hours}{400\ veh} = 12.2\ hours/day$$

BENEFITS OF A SIMPLIFIED MODEL

The model developed in this chapter is intentionally simplified, as compared with most of the comparatively complex SAV models detailed in Chapter 2. This section argues that such a broader, more flexible, and less-computationally intensive approach is best for addressing the research questions laid out in this dissertation. A simpler model is better able to study a wide variety of city types, as it is not tied to a specific city's geometries, street network, travel demand, etc. Lower computational demands are themselves a benefit, as they allow for a sensitivity analyses with large numbers of different simulation runs. This section first provides a more extensive justification for simpler models. It then details the various simplifications that are made, and also includes a few choices where simplification was NOT possible. It concludes by comparing the results of this simplified model to those of previous SAV models as described in Chapter 2, showing how the simpler model provides equivalent results.

Broad Justification for Simpler Models

In general, complex systems theory advocates for a type of Occam's Razor with respect to modelling: When given a choice, the simplest model is usually the best (Railsback & Grimm, 2011). Simpler models are especially useful in understanding new systems with many unknowns;

they are very good at showing and identifying patterns of behavior, though at the cost of less precise quantitative predictions. A truism of modelling is that they are ALL incorrect. It is the responsibility of the modeler to choose the right model for the task, the one that will be the least wrong in the ways that matter, and this dissertation argues that understanding the behavioral patterns of SAVs is of paramount importance at this stage of the technology. Automation is an early stage innovation, a technology that has the potential to be massively disruptive, but as past work has shown, the greater the potential disruption, the harder its effects will be to predict in general (Rogers, 2003). By focusing on a few key components — fleet size, urban density, urban form, and mode choices — a simple model allows for this dissertation to provide deep and detailed descriptions of how each of these can affect overall SAV performance. This dissertation does not pretend to predict the exact optimal fleet size for a given city, or the total expected environmental impact, or the market share in 2030, as a McKinsey Consultancy's report did (Kaas, 2016). However, by describing the overall behavior of SAV systems, and how they could work across different types of cities, this dissertation hopes to provide generalized knowledge on the factors that matter most, the conditions under which SAVs are most likely to succeed.

Additionally, simple agent-based models (ABMs) are especially useful at showing emergence — how a relatively narrow set of rules/behaviors for the actors can lead to large-scale systematic behavior (Axelrod, 1997). One subset of this behavior above are tipping points. These occur when a small shift in input values leads to a huge change in the outputs; the tipping point is the input value — or narrow range of values — around which these large changes in system behavior occur (Granovetter, 1978). A classic example is "Schelling's Segregation Model," which uses a very simple ABM to show that the addition of even very slight preferences to live near people of a similar race (e.g. a person is satisfied as long as they have at least 2 of 8 neighbors of the same race) is sufficient to yield highly segregated housing patterns (J. Zhang, 2011). The extreme version of tipping points are “phase shifts,” which occur when even a miniscule change in inputs leads to a totally different outcome, as with water going from 31 to 33 degrees changing from solid to liquid (Solé, 2011). Identifying tipping points is thus an empirical question, and one that can only be answered via sensitivity analyses. For example, the following chapter examines the effects of varying fleet size. Previous SAV modelling work that looked at a few different fleet sizes indicates that decreases in fleet size from a given optimal cause huge decreases in system performance, but that increases in fleet size from the optimal only bring small increases in performance (D. Fagnant, 2014; Martinez & Viegas, 2016; Spieser et al., 2014). These findings indicate that tipping point behavior *might* exist, but the only way to definitively identify the tipping point and obtain a detailed understanding of the relationship between fleet size and system performance is by testing a wide variety of different sizes (i.e. a parameter sweep). The work shown in Chapter 4 successfully quantifies this pattern, identifying an effective tipping point where queues start to form (i.e. more trip requests than available vehicles).These necessary sensitivity analyses are enabled by the lower computational demands of the simplified models. The fewer choices an agent needs to make, the faster the simulation will run.

Furthermore, this dissertation is specifically interested in showing how SAV performance could vary across different types of cities. The simple model presented here is not directly tied to any one city in particular, so is thus able to better simulate a wide variety of different cities, sacrificing specificity towards a given city for more generalizability. The complex models detailed in the previous chapter often include the specific street network of the city and use existing travel demand models (or modify them) to yield the origin-destination matrices and link-by-link travel speeds. By eschewing a specific street network, or specific travel demand patterns, or specific congestion profiles, the simple model used in this dissertation is not handcuffed to a specific place, and can thus address the issues of varying city type and density that other SAV models have so far been unable to attempt.

Key Simplifications

The model construction and usage sections describe a number of the model simplifications, which are summarized here. Key simplifications include no unassigned vehicle relocation (vehicles moving not in response to a trip request but in expectation of where these trips will be) and no ridesharing (single occupancy vehicles only); both of these are computationally intensive, the latter so especially. Additionally, travel within the system occurs at fixed speeds in open space over straight-line distances, as opposed to on a fixed street network. Requiring a street network would also add computational-intensity, as it would require some amount of routing logic. Additionally, a street network, even a basic gridded one, would make the model represent one type of city more than others. Overall, a model operating in open space can effectively answer all of the research questions. One difference that open space operation does create is direct routing; instead of having to make turns, the vehicle travels directly to the passenger. To accommodate for this, this model uses a somewhat lower speed of 28 kph; if vehicles were operating on a basic gridded network, this would result in an equivalent speed of 34 kph, or 28 multiplied by $(\sqrt{2} + 1)/2$. Ultimately, as shown in the following section, these simplifications do not create markedly different results from those of the more complicated models in Chapter 2.

One simplification that was attempted and then discarded was constant trip generation rates. Instead, the model varies travel demand throughout the day according to a New York City daily taxi travel demand profile. Admittedly, this may reduce generalizability somewhat, especially if someone wants to examine different demand profiles, such as total daily travel demand. However, some variable daily demand is necessary to ensure that the generalized model developed here shows agreement with other more complex models, as a fixed demand profile would have allowed for artificially high vehicle utilization rates. If the vehicles always need to serve an identical demand level, it is easy for them to almost always be in service; modelling runs not presented here bear this out. Using a variable demand profile leads vehicle utilizations in the 7-14 hour range (6-12 revenue-hours when actively serving passengers), which is in line with previous modelling efforts.

While this model varies demand *temporally* in this way, it does not have a corresponding *spatial* variation, where there would be proportionately more trips from the city outskirts to the city center during the morning peak, and more trips from the city center to the outskirts during the evening peak. This is a necessary approach to allow for easier comparison between cities of different centralities; for example, a sprawling city with a fully random trip distribution does not have a true "center."

Comparing the Results of the Simplified Model with Previous SAV Models

Sacrificing specificity for generalizability across a range of different environments is a reasonable tradeoff, especially since the more complex, detailed models described in the previous chapter can have reliability issues do to the high degrees of uncertainty inherent in modelling such an unknown and non-existent system of automated vehicles. Still, these models represent the best currently available understanding of SAV systems. As such, the model described here attempts to provide reasonable and believable results via careful tuning of the parameters in this simplified model, which enables for the end results to be in line with what the more complex models have been providing. Given below, Table 9 takes the summary from the previous chapter that described the general service, context, market, and performance characteristics of SAV models to date (Table 1), and adds the relevant characteristics for the simplified SAV model developed in this chapter. Note that the match-up is relatively good throughout, though since this work treats fleet size and urban density as independent variables, these show a wider range than all of the combined work to date.

Table 9: Comparing Previous SAV Models with Simplified Model Developed Here

SERVICE	Previous SAV Models	Simplified Model
Service Type	Automated Taxi, Automated Shared Taxi, and Automated Shuttle.	Automated Taxi
Vehicle Allocation	Most use nearest available vehicle and satisficing rules for sharing. Some optimization of assignment with respect to different objectives: wait times, profit, relocation distance, etc.	Nearest available, no sharing
Behavior When Not Serving Customer	Most have parking in place, with relocation to serve expected demand where appropriate. Some relocation to dedication stations.	Park in place
Connections with Transit	Mostly none, though some shuttles connect with transit as a first-mile solution. One study explicitly allows for one mode switch for AV to/from transit, and another, non-automated study revamps the transit network to high-frequency bus routes supported by taxis.	None (indirectly addressed in Chapters 7 and 8 on mode choice)

CONTEXT	Previous SAV Models	Simplified Model
City Population and Density	~80,000 to 8 million people. Most used entire city, metro area, or equivalent. ~100 to 27,000 ppl/km^2, most in 1000 to 6000 range.	Varies from 1,000 to 1,250,000 people. Varies from 10 to 12,500 ppl/km^2 (both independent variables).
Street Network and Speeds	Most use either abstract or real street networks, with speeds either fixed or given by a travel demand model, ranging from ~20-60 kph, typically assuming present-day roadway capacities.	No street network, 28 kph
Marginal Cost of Vehicle Automation	Either not reported or a one-time cost of ~$2,000-15,000, with some allowing for higher figures.	Not stated directly; use $0.15 per km profit for passenger distance, $0.35 per km cost for relocation distance.
Average Trip Length	Either not reported or from ~4-10 km, with a min of <2.4 km and a max of 34 km.	5.23 km
Model Type	Aggregate, Network, Agent-Based, or Hybrid.	Agent-Based

MARKET	Previous SAV Models	Simplified Model
Daily Baseline Market Size and Density	Wide ranging, from ~20,000 to 11 million trips, ~1% to 100% of all trips in the region. Densities also wide ranging, from ~ 100 to 10,000 trips/km^2.	No Baseline. SAVs are 2% of total trips. SAVs vary as independent variable from 80 to 100,000 trips, 0.8 to 1,000 trips/km^2.
Alternate Market Sizes	Most use a fixed market size, but some do perform sensitivity analyses using selected different market sizes (i.e. double and half of baseline).	Variable market size; see above.
Fleet Size and Determination	From ~50 to 300,000. Typically determined by a satisficing approach. A few studies vary fleet size from the baseline.	From 2 to 2,800, treated as independent variable.
Mode Choice	Generally not considered. Where mode choice is included, it is typically as a rule-based approach as opposed to a logit model or similar method to estimate user behavior.	Considered in Chapter 7 and 8 as integrated part of ABM.

PERFORMANCE	Previous SAV Models	Simplified Model
Average Passenger Wait Time	Typically under 3 minutes, often under 1 minute. Some evidence of substantially higher peak wait times (i.e. over 15 minutes).	~0-20 minutes, with a max of ~60 minutes.
Cost per Passenger KM	Where reported, from approximately $0.20 to $0.65/km, with some higher outliers.	Varies based on assumptions and measured relocation percentage.
Daily Vehicle Utilization	Taxis and shared taxis generally in motion from ~6-12 hrs/day, most frequently from ~7-8 hrs/day, and serve from ~10-40 trips/vehicle, mostly centering around 30 trips/vehicle. More sharing leads to more utilization to varying degrees; shuttles can serve up to 360 trips/vehicle.	~6-12 hrs/day transporting passengers, ~6-14 hrs/day in motion. Serve ~25-65 trips/vehicle
Vehicle/Passenger Distance Ratio	For automated taxis, generally ~1.05-1.2, with some higher outliers if required relocation to depots. For shared taxis, typically slight reductions from sharing (i.e. from 1.11 to 1.07). Substantially lower values for automated shuttles.	~1.05-1.25
Vehicle Replacement Ratio	Mostly 1 AV replacing ~6-12 traditional vehicles, but numerous higher and lower outliers.	Not measured
Economies of Scale	Only a minority of studies report, showing a wide range from miniscule to moderate economies of scale, including smaller relative fleet sizes as market size increases.	Yes; non-consistent economies due to tipping point behavior.
Environmental Effects	Most report only small reductions in emissions of 0-15%, though a few outliers report massive improvements, especially if SAVs allow for switching to an EV fleet.	Not measured.

Chapter 4 — Effects of Varying Fleet Size

This chapter uses the agent-based model detailed in Chapter 3 to model the effects of varying fleet size for three different urban densities, effectively performing a sensitivity analysis for fleet size. It examines a variety of different performance metrics, including average wait times, relocation percentage, and a holistic metric denominated in dollars per vehicle. It provides a detailed analysis of the phenomenon of queuing and how this can affect performance. It also introduces the concept of density of unassigned vehicles: the average density of vehicles with an "available" status that are free to be assigned to new trip requests. A separate set of modelling runs examines the effects of the density of unassigned vehicles on relocation percentage and wait times, and uses these findings to help explain the patterns seen in the main model runs on the relationships between fleet size, density, and wait time. Overall, the findings of this chapter provide context for and lead to the support of Hypothesis #1, which states that SAV performance increases at a decreasing rate with fleet size.

IMPORTANCE OF MODELLING FLEET SIZE

Setting the appropriate fleet size is a necessary step for all deployment efforts. In light of the inherent uncertainty around the deployment of SAVs, and consistent with this dissertation's goals of broad applicability, the results of this chapter seek to provide a broad guidance of the effects that different fleet sizes can have on system performance. Practically, this is especially important if the initially chosen fleet size is inappropriate, showing the relative risks of both a slightly too high and too low fleet size. As detailed in Chapter 2, most previous modelling efforts used a satisficing approach to choosing a single fleet size, such as determining the smallest possible number of vehicles necessary to ensure that the 95^{th} percentile wait time is less than 5 minutes. Those that do vary fleet size take a quantized approach, picking a few values above and below the baseline fleet size. By performing a sensitivity analysis of a broad range of fleet sizes, the results presented here can investigate the emergence of "tipping points," or narrow ranges across which small changes in fleet size result in large changes in system performance.

This chapter concludes that such tipping points exist with respect to fleet size, and that they can be attributed to queueing. As long as queues are avoided, wait times can remain moderate, and are not lowered much by increases in fleet size. However, wait times can massively increase when queues exist; if the fleet size is small enough to create the presence of even mild queueing, wait times can rapidly become unacceptable. Furthermore, the sensitivity analyses allow for a quantitative investigation of the relationship between fleet size and density, which is continued in Chapter 5. For example, this chapter examines the different shapes of tipping point behavior for different urban densities, showing that the transition between acceptable and unacceptable performance as a function of fleet size is sharper for higher densities. It also shows how the minimum necessary relative fleet size to avoid queueing decreases as density increases. Lastly, this work identifies the density of unassigned vehicles as the main determination of relocation wait time and percentages — the time and distance that the vehicle must travel to get to the passenger after being hailed. It shows that increasing the density of unassigned vehicles decreases these relocation wait times/percentages, but at sharply reducing marginal returns. This phenomenon explains why increasing the fleet size once queueing has been avoided does not substantially increase performance, and also helps justify the main finding of Chapter 5, that overall SAV performance as a function of density also increases with decreasing marginal returns.

Fleet Size Model Runs

This section describes two sets of model runs. The first treats fleet size as the independent variable at three different baseline densities. The second uses an artificially high fleet size to model the effects that the density of unassigned vehicles has on relocation wait times and distances. For both sets of model runs, trip origins and destinations are randomly assigned throughout the 10 x 10 km city.

Simulating the Effects of Varying Fleet Size

Table 10 describes the input values for the primary simulation runs that treat fleet size as the independent variable. Fleet size varies by increments of 0.05 vehicles per 100 trips; for densities of 500 and 2,500 people per km^2, each combination of density and fleet size has 10 simulation runs, and for densities of 12,500 people per km^2, 5 simulation runs. Table 10 also provides the equivalent absolute fleet size in number of vehicles for each of the densities. For these fleet size model runs, runs abort early if the 95th percentile wait time exceeds 90 minutes.

Table 10: Varying Fleet Size Parameters

Density ($\frac{ppl}{km^2}$)	Total Trips/Day	Fleet Size Range ($\frac{vehicles}{100\ trips}$)	Equivalent Number of Vehicles
500	4,000	2.0 – 3.0	80 – 120
2,500	20,000	1.7 – 3.0	340 – 600
12,500	100,000	1.5 – 3.0	1,500 – 3,000

Simulating Effects of the Density of Unassigned Vehicles

Additionally, this section seeks to isolate the effects of fleet size with respect to the density of unassigned vehicles. At any given moment, a given number of vehicles in the fleet are “unassigned” and thus available to be hailed. Logically, the higher the density of unassigned vehicles, the lower the relocation percentage, increasing the likelihood that an available vehicle will be close to the passenger request. To explicitly model this trend, this chapter performs an additional set of simulations, which focus on the effects of the density of unassigned vehicles of by ensuring that at any point in time, almost all of the vehicles in the fleet are unassigned. The results would become messy if the density of unassigned vehicles varied greatly throughout the day, such as by dropping substantially during peak demand. Therefore, these model runs use a very high relative fleet size of 25 vehicles per 100 trips. Trip densities vary to yield an absolute vehicle number from 1 to 1,000 vehicles, or 0.01 to 10 vehicles per km^2. 100 distinct simulations are run and averaged together for each absolute vehicle number. For each passenger trip, when the passenger is assigned to a vehicle, the passenger records the number of available vehicles at the time. These values are then averaged together for all assigned trips to yield the average number of unassigned vehicles, and divided by 100 to give the density of unassigned vehicles for a 10 x 10 km city.

RESULTS OF VARYING FLEET SIZE

This section provides and analyzes the results of the model runs for varying fleet size, including how this affects the three metrics of wait time, relocation percentage, and per vehicle holistic performance. It also uses an in-depth analysis of the wait time results to show the emergence of queueing when there are not sufficient vehicles to meet demand.

Wait Times

The relationship between average wait time and fleet size is shown below in Figure 5.

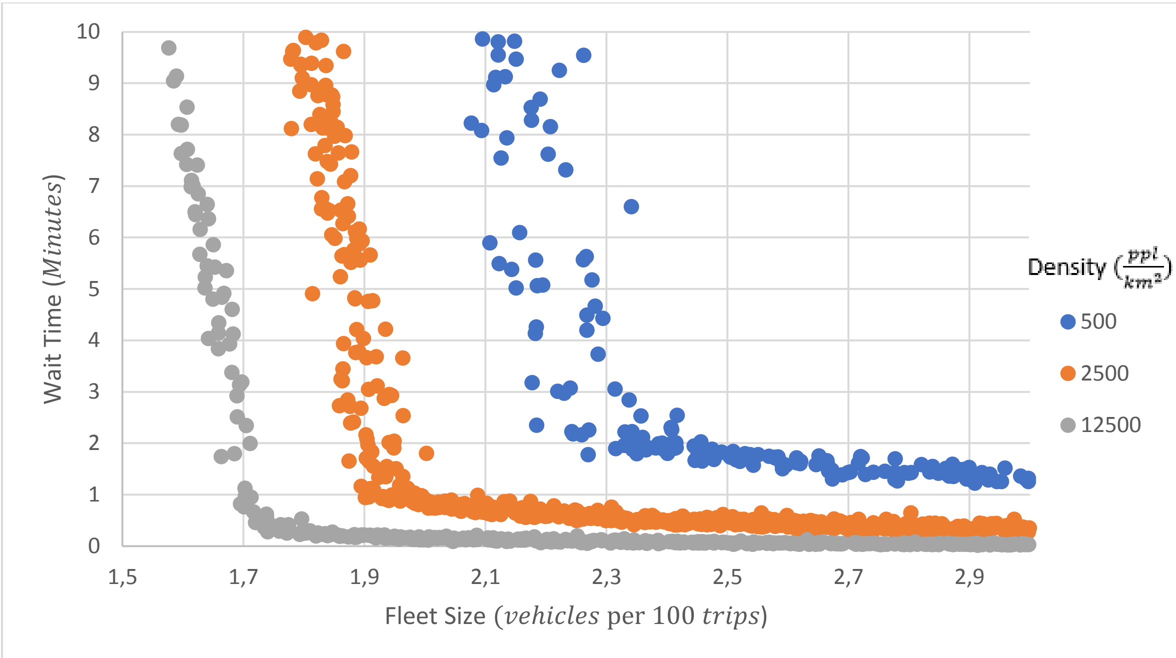

Figure 5: Average Wait Time vs. Fleet Size

What immediately stands out is the extreme non-linearity of performance as fleet size changes; this is tipping point behavior. To take the 2,500 people per km^2 example (orange), an inflection point occurs at a fleet size of about 1.95 vehicles per 100 trips, where wait times are about 1 minute. Fleet sizes smaller than this lead to massive increases in wait time, such as about 10 minutes for 1.8 vehicles per 100 trips. However, increases in fleet size lead to only small reductions in wait time, such as about 0.5 minutes at a 3.0 vehicles per 100 trips.

Effects of Queueing

The tipping point behavior of wait times with respect to fleet size is due to the emergence of queuing, which occurs when there are more passenger requests than sufficient available vehicles to serve them. As is familiar in transportation studies (May, 1990), queuing and the attendant delays will persist until the number of available vehicles is greater than passenger demand. Splitting wait time into its two constituent parts helps show the phenomena of queueing: hailing wait time — the time a passenger waits for a vehicle to be assigned to them —

and relocation wait time — the time the passenger is waiting while the vehicle is driving to them. Figure 6 shows the pattern for both hailing and relocation wait times over the course of a day at a density of 2,500 people per km^2. Two separate simulation runs using different fleet sizes of 1.6 and 1.8 vehicles per 100 trips allow for relatively high and low levels of queuing respectively. The wait times are averaged over ten minute intervals for more consistent results, and the daily travel demand profile is also reproduced in Figure 6 for reference.

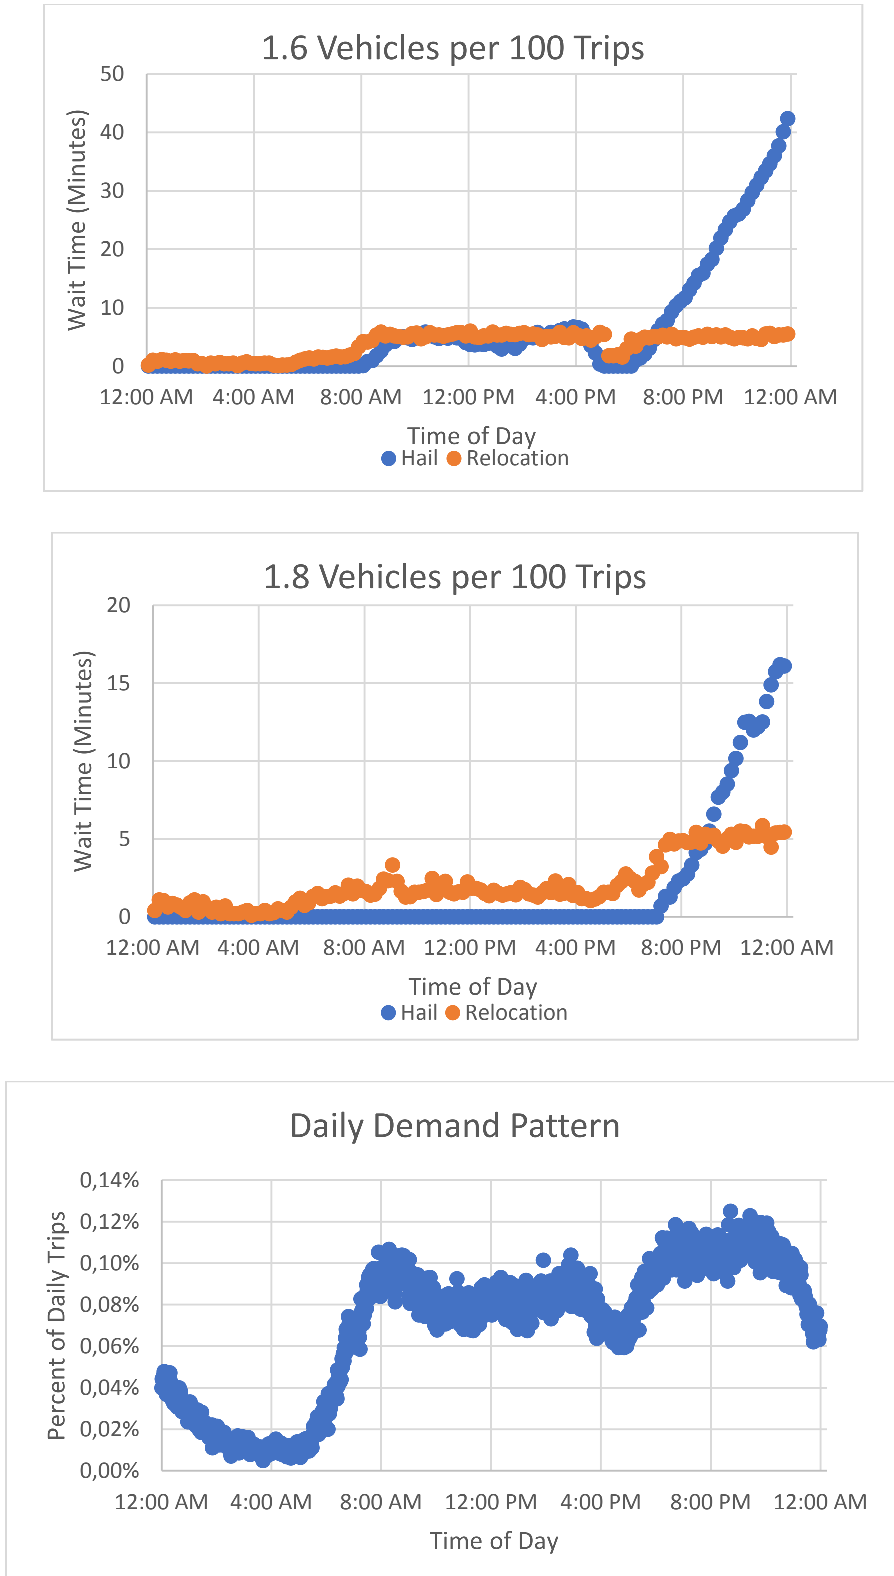

Figure 6: Queueing over Simulation Runs

When no queues exist, hailing wait times are zero; there are more available vehicles than there are hailing passengers, so vehicles are assigned to passengers immediately. Queues will first emerge during peak demand levels, and with fewer vehicles, queues will occur over longer periods. As seen above, the highest demand starts at about 7 pm, which is consistent with where queueing begins with a fleet size of 1.8 vehicles per 100 trips. The smaller fleet size of 1.6 vehicles per 100 trips sees queuing emerge earlier, corresponding with the local demand peak at about 8 am. Queues take time to dissipate, so even when travel demand declines towards midnight, queuing persists. For both cases, when queueing emerges, relocation wait times increase as well, but have a maximum of about 5 minutes; this phenomenon will be explained in the following section on the density of unassigned vehicles.

Queueing can also be seen by splitting the total wait times as a function of fleet size, as presented in Figure 5, into their discrete components of hailing wait times and relocation wait times, as given below in Figure 7.

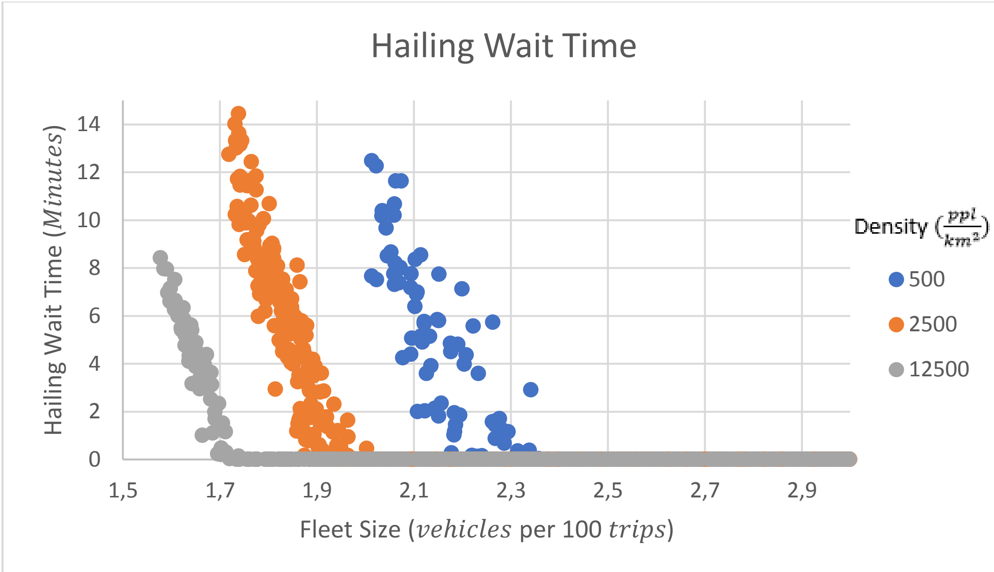

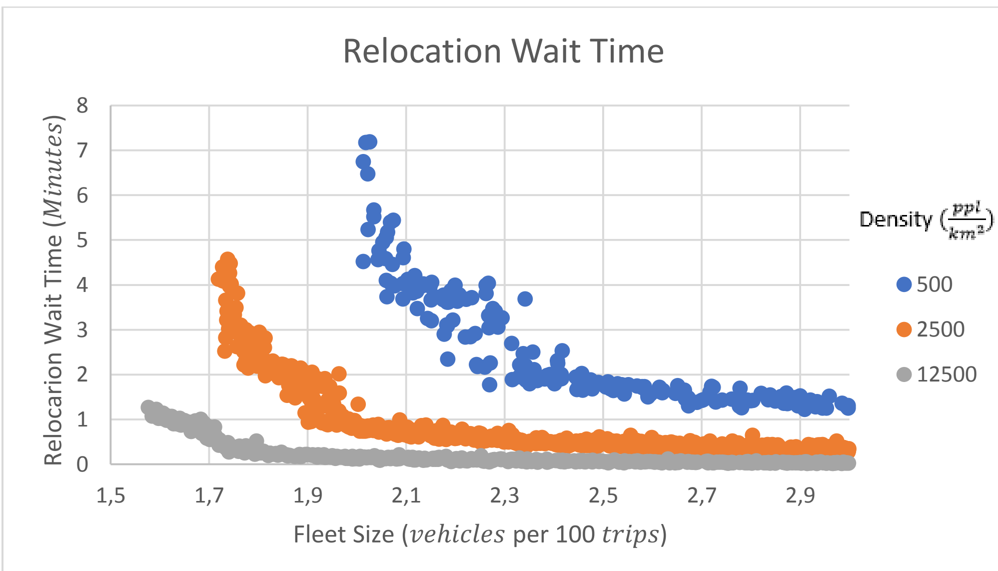

Figure 7: Separating Hailing and Relocation Wait Time

The onset of queuing is marked by the point when hail wait times become non-zero. As seen in Figure 7, queuing begins at fleet sizes of approximately 1.7, 1.9, and 2.25 vehicles per 100 trips for corresponding densities of 12,500, 2,500, and 500 people per km^2. Under conditions of queuing, hail times sharply increase in a similar pattern for the different densities as fleet sizes decrease. Queuing occurs at relatively smaller fleet sizes at higher densities; this pattern is addressed in greater detail in the following chapter. Relocation wait times do not have the same extreme inflection point as hail times, since relocation wait times are still non-zero without queueing.

Relocation Percentage

Figure 8 shows the second performance metric, relocation percentage, and how it varies as a function of fleet size for the three density levels.

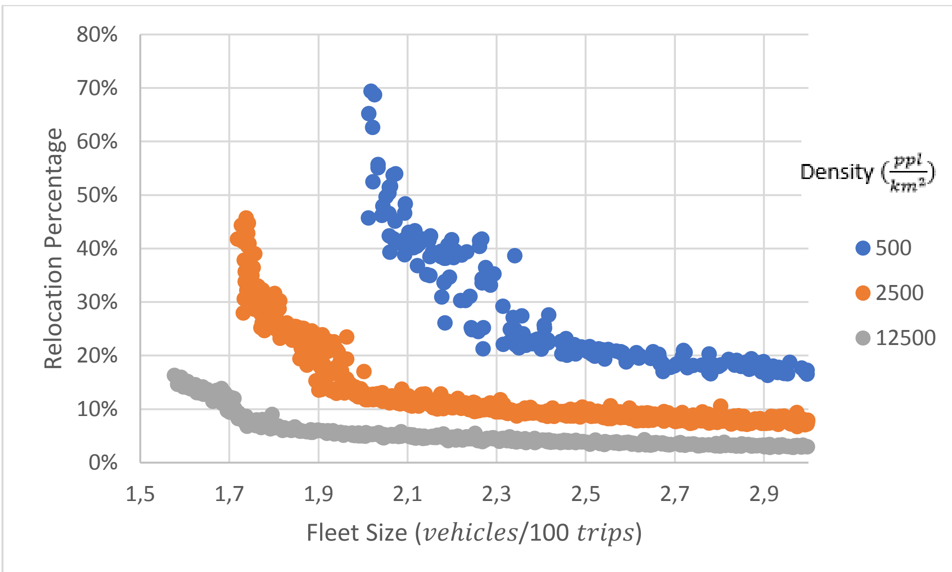

Figure 8: Relocation Percentage vs. Fleet Size

Relocation wait times and relocation percentage are two sides of the same coin, looking at the additional time/distance driven by the vehicles while empty and going to pick up the passengers. This work assumes fixed travel speeds, which create an equivalence between time and distance, hence the identical shape of the curves in Figure 8 for relocation percentage and the curves in Figure 7 for relocation wait time. A key difference though, is that relocation percentage is a scale-free metric, and thus can reasonably be compared against different cities and travel patterns.

Two things stand out in the relocation percentage patterns. First, maximum relocation percentage is lower for higher urban densities; at 500 people per km^2 it approaches 70%, but for

12,500 people per km^2, it never gets higher than 18%. Second, the inflection point for relocation percentage as a function of fleet size is sharper for higher densities. These two phenomena are linked, and can best be understood through the concept of the density of unassigned vehicles, which is addressed in detail in the following section. Briefly, higher densities correspond with higher densities of available vehicles, so that an available vehicle is almost always fairly close to a hailing passenger; this is true even during queueing due to vehicles constantly dropping off vehicles. The main takeaway is that relocation matters less the higher the density; at a density of 12,500 people per km^2, relocation percentages/wait times are always fairly small, so the biggest worry for an effective system is avoiding queues.

Holistic Performance Metrics

As indicated in Chapter 3, a satisficing rule for determining a fleet size to match a given urban density and size would be one that avoids the formation of queuing, which is quite bad for both relocation percentage and wait times. However, a more direct approach is to use a third, holistic performance metric. Two related metrics are used here. The first is profit-maximizing, and comes from the point of view of the system operator, so it ignores the wait time costs. The second is societally optimal, taking a broader view that includes the needs of the travelers and thus includes wait time costs. This societally optimal measurement is intended to isolate the effects of considering wait time, and its limited focus is in line with the goals of using an abstract and simplified model as described in Chapter 3. The true societal cost would likely include other factors, such as congestion and environmental impacts. Though not explicitly modelled here, these and other externalities are qualitatively discussed in the conclusion, Chapter 8.

Figure 9 shows holistic performance as a function of fleet size for two situations. For the societally optimal case, total costs can explode under situations of queueing; both densities of

500 and 2,500 people per km^2 have losses greater than \$250 per vehicle under queueing, but the corresponding figure is abbreviated at a loss of \$100 per vehicle for ease of presentation.

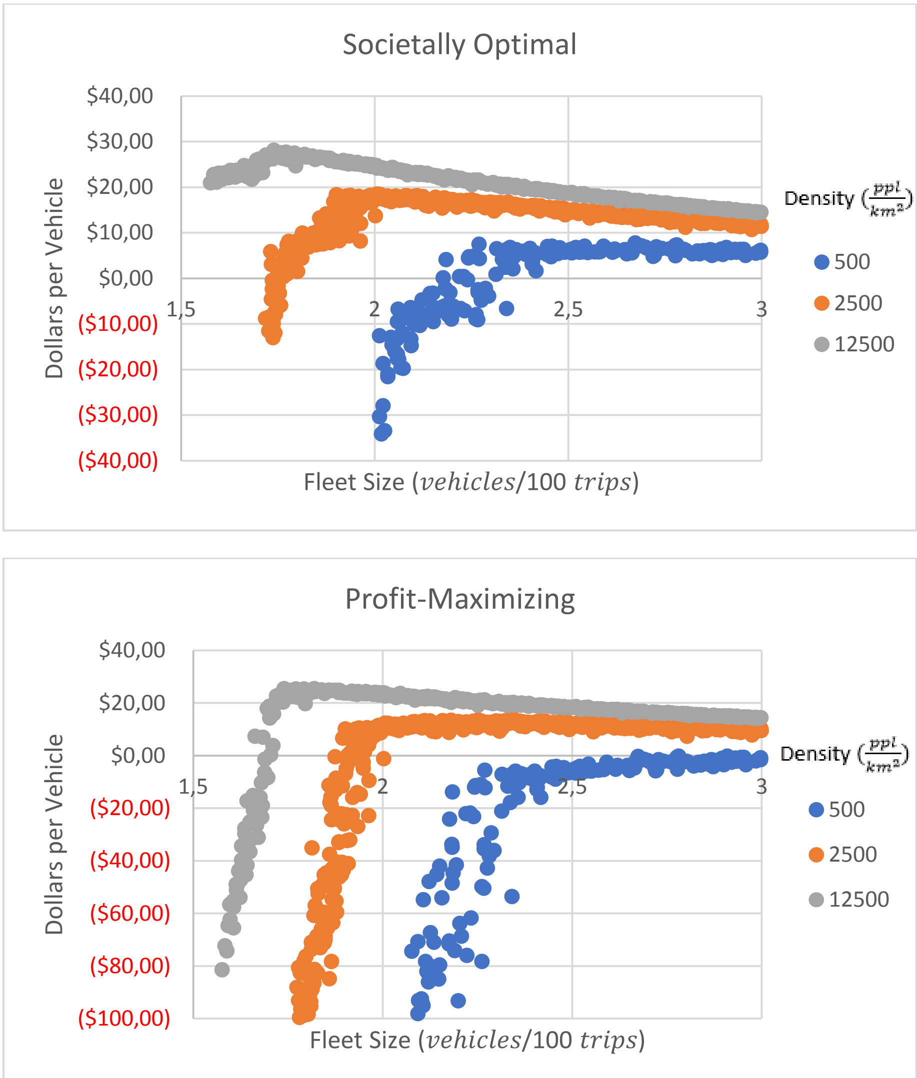

Figure 9: Holistic Performance vs. Fleet Size

Similar to the other performance metrics, there is an inflection point about where queuing begins. Note though, that even in the highest density case, the maximum performance occurs at a density just above that of where queuing begins to occur: approximately 1.73 for both holistic metrics vs. 1.7 vehicles per 100 trips for the emergence of queueing. Lower densities benefit from higher relative fleet sizes. For example, at 500 people per km^2, when waiting costs are considered, the system reaches its profit-maximizing point (this equates to breaking even with these cost assumptions) at its maximum fleet size modeled, 3.0 vehicles per 100 trips. One takeaway from these findings is that if wait times are ignored, and only the profit of the system operator is considered, then fleet size will be set closer to the point where queueing emerges. Therefore, heavier public sector involvement, or other inducements towards a societally optimal solution, would have slightly larger fleet sizes and better at avoid queuing. This phenomenon is also addressed in the following chapter.

Overall, lower densities offer lower holistic performance across all fleet sizes. As will be explained in greater detail in the following section on the effects of density, lower densities require higher relative fleet sizes and by extension, the vehicles in these systems will be operating fewer hours a day. One way to think about the different optimal systems for different densities is that in the profit-maximizing system for a density of 12,500 people per km^2, the vehicles will be ferrying paying passengers on average about 11 hours a day, but for a density of 500 people per km^2, only for about 6 hours a day. This can be explained by two reasons. First, the lower density cities has higher relocation percentages, and thus more time the vehicle must be travelling while empty. Second, the less dense city requires relatively higher fleet sizes to reach an equivalent level of performance.

RESULTS OF VARYING THE DENSITY OF UNASSIGNED VEHICLES

This section provides and analyzes the results of the model runs for varying the density of unassigned vehicles. Since fleet size is set very high for these runs, 25 vehicles per 100 trips, hail times are always zero, and vehicles are sitting empty and available to be hailed over 95% of the time. As such, any wait times are relocation wait times. To focus on the relative degrees of relocation, Figure 10 gives the relocation percentage as a function of the density of unassigned vehicles.

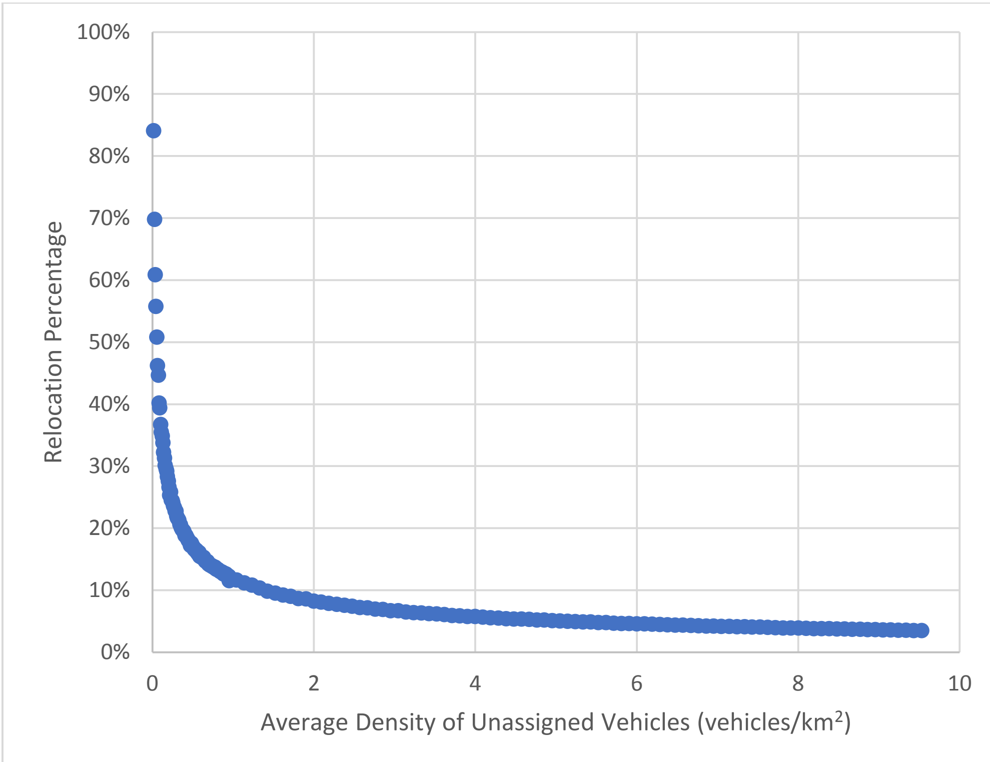

Figure 10: Relocation Percentage vs. Density of Unassigned Vehicles for High Fleet Size Runs

The higher the density of unassigned vehicles, the lower the relocation percentage, but at sharply decreasing marginal returns. Very low densities, less than 0.2 vehicles per km^2, can have relocation percentages above 50%. However, there are very small reductions in relocation percentage after 2 vehicles per km^2; going from 4 to 8 vehicles per km^2 only reduces relocation from 6% to 4%. This relationship between relocation percentage and the density of unassigned vehicles explains some of the patterns seen in the previous section on the relationships between fleet size, density, and wait time as follows:

Average Wait Times are lower for Higher Densities under Non-Queueing Conditions (Figure 5): At low densities, even if there are a sufficient number of vehicles to serve demand, those vehicles will on average have to travel further to the passengers once they are assigned; i.e.

the density of unassigned vehicles is higher. The sharply reducing marginal returns for increasing densities of unassigned vehicles indicates that there are reducing economies of scales for increasing trip density. For example, going from a density of 100 to 200 people per km^2 could decrease average wait times by a lot under non-queueing conditions, but going from 10,000 to 20,000 people per km^2 only by a little. Chapter 5 looks in-depth at the effects of varying density.

Relocation Wait Times Reach a Maximum under Queueing (Figure 6): Recall that at a fleet size of 1.6 and 1.8 vehicles per 100 trips and a density of 2,500 people per km^2, no matter how large the queues became (i.e. how high queueing wait time was), relocation wait time stayed at a maximum of about 5 minutes. Under queueing, all vehicles are effectively in service, and the vehicles will drop off passengers and become available at a relatively constant rate; a given percentage of the fleet becomes available at every minute time step, irrespective of the size of the queue. Therefore, the density of available vehicles, and thus the relocation times/percentages, stays at a relatively consistent maximum under conditions of queueing.

The percentage of vehicles released every minute varies for different cases; it is all the vehicles that are not either dropping off a passenger (fixed at one minute), picking up a passenger (fixed at one minute), transporting a passenger (averages 11.3 minutes) or relocating to a passenger (varies depending on case). Therefore, as a rough guide, for a fleet size of 1.6 vehicles per 100 trips and a density of 2,500 people per km^2, there is a relocation wait time under queueing of about 5 minutes, so the percentage of vehicles that become available every minute is approximately 5.5%, or 1 / (1+1+11.3+5). This in turn corresponds to an average of about 20 vehicles being released every minute and an average density of unassigned vehicles of 0.2 vehicles per km^2.

Relocation Wait Times as a Function of Fleet Size Have Lower Maximums at Higher Densities (Figure 7): As described above, there is a relatively constant density of unassigned vehicles under queueing. Higher urban densities correspond to higher absolute fleet sizes, and thus higher densities of unassigned vehicles and lower relocation wait times. Relocation wait times reach a maximum under queueing, and this maximum is therefore smaller for higher densities and their higher absolute fleet sizes. For example, in Figure 3, relocation wait times reach a maximum of approximately 1.1, 4.6, and 7.2 minutes respectively for densities of 12,500, 2,500, and 500 people per km^2.

Relocation Wait Times as a Function of Fleet Size Have Sharper Inflection Points at Higher Densities (Figure 7): At higher densities without any queueing, adding more vehicles does increase the density of unassigned vehicles. However, since this density of unassigned vehicles is already fairly high, increasing it further only brings very slight improvements in relocation percentage/wait times. By comparison, at lower densities, even if queues are already avoided, adding more vehicles can still decreases relocation percentage/wait time. For all cases, increasing fleet sizes under conditions of queuing decreases wait times by reducing/eliminating the size of the queues. Therefore, the distinction between queueing and non-queuing relocation wait times is more binary for higher densities, resulting in sharper inflection points for relocation wait time as a function of fleet size. For example, for the low density case of 500 people per km^2, even if fleet size is already sufficient to avoid queues, adding more vehicles will have a non-negligible effect on average wait times, thus making the relocation wait time transition less abrupt. By comparison, at a density of 12,500 people per km^2, as long as queueing is avoided, wait times will be close to zero, so adding more vehicles will not produce a substantial benefit. Figure 11 gives the relocation percentage as a function of the average density of unoccupied vehicles for each of the three density levels tested in the main experiment. Note that since these experiments did not have ultra-high starting fleet sizes, the density of unoccupied vehicles varies substantially throughout the day, and the number reported here is simply the average value as experienced by the passengers. Additionally, for consistency, these values are only calculated for situations without queuing: hail times are zero, so all wait times are solely relocation wait times. This overall pattern matches that in Figure 10, but the results appear as truncated because of the wide spread in the three density cases. For example, the average density of unassigned vehicles for the 2,500 people per km^2 case with the highest fleet size is about 4 vehicles per km^2, which is still less than the approximately 5 vehicles per km^2 for the 12,500 people per km^2 case with the lowest fleet size.

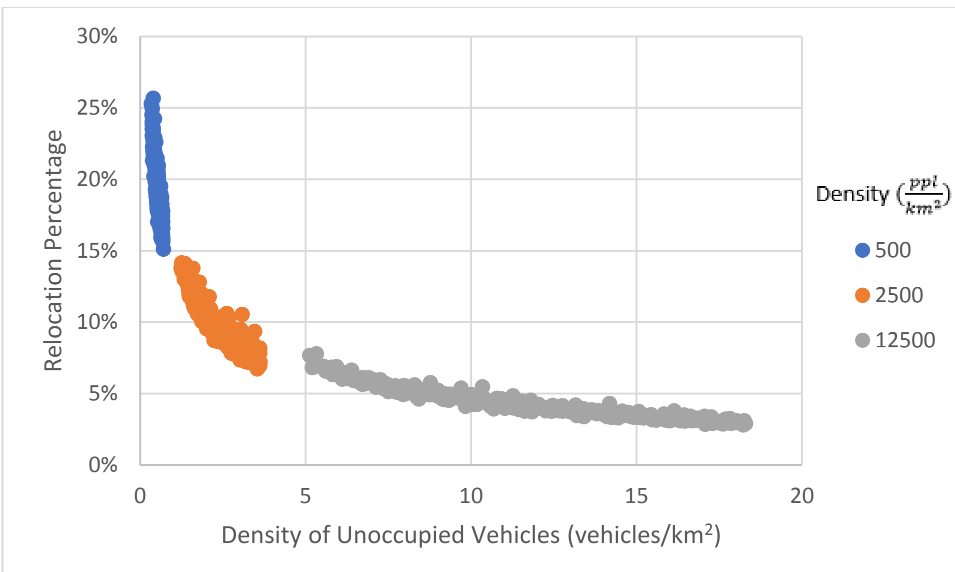

Figure 11: Relocation Percentage vs. Density of Unassigned Vehicles for Primary Runs

CONCLUSION

First, the primary hypothesis for this chapter is supported; SAV performance increases at a decreasing rate with fleet size. More specifically, the emergence of queueing explains the highly non-linear, tipping point behavior of SAV performance as a function of fleet size. Once fleet size becomes small enough for queues to emerge, wait times can skyrocket, becoming worse the longer the queues. Fleet sizes should always be set large enough to avoid queuing; this is the simplest and most important advice that can be gleaned from this chapter. The minimum vehicle limit for any given scenario is more stringent than the maximum vehicle limit; if the former is broached, wait times become excessive, as does relocation percentage to a lesser extent. By comparison, too many vehicles represents an inefficiency; i.e. 100 vehicles don't do a substantially better job than 90.

The results presented by this chapter also show the effects that density has on SAV performance as a function of fleet size; Chapter 5 provides an in-depth look at the effects of density on overall system performance. Most notably, as urban density increases, the required fleet size to avoid queuing decreases. Additionally, at higher densities, the change between queueing and non-queuing wait times is abrupt, representing more extreme tipping point behavior; under queueing, wait times are quite high, but without queueing, wait times are near zero. At low densities however, even when queueing is avoided, there is still a decrease in wait times once fleet sizes are increased. This behavior also explains why there is a narrower range of near-optimal fleet sizes for higher density cases using the holistic performance metric. As a positive, adding vehicles in the lower density cases decreases relocation times and distances, but as a negative, it decreases vehicle utilization. At the higher densities, the effects of decreased vehicle utilization stay the same, but the relocation benefits are much smaller.

Lastly, this chapter introduces the concept of the density of unassigned vehicles; these are the vehicles that are available to be hailed. The density of unassigned vehicles defines the average relocation percentage/wait time. The higher the density of unassigned vehicles, the more likely a vehicle will be near a passenger request, and thus the shorter amount of distance and time it must travel while empty to pick up the passenger. Low densities of unassigned vehicles correspond to high relocation percentages, such as 50% at 0.055 vehicle per km^2, but increasing the density of unassigned vehicles brings sharply decreasing marginal returns; relocation percentages are always less than 10% for densities of unassigned vehicles of more than 1.8 vehicles per km^2.

Chapter 5 — Effects of Varying Urban Density

This chapter uses the agent-based model detailed in Chapter 3 to model the effects of varying urban density, performing a sensitivity analysis of densities ranging from 10 to 12,500 people per km^2. For each density level, it also simulates a range of fleet sizes. The results are split into three groups. First is an examination of the required fleet size for different density levels to achieve a given wait time. Second is an examination of the relationship between relocation percentage and density, using a number of different wait time constraints. The final group looks at how holistic performance, denominated in dollars per vehicle, varies as a function of density. This includes an examination of the effects of varying assumptions used to calculate holistic performance — different costs per relocation km, profit per revenue km, and consideration of passenger wait time — and the identification of "break-even points:" the fleet size and density combination where SAV systems start to give positive dollar values. Ultimately, the findings presented in this chapter help explain and support Hypothesis #2, which states that SAV performance increases at a decreasing rate with urban density.

UNDERSTANDING AND MODELLING URBAN DENSITY

This section is split into two parts. The first describes how the urban density used in this abstract model compares to that of real world cities, and the second gives a broad justification of the importance of urban density as a key parameter to understand SAV performance.

Comparing Density in This Abstract Model to Real World Cities

This chapter treats density as an independent variable, varying it from very low (i.e. 10 people per km^2) to very high (i.e. 12,500 people per km^2) in a hypothetical and abstract city with constant density throughout and an area of 100 km^2 (i.e. 10 x 10 km in size). Chapter 6 looks at another component of density — urban form — seeing how SAV system varies as urban density goes from constant, as described in this chapter, to more concentrate around the city center. Isolating these two effects is necessary, as density is an inherently heterogeneous term and difficult define term, especially with respect to the scale of measurement. For example, in the US, New York is typically held as the paradigm of dense urban living, and Los Angeles as the quintessential sprawling metropolis. Yet measured across the entire metropolitan statistical area, the average population density of Los Angeles at 1,020 people per km^2 is only slightly less than New York's 1,090 people per km^2. Looking at a slightly smaller metropolitan region, LA can actually have a higher density, at 3,090 vs. 4,250 people per km^2. While at first glance this is strange, it occurs because LA has fairly consistent population medium density within its city limits: lots of single family houses on micro-lots or multiple apartments in 2-3 story buildings. By comparison, though Manhattan itself is super dense 27,812 people per km^2 the far larger Staten Island is very suburban (3,080 people per km^2). In the same way, sprawl cannot just be defined as low-density. A lack of clustering, as implied by LA's more consistent urban form, can be one aspect, but far from the only one, and current researchers agree that sprawl is at best multi-dimensional and at worst a fully fungible concept (Hamidi, Ewing, Preuss, & Dodds, 2015).

This uncertainty creates practical problems for planners. Saying that denser cities are better able to support public transit is fine, or even broader, saying that a city should encourage more density. However, specific blanket statements are inherently insufficient. For example, it is impossible to accurately claim, as some have done (Demery, Higgins, & Setty, 2007), that cities denser than a given threshold can support subway service, since the type and distribution of density clearly matters; Cervero & Guerra (2011) provide a good overview of this topic. One possible solution is the use of weighted density, a measure of the average experienced density by people in a region; it is obtained by measuring the density of each population tract (a unit of measurement in the US census, each comprising about 4,000 people) and weighting its overall effect by population. This can help provide better comparisons between cities of different types, but requires access to the raw census data. Therefore, instead of seeking a holistic metric combining both density and urban form, this dissertation splits them into two variables.

Chapter 3 addressed and justified the decision of this dissertation to forgo specificity to a given city in order to make broader claims on the variations in SAV performance. However, it is still necessary to provide a means to compare the density results of this dissertation with those of real world cities. Importantly, as a relatively small service area (100 km^2) the findings here are best compared against roughly equivalent areas; as described above, moving to larger regions and considering the more sprawling regions can have an outsized effect on overall city density. For example, the two lowest density areas studied by other automated taxi (i.e. SAV) models reviewed in Chapter 2 are 675 people per km^2 for Sioux Falls (Horl, Erath, & Axhausen, 2016) and 467 people per km^2 for the entire state of New Jersey (Zachariah, Gao, Kornhauser, & Mufti, 2014). It is hard to say that New Jersey is actually less dense that Sioux Falls — it is the densest state in the nation, in fact — but it includes rural farmland in addition to cities like Hoboken and Newark. As a means of comparison, Table 11 shows the densities of some selected cities, focusing on areas that have sizes somewhat close to 100 km^2 and that were modelled by previous

SAV work as detailed in Chapter 2. Note that the range studied here goes far lower than the lowest density areas previously modelled, but its maximum density is not quite as high as for Manhattan. As addressed below, this is a reasonable accommodation, since ultra-high densities are computationally intensive and because there are not huge differences in SAV performance once density gets to these density levels.

Table 11: Density of Selected Real World Cities

City	Density (ppl/km^2)	Size (km^2)
Abstract Model	$10 - 12{,}500$	100
Sioux Falls	675	125
Ann Arbor (Region)	846	337
Atlanta	1,288	348
Austin Region	1,297	718
Ann Arbor (City)	1,500	74
Minneapolis	2,959	151
Miami	4,866	145
San Francisco	6,226	120
Lisbon	6,678	85
Manhattan	27,812	59

Cities that were considered by other SAV models are highlighted in yellow.

The Importance of Urban Density

That shared mobility services perform better in higher density areas is an eminently reasonable assumption. Taxis and transit certainly benefit from higher densities, as having more people near the transit routes or taxis reduces inefficiencies and the need for detours and deadheading, ensuring that the people are already somewhat close to the vehicle and encouraging higher utilization of these shared services (Cervero & Guerra, 2011). SAV research to date, as detailed in Chapter 2, has generally backed this claim that higher densities improve SAV performance. For example, some papers show moderate economies of scale with increasing demand levels in a given city. Such economies were shown to take various forms, such as reduced wait times or smaller relative fleet sizes — e.g. a doubling of travel demand only requires 90% more vehicles to provide an equivalent level of service. However, the detailed and city-specific models of these previous efforts have prevented an in-depth consideration of the effects of urban density on SAV performance; such a study is described in this chapter.

In particular, this chapter is interested in showing the relative differences in performance of SAV systems across a wide range of different densities. It looks at ultra-high density areas, as have previous studies (e.g. with Manhattan or Singapore as modeled cities), but also looks at far lower densities, down to 10 people per km^2. By looking at densities from very low to very high, this chapter shows the types of cities where SAV systems can reasonably perform. Notably, the findings indicate that very low densities will struggle to support SAVs. However, due to the decreasing economies of scale, there is not nearly as big a benefit in going from medium density to high density as there is going from low to medium. This represents a form of tipping point behavior, albeit not as extreme a version as shown with fleet size in the previous chapter. Roughly stated, the tipping point occurs over the range of approximately 100 to 400 people per km^2 (equivalent to 8 to 32 trips per km^2 per day); below this range, performance drops off precipitously, and above this range, improvements in performance slows down. These findings could help aid cities that are planning to develop shared mobility systems, showing that while shared mobility in rural communities may not be feasible, they can function reasonably well in both cities like Ann Arbor and Manhattan, and that the difference between those two is not be as high as some would suspect.

URBAN DENSITY MODEL RUNS

Table 12 provides the input values for the urban density model runs, specifying both the density and the fleet size. As with Chapters 4, trip origins and destinations are randomly assigned throughout the 10 x 10 km city. The runs abort early if the 95th percentile total wait time exceeded 30 minutes; this is a sharper cutoff than for the 90 minutes for the fleet size runs. High wait times and the queues they represent are computationally intensive, and since the previous chapter established that very high wait times are undesirable, a sharper cutoff is a reasonable accommodation to allow for a larger number of model runs.

The general approach taken in these runs is to test a wide range of different densities, each density level with given range of fleet sizes, and each density-fleet size pair with a given number of identical simulation runs. The urban density is split into three basic groups: from 10 to 200, 200 to 2,500, and 2,500 to 12,500 people per km^2. For the sensitivity analysis the lower density group varies the density by smaller steps of 10 people per km^2, and the medium and higher density groups by larger steps of 100 and 500 people per km^2 respectively. The fleet size range for each density group have values informed by the modelling results in the previous chapter, so that the emergence of queuing always occurs within the given range. Fleet size varies in steps of 0.05 vehicles per 100 trips for all cases. Due to computational demands, the higher density groups have fewer identical simulations: two simulation runs for each density-fleet size pairing from 2,500 to 12,500 people per km^2, but ten runs for 10 to 200 people per km^2. All of these identical simulations are averaged together to yield more consistent results. In total, Table 12 represents 12,549 simulation runs, which together took approximately 60 hours to run. Lastly, recall that Chapter 3 provides equations to convert between the different values presented here, such as between the fleet size in vehicles per 100 trips to absolute fleet size, or between urban density, trips per day, and trip density. For the latter, the total trips per day equivalent to the urban density is also provided in a column in Table 1. These trips per day, divided by 100, also provides the trip density in trips per km.

Table 12: Varying Urban Density Simulation Runs

Urban Density Range (ppl/km^2)	Equivalent Trips/Day	Vary Density by (ppl/km^2)	Fleet Size Range $(\frac{vehicles}{100\ trips})$	Total Runs
$\mathbf{10 - 200}$	$80 - 1{,}600$	10	$2.1 - 4$	7,800
$\mathbf{200 - 2{,}500}$	$1{,}600 - 20{,}000$	100	$1.6 - 3.5$	3,625
$\mathbf{2{,}500 - 12{,}500}$	$20{,}000 - 100{,}000$	500	$1.5 - 2.8$	1,124

RESULTS OF VARYING URBAN DENSITY

This section provides and analyzes the results of the model runs for varying urban density, including how these variations affect both hailing wait time (the amount of time a person waits to be assigned to a vehicle) and total average wait time (the combination of hailing wait times and relocation wait times, or the amount of time it takes for the vehicle to travel to the person. The additional consideration of hailing wait times to total average wait times helps answer questions regarding the emergence of queuing, which as detailed in the previous chapter is a major factor in determining the performance of SAV systems. For example, the fleet sizes where queuing just starts to emerge for a given urban density can be identified as where average hail time starts to become non-zero. As with the previous chapter, these results also include analyses of how changing urban density affects relocation percentage and per vehicle holistic performance. Since the model runs vary both fleet size and urban density, the basic approach for presenting the results is to show the density-fleet size pair that corresponds to a given performance metric, such as the range of pairs that result in total average wait times of around 1 minute. Ultimately, the combination of these results shows that changing urban densities exhibits tipping point behavior, albeit not as extreme as varying fleet size.

Fleet Size

Figure 12 gives the relationship between density and fleet size for some hail time intervals, indicating the fleet sizes at which queueing just begins for different density levels. Notably, there is heavy overlap between the intervals (average hail times of less than 0.1 minutes, 0.1 to 0.5 minutes, and 0.5 to 1.0 minutes), since as would be expected from the fleet size analysis, relatively small changes in fleet size will bring about outsized changes in hail times, and by extension, total wait times. Note that the large marginal changes in performance at low densities has a strong tradition in queuing theory. The queueing phenomenon described here and in the previous chapter can be seen as a form of resource pooling, where "customer delay is distributed as in a system where there is a single queue with multiple servers." (Kelly & Laws, 1993) In general, the non-linear performance of systems with congestion is well established in queueing theory, but as described in Chapter 2, prior SAV work has not focused on queuing.

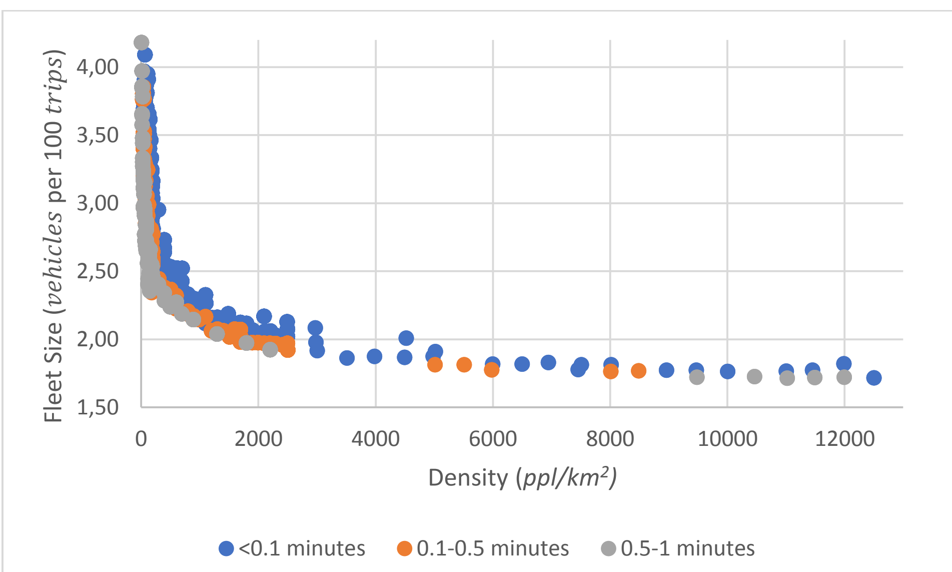

Figure 12: Fleet Size vs. Density for Hail Time Intervals

Overall, these findings indicate that some degree of tipping point behavior exists at densities of around 100 to 400 people to km^2; required fleet sizes to avoid queuing massively increases towards the lower end of this range, but the required fleet sizes only increase slightly for densities greater than this range. For example, from 20 to 400 people per km^2, the fleet size at which queueing occurs drops massively from about 4.3 to 2.4 vehicles per 100 trips, but with a much larger density increase of 400 to 12,500 people per km^2, the queueing fleet size only drops from 2.4 to 1.7 vehicles per 100 trips. As established in the previous chapter, while SAV systems may want a fleet size somewhat larger than the one for which queueing emerges, they definitely do not want one smaller.

Figure 13 looks at the same relationship between fleet size and density for total wait time intervals. The same basic trend of decreasing marginal returns at increasing densities persists, but it is interesting to see the differences between the different intervals. The high wait time interval (5-10 minutes) shows a similar pattern to the hail time one in Figure 12. This makes sense, since high wait times tend to occur under queuing situations. By comparison, the smallest interval, for wait times around 15 seconds (between 0.225 and 0.275 minutes), where throughout the range of densities considered, increasing density allowed for non-negligible reductions in effective fleet size. Therefore, tipping point behavior becomes more extreme the less importance society (or the users, or the system operators, etc.) places on wait times. For example, if wait times of five minutes are acceptable, as the system is seen to function at an acceptable level as long as queues are avoided, then the tipping point behavior is fairly extreme. Density here is something a binary question; as long as the density is greater than 600 people per km^2 or so, then the system will be seen to operate as a reasonable level, and further increases in density will not improve system performance by much. However, if society places a high value on wait times, such as by expecting average waits to be less than one minute, then tipping point behavior is less evident, and increasing density brings steady improvements in SAV system performance.

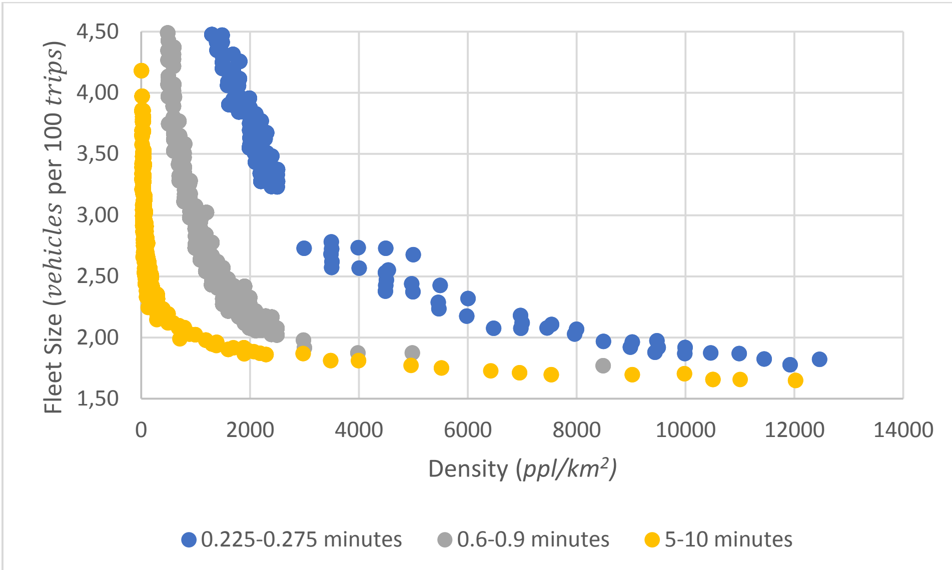

Figure 13: Fleet Size vs. Density for Total Wait Time Intervals

Relocation Percentage

Figure 14 gives the relationship between relocation percentage and density for hailing time. In general, except at densities below 500 people per km^2, additional relocation percentages tend to be from about 7% to 25% more than the passenger distances driven, with more of the relocation differences falling towards the lower end of that range. This matches up with the results of the previous section, showing that tipping point behavior occurs roughly over the range of 100 to 400 people per km. As with Figure 12, which looks at the relationship between fleet size and density for hail time intervals, there are not major differences between the different hailing wait time intervals; i.e. the measured relocation percentages at a given density are only slightly higher when average hailing wait time is between 0.5 and 1 minute as when it is non-zero but less than 0.1 minute.

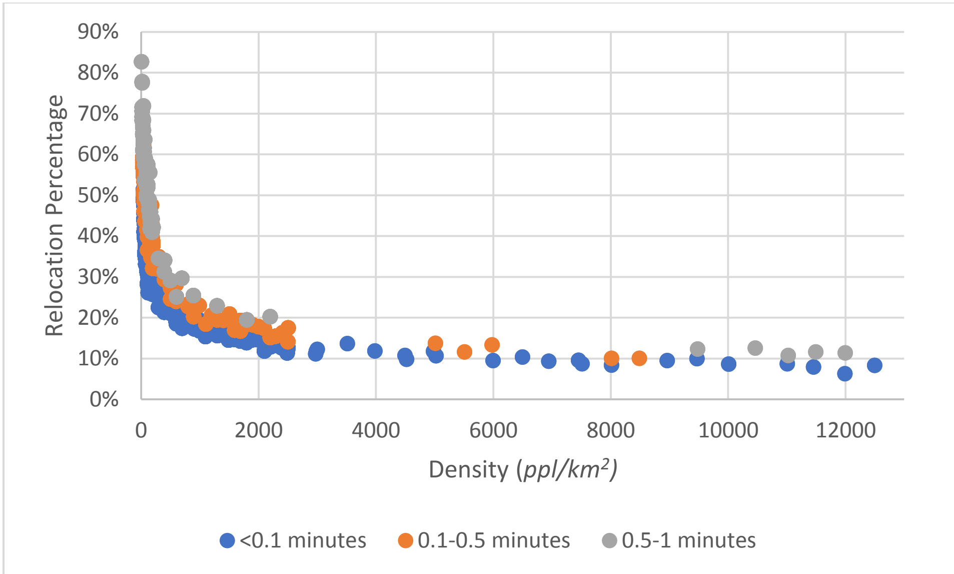

Figure 14: Relocation Percentage vs. Density for Hail Time Intervals

Figure 15 gives the relationship between relocation percentage and density for average wait time respectively. For the highest wait time interval of 5-10 minutes, there is an inverse relationship between relocation percentage and density. However, for the smaller intervals of 0.225-0.275 and 0.6-0.9 minutes, the curves are almost flat; corresponding to about 8% and 11% respectively for all densities. Additionally, wait time intervals of 0.225-0.275 and 0.6-0.9 minutes can only be achieved with densities greater than approximately 400 and 1400 people per km^2 respectively. These patterns can be explained through the dual concepts of unassigned vehicle density and hailing vs. relocation waiting times, both of which are described in greater detail in Chapter 4. While the higher wait times will include some hailing times, these sub-minute wait times mean that there are enough vehicles present so that no queueing occurs. As such, all of the wait times are relocation times, determined by the density of unassigned vehicles. Relocation wait time is thus fully determined by the relocation percentage (recall that they are two sides of the same coin), so is unaffected by density. As long as no queueing occurs, and all wait times are relocation wait times, relocation percentage will be constant as a function of density for given wait time intervals. Ultimately, at higher densities, as long as queueing is avoided, there are typically enough available vehicles so that one is on average quite close to any trip request. However, at lower densities, even when queueing is avoided, there are simply not that many vehicles available. The fleet sizes necessary to achieve these wait time intervals can and do vary as a function of density, such as evidenced by Figure 13.

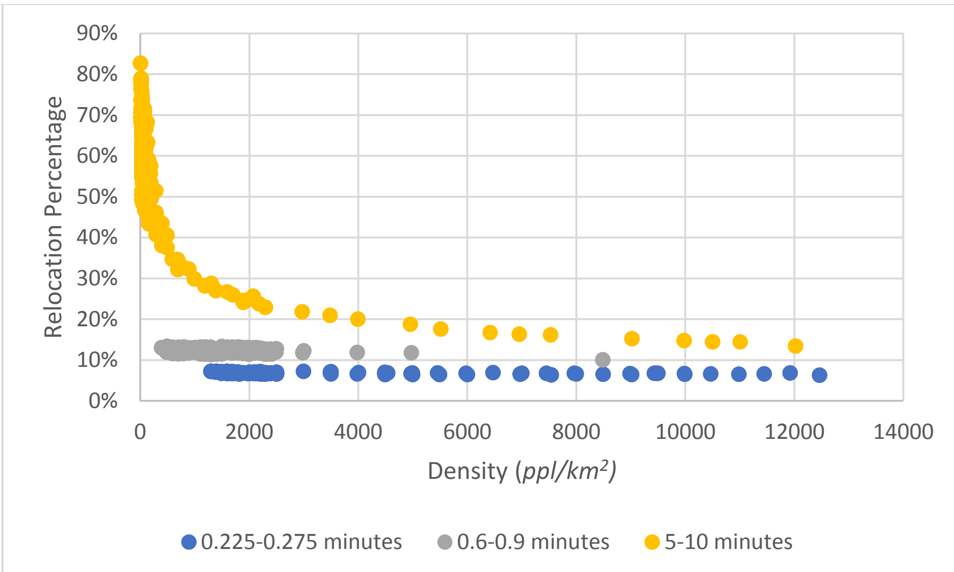

Figure 15: Relocation Percentage vs. Density for Total Wait Time Intervals

Holistic Performance Metrics

This analysis has the same baseline cost assumptions as for the fleet size holistic performance analyses in Chapter 4, with the addition of considering cases of halving and doubling profit per revenue km while keeping the cost per relocation km constant at $0.35, and the reverse of halving and doubling the cost per relocation km while keeping the profit per revenue km constant at $0.15. The goal of these variations is to show how different cost assumptions can vary the holistic performance of the SAV system. These values are given in Table 13, with the base case values given in bold.

Table 13: Varying Holistic Metric Assumptions

Waiting Time	**$10/hour**
Vehicle Fixed Cost	**$10/day**
Profit per Revenue Kilometer	$0.075, **$0.15**, $0.30
Cost per Relocation Kilometer	$0.175. **$0.35**, $0.70

For each density level modeled, from 10 to 12,500 people per km^2 as laid out in Table 7, the fleet size that gives the maximum performance per vehicle is chosen. As such, at each density level, the performance per vehicle represents the best performing system with the optimal fleet size, based on the above cost assumptions. All figures include both the societally optimal and profit-maximizing and societally optimal metrics, i.e. where wait times are considered and excluded as a cost respectively. As with the holistic performance figures in Chapter 4, for ease of presentation some of the charts are cut off for heavy losses which occur at very low densities. For example, in the base case for societally optimal performance (i.e. considering wait times as a cost), the highest performance that could be obtained for a density of 10 people per km^2 is actually a loss of $67.52 per vehicle.

Three related analyses are presented below. The first looks at the base case as a function of different densities, including a quantification of the optimal fleet sizes for the different densities. Second is an investigation into the "break-even" point: the combined density and fleet size where per vehicle SAV performance goes from negative to positive dollar values. The third looks at the effects of varying the cost and profit assumptions from the base case.

Base Case Holistic Performance

Figure 16 gives the per vehicle holistic performance as a function of density for the base case values. It provides both holistic performance metrics: profit-maximizing and societally optimal. The general trend of performance as a function of density are similar for these holistic metrics as those shown above for fleet size and relocation percentages: a sharp increase in performance as we increase from very low densities (i.e. from 0 to 200 people per km^2), smaller increases for moderate densities (i.e. from 200 to 2,000 people per km^2), then a relative leveling off at higher densities (i.e. over 2,000 people per km^2). Where wait time costs are included (societally optimal), there is a more consistent increase in holistic performance as a function of density (i.e. less extreme tipping point behavior) than when wait time costs are excluded (profit-maximizing), though the differences are slight. For example, going from 1,000 to 2,500 people per km^2 yields an increase from \$5.03 to \$11.41 in the societally optimal metric, but only from \$12.44 to \$17.41 for profit-maximizing.

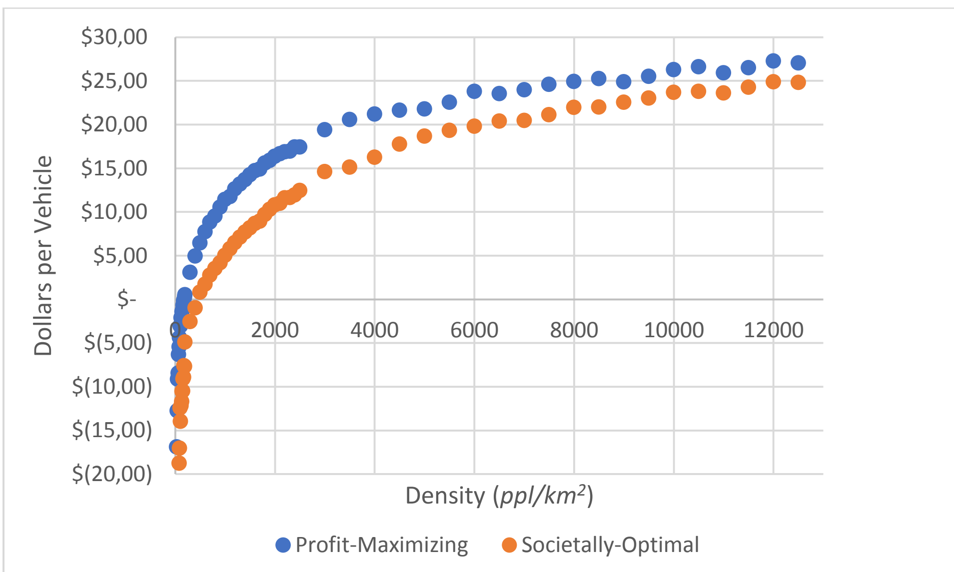

Figure 16: Holistic Performance vs. Density for the Base Case

Figure 17 shows that the optimal fleet size decreases as a function of density; recall that fleet size for these model runs is never greater than 4.0 vehicles per 100 trips, so that is the upper bound of the figure. The curves are not perfect for two reasons. First are the inherent variations between simulation runs with identical inputs and second are the fairly gentle slopes around the maximum holistic performance as a function of fleet size (i.e. relatively large changes in fleet size bring only small changes in the dollars per vehicle). Additionally, though not directly shown in the figure, for higher densities optimal fleet size tends to be relatively close to the point where queueing starts to emerge, and for lower densities, the optimal fleet sizes tend to be higher than fleet size where queues emerge. Lastly, note that the societally optimal fleet size is always greater than the profit-maximizing fleet size, and that the difference is greater for lower densities. Therefore, planners should be careful of wholly privately owned and operated systems using sub-optimal fleet sizes, especially in lower density areas (i.e. less than 1,000 people per km^2).

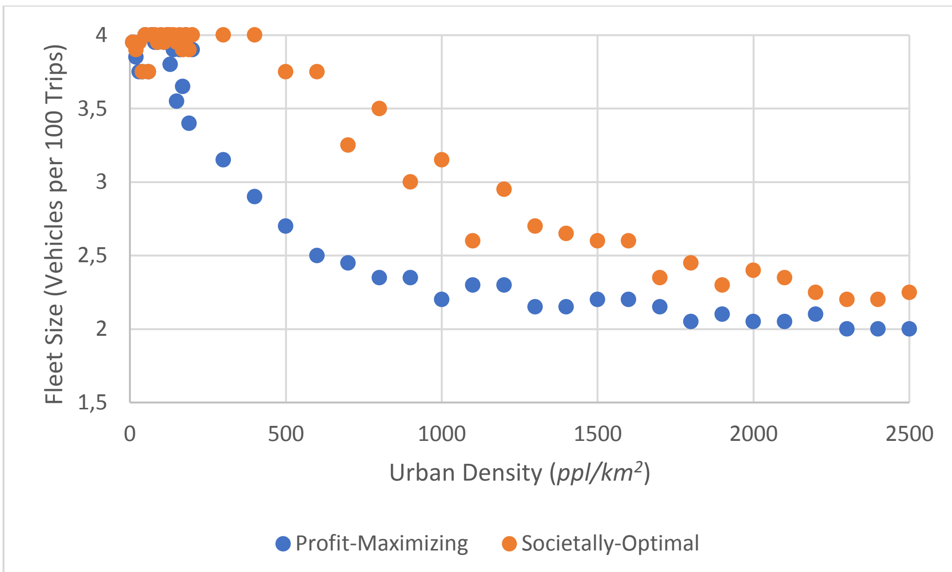

Figure 17*: Optimal Fleet Sizes for Different Densities*

Minimum Viable Fleet Size and Density

There is no definitive minimum viable fleet size and density, as these values will depend on the assumed costs and profits of SAV service. However, broad generalizations can be made. As shown in Figure 16, for the base case, the “break even” point occurs at approximately 150 people per km^2 when using the profit-maximizing metric and 500 people per km^2 with the societally optimal metric. For each of these two cases, Table 14 gives the equivalent absolute fleet size, fleet size in vehicles per 100 trips, total daily trips, and daily trip density, using Figure

17 to provide the optimal fleet size (i.e. the fleet size that produces the highest dollars per vehicle).

Table 14: Break-Even Fleet Sizes and Densities

	Profit-Maximizing	Societally Optimal
Urban Density (ppl/km^2)	150	500
Fleet Size $(\frac{vehicles}{100\ trips})$	4.0	3.75
Fleet Size $(\#vehicles)$	48	150
Total Daily Trips	1,200	4,000
Daily Trip Density $(trips/km^2)$	12	40

The following subsections show how the break-even density can vary with cost assumptions, as laid out in Table 13. However, the results presented in Table 14 provide reasonable worst-case and best-case scenarios. It is highly unlikely that fleet sizes much below 48 vehicles and densities below 150 people per km^2 can provide a sustainable business model. By comparison, fleet sizes greater than 150 vehicles and densities greater than 500 people per km^2 are much more likely to allow for a profitable and effective service. These are smaller fleet sizes than most previous SAV models have used, as detailed in Chapter 2, but will likely be a reasonable hurdle for any prospective system operator; one would not be able to simply purchase a handful of automated vehicles and effectively run an SAV system.

Effects of Varying Cost and Profit Assumptions

Figure 18 and Figure 19 give the holistic performance metrics as a function of density for three different relocation costs and three different profits per revenue mile respectively. The x-axis is truncated at 2,500 people per km^2 to better focus on low-density scenarios. Comparing these two figures, the largest takeaway is that profit per revenue mile (i.e. how much money a vehicle makes for each km they take a passenger) has a larger effect on holistic performance than does relocation costs (i.e. how much it costs a vehicle to travel while empty to another passenger). Additionally, changing the relocation costs does not substantially change the shape of the overall holistic curve, as shown in Figure 18. Moreover, the effects are relatively small for variations in relocation costs. By comparison, as shown in Figure 19, changing profit per revenue km can both have a large effect on the magnitude of holistic performance (i.e. higher profits per revenue km shift the performance curve upwards) and change the shape of the curve itself. There is a greater leveling off of holistic SAV performance at higher densities when halving the profit per revenue km as compared to the base case. By comparison, there is more of a steady increase in holistic performance throughout all densities for the doubled profit per revenue km case.

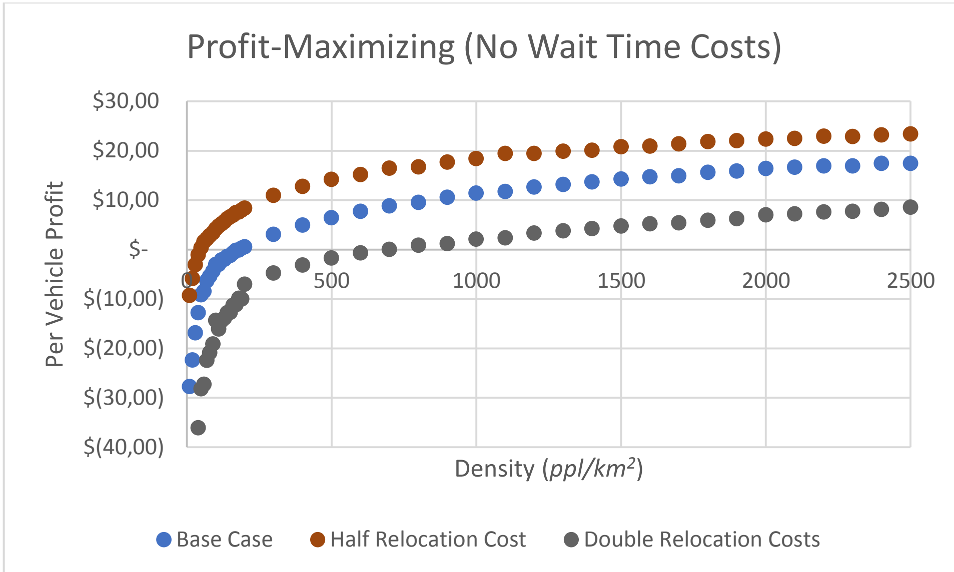

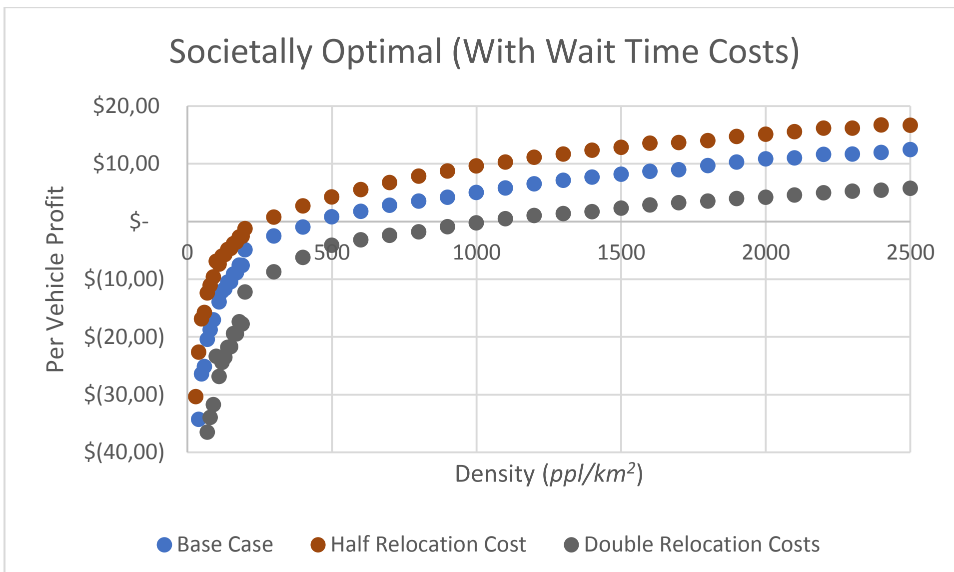

Figure 18: Holistic Performance vs. Density with Different Costs per Relocation Kilometer

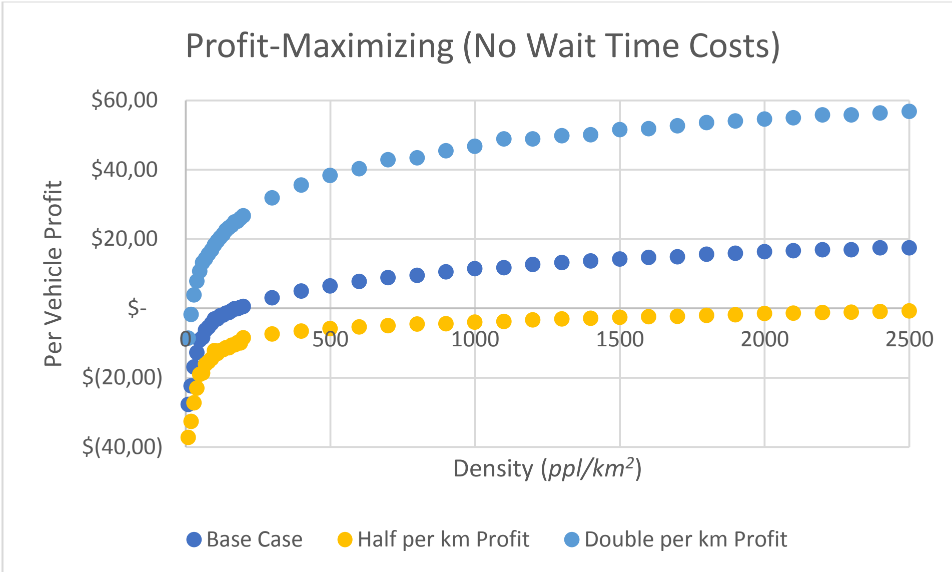

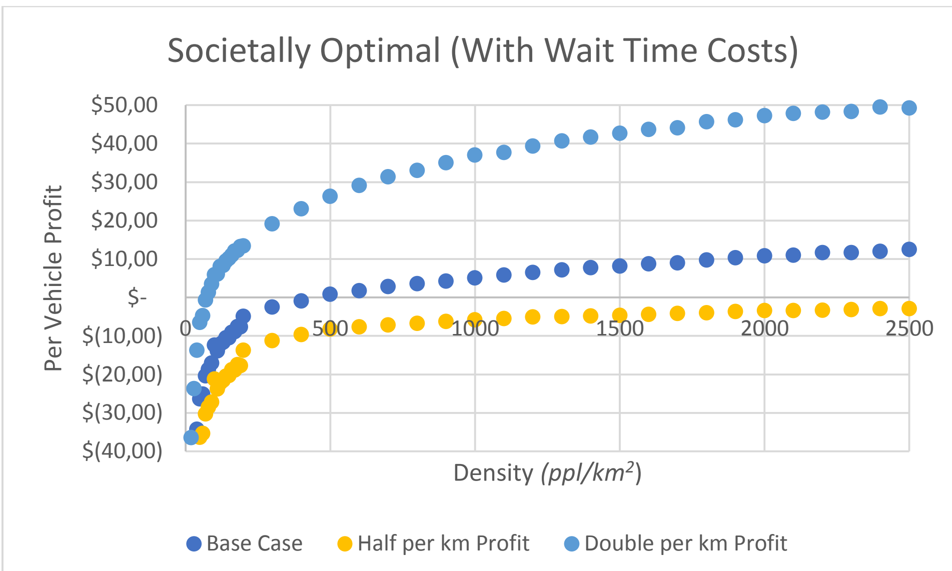

Figure 19: Holistic Performance vs. Density with Different Profits per Revenue Kilometer

Conclusion

The primary hypothesis for this chapter is supported; SAV performance increases at a decreasing rate with urban density. At ultra-low densities (i.e. less than 100 people per km^2), SAV performance is quite poor, with large wait times, high required fleet sizes, and low holistic performance (i.e. losses of more than $50 per vehicle per day). For these ultra-low densities from 0 to approximately 100 people per km^2, performances increases very sharply. Improvements in performance slows down somewhat, but are still very pronounced from approximately 100 to 500 people per km^2, then begin to level off for higher densities (i.e. comparatively small marginal improvements in performance). These variations in performance indicate a "tipping point" form of behavior around the range of 100 to 500 people per km^2, in a 10 x 10 km city where SAVs carry 2% of all trips. Note though, that this is not as strong a tipping point as for variation in fleet size, as described in the previous chapter.

With the cost assumptions made here, this tipping point range coincides with the break-even point: the combination of density and fleet size for which SAV systems start to become profitable. This chapter gives a primary range for this break-even point from 150 people per km^2 and 48 vehicles to 500 people per km^2 and 150 vehicles. Importantly, the break-even point occurs roughly across the tipping point, which should give pause to planners of SAV systems. Since across this tipping point range, small changes in density could result in large changes in performance, slightly overestimating the density, and by extension the travel demand, could have outsized negative effects. On the other hand, the relatively low densities for the tipping/break-even point indicates that SAV systems can be implemented over a wide range of different cities. SAV systems can be successful in medium-density areas as well as dense downtowns, and the overall performance may not differ that much between the two.

Overall performance is determined by a number of assumptions on what constitutes acceptable performance and the various costs of the SAV system. This chapter also investigates the effects of varying these assumptions from the base case, which create a few interesting findings. First, tipping point behavior becomes more extreme with higher acceptable wait total times. For example, if ultra-low average wait times of around 15 seconds are required, then there is a fairly steady decrease in required fleet sizes for higher densities. However, if longer wait times of 5 to 10 minutes are tolerated, then behavior becomes more binary; for all densities greater than 500 people per km^2, fleet sizes of approximately 2 vehicles per 100 trips are sufficient. Second, the optimal fleet size decreases with density, but at decreasing marginal rates. For example, for all densities greater than 1,750 people per km^2, the optimal fleet size is around 2.4 vehicles per 100 trips when using the societally optimal metric and 2.0 vehicles per 100 trips with the profit-maximizing metric. Thirdly, varying the cost of relocation distances does not substantially affect the shape of the holistic performance curves as a function of density; higher costs shift the curve up slightly, and lower costs shift it downwards. Lastly, varying the profit per revenue km does affect the shape of the cost curves; higher profits per km mean more consistent increases in system performance as density increases, while lower profits mean a more binary relationship, where system performance starts to plateau after approximately 300 people per km^2.

Chapter 6 — Effects of Varying Urban Form

This chapter models the effects that different urban forms have on SAV system performance. While the previous chapter examined how the magnitude of density affects performance, this chapter examines the effects of the distribution of that density. More specifically, it takes the city used in the previous two chapters, where people are distributed evenly throughout the region, and relaxes the assumption that people are evenly distributed throughout the city. Specifically, it examines how performance differs where more people are concentrated towards the center of the city and less towards the outskirts. A major challenge of this effort is that changing the distribution of density can affect the overall magnitude of travel demand. All else being equal, if two cities have the same overall density, the city with more people concentrated in the center will have shorter average trips. Therefore, this chapter takes two separate modelling approaches, both of which address this issue of different travel demand magnitudes with different urban forms. The first is a sensitivity analysis, varying this density distribution and quantifying how it affects relocation percentage using the density of unoccupied vehicles concept described in Chapter 4. The second compares three different cities: a base case city with people evenly distributed throughout, a “compact” city with more people concentrated in the center, with same density but shorter average trip length as compared to the base city, and a “centralized” city, also with more people concentrated in the center but with lower density and the same average trip length as compared to the base city. Note that the centralized city is larger than the base and compact cities, thus relaxing the assumption of 10 x 10 km service area used elsewhere. Overall, the findings of this chapter provide context for and support of Hypothesis #3, which states that more compact cities lead to higher SAV performance and more centralized cities lead to lower performance as compared to a city with constant density throughout.

IMPORTANCE OF VARYING URBAN FORM

Urban form is a necessary component in understanding the overall effects that different densities have on SAV performance. Just as the previous chapter investigated the ramifications of different magnitudes of density, this chapter examines how the distribution of that density affects performance. Understanding which types of cities are most amenable to SAV systems is a primary goal of this dissertation work, and studying the effects of varying urban form are a vital part of this effort. As described in Chapter 2, no SAV modelling efforts to date have engaged in detailed comparisons between different cities, a gap this chapter partially fills. Continuing with the general approach of this dissertation to sacrifice specificity for broad applicability, these variations in urban form are kept as simple as possible, going from people evenly dispersed throughout the city to highly concentrated along the center. This paper uses the degree of urban centrality, which gives the normal distribution of people from the city center, as a single parameter to vary urban form. Higher degrees of centrality (lower standard deviation values) correspond to more centralized cities; very high standard deviation values correspond to cities with more even densities throughout. The findings presented here show the difference in performance across different urban forms, such as the more consistent density of sprawling Los Angeles vs. a more centralized San Francisco.

A major challenge in this work is isolating and identifying the different effects of urban form. Cities with more people concentrated along the center are likely to have lower average trip lengths than a city of equivalent size and density with people evenly spread throughout the region. If people are more concentrated along the center, the trip origins and destinations are likely to be closer together (i.e. closer to the center), thus reducing average trip length. Shorter average trip lengths correspond to a lower total travel demand. For example, assume that two cities have 1,000 total trips a day, but the first had an average trip length of 11.3 minutes (as with the city in Chapters 4 and 5), while the latter had an average trip length of 5 minutes. Further assuming 1 minute for both boarding and alighting, vehicles need to be servicing passengers for 222 hours in the first city but only 117 hours in the second city. Quite reasonably, this lower demand level can be served by fewer vehicles, though this in itself is an important finding; more compact cities can be better served by SAV systems due to smaller average trip lengths.

This research seeks to isolate the effects of the urban form, separate from differences in trip length, in two ways. First, by performing the density of unoccupied vehicles analysis with ultra-high fleet sizes, it avoids the variations in travel demand issue, and looks at how relocation percentage varies as a function of the degree of urban centrality. Second, in addition to looking at a more "compact" city with the same density but a more concentrated distribution of people, and thus shorter trip lengths, it also looks at a more "centralized" city. Here, the centralized city itself is larger than the base city, so that the overall density of the city goes down but the average trip length stays the same as compared to the base. While more trips are concentrated along the center, as with the compact city, the larger city size also allows for longer trips. Effectively, as compared to the base city, the centralized city has both longer and shorter trips than the base city, and their effects cancel each other out so that the average trip length is the same as for the base city.

Lastly, as with the previous two chapters, this work is interested in the presence of tipping point behavior, where relatively small shifts in urban form could have outsized effects in system performance. As described below, while performance is not strictly linear with respect to urban form, tipping point behavior is not evident as it was for fleet size and the magnitude of urban density.

URBAN FORM MODEL RUNS

This section describes the two sets of model runs on urban form. The first treats degree of urban centrality as an independent variable, varying from more to less centralized cities. The second compares three cities: a base case, a compact city, and a centralized city.

Simulating the Effects of Varying the Degree of Urban Centrality

These simulation runs use a high relative fleet size of 15 vehicles per 100 trips, and fixes the population density at 15 people per km^2, such that there are always 60 vehicles. By using such a high fleet size, ensuring that there is always a surfeit of vehicles, these model runs separate the effects of the degree of urban centrality on relocation percentage from that of more compact cities having smaller average trip lengths and thus lower total travel demand. The degree of urban centrality is given as standard deviation of population distance from the center of the city. Standard deviations vary from 5 to 99 km in 1 km increments. Both origins and destinations are assigned randomly as according to these standard deviations. The city size is set at 10 x 10 km; if an initial trip origin or destination is selected outside of the city size range, it is reselected as fully random within the city. Each standard deviation value has 1000 simulation runs whose results are averaged together; the high number of simulations is necessary due to naturally high variance.

Figure 20 provides an example of vehicles distributed throughout the city with standard deviations of 10, 18, and 26 km. It uses a fleet size of 400 vehicles to make the differences between the three cases easier to see.

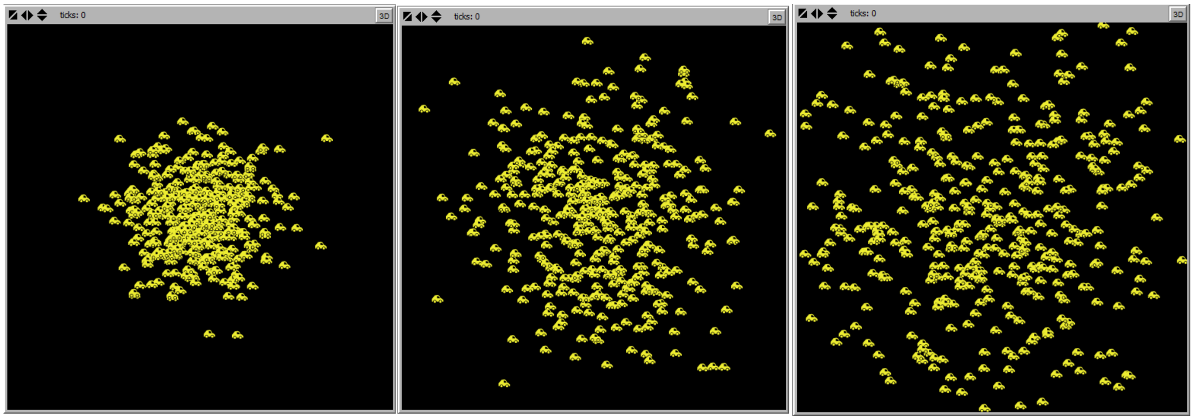

Figure 20: Vehicle Distribution for Standard Deviations of 10, 18, and 26 km

Comparing Base, Compact, and Centralized Cities

These simulation runs compare three cities: one a base case with a random density distribution, the second a compact city at the same density as the base case but lower average trip lengths, and the third a centralized city with lower density but the same average trip length as the base case. Figure 21 provides an example of these three city types, again using a fleet size of 400 vehicles to make the differences between the three cases easier to see.

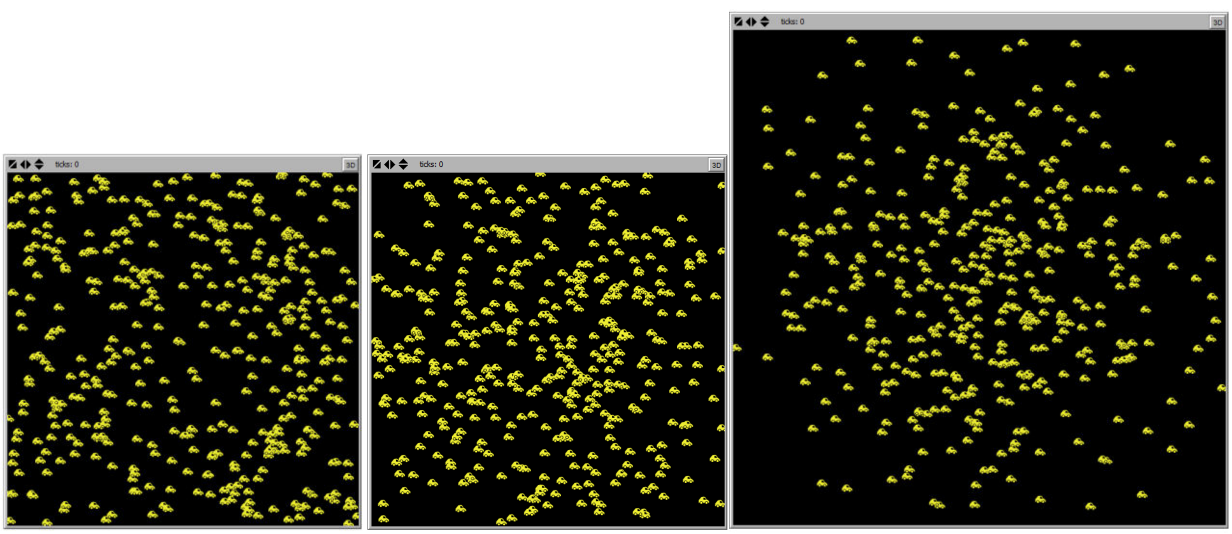

Base City *Compact City* *Centralized City*

Figure 21: Base, Compact, and Centralized Cities

Each city type is run across a range of fleet sizes, from 1.0 to 3.0 vehicles per 100 trips (from 200 to 600 total vehicles) with gaps of 0.05 vehicles per 100 trips and a 95th percentile wait time cutoff of 30 minutes. Each fleet size and city type combination is run 10 times, with the results averaged together. The model parameters are summarized in Table 15. Note that the average trip length is an output, rather than an input; it is determined by the combination of city size and standard deviation. Also note that the centralized city is larger than the base and compact cities, at 199 km^2 in area, or 14.1 x 14.1 km.

Table 15: Comparing Different Urban Density Types

City Type	Degree of Centrality: Standard Deviation (km)	City Size (km^2)	Total Daily Trips	Density $(\frac{ppl}{km^2})$	Average Trip Length $(minutes)$
Base	∞	100	20,000	2,500	11.3
Compact	32	100	20,000	2,500	10.1
Centralized	32	199	20,000	1,257.5	11.3

Results of Varying Urban Form

All of the cases above had a random distribution throughout the city: i.e. a fully decentralized city. This section examines the effects of varying the centrality of urban density, first by looking at how relocation percentage varies as a function of this degree of urban centrality. Second is an examination of three different city types: one base case with a random distribution, one compact city with the same overall density as the base case, and one centralized city with the same average trip length as the base city.

Effects of Varying the Degree of Urban Centrality on Relocation Percentages

The results of varying the degree of urban centrality on relocation percentage for ultra-high fleet sizes are presented below in Figure 22.

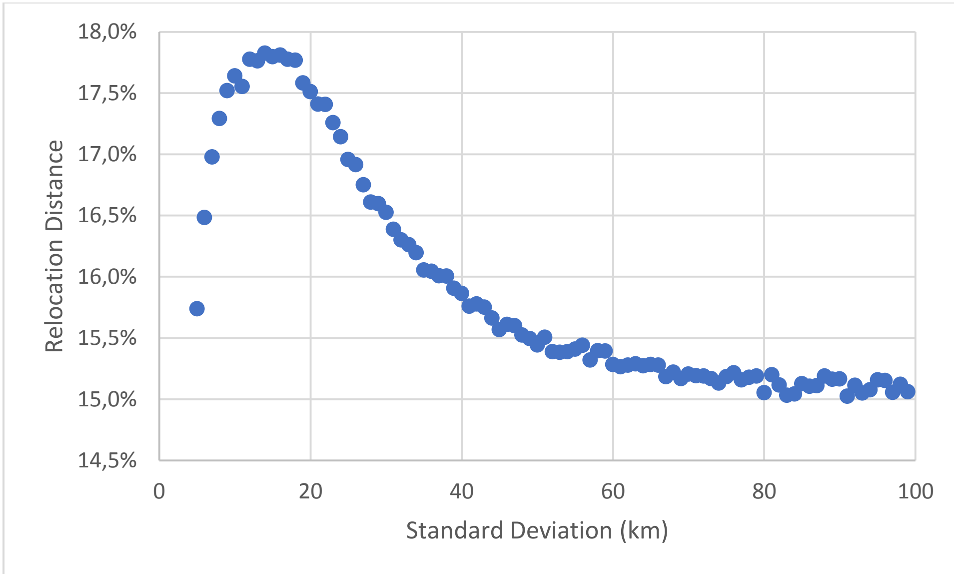

Figure 22: Standard Deviation vs. Relocation Percentage

Relocation percentages increase up to a standard deviation of 18 km with a maximum of about 17.9%, and decrease in an asymptotic fashion as distribution becomes more random to a minimum of about 15.1%. This study tentatively disregards standard distributions of less than 18 km, since these are relatively unrepresentative of real cities, as the density on the outskirts borders on zero; this is exhibited by Figure 20, where with a standard deviation of 10 km, no vehicles come close to the city limits. Therefore, the findings of these model runs indicate that more centralized cities have slightly higher relocation percentages than less centralized cities, though the distinction is not huge.

The reason that more centralized cities have higher relative relocation percentages comes from the shape of the "density of unoccupied vehicles" curve described in Chapter 4. Recall from those findings that low densities result in very high relocation percentages, and only small increases in the density can drastically reduce the relocation requirements. However, if the density is already fairly high, a further increase in density will not have much of an effect. As such, a set of trips with a low variance in trip distance will have lower relocation distances than one with a higher variance, even if the two have the same average, since the losses from the longer trips (higher relocation distances) more than outweigh the gains from the shorter, more concentrated trips.

Results of Comparing Base, Centralized, and Compact Cities

Figure 23 shows how SAV performance varies as a function of fleet size for the three city types: base, centralized, and compact. It gives two separate performance metrics: wait time and relocation percentage.

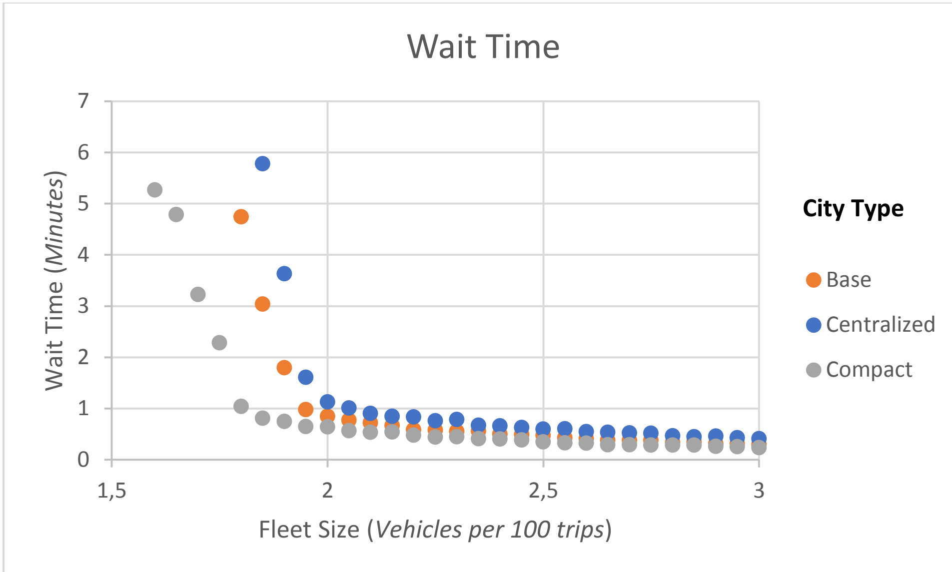

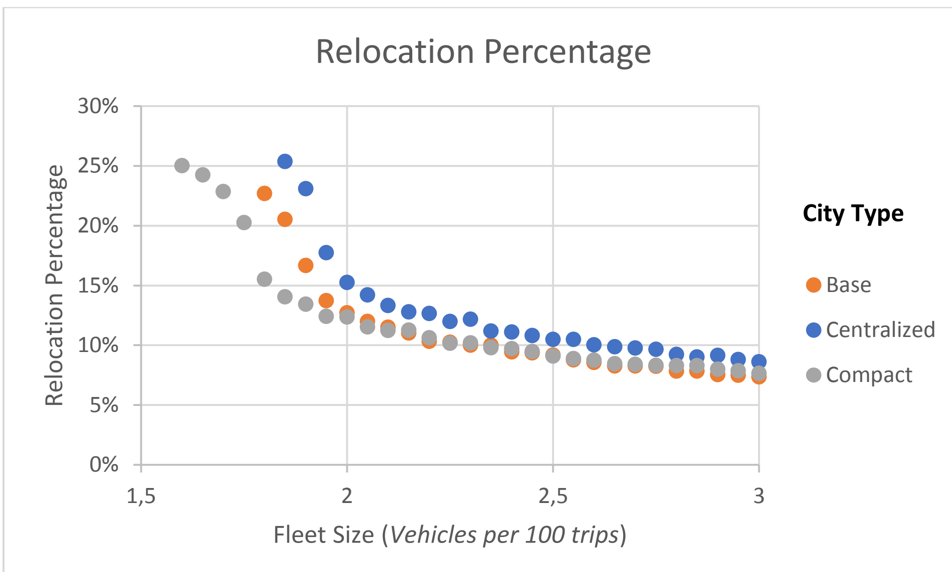

Figure 23: SAV Performance for Base, Centralized, and Compact Cities as a Function of Fleet Size

These results help confirm Hypothesis #3. Compared to a base city with constant density throughout, a centralized city has worse performance, and a compact city has better performance. For a given fleet size, a centralized city has higher wait times and relocation percentages than a base city, which in turn has higher wait times and relocation percentages than the compact city. As described in the previous section, compared with the base case, a compact city has shorter average trips (10.1 vs. 11.3 minutes), reducing total travel demand. Therefore, fewer vehicles are needed for the compact city's SAV system to avoid queuing. Interestingly, once queuing is avoided, performance becomes roughly the same for both the compact and centralized cities. In this case, the higher availability of vehicles with the compact city cancels out the increased relocation percentage from the higher degree of urban centrality. By comparison, this increased relocation percentage is a consistent feature of all fleet sizes for the centralized city. Since the base and centralized city have the same overall trip length of 11.3 minutes, the total travel demand is the same, so the need for greater relocation distances also causes for a slightly increased fleet size necessary to avoid queuing.

Figure 24 provides a closer, more detailed view of this tipping point when queuing starts to emerge for the base, centralized, and compact cities described above. The same modelling approach is taken (i.e. varying the fleet size, with 10 identical fleet-size and city combinations), but with smaller steps between fleet sizes of 0.01 vehicles per 100 trips and a narrower fleet size range of 1.6 to 2.0 vehicles per 100 trips.

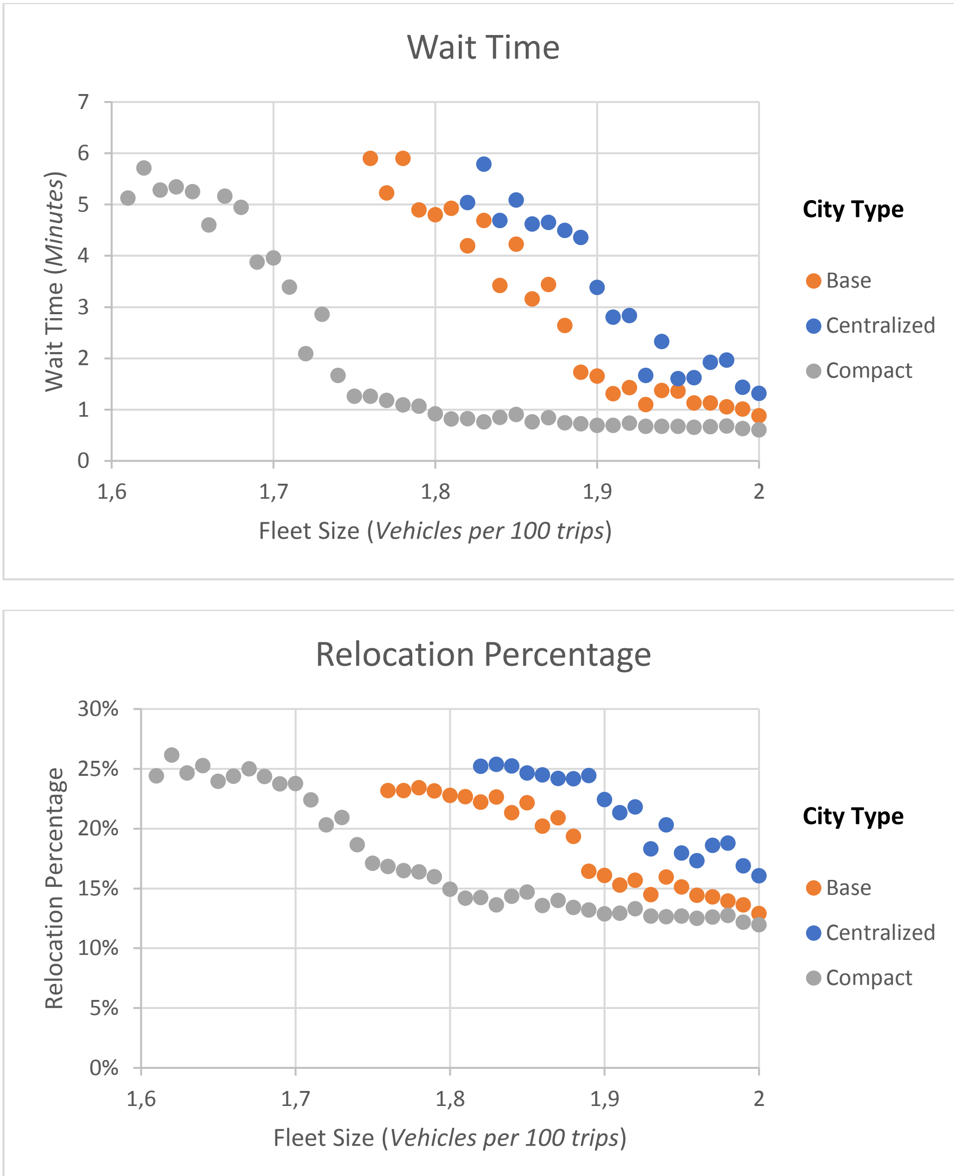

Figure 24: Tipping Point Behavior for Base, Centralized, and Compact Cities

Conclusion

The primary hypothesis from this chapter is supported; compared to a base city with constant density throughout, a compact city with the same overall density throughout but more people concentrated in the center will have higher overall performance. Since the average trip length is lower with a compact city compared with base city, there is effectively lower overall travel demand (i.e. less total passenger hours traveled), so that a smaller number of vehicles could provide an equivalent level of service. However, a centralized city leads to worse performance than with the base city. In a centralized city, more people are concentrated around the center of the city as with the compact case, but people also live further out, so that overall density is lower than and the average trip length the same as the base city. In other words, correcting for trip length, SAV performance improves for more sprawling cities with lower degrees of centralization. This finding is also confirmed by the model runs varying the degree of urban centrality with ultra-high fleet sizes, thus avoiding the lower utilization issue from shorter trips with more compact cities. These runs show that for the studied city with a population density of 15 people per km^2 and fleet sizes of 15 vehicles per 100 trips, relocation percentage decreases from about 17.9% in cities with the highest degree of centralization to about 15.1% for fully decentralized cities. The reason for lower performance with higher degrees of centralization, assuming that trip length issues are avoided, is that trips to less dense areas provide more negatives (i.e. are harder to serve) than the denser trips provide positives. As described in Chapters 4 and 5, since SAV performance increases with decreasing marginal returns at higher densities, from a base case, a doubling of density provides fewer benefits than a halving produces losses. Higher degrees of centrality correspond to a wider range of different trips, even if the average trip length stays the same: more and shorter trips about the city center, and more and longer trips to and from the outskirts.

Beyond the more technical concerns, differences in performance between the city center and the outskirts can have large policy ramifications. Especially combined with the established lower performance at lower densities, there is a risk of existing and future shared ride services not providing accessibility to more remote parts of the city. Already in New York City, over 90%

of all yellow cab trips occur in Manhattan and or one of two airports (New York City Taxi and Limousine Commission, 2016). Uber has released some data that they provide more trips to the outer boroughs, but are still concentrated in Manhattan and the densest parts of Brooklyn. The expected rational choice of private companies to avoid the outskirts of cities must be considered by the cities themselves, especially if SAVs become a substantial proportion of total trips in a region. Providing transportation access to the whole of a city, especially the poorer communities, is a moral imperative, and not one that will happen on its own.

Chapter 7 — Effects of Considering Mode Choice

This chapter extends the core agent-based model described in Chapter 3 to include and consider mode choice, thus relaxing the assumption used in the previous three chapters that SAV mode share is fixed at 2% of all trips, as used in the previous three chapters on the effects of varying fleet size, urban density, and urban form respectively. In this chapter, by implementing a mode choice component into the ABM, travel demand served by the SAV system becomes an independent variable. Total travel demand remains a fixed input, but the passenger agents are given a choice on which mode to take: SAVs, private autonomous vehicles (PAVs), and transit. The chapter first describes the implementation of mode choice into the core model, then provides the results of three main modelling run sets. The first set takes a fixed density and shows the effects of varying fleet size on mode choice (i.e. what percentage of the trips occur via SAVs). The second shows how SAV mode choice varies as a function of density when the traveler is choosing between all three modes. The third allows looks as SAV mode share as a function of density, but adds the possibility for multimodality, simulating an environment where a combination of transit and SAVs allows for a person to avoid the need to purchase a vehicle. Overall, the findings of this chapter provide context for and lead to the support of both Hypotheses #4a (with provisos) and #4b. For #4a, SAVs show the greatest improvement over other modes of travel for medium densities, though the range as measured here is from 500 to 4,000 people per km^2, lower than the approximately 2,000 to 8,000 people per km^2 as initially hypothesized. For #4b, SAVs alone are not enough to break the dominance of private vehicle ownership, but when combined with transit, they lower the densities at which people can feasibly go car-free.

IMPORTANCE OF MODELLING MODE CHOICE

This chapter examines the conditions where SAV systems will thrive compared to other modes of travel. This expands on the work of the previous three chapters, which identify the conditions under which SAV systems will be more or less successful irrespective of other modes. Chapter 4 shows how fleet sizes should be large enough to avoid queuing during hours of peak demand in order to avoid ultra-high wait times. Chapter 5 identifies that the increasing urban densities increases SAV system performance, but with decreasing marginal returns. Chapter 6 shows that all else being equal, more centralized cities will have higher relocation percentages than less centralized cities. Such findings are inherently useful for planning and deploying an SAV systems (e.g. such system designers should know how to avoid ultra-high wait times), yet more so for identifying if such systems *could* work rather than if they *will* work. Mode choice is needed to help answer the latter question; SAV systems will only succeed if they can provide a superior option to other modes of travel. And if the advance of autonomous vehicle technologies allows for SAVs to become reality, it will likely likewise allow for PAVs, so they are included as a viable mode here.

A flexible and a simplified model, as used in this dissertation, is especially merited for investigating mode choice. In traditional four-step travel demand forecasting, the mode choice step is deeply dependent on calibration; utility functions for the different modes are varied so that the results of the mode choice model match the observed, real-world mode share patterns — how many people take which modes for different trips. Such calibration is impossible for nonexistent modes like SAVs and PAVs. However, as done here, one can obtain reasonable estimates of the utility functions of these new modes by applying the costs/values of know attributes (e.g. wait time costs) and make reasonable assumptions for unknown attributes (e.g. cost per km for SAV operation). Still, as described in Chapter 3, the more transformative a technology, the harder it is to precisely predict its effects. As a result, even more so than with the previous chapters, the results presented here should be taken as approximations, with the overall trends more important than point predictions. For example, with Hypothesis #4a, the general trend is that SAVs show the greatest benefit over other modes at medium densities. High wait times and relocation requirements for SAVs make private vehicles more attractive in lower density cases, and transit becomes more attractive in higher density places. Results presented here offer confidence about the validity of this overall trend, but less so about the specific definition of medium density, which is given here as from 2,000 to 8,000 people per km^2.

However, while the model runs testing Hypothesis #4a do confirm that SAVs are comparatively more desirable at medium densities, they still find that private automobiles are likely to predominate throughout the vast majority of urban densities, assuming that people can afford to purchase them in the first place. Hypothesis #4b, and its attendant model runs, is an attempt to identify the conditions under which SAVs could attack the primacy of the personal vehicle. Neither transit nor SAVs on their own would be enough to attract the majority of people to forego PAVs. However, if people use SAVs to supplement transit trips, then this SAV-transit combination can compete against PAVs and obtain a sizable mode share. Using transit as a primary mode of travel, with SAVs as a supplement for trips where transit is undesirable, can effectively lower the feasible density for people to go car-free.

IMPLEMENTATION OF MODE CHOICE

This section describes the extension of the core ABM detailed in Chapter 3 to include mode choice. It first describes the three modes under consideration — SAVs, PAVs, and Transit — and provides initial utility functions for each. It then describes how the agents in the model chose between these different modes.

Defining Different Mode Options

Following with the overarching approach of this dissertation's model to maintain simplicity and flexibility wherever possible, this chapter only considers three modes – SAVs, PAVs, and Transit. Only SAV service is explicitly modelled. The utility functions of both transit and PAVs are estimated from the origin and destination of the trips. These base utility functions for all three modes are presented below, considering only wait times, travel times, and travel costs. This is a highly simplified approach, and various policy and technical factors could affect the utilities of the different modes. For example, heavier transit investment would reduce the cost of transit (i.e. make the transit utility less negative), eliminating minimum parking requirements in zoning codes or otherwise reducing available parking would increase the cost of private vehicles.

$$U = a(Wait\ Time) + b(Travel\ Time) + c(Travel\ Distance/Cost) \qquad (9)$$

During the simulation runs, each traveler agent is presented with a choice between the three modes via a traditional logit model. If the agent chooses transit or PAV, then this choice is recorded, but no actual travelling is modelled, and the agent is removed from the simulation. The total travel demand profile is the same as with the previous three chapters. However, unlike these previous chapters where SAV demand was fixed at 2% of total travel demand (approximating most cities' taxi mode share), in this this chapter, SAV mode share is a model output. All other modelling assumptions remain the same, including the daily average of four trips per person and the variations in daily travel demand linked to the temporal ridership pattern of the New York City taxi system. Note that this daily travel demand pattern is not the same as for overall travel demand throughout a traditional city, but it is used in this chapter in order to maintain consistency with previous results. Once the trip request is created, the traveler agent estimates the utility for each of the three modes, with all values given in dollar terms.

SAV and Ridesourcing Utility Function

This chapter jointly considers ridesourcing services like Uber and Lyft with SAVs. Both services provide point-to-point on-demand service. The only difference between the two is given in this section; SAVs eliminate the labor cost of having a paid driver, and thus are less expensive than ridesourcing.

For the cost estimates, first, the expected wait time is given by the relocation wait time only; within the simulation, there is an open and available vehicle to be hailed. If there are no available vehicles (i.e. queueing exists) then travelling via a SAV is immediately dismissed as an option. This is a necessary simplification; estimating the relocation wait time is possible when a trip request is made, but not estimating the length of the queue, and thus how long the traveler will have to wait before they can hail a vehicle. As previous chapters stressed the importance of avoiding queuing for any SAV system, this is also a reasonable simplification. Importantly, this wait time is part of the utility function that is determined by a model output. Both travel time and travel cost are solely determined by the trip origin and the destination. Travel costs for SAVs are assumed to be $0.50 per km, in line with previous SAV estimates and far below the current costs of ridesourcing firms like Uber or Lyft. The cost of in-vehicle travel time is set at $6 per hour; the travel time includes both the vehicle travel time and a minute each for boarding and alighting. The cost of wait time is set higher at $10 per hour, as with the assumption made in previous Chapter 3. The difference between the wait and travel times is based on the established premise that people dislike waiting more than travelling (Walker, 2011). The SAV utility function thus becomes:

$$U_{SAV} = \left(\frac{\$10}{hour}\right)(SAV\ Wait\ Time) + \left(\frac{\$6}{hour}\right)(Vehicle\ Travel\ Time + 2\ min) + \left(\frac{\$0.50}{km}\right)(Travel\ Distance) \quad (10)$$

Note that the financial cost of this service, while higher than that for personal vehicle usage as given below, is still substantially above that of traditional taxi and ridesourcing services (e.g. Uber and Lyft). For such services, the basic utility function stays the same, but the cost increases to $1 per km, which is in line with current Uber and Lyft fares (Henao, 2017):

$$U_{RS} = \left(\frac{\$10}{hour}\right)(SAV\ Wait\ Time) + \left(\frac{\$6}{hour}\right)(Vehicle\ Travel\ Time + 2\ min) + \left(\frac{\$1.00}{km}\right)(Travel\ Distance) \quad (11)$$

PAV Utility Function

The same travel cost time is used for PAVs as with SAVs, and PAVs are assumed to be able to travel at the same speed as SAVs. There is no wait time, since the traveler is assumed to have immediate access to a vehicle; this assumption is relaxed in the multimodal modelling runs. Furthermore, variable travel costs are assumed to be less for PAVs than for SAVs, not having to account for the expenses of operating these services. These travel costs are assumed to be $0.30 per km, leading to an initial PAV utility function of:

$$U_{PAV} = \left(\frac{\$6}{hour}\right)(Vehicle\ Travel\ Time) + \left(\frac{\$0.30}{km}\right)(Travel\ Distance) \quad (12)$$

Transit Utility Function

First, all transit rides are assumed to have a fixed fare of $1.00. Obviously, transit includes some waiting time as well as travel time, but both are lumped together here, since transit rides are not explicitly modelled. The cost of this travel time is set at $6 per hour, the effects of wait times and the relative attractiveness of transit are considered to be contained within an adjustment factor j_t. Estimating the travel time is a harder task than for PAVs or SAVs. For every trip, this chapter creates an adjustment factor to show if the transit travel time will be lower or higher than the vehicle travel time. An adjustment factor is created for both the origin and the destination, and the two are averaged together to get the total trip adjustment factor.

This adjustment factor relies on two major assumptions, both on the relationship between city type and transit performance. First, it assumes that transit accessibility is not uniform throughout the city, but rather improves the closer one gets to the city center. Note that this assumption is made even though all of the model runs in this chapter use a uniform distribution of people throughout the 10 x 10 km city. This first assumption is that the center of the city will have three times as good a transit service than the furthest point in the city, which in the 10 x 10 km city is 7.07 km away. Second, it assumes that denser cities will have superior travel performance. It does this by stating that the cities' density, divided into a constant of 3,165 people per km^2 provides this second adjustment factor. Both of these assumptions are tentative; obviously, many other factors can and do affect transit system performance, such as existing infrastructure and government support. However, they are intended to show how transit performance can vary both as a function of urban density and a traveler's location within the city.   Together, they lead to an adjustment equation as follows:

$$j = \frac{3{,}165\frac{ppl}{km^2}}{density} * \left(1 + \frac{3 * (distance\ from\ center)}{7.07\ km}\right) \tag{13}$$

$$j_t = \frac{j_o + j_d}{2} \tag{14}$$

As an example, assume a city with a density of 8,000 people per km^2, an origin 2 km from the center, and a destination 5 km from the center. The total adjustment factor is:

$$j_o = \left(\frac{3{,}165 \, \frac{ppl}{km^2}}{8{,}000 \, \frac{ppl}{km^2}} \right) * \left(1 + 3 * \frac{2 \, km}{7.07 \, km} \right) = 0.731$$

$$j_d = \left(\frac{3{,}165 \, \frac{ppl}{km^2}}{8{,}000 \, \frac{ppl}{km^2}} \right) * \left(1 + 3 * \frac{5km}{7.07km} \right) = 1.235$$

$$j_t = \frac{0.731 + 1.235}{2} = 0.983$$

The transit utility function can now be given in terms of the total adjustment factor:

$$U_{Transit} = \left(\frac{\$6}{hour} \right) j_t(Vehicle\ Travel\ Time) + \$1.50 \tag{15}$$

Note that the transit and PAV utility functions are constructed in such a way so that for the model runs to follow, transit utility is approximately equal to PAV utility at a density of 8,000 people per km^2. This can be seen in Figure 27, where the PAV and transit mode share curves intersect.

Logit Model

Once the utility has been calculated, the traveler decides to make the trip based on a logit model, which gives the probability of the traveler choosing each of the three modes. The simulation then assigns the mode based on these probabilities.

$$P(U_i) = \frac{e^{-\alpha U_i}}{e^{-\alpha U_{SAV}} + e^{-\alpha U_{PAV}} + e^{-\alpha U_{Transit}}} \tag{16}$$

α is a calibration constant. The first modelling run determines its value such that a city with a density of 2,500 people per km^2 will have an SAV mode share of 2%. However, as an example of the logit model here, assume a α value of 1 and utility values for SAVs, PAVs, and transit respectively of \$5, \$4, and \$10. The likelihood that a traveler will take a SAV is given as:

$$P(U_{SAV}) = \frac{e^{-5}}{e^{-4}+e^{-5}+e^{-10}} = 27\%$$

Running the Model

The following four sections describes the mode choice model runs — calibration runs, varying fleet size runs, varying density runs without multimodality (which tests Hypothesis #4a), and varying density runs with multimodality (which tests Hypothesis #4b). Since each of these model runs builds off its predecessor, the run parameters and the results are given together within each section.

All of these model runs have an explicit cut-off for the emergence of queueing, i.e. when hailing wait times become non-zero. As discussed in previous chapters, queueing is both highly undesirable for SAV performance and computationally intensive for the model.

CALIBRATION RUNS

These runs determine the calibration factor α so that a ridesourcing service would serve 2% of all trips in a city with a density of 2,500 people per km^2 and a fleet size of 400 vehicles; this is the equivalent of 2.0 trips per 100 vehicles with SAVs comprising 2% of all trips as given in Chapter 4. This fleet size is roughly the "optimal" fleet size using the profit-maximizing holistic performance metric, as determined in Chapter 4, and is slightly higher than the bare minimum fleet size necessary to avoid queuing. Recall from above that in this model, SAVs and ridesourcing provide identical service. The only difference is in in their utility functions, specifically the cost per km to the traveler. The \$1.00 per km for ridesourcing represents services like Uber and Lyft, and the \$0.50 per km for SAVs represent SAVs and the cost savings that could come from eliminating the paid driver.

The calibration factor α is varied from 1.4 to 1.8 in increments of 0.01. α=1.62 provides a 2% mode share for ridesourcing/SAVs. Higher α values provide lower mode shares for ridesourcing/SAVs. Equation 8 thus becomes:

$$P(U_i) = \frac{e^{-1.62U_i}}{e^{-1.62U_{SAV}} + e^{-1.62U_{PAV}} + e^{-1.62U_{Transit}}} \quad (17)$$

VARYING FLEET SIZE RUNS

As with the above, these runs take a city with a density of 2,500 people per km^2. They test two different shared mobility services — ridesourcing and SAVs — against transit and PAVs across a range of different fleet sizes.

Ridesourcing Fleet Size Runs

The runs use the utility function for ridesourcing as given in Equation 3, the PAV and transit utility functions as given in Equations 4 and 5, and the adjusted logit model as given in Equation 9. They vary fleet size from 100 to 600 vehicles. Figure 25 shows the results of these runs, both for all three modes and for a zoomed-in view of ridesourcing mode share.

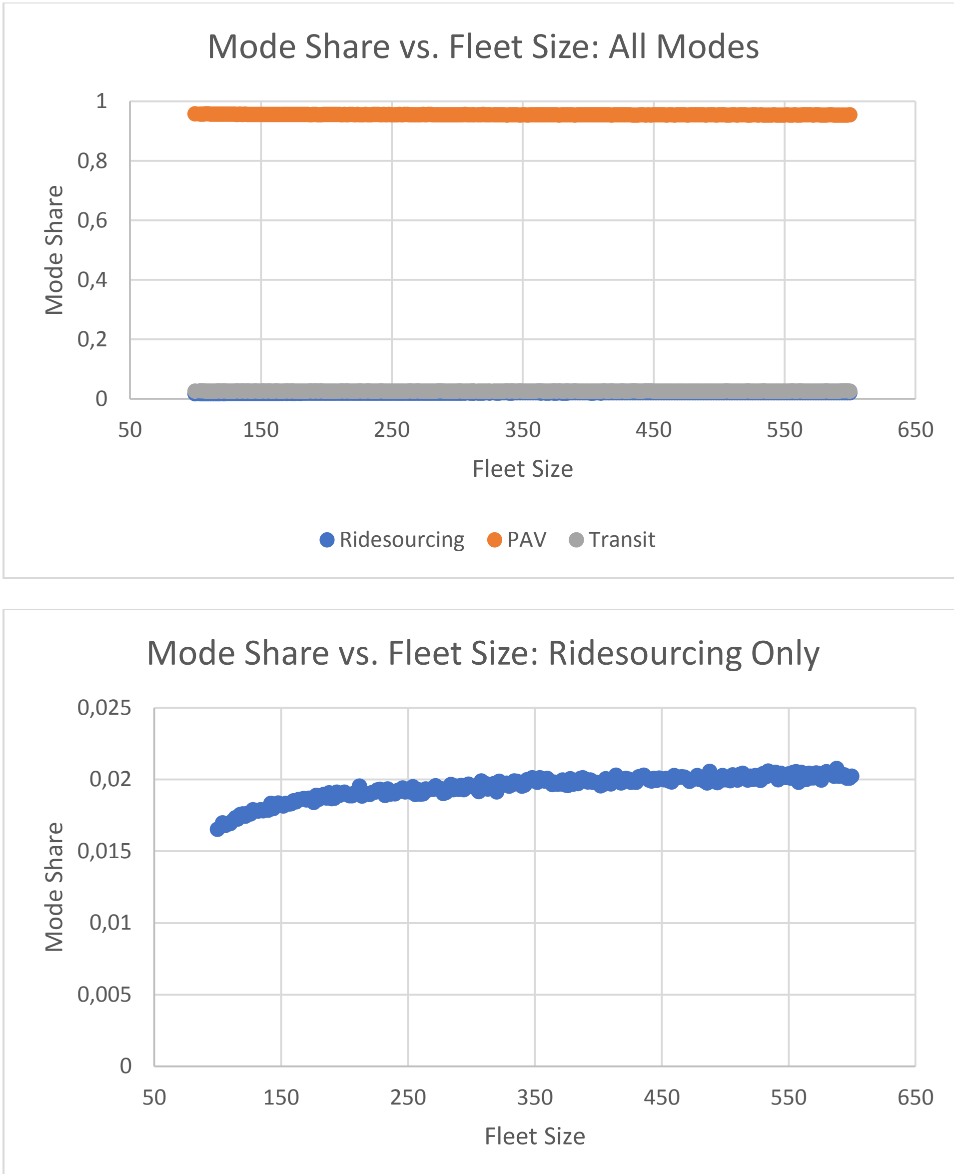

Figure 25: Mode Share vs. Fleet Size for Ridesourcing

For all fleet sizes, mode share remains relatively consistent, approximately 95% PAV, 3% transit, and 2% ridesourcing. Additionally, average trip length is relatively consistent at 1.64 km, or only 31.4% the average trip length from the previous three chapters of 5.23 km. This is because for longer trips, the logit model more heavily favors PAVs over ridesourcing/SAVs. Queuing eventually emerges, but only at fleet sizes fewer than 100 vehicles. Furthermore, mode share drops only slightly, i.e. from approximately 2% to 1.8% for a fleet size decrease from about 400 to 120. This is different from the pattern seen for varying SAV size when serving a fixed demand, such as given in Chapter 4. Those results show that queuing begins to emerge with fleet sizes smaller than about 360 vehicles in a city with a density of 2,500 people per km^2 and a SAV fixed mode share of 2%. This occurs for two reasons. First is the lower average SAV trip length. A related pattern is seen in the compact city evaluated in Chapter 6, where shorter average trips enable smaller fleet sizes. The second reason is that implementing mode choice as done here can effectively smooth peak ridesourcing/SAV demand. When there are fewer vehicles available, which would occur during peaks when SAVs are more heavily used, then wait times will increase, and overall ridesourcing/SAV utility will decrease. In other words, SAV utility is higher (i.e. lower cost) when there are more available vehicles, and lower when there are fewer.

For the varying density model runs in the following section, the relative consistency of mode share represents an opportunity to simplify the model runs by fixing ridesourcing fleet size at 100 vehicles for each density of 1,000 people per km^2. This equates to 250 vehicles for a city with a density of 2,500 people per km^2.

SAV Fleet Size Runs

SAVs have a higher mode share than ridesourcing due to their lower cost per km, and thus should require a correspondingly greater fleet size. The following model runs replicate the above for ridesourcing, but use the SAV utility function as given in Equation 2. These runs vary fleet size from 500 to 6,000 vehicles, and the results are shown in Figure 26.

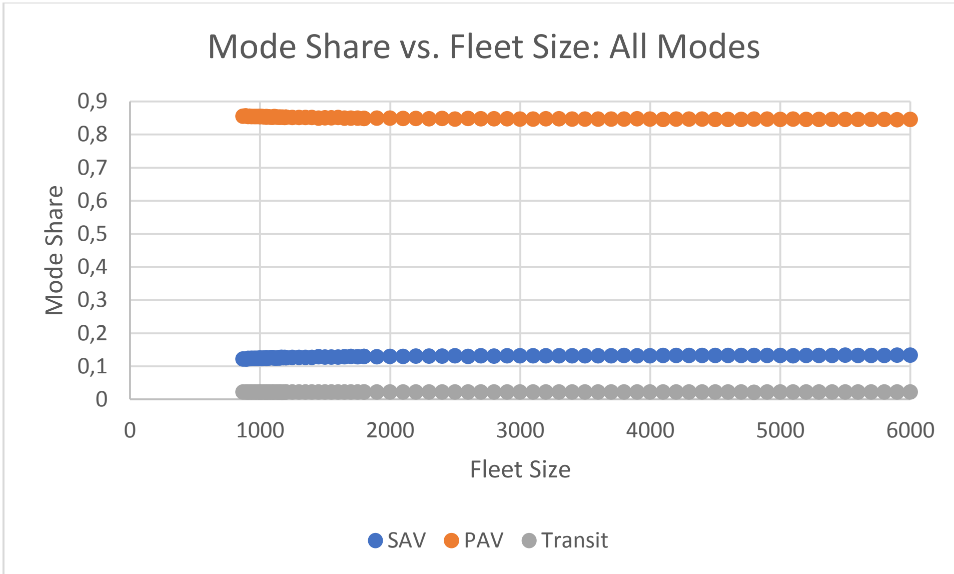

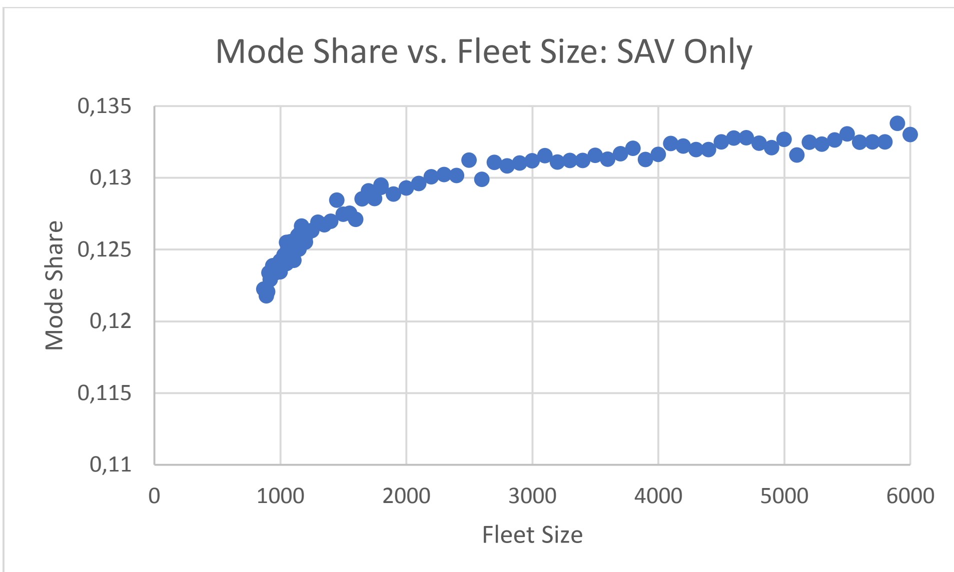

Figure 26: Mode Share vs. Fleet Size for SAVs

As with the above, for all fleet sizes, mode share remains relatively consistent. However, the lower cost of SAVs substantially increases their mode share, so that the split is approximately 85% PAV, 2% transit, and 13% SAV. Therefore, both transit and PAV mode share decrease with the advent of SAVs. Queuing emerges at fleet sizes less than 850 vehicles. Additionally, average SAV trip length is again consistent, but higher than for ridesourcing, at 3.73 km. For any trip, PAV utility is still higher (lower cost) than SAV utility, but the difference is smaller than with ridesourcing, hence the longer average trip length.

For the varying density model runs in the following section, the relative consistency of mode share again allows for fixing SAV fleet size, here at 800 vehicles for each density of 1,000 people per km^2. This equates to 2,000 vehicles for a city with density of 2,500 people per km^2. This fleet size for SAVs is equivalent to the same point in the curve for ridesourcing; compare the ridesourcing and SAV only curves in Figure 25 and Figure 26 respectively.

VARYING DENSITY WITHOUT MULTIMODALITY

These model runs focus on Hypothesis #4a, examining the mode share of shared mobility services relative to transit and PAVs as a function of density. As with the above section, these runs test two types of shared mobility service: ridesourcing and SAVs. Based on the above section, fleet size is fixed at 100 vehicles for each density of 1,000 people per km^2 for ridesourcing, and at 800 vehicles for each density of 1,000 people per km^2 for SAVs. Densities range from 10 to 10,000 people per km^2. Figure 27 shows the results for ridesourcing, and Figure 28 the results for SAVs. Both figures show an overview of the mode shares for all modes, and a zoomed-in image of the mode-share only for ridesourcing/SAVs.

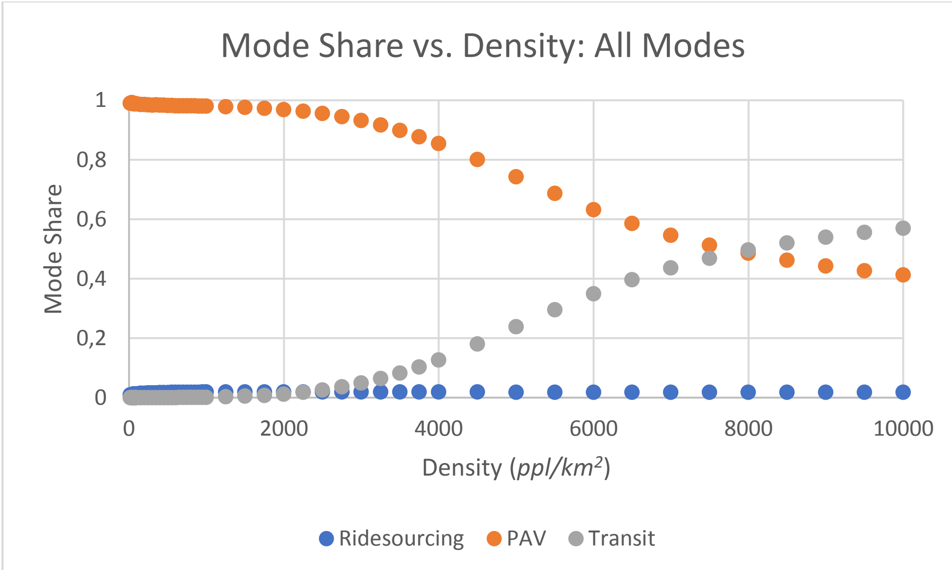

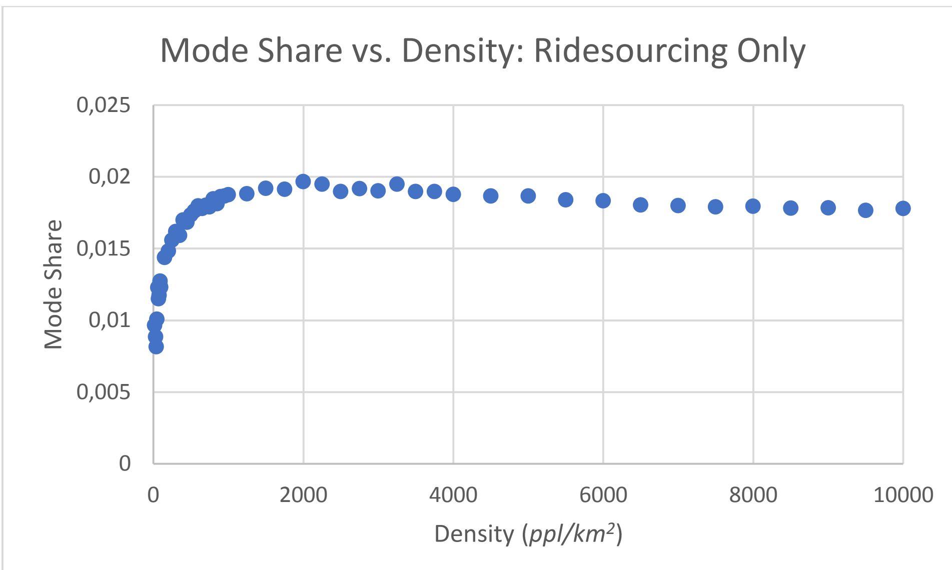

Figure 27: Mode Share vs. Density for Ridesourcing

For ridesourcing, with its higher per km cost, mode share is at a minimum for the lowest density (i.e. 0.7% at a density of 50 people per km^2. It sharply increases with density to a maximum of 2% at a density of 2,000 people per km^2. It then very slowly declines to 1.7% at a density of 10,000 people per km^2.

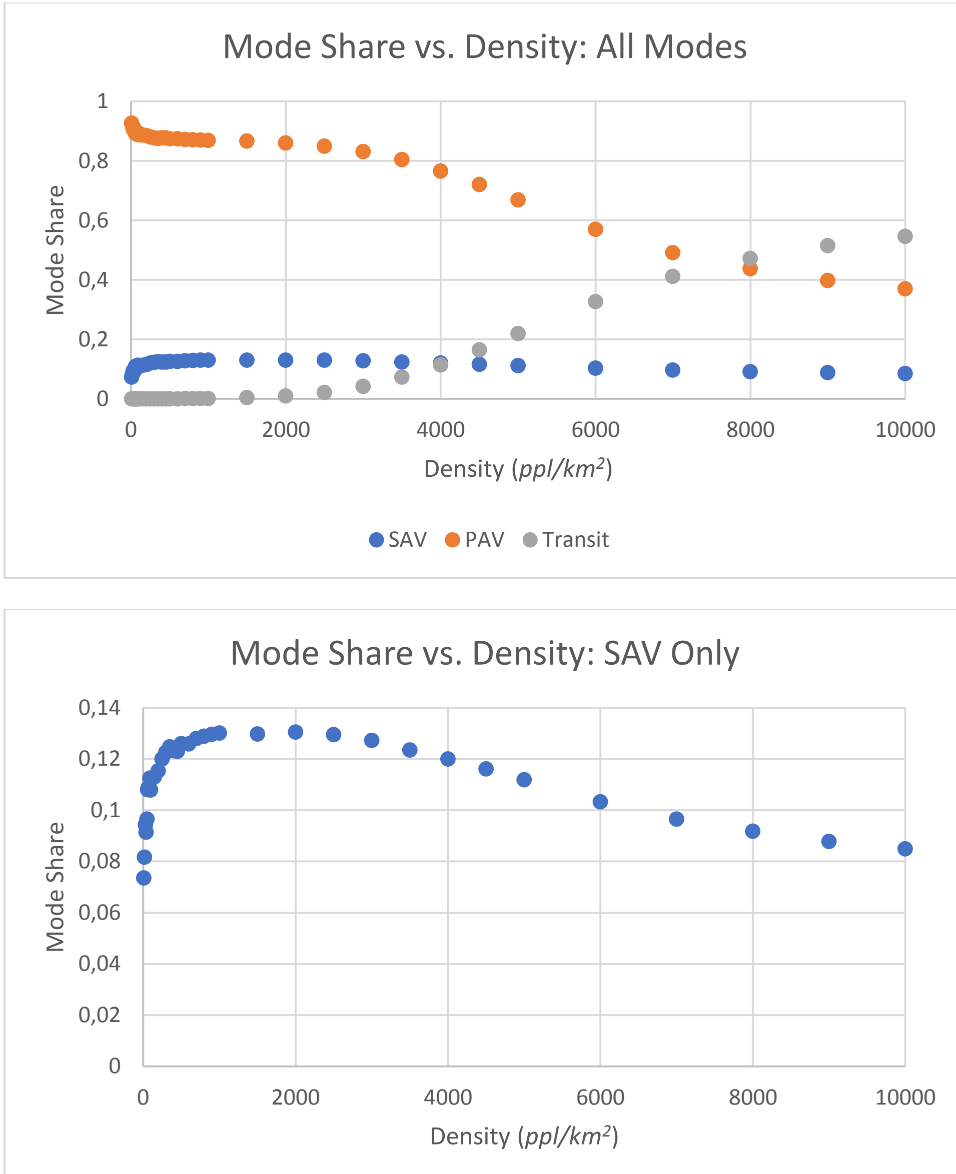

Figure 28: Mode Share vs. Density for SAVs

The lower per km cost of SAVs allow for them to capture a higher mode share, with a peak of 13% at 2,000 people per km^2. As city density increases from this point, SAV mode share decreases in a sharper fashion than with ridesourcing, but is still not as sharp as the decrease with lower densities. SAV mode share reaches a minimum of 7% at a density of 50 people per km^2, while going to a density of 10,000 people per km^2 only decreases mode share to 9%. Overall, Hypothesis #4a can be supported, albeit with some provisos. SAVs do perform better at medium densities, capturing a higher mode share, but these model runs indicate that the definition for "medium density" should be smaller. SAV mode share again peaks at about 13% at a density of 2,000 people per km^2, but remains above 12% for densities from about 500 to 4,000 people per km^2. SAV mode share for densities less than this range drop sharply, and drop more slowly for densities greater than this range.

Lastly, Figure 29 shows the effects of both ridesourcing and SAVs on transit and PAV mode shares. Since total travel demand is fixed (i.e. there is no induced demand), the rides provided by ridesourcing and SAVs come from both transit and PAVs, leading to a drop in both of their mode shares compared to if these shared mobility services did not exist. However, these drops are not consistent. Since ridesourcing never carries a large percentage of trips, SAVs only slightly reduce the mode share of PAVs and transit. By comparison, since SAVs are more attractive, they have a greater effect on transit and PAV mode share. The effects on transit are especially notable for lower densities. For all cities with densities at or lower than 4,500 people per km^2, SAVs can reduce transit mode share by more than 10 %.

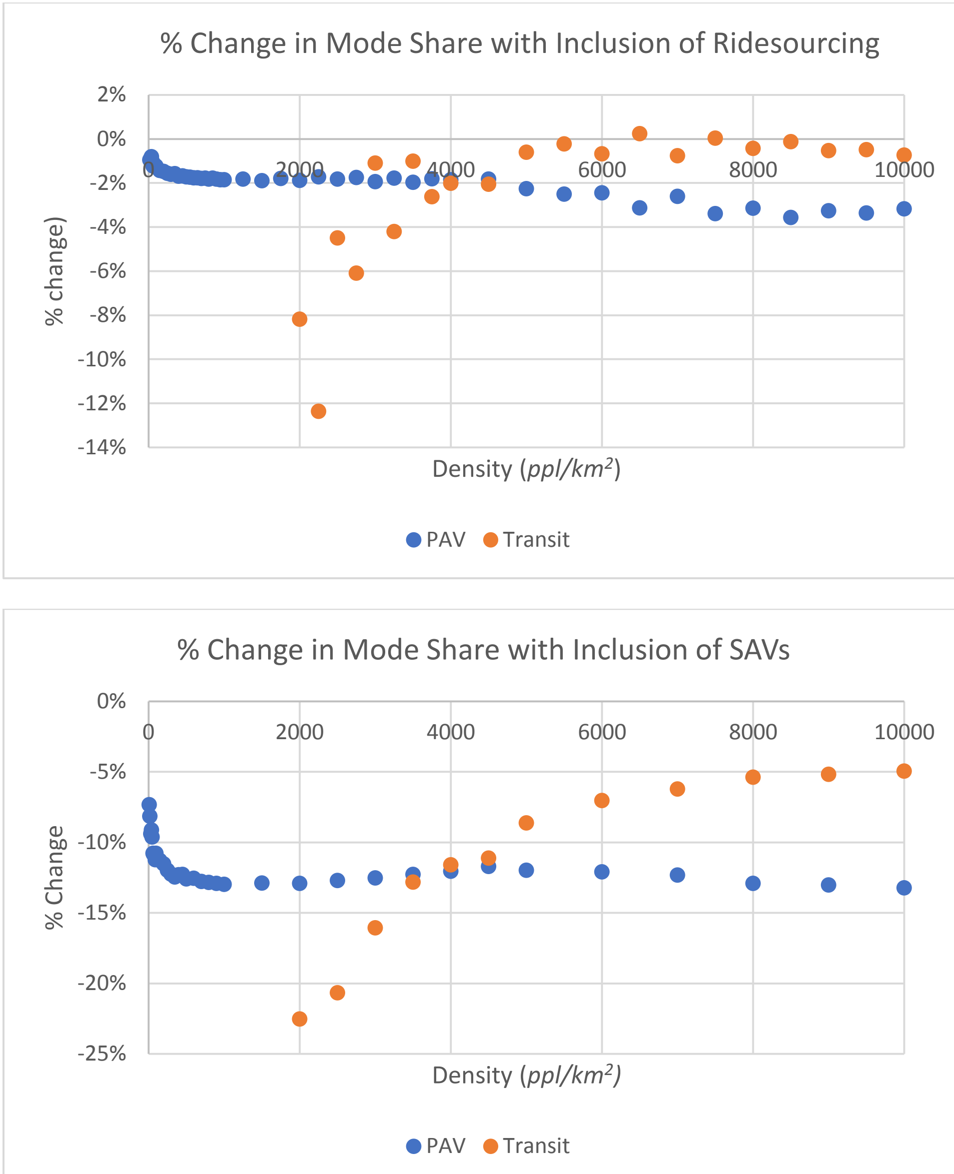

Figure 29: Effects of Ridesourcing and SAVs on Transit and PAV Mode Share

Ultimately, the findings of this section show the limitations of SAVs. Even assuming full automation and low per-km costs, SAVs never obtain greater than a 13% mode share. This is not totally surprising, as private vehicles are inherently low variable cost and high fixed cost, so that if one owns and has access to a vehicle, they are likely to use it. This is the norm for transportation in the US for all but the densest cities, where vehicles are simply too inconvenient to use and transit provides good accessibility and rapid travel. Moreover, especially as lower densities (i.e. less than approximately 4,000 people per km^2), the addition of SAVs can have a sizable negative effect on transit ridership, a greater reduction in percentage terms than the reduction of private vehicle (i.e. PAV) mode share.

VARYING DENSITY WITH MULTIMODALITY

This final set of model runs address Hypothesis #4b by relaxing the assumption that everyone owns a vehicle. They seek to show that a combination of two modes, SAV and transit, can offer substantially better service than either of the two on their own, and effectively lower the densities where people can feasibly forgo owning a vehicle. This section first describes the parameters of these multimodal model runs, then gives the results, and concludes with a discussion of the results.

Multimodal Model Runs

These model runs look at the same density range as the previous section, from 10 to 10,000 people per km^2. As these runs show, PAVs are likely to predominate over both SAVs and transit if everyone owns a PAV. These relax the ownership assumption by allowing for a choice to own a vehicle or not. This requires three separate additions to the above model.

First is calculating the average utility per trip as a function of density for thee mode mixes — transit only, ridesourcing/SAV only, and a combination of transit and ridesourcing/SAVs. For the combination of SAVs and transit, since the relative mode share between the two vary as a function of density (e.g. in Figure 33 below), these runs must vary both fleet size and density. For every density level, the smallest fleet size sufficient to avoid queuing is chosen. For all of these runs, the ridesourcing/SAV combination serves 10% of all trips. Serving 100% of trips would be too computationally intensive, and 10% provides effectively equivalent results. Second is calculating the average utility for a fourth mode mix — PAV only. Within this model, the PAV is constant throughout densities, at $2.73/trip

Lastly, at each density level, the average PAV utility is compared individually with each of the other three mode mixes via the logit model in Equation 9, resulting in a modal mix for each of the three mode mixes: PAV and Transit, PAV and ridesourcing/SAV, and PAV with a combination of transit and ridesourcing/SAV. As with the above model runs, ridesourcing and SAVs are considered separately to show the potential gains that can come with going from a non-automated ridesourcing service like Uber to lower cost SAVs.

Multimodal Model Results

Figure 30 shows the mode share for the PAV and transit only case. These results bear strong resemblance to the results in Figure 27, which showed the mode share when ridesourcing, PAVs, and transit are considered at the same this. This is to be expected, as ridesourcing only carried about 2% of all rides in Figure 27, so the PAVs and transit carry the vast majority of all trips. Note that as before, PAV and transit mode share are equal at a density of 8,000 people per km^2.

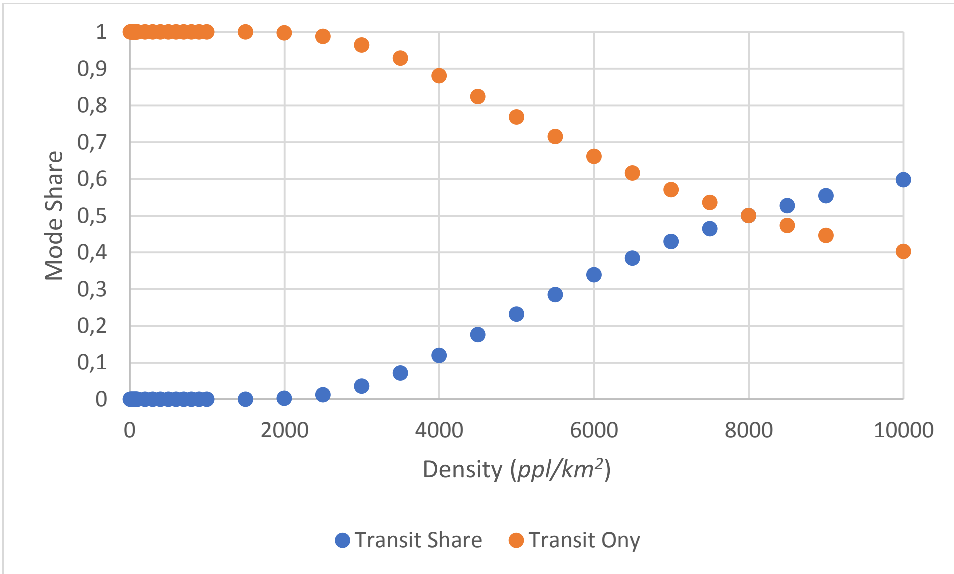

Figure 30: Transit vs. PAV Mode Share

The following subsections split the results into two groups, for ridesourcing and SAVs alone with PAVs and for ridesourcing and SAVs when combined with transit.

Ridesourcing and SAVs Only

Figure 31 shows the PAV mode share when these private vehicles are only competing with shared mobility services: SAVs and ridesourcing. Clearly, PAVs dominate across all densities; if someone owns and has access to a private vehicle, it will always have a lower cost than SAVs or ridesourcing. The lower cost of ridesourcing does allow it to capture a non-negligible mode share, maxing at about 12 % (PAV mode share). These results show that, on their own, even with the lowered cost, SAVs do not appear to be sufficient to supplant the dominance of private vehicles.

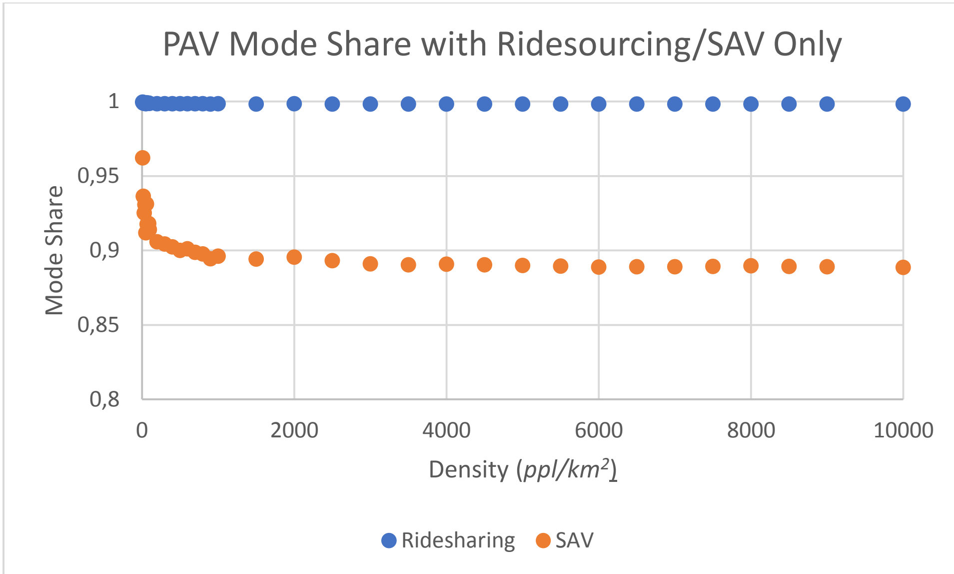

Figure 31: PAV Mode Share with Ridesourcing/SAV Only

Combining Ridesourcing/SAVs with Transit

Figure 32 shows the results of the multimodal runs for a combination of ridesourcing and transit against PAVs. These results are very similar to those of Figure 27, which showed the mode shares as a function of density without multimodality for ridesourcing, transit, and PAVs. Put simply, adding ridesourcing to transit does not substantially increase its attractiveness to users. While it can supplement transit in certain cases, ridesourcing is simply too expensive to act as a regular means of travel. Even Uber has admitted as such, albeit indirectly. They once defined "heavy users" as those who took Uber eight or more times in a five and a half month period – less than twice a month (Cohen, Hahn, Hall, Levitt, & Metcalfe, 2016).

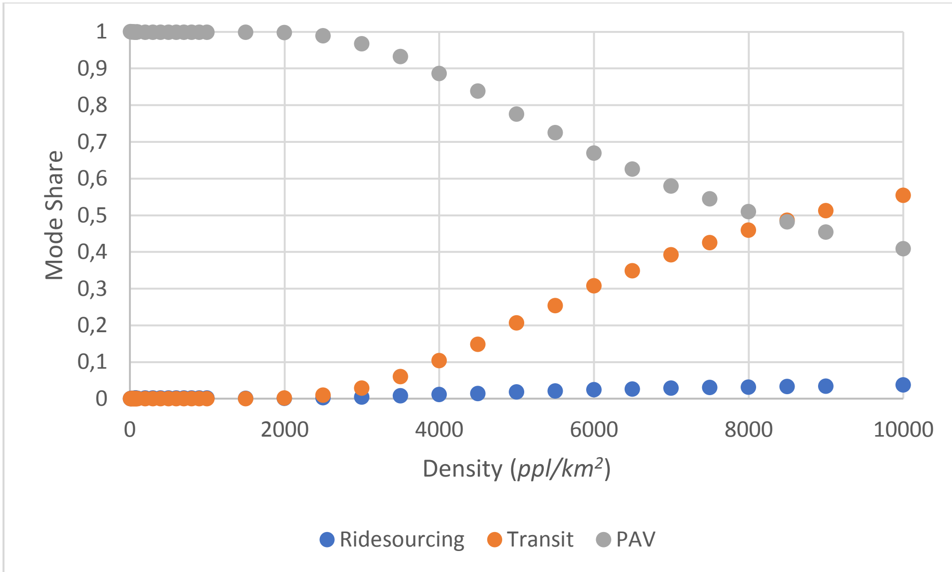

Figure 32: Ridesourcing, Transit, and PAV Mode Share with Multimodality

Figure 33 shows the results of the multimodal runs for a combination of ridesourcing and transit against PAVs. Unlike the above for ridesourcing, these SAV results are substantially different from those of Figure 28, which shows the mode shares without multimodality for SAVs, transit, and PAVs. The lower cost of SAV service enables a SAV-transit hybrid to provide superior service than either of the two on their own. This substantially lowers the density for people to go car free. For example, with transit alone, PAVs carry 50% of all trips at a density of 8,000 people per km^2. Without multimodality, SAVs and transit as separate modes reduce this 50% PAV density to 7,000 people (see Figure 28) , but treating SAVs and transit together produces a much larger effect, reducing the 50% PAV density to about 5,500 people per km^2. Beyond just reducing the feasible density for going-free, SAVs and transit together increase the attractiveness of forgoing a private vehicle across all densities.

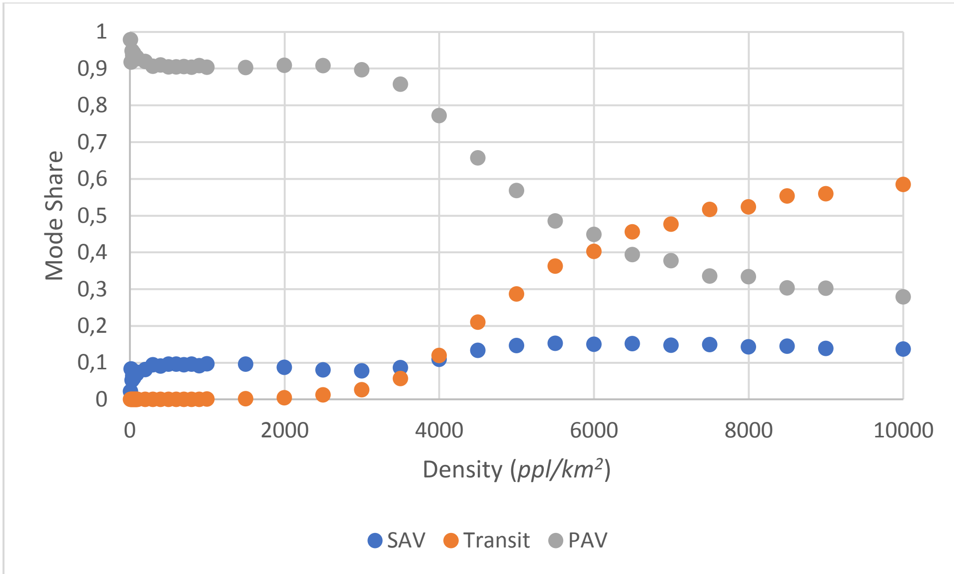

Figure 33: SAV, Transit, and PAV Mode Share with Multimodality

Discussion of Multimodal Results

Though this model did not show ridesourcing as having a substantial effect on lowering the attractiveness of going car-free, this does not mean they are currently having no effect. Some current work has looked at the effects of ridesourcing services like Uber and Lyft on transit usage (Rayle, Dai, Chan, Cervero, & Shaheen, 2016). Though the findings are inconclusive, there is evidence that for some people, in some situations, ridesourcing does encourage more transit usage.

The results show that SAVs can have a large effect in encouraging people to go car-free, but they are not enough by themselves. Even with the lowered cost of SAVs compared with ridesourcing, SAVs alone are not sufficient to supplant the dominance of private vehicles across all densities. However, by acting as a supplement for transit, the inclusion of SAVs can substantially increase the attractiveness of going car-free. SAV mode share never actually rises above 13%, but it helps lower the cost compared to if a person had to rely only on transit. Effectively, SAVs can provide transportation when transit is not feasible or desirable. In this model, transit trips become more expensive (i.e. longer and less desirable) when they have origins or destinations towards the outskirts of the city, so SAVs can help serve these trips. The need for at least occasional alternatives to transit is vital for any shared mobility future. For example, carsharing has already helped provide this role, allowing for people living in the city to predominantly rely on transit but still have access to a vehicle for a trip to Ikea or the countryside (Shaheen & Cohen, 2007).

CONCLUSION

Both of the hypotheses from this chapter are supported, though Hypothesis #4a comes with some corrections. SAVs do provide the biggest improvement over both transit and private automated vehicles (PAVs) for medium densities, but this model shows this medium density range to be between approximately 500 and 5,000 people per km^2, as opposed to the 2,000 to 4,000 people per km^{2}as initially hypothesized. For Hypothesis #4b, across all densities, SAVs never obtain a mode share much above 13%. Therefore, on their own, SAVs are not enough to break the primacy of private vehicles in places where they currently dominate. However, SAVs in tandem with transit can substantially reduce overall travel costs below what either of the two are on their own. Therefore, treating transit and SAVs as a combined mode can lower the feasible density to go car-free. As modelled here, non-PAV trips reach a 50% mode share at a density of 8,000 people per km^2 when transit is the only non-PAV option, but at 5,500 people per km^2 when SAVs and transit are treated as a combined mode. Ridesourcing services like Uber and Lyft do not have the same effect; their relatively high costs preclude them from being a regular replacement for transit. By comparison, the relatively low cost of SAVs enables them to supplement transit and provide less expensive and more convenient trips when transit would not be a desirable option.

Even more so than the results of the previous three chapters, the findings of this chapter should be treated as rough estimates more than reliable predictions. This is due to the cost assumptions made in this chapter with respect to the utility functions of the three modes, combined with the overall approach of this dissertation to develop a simple and flexible ABM that evaluates SAV performance in hypothetical cities. Still, the findings robustly show that SAVs have the potential to encourage far more shared mobility, enabling more people to go car-free. Any future will likely contain a combination of shared and private transportation, but there are many policy choices that will encourage greater amounts of sharing. The most important of these, and the one directly modeled here, is that automated vehicles must partner with transit. In particular, SAVs show the greatest potential in allowing people to go car-free by providing lower cost trips to places that transit does not reasonably serve. As such, policy can and should help encourage this sort of behavior, such as by encouraging/subsidizing SAV system operators to serve areas that transit does not effectively reach and/or coordinating SAV service with transit service. Similarly, policies that discourage private car usage would also push more people towards transit and SAVs. Such examples include making parking more difficult by reducing parking minimums and other free parking opportunities, enforcing higher fuel taxes, implementing vehicle kilometers traveled fees, or providing dedicated lanes for buses and SAVs. These and other policy considerations are addressed in greater detail in the following, final chapter.

Chapter 8 — Conclusion

Vehicle automation is not a panacea. It is a new technology, potentially a transformative one, but no new technology is a universal good. For example, it could eliminate the need for paid drivers, making ridesourcing services like Uber and Lyft massively cheaper, but at the same time put millions of people out of work.

Nor is automation imminent. Tens of billions of dollars have been invested in research, but fully automated vehicles are not commercially available, and no one can say when they will be. Building a car that will drive itself under all conditions, or even most conditions, is extremely difficult, especially if the automated vehicle must interact with traditional, human-driven cars.

Nor is automation unitary. Just as the computer was more than just a better electric typewriter, or the smartphone more than just a better phone, the automated vehicles will almost certainly look and be used differently than the cars of today – or at least some of them will. A pleasant autonomous luxury sedan is certainly one option. It could be parked in a suburban garage, alternately picking up the kids from school, or the groceries from the store, or the parent from the bar. Small, low-speed electric pods shuttling people through the city is another option. “Sleeper” cars for long-distance inter-city travel are a third option. Automated buses or shuttles is a fourth, and there are many, many more.

What automation does offer is a grand opportunity to reshape transportation and land use like nothing since the automobile itself. Yet unlike the streetcars creating the first bedroom communities miles from the city center, or the interstates encouraging the endless exurbs, automation has the potential to encourage denser living and discourage sprawl. Of course, this is only a potential, and the future remains uncertain. Still, this dissertation argues that urban planners, and society at large, should try to encourage higher degrees of sharing. Private automated vehicles (PAVs) will almost certainly be part of the vehicle mix, but so will shared automated vehicles (SAVs), similar to self-driving Ubers or taxis ferrying people throughout the region.

This final chapter provides guidance on how to encourage higher degrees of sharing and enable the successful deployment of SAVs to planners, government officials, potential SAV system operators, and other institutional stakeholders. It first summarizes the findings presented in the previous chapters, which rely on a flexible and simple agent-based model to study the conditions under which SAVs are most likely to succeed. Specifically, it reviews the modelling results of testing four interconnected hypotheses, each with their own chapter. The chapter then synthesizes the findings to show the multiple potential futures of SAVs, showing under what conditions they will be more or less successful. This is followed by a brief summary of opportunities for future work. The dissertation concludes with prescriptive policy advice in planning for SAVs, including both the opportunities and major risks.

CORE RESEARCH FINDINGS

This section summarizes the results of Chapters 4 through 7, each of which uses the developed ABMs to test its own hypothesis. The first three hypotheses look at the effects of varying fleet size, urban density, urban form respectively, and the fourth includes mode choice while varying density. Since all four hypotheses look at various aspects of SAV performance, an initial subsection summarizes the different performance metrics used in this dissertation. Also, these hypotheses results are presented as technical findings; the following sections of this conclusion more directly address policy-relevance.

Review of Performance Metrics

This dissertation uses three main performance metrics. The first is wait times, or how long passengers wait for a vehicle to pick them up. The second is relocation percentage, or how much the vehicles must travel while empty while driving to pick up a passenger (i.e. a relocation percentage of 10% means that for every 100 km a SAV takes a passenger, it travels 10 km while empty). The third is a holistic performance metric, with two variants, both denominated in dollars per vehicle. The "profit-maximizing" holistic metric considers performance from the point of view of the system, considering on a per-vehicle basis profit per revenue km (i.e. much money an SAV could make carrying a passenger), costs per relocation km (i.e. how much it costs to travel while empty), and daily fixed costs. The "societally optimal" variant is the same as profit-maximizing, with the addition of considering wait time costs. By measuring performance on a per-vehicle basis, the holistic metrics can balance the first two metrics – wait time and relocation percentage – with the effects that fleet size have on vehicle utilization. A system with an ultra-high fleet size would have very low wait times and relocation percentages, but the vehicles would hardly be used. Carried to an extreme, an SAV system with a high enough fleet size could approximate a system of private ownership, with one or more vehicles for every person. In addition to these holistic metrics, the vehicle utilization issue is also addressed by setting wait times as a constraint; i.e. what fleet size for a given density is necessary to ensure that average wait time is less than one minute. However, the holistic metrics have the added benefit of showing how perceptions of performance can vary depending on the stakeholder. For example, the work around Hypothesis 2 and 3 show that system operators will prefer smaller fleet sizes and the resulting higher wait times than would society as a whole.

These three main performance metrics are used for the first three hypotheses, which look at SAV performance with a fixed demand profile. Here, demand is wholly determined by urban density; every resident is assumed to take four trips a day, and SAVs are assumed to take 2% of all trips. The fourth hypothesis relaxes this final assumption by introducing mode choice. Therefore, the percentage of trips served by SAVs (mode share) is determined by a logic model pitting the relative utility of an SAV trip to that of transit trips and PAV trips. As such, for the fourth hypothesis, mode share acts as the performance metric; SAV systems are considered to perform better in places where they carry a higher percentage of trips.

Hypothesis #1: Fleet Size

Modelling results on the effects of varying fleet size are presented in Chapter 4, and they show that SAV performance increases at a decreasing rate with fleet size. More specifically, Chapter 4 introduces the concept of queueing, which occurs when there are more passenger requests for SAV rides than there are available vehicles. The emergence of queueing leads to highly non-linear, tipping point behavior of SAV performance as a function of fleet size. Once fleet size becomes small enough for queues to emerge, wait times can skyrocket, becoming worse the longer the queues. Once queueing is avoided, further increases in fleet size can bring relatively small increases in performance by reducing the average relocation times/distances; this is the time and distance it takes for an empty vehicle to travel to a hailing passenger. These improvements in relocation are more pronounced for lower densities (e.g. 500 people per km^2), so that even if the initial fleet size is large enough to avoid queues, further increases in fleet size will bring non-negligible reductions in average wait time. By comparison, higher densities (e.g. 12,500 people per km^2) exhibit more binary tipping point behavior; as long as queues are avoided average wait time is less than a minute, so adding more vehicles has no real effect.

Hypothesis #2: Urban Density

Modelling results on the effects of varying urban density are presented in Chapter 5, and they show that SAV performance increases at a decreasing rate with urban density. Like with varying fleet size, varying urban density also brings about tipping point behavior for SAV performance, albeit not as extreme. It is more of a tipping range, from about 100 to 500 people per km^2. For densities at the lower end of this range and below, SAV performance drops off precipitously. For example, densities less than 100 people km^2 can bring average losses of more than $50 per vehicle per day; the vehicles simply cannot serve enough passengers, so that they are either sitting empty or travelling while empty for too great a portion of the day. By comparison, densities greater than 500 people per km^2 bring some improvements in overall SAV performance, but they are relatively small. For example, there is not a huge difference in SAV system performance between cities with densities of 2,000 and 10,000 people per km^2. Additionally, with the cost assumptions made in this dissertation, this tipping point range coincides with the breakeven point for system profitability. It presents a few possible breakeven points, but all occur within the range of 150 people per km^2 and 48 vehicles to 500 people per km^2 and 150 vehicles. Ultimately, the relatively low densities for the tipping/breakeven point indicate that SAV systems can be successfully implemented across a wide range of different types of cities. SAV systems will work in both medium-density areas and dense downtowns. In fact, the overall performance will not differ that much between the two.

Hypothesis #3: Urban Form

Modelling results on the effects of varying urban form are presented in Chapter 6, and they compare three different types of cities. The first is the base case, a city with constant density throughout. The second is a compact city with the same overall density throughout but more people concentrated in the center. The third is a centralized city, where more people are concentrated around the center of the city as with the compact case, but people also live further out, so that overall density is lower than and the average trip length the same as the base city. The results show that an SAV system in a compact city has better performance than in the base city. In turn, an SAV system has better performance in the base city than in a centralized city. Since the average trip length is lower with a compact city compared with the base city, there is effectively lower overall travel demand (i.e. fewer total passenger hours traveled), so that a smaller number of vehicles could provide an equivalent level of service. However, a centralized city has worse performance than the base case, even though the overall travel demand is effectively the same. This is because the centralized city has more and longer trips, as well as more and shorter short trips. The short trips are easier to serve than the long trips, but the additional cost and difficulty of serving more and longer trips outweighs the benefits from more and shorter trips. Therefore, more compact cities will generally have better SAV performance than cities with constant density throughout. However, if correcting for trip length, SAV performance then improves for more sprawling cities with lower degrees of centralization. Note though, that these effects are relatively small (i.e. a centralized city would require a five to ten percent increase in fleet size compared to the base city), so varying urban form does not bring about the same tipping point behavior as does changes in fleet size or urban density.

Hypothesis #4: Mode Choice

The modelling results of adding a mode choice component to the core agent-based model are presented in Chapter 7. These results show how mode share varies as a function of density for three different modes: SAVs, PAVs, and transit. This chapter evaluated two related hypotheses; #4a looks at the relative desirability of SAVs compared to other modes of travel as a function of density, and #4b looks at how SAVs in partnership with transit could work to enable more people to go car-free. For #4a, the results show that SAVs provide the biggest improvement over both transit and PAVs for medium densities between approximately 500 and 4,000 people per km^2, as opposed to the 2,000 to 8,000 people per km^2 as initially hypothesized. Note though, that this is a relative measure. SAVs reach their maximum mode share over this medium density range, but the maximum is only about 13% of all trips. The maximum is even lower for ridesourcing services like Uber and Lyft; the high cost of these services means they never serve more than about 2% of all trips.

For Hypothesis #4b, since SAVs never carry more than 13% of all trips, on their own they are not sufficient to break the dominance of private vehicles. However, a combination of transit and SAVs provides superior service than either of the two modes on their own. SAVs act as a useful supplement to transit, helping to provide rides that would be too difficult or expensive for transit to serve. Therefore, the inclusion of SAVs can help encourage people to forgo owning a vehicle. For example, the results show that PAVs would maintain a 50% mode share at 8,000 people per km^2 if transit were the only other option, but a combination of transit and SAVs could reach the same 50% mode share at a lower density of about 5,500 people per km^2

THE FUTURE OF SAVS

This synthesis section seeks to use the research findings to answer the primary question initially proposed in the introduction: "Under what conditions will SAVs be more or less successful?" It divides these conclusions into a number of subsections. The first looks at fleet size, showing the importance of queueing and how the number of vehicles affects system performance. The second examines system placement, or the type of cities that are most amenable to SAV systems (e.g. identifying minimum density levels below which SAVs struggle to operate).The third shows the importance that transit can have in working with SAV systems and encouraging people to go car-free. The fourth makes the broader point that different stakeholders have different needs, meaning that there is never a true optimal solution for deploying SAV systems, since tradeoffs must always be made. The fifth and final subsection is the dourest, describing some of the limitations of SAV systems, such as potential increases in vehicle distance traveled and relatively low mode shares across all urban densities. Key findings, which should be of use to both researchers and potential SAV system developers, **are indicated in bold**.

Fleet Size

SAV fleet size must all be sufficiently large to avoid the emergence of queueing during peak demand. This should be an inviolate rule for any SAV system. As addressed in Chapters 4 and 5, queueing occurs when there are more passenger requests than available vehicles. The emergence of queueing represents a tipping point for SAV system performance. Once fleet size is small enough for queuing to occur, performance degrades rapidly (i.e. wait times can grow to over thirty minutes and required vehicle relocation can multiply). Once fleet size is large enough to avoid queueing, then further additions of vehicles bring relatively small improvements in performance. Therefore, fleet sizes should be slightly higher than the point at which queues emerge. **The higher the urban density, the lower the relative fleet size (i.e. vehicles per 100 trips) required to avoid queuing**. This is because lower density areas are harder to serve, especially due to higher relocation percentages. It is far more likely for an SAV to have to travel a longer time in a rural area with a density of 100 people per km^2 than a city with a density of 10,000 people per km^2. Also for this reason, the optimal fleet size as according to holistic performance metrics is closer to the point where queues emerge for higher density areas, and higher than this queue-emerging point for lower density places.

Additionally, note that **allowing for ridesharing should effectively enable dampening of peak demand, so that fewer vehicles are needed to provide an equivalent level of service.** This dissertation did not consider ridesharing in the interest of preserving a simple and flexible model, and other SAV models that include ridesharing do not explicitly address queueing. However, these other models do show that while ridesharing tends to have limited effects on both average wait times and vehicle utilization (i.e. how many hours a day a vehicle is in operation), ridesharing can be substantially more successful at reducing necessary fleet sizes. This is because, while a SAV with ridesharing may sit empty the majority of the time, by serving multiple trips simultaneously during peak demand, it can prevent queues from forming. Effectively, ridesharing acts as a source of excess capacity that can be used when demand is high.

System Placement

Overall, SAV performance increases with density. Denser cities have low wait times and relocation percentages, lower relative fleet sizes to achieve the same wait times, and higher holistic performance (i.e. dollars per vehicle for both the profit-maximizing and societally optimal variants of holistic performance). **Although cities with ultra-low densities (i.e. less than 500 people per km^2) are unattractive opportunities for SAV service, there is not a huge difference in SAV system performance between medium (e.g. 2,000 people per km^2) and very high (e.g. 10,000 people per km^2) density cities.** Therefore, smaller and medium density areas appear to be attractive opportunities for initial SAV system deployment, providing similar levels of service to SAV systems in denser service areas but requiring substantially smaller fleet sizes. Other policy concerns would also encourage such smaller-scale initial deployments; it would be easier to map smaller developments, medium-density places would avoid some of the traffic congestion concerns that plague the densest cities, and less-busy places could enable an easier environment for the operation of AVs (i.e. gridlock, constant double-parking, and heavy pedestrian loads could be difficult problems for virtual driving systems to solve). For example, college campuses and the equivalent could be attractive early options; the presence of car-free students provides a sizable demand base, while avoiding some of the challenges of the densest city centers.

However, if the density is too small, or the fleet size too low, then the SAV system cannot function effectively**. There is no true "minimum density" or "minimum fleet size," as it is highly dependent on cost assumptions of vehicles, relocation percentage etc., but as indicated by this work, 500 people per km^2 and 180 vehicles is a conservative and reasonable starting point.** A large number of cities meet this density requirement (e.g. Ann Arbor's density is 1,500 people per km^2), but a guideline of 180 vehicles indicates that any SAV system provider will need to provide heavy up-front investment for the vehicle procurement alone. This presents two major issues for the successful deployment of SAV systems. First, such an SAV cost hurdle itself could present an advantage to private ownership (i.e. PAVs) over SAVs. Second, a ridesourcing firm like Uber or Lyft, which already has access to a fleet, would have an easier pathway to deployment by slowly integrating SAVs into their existing service. However, the fully private operation of SAV systems brings up a number of troubling concerns, which are addressed in further detail in the following section on planning for SAVs.

Role of Transit

Chapter 7 explicitly models the relationship between traditional transit and SAVs. By partnering with transit, SAVs can help society move away from a private ownership model of transportation to one of shared mobility services. This effect is substantially muted if SAVs and transit are considered as separate and competing modes of travel. **SAVs on their own are not enough to supplant private vehicles.** In no cases modelled here does PAV mode share go below that of SAVs. However, a combination of transit and SAVs can achieve what neither of the two can do on their own. **Combining transit with SAVs encourages people to go car-free by enabling occasional trips to places that transit does not reasonably serve.** Though a person might rely on transit or other modes for the majority of the trips, SAVs can substantially lower the costs and difficulties of going car-free by effectively serving places that transit cannot. Therefore, one of the strongest pieces of advice that this dissertation can provide is that **SAVs and transit must work together to help encourage a shared mobility future; SAVs are NOT area replacement for traditional transit**.

Furthermore, **SAV and transit partnerships are most attractive in places with a decent well-functioning existing transit system (i.e. places where transit already maintains a non-negligible mode share).** For example, SAVs may work well in a city like Los Angeles, which has a good and rapidly growing rail network and an expansive bus system. is already fairly dense at approximately 3,300 people per km^2 for the entire metro region A person could use the trains for their regular commute, and SAVs for times when they want to get to a place that is not transit accessible. A city like Detroit looks less attractive. Since there is no real high-speed transit system, a person without a car might be forced to rely on SAVs for all or almost all trips, which is simply not feasible. Note though, that while SAV-transit partnerships may not work in places without well-functioning transit systems, SAVs can improve transportation options in places where transit is not feasible, such as very low density areas (i.e. less than 700 people per km^2).

Different Stakeholders have Different Needs

All transportation systems, and the cities they serve, are inherently complex systems. **There is never a true "optimal" solution, as these systems are comprised of numerous different stakeholders with different needs.** For SAVs, we can consider three broad groups of stakeholders: the SAV system operators, the SAV travelers, and society as a whole. Tradeoffs must always be made between these groups. For example, higher fleet sizes correspond with lower wait times (better for the SAV travelers) and lower relocation distances (better for the travelers and society as a whole, due to fewer empty vehicle kilometers traveled) but also lower profits per vehicle (worse for the system operator) and lower per-vehicle utilization (worse for society as a whole, especially considering the resources needed to build and store these vehicles).

This dissertation touches on potential conflicts between stakeholders in a few ways. One is the use of two holistic performance metrics: profit-maximizing, which considers only the desires of the system operator, and societally optimal, which considers both the system operator and the SAV travelers. **If only the needs of the system operator are considered, fleet sizes will be lower and wait times higher than if the needs of the travelers are also included.** Additionally, as stated earlier, SAVs show the greatest potential in partnering with transit in providing access to places that transit does not reasonably serve. Note though, that these trips will likely be in the less dense parts of a city, away from the areas that high-quality transit operates. This represents an inherent tension, since system operators would prefer to operate in denser areas, as the costs are lower and the revenue higher (i.e. the profit-maximizing metric is higher in denser areas). **Therefore, society would prefer SAV service that supplements transit in lower density places where transit does not serve, while system operators would likely prefer to operate in the densest parts of a city.** Socio-economic and equity concerns could only exacerbate the issue of society and system operators preferring to operate in different areas. For example, low income areas may have low overall demand for SAVs, since people there cannot afford to take this more expensive service on a regular basis. However, they are also the areas with the greatest need for and potential reliance on rapid shared mobility option. Since they have the greatest percentage of people without access to a vehicle, people require a taxi or an SAV (in the future) to perform vital trips such as to a job interview or to day care when transit is running late or otherwise unavailable. Already though, studies have indicated that ridesourcing firms are less likely to serve such areas in particular and to pick up women and people of color in general (Ge, Knittel, MacKenzie, & Zoepf, 2016).

The phenomenon of queueing brings up another example of where the needs of society and system operators could diverge. Avoiding queueing during peak hours is expensive for system operators, as they effectively need to procure a marginal vehicle that is only needed during peak times. As mentioned before, allowing for ridesharing is one way to dampen peak demand and reduce the need for extra vehicles. However, another would be to simply encourage peak hour travelers to find a different mode, such as transit. **Ultimately private operators would prefer to "cherry-pick" the most efficient and profitable routes and times of traditional transit.** Transit could effectively subsidize SAV operators by serving the most expensive and difficult trips, but this could be disastrous from a public good perspective, especially if it put public transit budgets into a tailspin. Such cherry-picking of transit is a well-observed phenomenon(Walker, 2011). Over 90% of taxi trips originate in Manhattan south of approximately 90th street or in one of the two airports (New York City Taxi and Limousine Commission, 2014). In the 1920s, jitneys, which are van-like vehicles operating on semi-fixed routes, would "steal" fares from the trolleys, often racing ahead of them to pick up people waiting at the trolley stop. Trolley operators, who had government-granted monopolies to operate their service, were able to help enact regulations that effectively put jitneys out of business (Hodges, 2006).

Furthermore, **SAV operators have large incentives towards monopolization,** which society as a whole would not prefer, at least without heavy public sector involvement. Multiple providers must effectively split the trip density, thus reducing overall system performance. The tendency towards monopolies can already be seen in ridesourcing companies. In China, Didi merged with Uber to gain dominance over the ridesourcing industry, and there are only one or two ridesourcing services in most US cities. Even this understates the degree of concentration and lack of competition, as Uber and Lyft rely on the same driver base to a large extent (i.e. most people who drive for Uber drive for Lyft at the same time).

Limitations of SAVs

Society must be cautious about the promises of shared automation. It shows great promise to improve transportation, but a future where all people travel via SAVs is both unlikely and undesirable. Overall, SAV service has a number of limitations that must be considered. For example, as confirmed by this work and shown in many previous models, **SAVs are likely to sit empty for the majority of the day; twelve hours of operation is definitely on the high side of operation.** This is due to the daily variations in travel demand; the fleet size should be large to avoiding queuing during the morning peak, which means there will likely be far more vehicles than are needed at 3:30 AM. Furthermore, the need for relocation should increase total distance traveled. **Relocation percentages in the range of 7% to 20% are reasonable for most of the SAV systems modelled here.** Other modelling work shows how ridesharing can reduce the additional driving distances somewhat, but almost all efforts show an increase in vehicle kilometers driven; in other words, the additional relocation distance is greater than the benefits gained from shared rides.

Additionally, as stated earlier, in no cases modelled do SAVs achieve a mode share greater than 13%, and this includes generous assumptions made on the cost savings from automation. The same modelling work for ridesourcing vehicles, which require a paid driver, never achieve a mode share higher than 2%. **SAVs can supplement transit, but they do not appear to be an attractive option for replacing most or all trips in cities.** At lower densities (i.e. less than 2,000 people per km^2), PAVs predominate, and at higher densities (i.e. greater than 5,000 people per km^2), the mode share of transit starts to exceed that of PAVs.

FUTURE WORK

Chapter 2 provides a fairly extensive review of gaps in the SAV modelling research literature. This final section focuses more on research opportunities specifically identified by this dissertation's findings. First and foremost, more work is needed to show potential interactions between transit and SAVs. While Chapter 7 showed the benefits of transit and SAVs working together, it did not explicitly model an intentional partnership. This includes, but is not limited to, first-mile and last-mile service and direction of SAV vehicles to areas of low transit. Private control of SAV services is also a concern here. Private companies may not take such direction, or be willing to share the detailed ridership data necessary to optimize the timing and location of connections between transit services and SAVs.

More critical and questioning research on SAVs is also merited. Many previous studies take an optimistic view of SAV operations, with correspondingly positive assumptions. This work seeks to be more neutral, but far more work is needed to show not just the conditions under which SAVs will succeed, but the conditions where they will not. More research efforts also need to focus on how SAVs can affect the city in which they operate, not just the other way around. Part of this work can and should be critical, including how PAVs and other forms of automation could create induced demand (i.e. more vehicle kilometers traveled) and more sprawling development. While this conclusion chapter makes the point that both SAVs and PAVs can change land use patterns, with SAVs encouraging denser development, no studies to date have examined this phenomenon.

PLANNING FOR SAVS: FAILURES AND POSSIBILITIES

This final section advises planners and policy makers on what they can and should be doing to prepare for SAVs, and is split into four parts. The first stresses the importance of establishing guidelines on the type of transportation future we want; without knowing and agreeing on our goals, we can never hope to achieve them. The second establishes the importance of collecting and maintaining access to the data necessary to help manage and regulate SAV systems, including current data. The third states a number of actions that society and the public sector should take to help achieve its goals for SAVs, especially for transit and SAV cooperation. Especially attractive here is the potential coordination between transit and SAVs that could offer superior service to either of the two combined. As discussed earlier in this dissertation, such an approach could imply a "hub-and-spoke" transit system, where high-frequency express transit serves the highest demand routes, and SAVs or other shared mobility vehicles serve the less well-traveled areas, including providing first-and-last-mile service to the express transit routes where needed. The fourth part of this section concludes this dissertation, arguing that urban planners and policy makers must start taking actions immediate to help prepare for and encourage a shared automated future.

Establishing SAV Guidelines

Automation promises the biggest revolution in transportation since the car itself. However, as with all disruptive technologies, the bigger the potential effects, the harder the future is to predict. Many will claim with certainty that 10 years hence, almost all private cars in urban areas will be replaced by fleets of electrified and platooning SAVs (Lutz, 2017). Others argue just as confidently for something totally different. The truth is that the future is unknowable, but that is not an excuse to do nothing. The actions we take as a society now will shape the transportation system of the future. One of the most important roles of transportation engineers, urban planners, and other transportation professionals is to take their expertise on and knowledge of the current transportation, and use those skills and understanding to declare what sort of transportation we as a society would like to create. Critical here is the recognition that neither cities, nor transportation, nor SAVs are problems to be “solved.” As this chapter has argued, almost every decision comes with tradeoffs in system performance, and there can never be a single true solution when there are multiple objectives and stakeholders.

This dissertation took a relatively narrow view of performance metrics and stakeholders, laying out three for each: wait time, relocation percentage, and holistic performance for the metrics, and system operators, travelers, and society as a whole for the stakeholders. The reality is far more complex, but a few simple goals should provide a reasonable starting point, four for SAV systems themselves and one for the relationship between SAVs and PAVs.

- ***Equity:*** SAVs should not make the poorest among us worse-off. Where possible, the needs of the poor should take precedence over those of the rich, since they already have good transportation options.
- ***Resiliency:*** SAVs should increase the resilience of a transportation system, encouraging it to better maintain its function in response to disruptions both small (e.g. a traffic jam) and large (e.g. a hurricane).
- ***Sustainability:*** SAVs should decrease carbon emissions and other deleterious environmental impacts. Of special concern is modulating the potential increases in vehicle kilometers traveled due to eliminating the need for a driver reducing travel costs.
- ***Accessibility:*** SAVs should improve the transportation accessibility of the cities they serve, making more jobs and other desirable destinations easier to reach.
- ***Sharing:*** Where feasible and reasonable, policies should encourage SAV usage over PAV usage; this should help the system achieve the other four goals.

Note that congestion was not included in this list. That is not to say that congestion is a good thing; no one enjoys traffic jams. However, being able to get to a place easily and quickly (i.e. accessibility) is more important than getting there without congestion; a commute of 10 minutes travelling 2 miles should be seen as preferable to one of 60 minutes travelling 60 miles. Moving towards accessibility-based metrics (i.e. people's ability to reach desired locations), and away from mobility-based metrics (i.e. people's ability to move) is an emerging focus of transportation planning (Levine, Grengs, Shen, & Shen, 2012).

All of above the goals combine politics, policy, and technical issues. However, the majority of SAV work to date has come from technical fields such as operations engineering, computer science, and transportation engineering. Urban planners are especially well suited to consider the complexities of eventual automated vehicle deployments, but they have played a relatively small role in discussions around SAVs to date.

Importance of Ridership Data

The following section details how the public sector can and should guide SAV deployment. However, such effective action is impossible without knowing how the SAVs are operating. One small solution, as called for in the literature review (Chapter 2), is to encourage more standardization and openness when publishing SAV model results. Far more important though, for both researchers and the public sector, is gaining access to real-world shared mobility ridership data, especially from ridesourcing firms like Uber and Lyft. Notably, though ridesourcing is the best extant analog for SAVs, no SAV models have relied on ridesourcing data due to the lack of availability. Most have either used all of the trips in a region and assumed that SAVs serve a percentage of all trips, or used only one of two readily available taxi data sets, from New York and Singapore. This dissertation used the New York City data to determine the daily variations in travel demand.

Obtaining ridership data will not be easy; as private companies are loath to share, as it represents one of their key competitive advantages. A new entrant could develop a reasonable app, reasonable routing algorithms, and provide equivalent service to Uber and Lyft; the temporary banning of these companies in Austin in 2016 and the emergence of new players helped prove this point. However, no one knows as well as Uber or Lyft how people use their services — the spatial and temporal variations in travel demand, the willingness to pay for rides, how long vehicles must wait for people, how long it takes to board and alight the vehicles, how travel demand varies when ridesharing is included, etc. Since these data are so important to these companies, they have not and will not release it without a fight. This is one of the major reasons why previous researchers have struggled to estimate the effects of Uber and Lyft on current transportation systems. If researchers do not know where they operate, when they pick people up, and who uses the services, it is near impossible to say if these services have encouraged or discouraged more transit use, increased or decreased roadway congestion, or aided or hampered parking issues.

More than understanding the effects of the services, access to data is also necessary to both improve service and appropriately consider externalities. If ridesourcing continues to provide a relatively small percentage of total trips, i.e. in the vicinity of 2% in most cities, then these externalities will [hopefully] remain relatively small. However, if SAVs reach their promise and start carrying substantially more trips, externalities could increase, both positive and negative. Fleets of thousands of SAVs, alternatingly ferrying passengers, traveling while empty, and sitting parked, as posited by this dissertation and other SAV models, could potentially wreak havoc on a city's transportation system. A government may want to do things such as charge a SAV service for the congestion that it creates, subsidize them to provide transportation to underserved communities, or arrange interactions with transit to both avoid cherry-picking time pickups and drop-offs to match the transit schedule; these and other potential government actions are addressed in the following section. However, if the government does not know where and how the SAVs are providing service, none of this is possible. Therefore, this dissertation argues that to empower researchers, improve current ridesourcing operation, and most importantly, establish a precedent for how they will treat eventual SAV operators, governments should require ridesourcing companies to release ridership data. Again, they will not do so willingly, but allowing them to operate on public roads is a valuable concession.

Public Sector Involvement with SAVs

Like all companies, private operators are profit-seeking entities. This is fully rational, but problems arise when the decisions to maximize individual profits run counter to the aims of society. Fully public ownership and control of SAV systems appears unlikely; the current norm of operation is private, as with taxis, Lyft, and Uber, and the private sector is massively investing in vehicle automation, estimated at $80 billion in total as of October 2017 (Karsten, 2017), and will reasonably expect some sort of return. However, the public sector is vital to help guide the deployment of SAV systems and avoid some of the dangers of privately led SAV systems and PAVs, as described earlier in this chapter and summarized here:

- “Cherry-picking” transit routes in the densest part of the city
- Avoiding poorer and less dense parts of the city
- Creating congestion, increased distance traveled, and other negative externalities
- Not effectively partnering and coordinating with transit
- Setting too-low fleet sizes, which increases average wait times and makes queueing more likely
- Tendency towards monopolization of operators
- Losing out in market share to PAVs
- Maintaining secrecy of ridership data

To help ameliorate the above dangers, governments can take a number of steps. The most attractive feasible future is for SAV systems to look something like a utilities and/or public-private partnership. Korea, Hong Kong, and Singapore all have heavy private sector transportation involvement, but transportation and access is a public good, so that major decisions are always emanating from the public sector. The same should be true with SAVs. A true public-private partnership would avoid most of the above dangers, since the goals of the private entity and the public should be mostly aligned, allowing for things like the development of a combined SAV and Transit hub-and-spoke system. A ridesourcing company's business model is to provide as many trips as possible, as profitably as possible. Without a substantial change in this model, they will always have different aims than society, as their core client is the individual traveler. They will not want to share their data, provide less profitable trips to less dense and poorer areas, or avoid causing traffic jams or additional driving. However, if a private company treated the city as their customer, then achieving broader societal goals is more likely. This does not preclude competition, as numerous suppliers could bid on given contracts. For example, a company like GM or Ford would be happy to sell a profitable service contract to cities, providing both the automated vehicles and the software to support SAV operation (e.g. modelling, real-time management, user apps, and more).

While this dissertation views SAV public-private partnerships as the goal standard, even without them, planners should advocate for heavy public sector involvement with SAVs in a variety of ways, starting with current regulations around ridesourcing. As stated earlier, releasing anonymized ridership data should be a prerequisite to operation, and cities should begin to both set goals for SAV operation. They also need to establish regulations and incentives to reach these goals, both for ridesourcing now and SAVs in the future. This must be an ongoing discussion and debate, but a few attractive possibilities do stand out. One is variable pricing to address externalities both positive and negative. For example, SAVs could be charged more heavily when operating in a dense central business district during rush hour, addressing the negative externality of worsening congestion, and subsidized to provide service to poorer communities, addressing the positive externality of improving accessibility and equity. A second is helping to provide dedicated pickup and drop-off points, especially useful to help prevent double-parking-related traffic jams; this and other curb management issues only become more important as shared mobility usage increases. A third is explicitly encouraging transit and SAV/ridesourcing integration. Short of explicit public-private partnerships, both externality pricing (i.e. incentivizing transit connections) and providing curb space show promise, as does data sharing efforts to help coordinate pickups and drop-offs. More intensive transit and SAV coordination efforts could come from the concept of mobility as a service, where users could have a single payment platform, and/or a subscription service, that would include a variety of shared mobility options: transit, carsharing, ridesourcing, etc. A fourth, and final suggestion, is to accept the uncertain future of shared mobility and automation and establish guidelines around a city's core transportation goals, such as described here. Examples of this would include establishing things like equity, accessibility, and sustainability as goals in their strategic transportation plans and setting rules whereby the city agencies have an obligation to develop regulations to support these goals.

A Call to Immediate Action

SAVs show a potential for encouraging more shared use, fewer vehicle miles traveled, denser living, but only a potential. Broader outcomes are always determined by existing land use patterns, existing politics and power structures, existing rules and regulations, and more. Inaction is not neutral, but rather advocacy for the status quo, which in most of the developed world is private vehicle ownership. By not playing an active role, and playing it now, planners and policy makers are encouraging PAVs, rather than SAVs. The fight for more sharing, more equity, more accessibility, more sustainability, and more resilience must be waged on multiple fronts. The public sector alone is not enough, nor is the private sector; SAV policy and planning will come from federal, state, and local governments, community organizations, traditional planning organizations, the tech industry, vehicle manufactures, academia, and more. Most importantly, we cannot wait for automated vehicles to hit the road before starting the planning process. We are already planning for automated vehicles. And we are doing it wrong. But we can do better.

Bibliography

Axelrod, R. M. (1997). *The Complexity of Cooperation: Agent-based Models of Competition and Collaboration*. Princeton University Press.

Bagg, S., & El-Geniedy, A. (2016). Riding with Robots, (January).

Beiker, S. (2014). History and Status of Automated Driving in the United States. In G. Meyer & S. Beiker (Eds.), *Road Vehicle Automation* (pp. 61–70). Springer International Publishing. https://doi.org/10.1007/978-3-319-05990-7_6

Bischoff, J., & Maciejewski, M. (2016). ScienceDirect Autonomous taxicabs in Berlin – a spatiotemporal analysis of service performance, *00*.

Boesch, P. M., Ciari, F., & Axhausen, K. W. (2016). Required Autonomous Vehicle Fleet Sizes to Serve Different Levels of Demand. In *Transportation Research Board 95th Annual Meeting*. Retrieved from http://trid.trb.org/view.aspx?id=1392994

Brown, A., Repac, B., & Gonder, J. (2013). *Autonomous Vehicles Have a Wide Range of Possible Energy Impacts*. NREL, University of Maryland. Retrieved from http://www.osti.gov/scitech/biblio/1220221

Brownell, C. K. (2013). *Shared autonomous taxi networks: An analysis of transportation demand in NJ and a 21st century solution for congestion* (Bachelor Thesis). Princeton University.

Burgess, E. W. (2008). The growth of the city: an introduction to a research project. In *Urban ecology* (pp. 71–78). Springer. Retrieved from https://link.springer.com/chapter/10.1007/978-0-387-73412-5_5

Burns, L. D., Jordan, W. C., & Scarborough, B. A. (2013). Transforming personal mobility. *The Earth Institute*.

Carbaugh, J., Godbole, D. N., & Sengupta, R. (1998). Safety and capacity analysis of automated and manual highway systems. *Transportation Research Part C: Emerging Technologies*, *6*(1), 69–99.

Cervero, R., & Guerra, E. (2011). Urban Densities and Transit: A Multi-dimensional Perspective. Retrieved from http://its.berkeley.edu/publications/UCB/2011/VWP/UCB-ITS-VWP-2011-6.pdf

Chambliss, J. C. (2008). A question of progress and welfare: The jitney bus phenomenon in atlanta, 1915-1925. *The Georgia Historical Quarterly*, 486–506.

Chase, R. (2014, April 3). Will a world of driverless cars be heaven or hell? *The Atlantic*. Retrieved from http://www.citylab.com/commute/2014/04/will-world-driverless-cars-be-heaven-or-hell/8784/

Chemla, D., Meunier, F., & Calvo, R. W. (2013). Bike sharing systems: Solving the static rebalancing problem. *Discrete Optimization*, *10*(2), 120–146.

Chen, T. D. (2015). Management of a Shared , Autonomous , Electric Vehicle Fleet : Vehicle Choice , Charging Infrastructure & Pricing Strategies.

Childress, S., Nichols, B., Charlton, B., & Coe, S. (2015). Using an Activity-Based Model to Explore the Potential Impacts of Automated Vehicles. *Transportation Research Record: Journal of the Transportation Research Board*, *2493*(2493), 99–106. https://doi.org/10.3141/2493-11

Christensen, C. (2013). *The innovator's dilemma: When new technologies cause great firms to fail*. Harvard Business Review Press.

Clayborne, D. D. (2012). *Owner-drivers in the tro-tro industry: A look at jitney service provision in Accra, Ghana* (Doctoral Dissertation). University of Los Angeles.

Cohen, P., Hahn, R., Hall, J., Levitt, S., & Metcalfe, R. (2016). *Using Big Data to Estimate Consumer Surplus: The Case of Uber* (Working Paper No. 22627). National Bureau of Economic Research. https://doi.org/10.3386/w22627

Dell'Amico, M., Hadjicostantinou, E., Iori, M., & Novellani, S. (2014). The bike sharing rebalancing problem: Mathematical formulations and benchmark instances. *Omega*, *45*, 7–19. https://doi.org/10.1016/j.omega.2013.12.001

Demery, L., Higgins, W., & Setty, M. (2007). *Traffic density thresholds for rail transit: A retrospective* (No. Special Report No. 2). Vallejo, CA: Public Transit Insititue.

Fagnant, D. (2014). *The future of fully automated vehicles: Opportunities for vehicle-and ride-sharing, with cost and emissions savings* (Doctoral Dissertation). University of Texas Austin.

Fagnant, D. J. (2014). *The Future of Fully Automated Vehicles: Opportunities for Vehicle- and Ride-Sharing, with Cost and Emissions Savings*. University of Texas at Austin.

Fagnant, D. J., & Kockelman, K. M. (2014). Development and Application of a Network-Based Shared Automated Vehicle Model in Austin, Texas. In *Transportation Research Board Conference on Innovations in Travel Modeling* (p. 15).

Fagnant, D., & Kockelman, K. M. (2014). The travel and environmental implications of shared autonomous vehicles, using agent-based model scenarios. *Transportation Research Part C: Emerging Technologies*, *40*, 1–13.

Febbraro, A., Sacco, N., & Saeednia, M. (2012). One-way carsharing: Solving the relocation problem. *Transportation Research Record: Journal of the Transportation Research Board*, *2319*, 113–120. https://doi.org/10.3141/2319-13

Feynman, R. P. (1998). The uncertainty of science. *Whitfield And*. Retrieved from http://maps.unomaha.edu/maher/GEOL1010/lecture1/feynman-1.pdf

Fishelson, J., Freckleton, D., & Heaslip, K. (2013). Evaluation of automated electric transportation deployment strategies: Integrated against isolated. *IET Intelligent Transport Systems*, *7*(3), 337–344.

Ford, H. J. (2012). Shared Autonomous Taxis : Implementing an Efficient Alternative to Automobile Dependency, (June), 104.

Ge, Y., Knittel, C. R., MacKenzie, D., & Zoepf, S. (2016). *Racial and gender discrimination in transportation network companies*. National Bureau of Economic Research.

Granovetter, M. (1978). Threshold models of collective behavior. *American Journal of Sociology*, 1420–1443.

Greenblatt, J. B., & Saxena, S. (2015). Autonomous taxis could greatly reduce greenhouse-gas emissions of US light-duty vehicles. *Nature Climate Change*, *5*(July). https://doi.org/10.1038/nclimate2685

Gucwa, M. A. (2014). *The mobility and energy impacts of automated cars*. Stanford University.

Hamidi, S., Ewing, R., Preuss, I., & Dodds, A. (2015). Measuring Sprawl and Its Impacts An Update. *Journal of Planning Education and Research*, *35*(1), 35–50.

Hars, A. (2015). Self-driving cars: The digital transformation of mobility. In *Marktplätze im Umbruch* (pp. 539–549). Springer. Retrieved from http://link.springer.com/chapter/10.1007/978-3-662-43782-7_57

Henao, A. (2017). *Impacts of Ridesourcing - Lyft and Uber - on Transportation Including VMT, Mode Replacement, Parking, and Travel Behavior* (Dissertation). University of Colorado Boulder.

Henderson, P. (2017, May 25). Some Uber and Lyft riders are giving up their own cars: Reuters/Ipsos poll. *Reuters*. Retrieved from http://www.reuters.com/article/us-autos-rideservices-poll-idUSKBN18L1DA

Hodges, A. (2006). 'Roping the Wild Jitney': the jitney bus craze and the rise of urban autobus systems. *Planning Perspectives*, *21*(3), 253–276.

Hodges, G. R. G. (2009). *Taxi!: A social history of the New York City cabdriver*. Baltimore, MD: Johns Hopkins University Press.

Horl, S., Erath, A., & Axhausen, K. W. (2016). *Simulation of autonomous taxis in a multi-modal traffic scenario with dynamic demand*. ETH Zurich. Retrieved from https://www.ethz.ch/content/dam/ethz/special-interest/baug/ivt/ivt-dam/vpl/reports/1101-1200/ab1184.pdf

Iglesias, R., Rossi, F., Zhang, R., & Pavone, M. (2016). A BCMP Network Approach to Modeling and Controlling Autonomous Mobility-on-Demand Systems. *ArXiv:1607.04357 [Cs]*. Retrieved from http://arxiv.org/abs/1607.04357

International Transport Forum. (2016). *Shared Mobility: Innovation for Liveable Cities*. Retrieved from https://www.itf-oecd.org/shared-mobility-innovation-liveable-cities

Javanshour, F., & Dia, H. (2016). Development and Evaluation of Models for Autonomous Shared Mobility-on-Demand Systems, (Unrsc), 1–12.

Kaas, H. W. (2016). *Automotive revolution–perspective towards 2030: How the convergence of disruptive technology-driven trends could transform the auto industry*. McKinsey & Company. Retrieved from http://www.voced.edu.au/content/ngv:74173

Kang, N., Feinberg, F. M., & Papalambros, P. Y. (2016). Autonomous Electric Vehicle Sharing System Design. *Journal of Mechanical Design*, *139*(1), 011402-011402-10. https://doi.org/10.1115/1.4034471

Karsten, C. F. K. and J. (2017, October 16). Gauging investment in self-driving cars. Retrieved January 3, 2018, from https://www.brookings.edu/research/gauging-investment-in-self-driving-cars/

Kelly, F. P., & Laws, C. N. (1993). Dynamic routing in open queueing networks: Brownian models, cut constraints and resource pooling. *Queueing Systems*, *13*(1), 47–86.

Kranzberg, M. (1986). Technology and History: “Kranzberg’s Laws.” *Technology and Culture*, *27*(3), 544–560. https://doi.org/10.2307/3105385

Lazarus, J., Shaheen, S., Young, S. E., Fagnant, D., Voege, T., Baumgardner, W., … Lott, J. S. (2018). Shared Automated Mobility and Public Transport. In *Road Vehicle Automation 4* (pp. 141–161). Springer. Retrieved from https://link.springer.com/chapter/10.1007/978-3-319-60934-8_13

Levin, M. W., & Boyles, S. D. (2015). Effects of autonomous vehicle ownership on trip, mode, and route choice. In *Transportation Research Board 94th Annual Meeting*.

Levin, M. W., Li, T., Boyles, S. D., & Kockelman, K. M. (2016). A general framework for modeling shared autonomous vehicles, (January), 1–23.

Levine, J., Grengs, J., Shen, Q., & Shen, Q. (2012). Does accessibility require density or speed? A comparison of fast versus close in getting where you want to go in US metropolitan regions. *Journal of the American Planning Association*, *78*(2), 157–172.

Levine, J., Zellner, M., Shiftan, Y., Arquero de Alarcon, M., & Diffenderfer, A. (2013). *Effects of automated transit, pedestrian and bicycling facilities on urban travel patterns* (p. 260). Federal Highway Administration.

Liang, X., Correia, G. H. de A., & van Arem, B. (2016). Optimizing the service area and trip selection of an electric automated taxi system used for the last mile of train trips. *Transportation Research Part E: Logistics and Transportation Review*, *93*, 115–129. https://doi.org/10.1016/j.tre.2016.05.006

Linares, M. P., Montero, L., Barceló, J., & Carmona, C. (2016). A simulation framework for real-time assessment of dynamic ride sharing demand responsive transportation models. In *Proceedings of the 2016 Winter Simulation Conference* (pp. 2216–2227). IEEE Press. Retrieved from http://dl.acm.org/citation.cfm?id=3042370

Litman, T. (2015). Autonomous vehicle implementation predictions: Implications for transport planning. In *Transportation Research Board 94th Annual Meeting*.

Lutin, J. M., & Kornhauser, A. L. (2014). Application of autonomous driving technology to transit-functional capabilities for safety and capacity. In *Transportation Research Board 93rd Annual Meeting*.

Lutz. (2017, December 11). Bob Lutz: Kiss the good times goodbye. *Automotive News*. Retrieved from http://www.autonews.com/article/20171105/industry_redesigned/171109944/industry-redesigned-bob-lutz

Maheo, A., Kilby, P., & Van Hentenryck, P. (2016). Benders Decomposition for the Design of a Hub and Shuttle Public Transit System.

Martinez, L. M., Correia, G. H., & Viegas, J. M. (2015). An agent-based model to assess the impacts of introducing a shared-taxi system in Lisbon (Portugal). *Journal of Advanced Transportation*, *49*(3), 475–495. https://doi.org/10.1002/atr.1283

Martinez, L. M., & Viegas, J. M. (2016). Concepts and impacts of new urban shared mobility alternatives: An agent-based simulation model for the city of Lisbon, Portugal. In *Transportation Research Board 95th Annual Meeting2* (p. 16).

May, A. D. (1990). *Traffic flow fundamentals*. Retrieved from https://trid.trb.org/view.aspx?id=356201

New York City Taxi and Limousine Commission. (2014). The New York City taxicab fact book.

Railsback, S. F., & Grimm, V. (2011). *Agent-based and individual-based modeling: a practical introduction*. Princeton university press. Retrieved from https://books.google.com/books?hl=en&lr=&id=tSI2DkMtoWQC&oi=fnd&pg=PP2&dq=Agent-Based+and+Individual-Based+Modeling.&ots=dQ8_Wz8NWJ&sig=MVDxJTfsvmwZgrIZrD7OhI4MhH4

Rayle, L., Dai, D., Chan, N., Cervero, R., & Shaheen, S. (2016). Just a better taxi? A survey-based comparison of taxis, transit, and ridesourcing services in San Francisco. *Transport Policy*, *45*, 168–178. https://doi.org/10.1016/j.tranpol.2015.10.004

Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.

Shaheen, S. A., Chan, N. D., & Micheaux, H. (2015). One-way carsharing's evolution and operator perspectives from the Americas. *Transportation*, *42*(3), 519–536. https://doi.org/10.1007/s11116-015-9607-0

Shaheen, S., & Cohen, A. (2007). Growth in worldwide carsharing: An international comparison. *Transportation Research Record: Journal of the Transportation Research Board*, (1992), 81–89.

Shen, W., & Lopes, C. (2015). Managing autonomous mobility on demand systems for better passenger experience. *ArXiv Preprint ArXiv:1507.02563*.

Shepherd, S., & Muir, H. (2011). *CityMobil D2.3.2 Strategic Modelling Results* (No. May 2006) (p. 42). European Commission.

Shladover, S. E., Desoer, C. A., Hedrick, J. K., Tomizuka, M., Walrand, J., Zhang, W.-B., … McKeown, N. (1991). Automated vehicle control developments in the PATH Program. *IEEE Transactions on Vehicular Technology*, *40*(1), 114–130.

Solé, R. (2011). *Phase transitions*. Princeton University PRess.

Spieser, K., Treleaven, K., Zhang, R., Frazzoli, E., Morton, D., & Pavone, M. (2014). Toward a systematic approach to the design and evaluation of automated mobility-on-demand systems: A case study in Singapore. In *Road Vehicle Automation* (pp. 229–245). Springer.

Sterman, J. D. (1991). A skeptic's guide to computer models. In G. O. Barney (Ed.), *Managing a nation: The microcomputer software catalog* (pp. 209–229). Boulder, CO: Westview Press.

Townsend, A. (2014). Re-programming mobility: The digital transformation of transportation in the United States. *New York: Rudin Center for Transportation Policy & Management Robert F. Wagner Graduate School of Public Service, New York University*.

Wachs, M. (2010). Futurama. *Issues in Science and Technology*, *26*(4), 90–92.

Walker, J. (2011). *Human transit: How clearer thinking about public transit can enrich our communities and our lives*. Washington, DC: Island Press.

Winter, K., Cats, O., Correia, G. H. de A., & van Arem, B. (2016). Designing an Automated Demand-Responsive Transport System: Fleet Size and Performance Analysis for a Campus–Train Station Service. *Transportation Research Record: Journal of the Transportation Research Board*, (2542), 75–83.

Yang, H., & Wong, S. C. (1998). A network model of urban taxi services. *Transportation Research Part B: Methodological*, *32*(4), 235–246. https://doi.org/10.1016/S0191-2615(97)00042-8

Zachariah, J., Gao, J., Kornhauser, A., & Mufti, T. (2014). Uncongested Mobility for All: A Proposal for an Area Wide Autonomous Taxi System in New Jersey. In *Transportation Research Board 93rd Annual Meeting*. Retrieved from https://trid.trb.org/view.aspx?id=1288288

Zhang, J. (2011). Tipping and residential segregation: A unified schelling model*. *Journal of Regional Science*, *51*(1), 167–193.

Zhang, R., Spieser, K., Frazzoli, E., & Pavone, M. (2015). Models , Algorithms , and Evaluation for Autonomous Mobility-On-Demand Systems, 2573–2587.

Zhang, W., Guhathakurta, S., Fang, J., & Zhang, G. (2015a). Exploring the impact of shared autonomous vehicles on urban parking demand: An agent-based simulation approach. *Sustainable Cities and Society*, *19*, 34–45.

Zhang, W., Guhathakurta, S., Fang, J., & Zhang, G. (2015b). The Performance and Benefits of a Shared Autonomous Vehicles Based 2 Dynamic Ridesharing System: An Agent-Based Simulation Approach. In *Transportation Research Board* (p. 15).

Zhang, W., Guhathakurta, S., Fang, J., & Zhang, G. (2015c). The performance and benefits of a shared autonomous vehicles based dynamic ridesharing system: An agent-based simulation approach. In *Transportation Research Board 94th Annual Meeting*.

Zhang, W., Guhathakurta, S., & Ross, C. (2017). *Parking spaces in the age of shared autonomous vehicles: How much parking will we need and where? 2*. Retrieved from http://docs.trb.org/prp/17-05399.pdf

Zoepf, S., & Heywood, J. B. (2012). Characterizations of deployment rates in automotive technology. *SAE International Journal of Passenger Cars-Electronic and Electrical Systems*, *5*(2012-01-1057), 541–552.