Plug-in Vehicles and Carsharing: User Preferences, Energy Consumption and Potential for Growth

by

Stephen M. Zoepf

B.S. Electrical Engineering and Computer Science, Massachusetts Institute of Technology (2001)

S.M. Technology and Policy, Massachusetts Institute of Technology (2011)

Submitted to the Engineering Systems Division
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Engineering Systems: Technology, Management, and Policy
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2015

Author . . . .
Engineering Systems Division
May 18, 2015

Certified by . . . .
Dan Frey
Professor of Mechanical Engineering and Engineering Systems
Doctoral Committee Chair

Certified by . . . .
John Heywood
Sun Jae Professor Emeritus of Mechanical Engineering
Thesis Supervisor

Certified by . . . .
Christopher Knittel
William Barton Rogers Professor of Energy Economics
Doctoral Committee Member

Certified by . . . .
Robert Hampshire
Research Assistant Professor
Transportation Research Institute, University of Michigan
Doctoral Committee Member

Accepted by . . . .
Munther A. Dahleh
William A. Coolidge Professor of Electrical Engineering and Computer Science, Acting Director Engineering Systems Division

Plug-in Vehicles and Carsharing: User Preferences, Energy Consumption and Potential for Growth

by

Stephen M. Zoepf

Submitted to the Engineering Systems Division
on May 18, 2015, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Engineering Systems: Technology, Management, and Policy

Abstract

Plug-in Electric Vehicles (PEVs) are seen as a key pathway to reduce fuel consumption and greenhouse gas emissions in transportation, yet their sales are under 1% of new cars despite large incentives. Carsharing, a market where consumers rent vehicles for short durations, is a low-risk way for consumers to use Plug-In Electric Vehicles for their travel needs without a large financial commitment. However, deployment of PEVs in carsharing depends on three key factors: (1) consumer acceptance of PEVs for rental trips, (2) the ability of carsharing providers to manage technical limitations of PEVs, and (3) that real-world energy consumption of PEVs meets expectations. To explore the feasibility of PEV deployment in carsharing, this dissertation incorporates a Mixed-Integer Programming and simulation of the assignment of trips and vehicles, and a Hybrid Choice Model of carsharing user preferences. This dissertation's primary contributions consist of the combination of Hybrid Choice Models with a Structural Topic Model to incorporate respondent comments, a two-level representation of the assignment problem faced by carsharing providers in allocating trips to vehicles and locating vehicles, a case study of PEV deployment in Boston, and analysis of real-world energy consumption of two fleets of PEVs. Results suggest that a large fraction of round-trip carsharing fleets could be converted to PEVs, simultaneously increasing profitability *and* reducing gasoline consumption, and some benefits can be captured using simple heuristics. However, current user attitudes towards PEVs in carsharing vary widely, and while carsharing exposes many users to hybrids, few have tried PEVs.

Doctoral Committee Chair: Dan Frey
Title: Professor of Mechanical Engineering and Engineering Systems

Thesis Supervisor: John Heywood
Title: Sun Jae Professor Emeritus of Mechanical Engineering

Doctoral Committee Member: Christopher Knittel
Title: William Barton Rogers Professor of Energy Economics

Doctoral Committee Member: Robert Hampshire
Title: Research Assistant Professor
Transportation Research Institute, University of Michigan

Acknowledgments

This thesis would not have been possible without the help and support of my family, advisor, colleagues, sponsors, and the community upon which we depend as graduate students.

First and foremost I would like to offer my sincere thanks to John Heywood for his guidance and support these past six years as my advisor. I am humbled by his immense knowledge, talent and kindness. Thank you also to my doctoral committee– Dan Frey, Chris Knittel and Robert Hampshire–for their ideas and advice these past four years.

Much of this work is based on research that is the product of collaborative efforts and has been (or will shortly be) published separately: Chapter 2 with David Keith, Chapter 3 with Lennon Rodgers, Johann Prenninger, Don MacKenzie, David Keith and William Chernicoff, and Chapter 4 with Alex Jacquillat. I thank them for both their enormous contributions and their blessing to include this work.

I would like to thank many current and former students from the Sloan Automotive Laboratory and ESD: Jake MacKenzie, Valerie Karplus, Fernando de Sisternes, Vivek Sakhrani, Ta Chiraphadhanakul and others. My work has been improved greatly by their collaboration and feedback.

I would also like to thank Todd at Zipcar and Jose and Johann at BMW for their willingness to support this work with data and their own insights into this field of research. Thank you also to those who have sponsored my work financially: Ford Motor Company, MIT Portugal Program, Chevron, ENI, and the Martin Family fellowship. This work would not have been possible without your contributions.

I have been blessed by the unwavering support of my wonderful wife, Irma. I cannot thank her enough for accompanying me on this academic adventure. I am also fortunate to have two wonderful children, Mia and Oliver, who continually remind me what is most important in life, and my parents and sister who set a high standard in their own intellectual pursuits and have always supported my endeavors.

Contents

List of Figures

List of Tables

Foreword

As I began my research at MIT by studying technology deployment in the automotive sector, I was struck by how many years it took for certain technologies to gain market acceptance. In particular, the automotive market seemed particularly slow to adopt fuel efficient technologies– even those which resulted in a net savings for the purchaser over the life of the car. The literature was full of explanations for why consumers might be slow to purchase energy-saving technology.

Almost by accident, I discovered a niche–carsharing– in which it seemed that technology adoption was progressing far faster. By my estimation, penetration rates of a wide variety of automotive powertrain technologies in the carsharing fleet were 3-5x higher than the average new car sold in the same year. This was true of technologies that were more apparent to the consumer (plug-in hybrids and electric vehicles) and those that were more transparent (direct fuel injection and advanced transmission types).

Because carsharing vehicles are driven by many different users, these relatively few cars seemed to be exposing large numbers of consumers to new vehicle technology. Furthermore, this technology exposure was happening in a low-risk environment: rather than spend tens of thousands of dollars to purchase a car, consumers needed only to invest a few dollars and an hour to use a new vehicle.

I decided to gain firsthand experience, and I leased a Chevrolet Volt Plug-in Hybrid vehicle and enrolled it in a peer-to-peer carsharing service. Over the course of the subsequent three years dozens of people inquired about the car and rented it. Many people were interested in renting the car specifically because it was a PHEV, but there was widespread misunderstanding about the capabilities and limitations of the vehicle. I eventually wrote an FAQ section on the listing of the vehicle to answer basic questions such as “What do you have to do when the battery is exhausted” and “How do you recharge the vehicle”

This realization led to the framing of the first two key research questions of this thesis. Does vehicle technology play a role when consumers rent a vehicle for a few hours, or are consumers interested only in finding an available vehicle near them? And how do consumers drive these vehicles, especially with regard to energy consumption?

Finally, as I considered the introduction of electric vehicles I realized that there was a fundamental tension. On one hand, carsharing operators wish to keep vehicles in operation and generating revenue as much of the day as possible. Yet electric vehicles face range constraints and require substantial time to recharge, even with 240V or DC fast charging. The final research question of this thesis investigates this tradeoff: how does the introduction of electric vehicles affect energy consumption in carsharing fleets, and how many can be deployed before service quality is diminished?

Chapter 1

Introduction: The Rise of Carsharing and Plug-in Vehicles

As a form of private transportation, passenger cars and light trucks present both a tremendous asset and daunting challenge. While they are an attractive a form of safe, convenient and comfortable transportation, they also account for nearly half of petroleum consumption and nearly a quarter of energy usage in the United States. (EIA, 2010; Kromer and Heywood, 2007; EPA, 2010)

In the past decade, an increasing number of businesses have started to offer consumers fractionally-owned, leased and rented products in the place of purchased goods. In the domain of personal transportation this trend has given rise to carsharing, a business model in which consumers rent cars by the hour or by the minute. Unlike traditional car rentals, vehicles operated by carsharing providers are typically distributed throughout a metropolitan area. Reservations are typically made via a web site or smartphone application, and the vehicles are access by an application or smart card.

Carsharing in some form has existed for decades, with early systems such as the Dutch Witkar being launched in 1974. Dozens of carsharing operators now exist, typically based in large cities in the U.S. and Europe. According to Shaheen and Cohen (2014) the carsharing market in the United States now consists of nearly 20,000 vehicles, still a small fraction of the overall U.S. light-duty vehicle fleet, which includes approximately 230 Million passenger cars and light trucks. (Davis et al., 2012) However, carsharing members now number approximately 1.3 Million members, a small but substantially more significant percentage of the U.S. driver population.

Since carsharing is such a small fraction of the fleet, it its direct impact on the overall fuel consumption and emissions levels in the near-term is small. However, carsharing services are growing rapidly. Levine (2012) claims that in many cities carsharing membership levels are growing at greater than 10%. Such growth suggests that carsharing services will soon have a greater impact than their existing membership numbers suggest.

1.1 Carsharing Business Models: Competing and Co-existing

Multiple business models compete in the carsharing market. Traditional carsharing business models have been an evolution of the car rental business in which users pickup and drop off vehicles in the same location. This business model, frequently referred to as "round trip carsharing," typically allows rentals in half-hour increments of at least one hour. Fuel is generally included up to a certain maximum daily mileage allowance, and users refuel vehicles when the fuel level is low using a credit card stored in the vehicle.

While round trip carsharing represents a majority of operations in the United States, other carsharing business models also exist. One-way carsharing is a business model in which users rent a vehicle in one location and can drop it off at a different location at the end of their rental period. Within one-way carsharing, there are two common implementations. Some are *station-based carsharing* systems, in which a user must return a vehicle to one of a specific set of permitted locations in an urban area. Others are *free-floating carsharing* systems, in which a user may leave a carsharing vehicle in any legal public parking spot within a defined perimeter in a metropolitan area. As with round-trip carsharing systems, fuel is also typically included in one-way carsharing rentals up to a daily cap.

A third type of carsharing operation is the peer-to-peer model. Peer-to-peer carsharing services do not actually own or operate a vehicle fleet. Rather, they rely on members to enroll their privately-owned or leased vehicles for rental to others and act as a marketplace for the service. Because vehicles are purchased by members, the variety of vehicles in the service is generally much greater, and users typically set the rental prices for their own vehicles. In some systems peer-to-peer rentals are available only on a daily basis, while others allow hourly rentals. Since vehicles are owned by individuals, peer-to-peer rentals typically require that the renter refill gasoline and also pay fees for mileage in excess of a pre-set limit.

In many major metropolitan areas in the United States, several of these carsharing providers co-exist. Because of their pricing models, the total cost a carsharing user pays for a reservation depends on both the length of time and the distance driven. However, since mileage limits, overage charges, and gasoline costs (if any) vary greatly, certain business models become more attractive to certain types of usage.

Figure 1-1 illustrates the intensively competitive nature of the carsharing business at present. These figures compare the prices of three large carsharing providers under their current price structure for a user taking a round-trip reservation.

Figure 1-1a shows a comparison of Car2Go (D.C. pricing), a Zipcar at $10/hour vs. a RelayRides vehicles at $50 per day with 250 miles allowed and refueled at $2.50/gallon and with an efficiency of 20mpg. In this scenario, Car2Go is the cheapest option only for users with very brief trips, while Zipcar is the cheapest for users traveling farther in a shorter period of time, while RelayRides is the cheapest for a longer reservation.

However, if the pricing changes slightly, the competitive field is entirely different. Figure 1-1a shows a scenario in which a user compares Car2Go, a Zipcar vehicle at $14/hour, and a RelayRides vehicle ($60 per day with 100 miles allowed and refueled at $4.50/gallon) and a Zipcar at $10/hour. In this scenario, Car2Go becomes the cheapest option in vastly more distance-time combinations, while Zipcar and RelayRides are the most attractive option for users with reservations that are long in time but short in distance, with Zipcar the cheapest under exactly the opposite scenario.

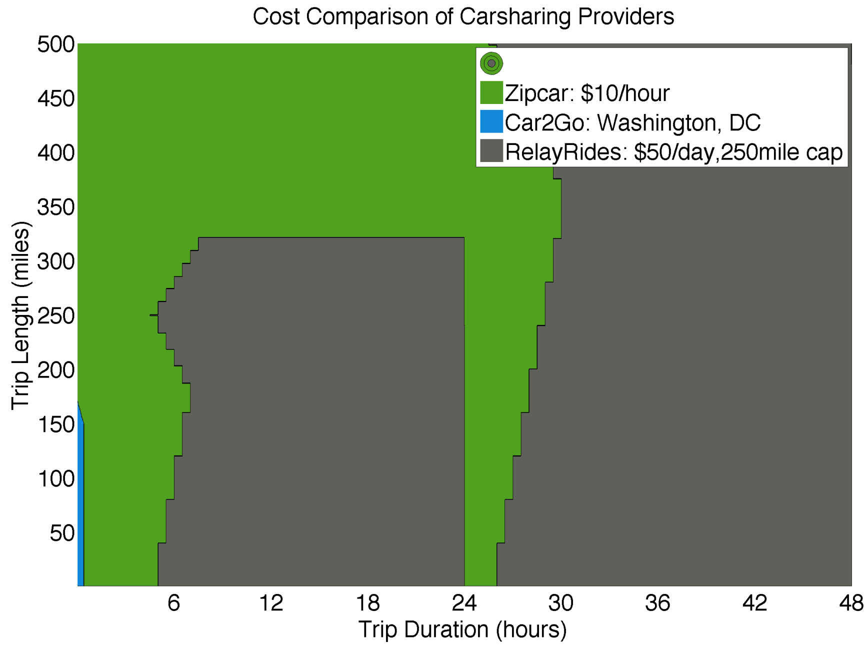

(a) Car2Go, a Zipcar vehicle at \$10/hour, and a RelayRides vehicle (\$50 per day with 250 miles allowed and refueled at \$2.50/gallon

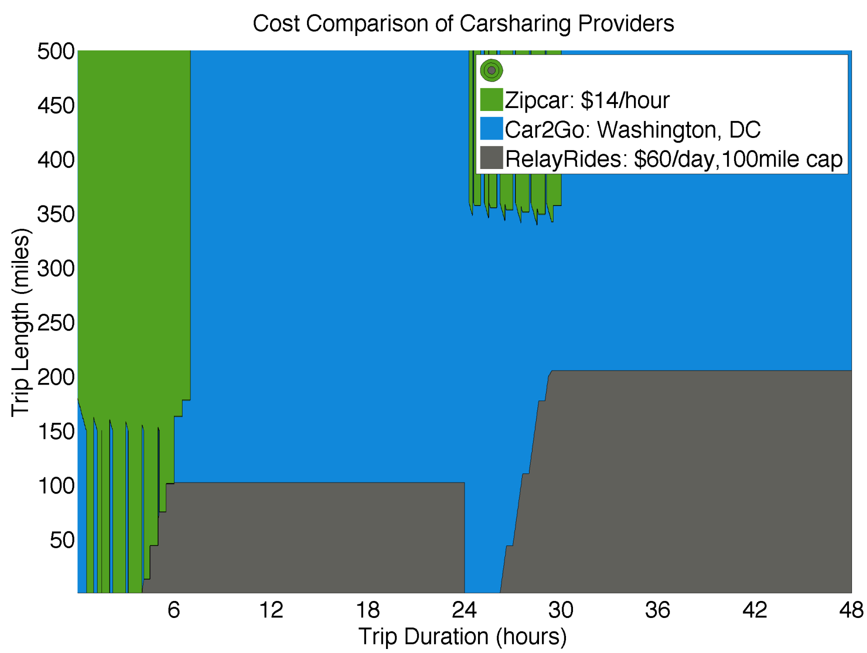

(b) Car2Go, a Zipcar vehicle at \$14/hour, and a RelayRides vehicle (\$60 per day with 100 miles allowed and refueled at \$4.50/gallon

Figure 1-1: Price comparison of a Car2Go rental, a RelayRides vehicle and a Zipcar under different pricing conditions.

These comparisons help explain why so many carsharing providers continue to co-exist: the business models are different enough in pricing alone to offer a competitive advantage under a wide variety of scenarios. No matter how extreme the pricing, no scenario resulted in a situation where a single provider held a pricing advantage to all price-distance combinations.

1.2 Plug-in Electric Vehicles

PEV is are a general term which is used to describe vehicles that source some or all of their motive energy from electricity. While there is some variation in precise definitions, PEVs are broadly divided into two categories: Plug-in Hybrid Electric Vehicles (PHEVs) and Battery Electric Vehicles (BEVs), which are sometimes abbreviated further to EVs (Electric Vehicles). Figure 1-2 highlights the naming conventions used in this work. Note that hybrid vehicles are generally not considered PEVs since none of their energy comes from the electrical grid.

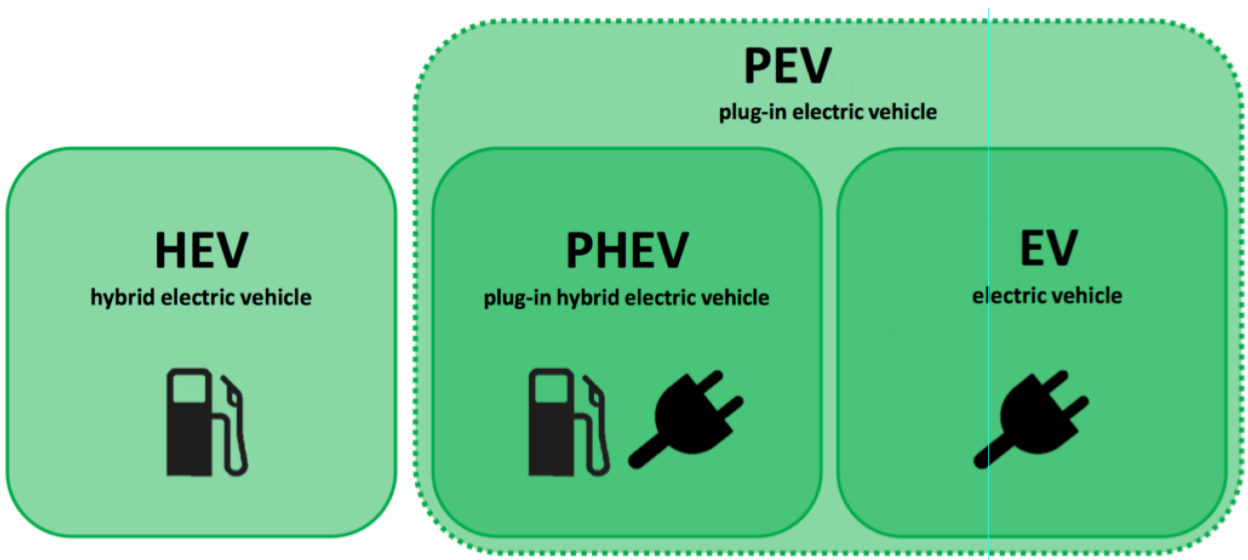

Figure 1-2: Types and names of PEVs. Figure adapted from U.S. Department of Energy.

BEVs are powered exclusively by electricity, and as a result they are typically characterized by a shorter range and longer refueling (charging) times than gasoline vehicles. Once their batteries are exhausted they must be connected to a source of electricity and recharged to replenish the battery. An alternative to long recharging times is "battery swapping" in which a vehicle's battery is disconnected and replaced by a full battery, but such systems are relatively rare in the transportation environment today. Because these vehicles cannot be run on gasoline, they are sometimes known as "pure electric vehicles."

PHEVs are vehicles that combine a gasoline engine with a battery pack and electric motor. Unlike a traditional hybrid, PHEV batteries can be recharged from the electric grid in addition to storing vehicle kinetic energy through regenerative braking. PHEVs are usually designed to favor electric operation over gasoline operation. In some cases this is strictly true, with the gasoline engine activating only when the battery is completely depleted. Other PHEVs choose to engage the gasoline engine under situations of high power demand, such as high speed, fast acceleration, or in very high or low temperatures when climate control is used aggressively.

The words BEV, EV or PHEV are sometimes appended with a number (e.g. EV100 or PHEV30). This number refers to the electric range of the vehicle in miles when the battery is fully charged, under nominal driving conditions. In the case of BEVs, this range is the total range of the vehicle. In the case of PHEVs, the number refers to the range of the vehicle on electricity and does not include gasoline driving range, which is typically substantially longer.

1.3 The Market for Plug-In Vehicles

Automotive development, production and usage and is characterized by life-cycle properties of a large-scale, complex, and socially embedded system. Vehicles are designed with long service lives, and according to the National Automobile Dealers Association (NADA), the average on-road vehicle in the U.S. is now nearly 12 years old.

Existing research has shown that new technologies take decades to form a substantial fraction of the automotive fleet. Bandivadekar (2008); Zoepf and Heywood (2012); Chon and Heywood (2000) Similarly, consumers have shown themselves to be risk averse regarding new powertrain technologies: conventional wisdom holds that mainstream buyers expect a three-year payback period before purchasing a vehicle that incorporates fuel-saving powertrain technologies. Cardell and Dunbar (1980) and Jaffe et al. (2000) point to a variety of reasons for this unwillingness to invest in energy efficiency technology, including market failure and uncertainty of future energy prices which result in a relatively high discount rate for future benefits of efficiency.

This unwillingness to invest in efficient technology has manifest itself in the slow adoption rate of hybrids and Plug-in Electric Vehicles (PEVs). Hybrid vehicles, fifteen years after their introduction to the U.S. market, still make up just 3% of new vehicles sales (Cobb, 2015) and less than 1% of the in-use fleet. PEVs have a smaller market share, with PHEVs accounting for 0.34% of new vehicle sales, and BEVs 0.39% of new vehicle sales despite subsidization by both U.S. government programs and automobile manufacturers. PEVs which were leased in 2011 and 2012 are now ending their contracts, and industry observers (Rogers, 2015) have noted that their resale values are far lower than predicted even a few years ago, with the average 2012 Nissan Leaf selling at auction for 25% of its original retail price.

Consumer adoption theory has a rich body of literature surrounding the adoption of products and services. Existing work by Keith and others have shown that consumer adoption of hybrid vehicles can be modeled as a contagion process based on the work of Everett Rogers, who cites five factors for the adoption of new technologies. (Rogers, 2003; Keith, 2012) One of these five factors, *trialability*, describes the ease with which a potential adopter can try a new technology with limited risk. For many potential purchasers, visiting a car dealer may have a real cost (time) or psychological cost, while carsharing offers the opportunity to try a plug in vehicle independently for a few hours and a few dollars. Thus carsharing may be viewed as a way to increase the trialability of new powertrain technologies.

1.4 Plug-in Vehicles and carsharing

A number of researchers have begun to highlight potential for benefits for the application of new vehicle powertrain technology in carsharing systems. Green et al. (2014) criticize what they view as a "mainstream market bias," or a tendency to want to push for mainstream adoption of new technology, when in reality it may be a better fit for smaller subsets of the market. The authors identify carsharing as one such well-suited market.

Vehicles in carsharing systems are typically small, fuel efficient vehicles that are characterized by low operating costs and compact size that makes them easy to drive and park in urban environments. Unlike traditional rental cars, which are frequently thought of as being bare-bones vehicles for fleet use, carsharing vehicles have a higher penetration of fuel-efficient technology than their privately purchased counterparts. Figure 1-3 shows a comparison of the approximate technology penetration of the shared fleet in 2012 (which included a variety of vehicles from the previous few years) with the penetration of the same technologies in the new vehicle fleet over the past three years. The usage of hybrids and plug-in vehicles was substantially higher, but there was also greater usage of other fuel-saving technology that is more transparent to users, such as continuously variable transmissions (CVTs) and direct fuel injection (DI).

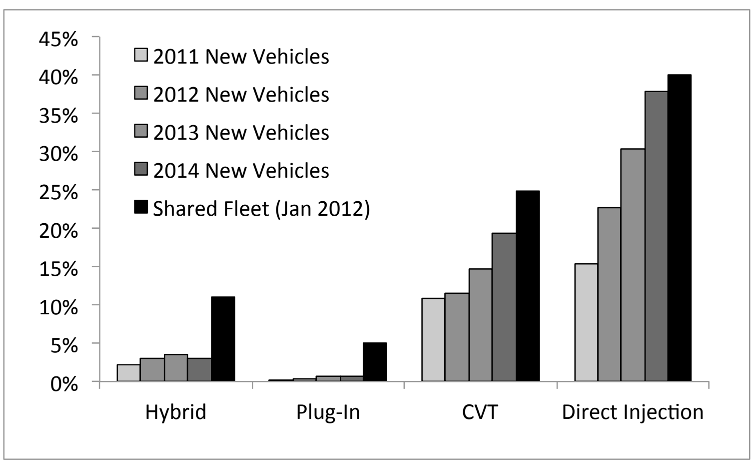

Figure 1-3: Adoption rates of fuel efficient technology are higher in carsharing fleets than the new vehicle fleet.

Researchers have identified a long and growing history of use of PEVs in carsharing Shaheen and Chan (2015). The behavior of members when selecting vehicles offers the opportunity to study consumer adoption of hybrids and PEVs under different circumstances than at the time of vehicle purchase. Carsharing users can select to drive specific vehicles including hybrids and PEVs with very low investment of money or time, lowering the barrier to entry. Additionally, since carsharing users choose a car for a specific purpose, they can select a gasoline vehicle for long trips while using an electric vehicle for shorter trips, choosing the best vehicle for their immediate needs.

A number of studies have noted that the introduction of PEVs into carsharing also alleviates the chicken-and-egg problem of Alternative Fuel or Plug-in Vehicles and their refueling/charging infrastructure. (Kriston et al., 2010; Green et al., 2014) In the private market, prospective plug-in vehicle buyers may be reluctant to purchase PEVs due to lack of sufficient density of infrastructure. Similarly, the earning potential of infrastructure depends on presence of vehicle owners who will use the recharging stations. In shared vehicle environments, the operator both bears the costs of infrastructure installation and appropriates the benefits of their usage. Similarly, the operator can also control access to the infrastructure to make sure that it is used effectively.

While the use of electric vehicles in carsharing seems highly attractive in theory, there are numerous practical challenges associated with the conversions of fleets from gasoline to alternative fuel vehicles. Infrastructure must be purchased, installed and maintained. Users must also be willing to use non-gasoline vehicles for their travel. Users must be educated to understand the benefits and limitations of new vehicle technology. Fundamentally, users must also be aware of their travel behavior and plan ahead. Since electric vehicles are characterized by shorter ranges and longer refueling (recharging) times than their gasoline counterparts, users must be able to distinguish whether specific PEVs will be adequate for a specific trip.

These questions of the integration of PEVs into carsharing services are explored here in greater detail. What are user attitudes towards PEVs and willingness to use them in the carsharing environment? What do we know about the performance of PEVs in the carsharing environment, and how can that shape our plans to use them? And finally, what is the feasibility for cost savings and environmental benefit for the greater integration of PEVs into the carsharing environment, and how can we use PEVs in an effective way? These questions are discussed in detail in the following sections.

1.5 Overview of the dissertation

This dissertation is organized into five chapters addressing a series of related topics at the intersection of vehicle technology adoption, energy consumption and alternative fuel infrastructure. This document is connected largely in a thematic fashion, and while the documents build on and develop results from prior chapters, each is written so as to to be largely self contained.

Chapter 2 investigates the question of how consumers choose vehicles to rent within a carsharing service. This chapter presents the results of a discrete choice survey that asks consumers to tradeoff amongst different vehicle and service attributes, and attempts to evaluate the role of demographics and attitude in these decisions.

Chapter 3 addresses the question of how consumers drive by analyzing the energy consumption of both PHEVs and BEVs in in two fleets of vehicles, one of which included vehicles in a carsharing service.

Chapter 4 investigates the potential for future deployment of EVs and other range-limited vehicles, analyzing the tradeoffs that shorter range and longer refueling time present when introduced in the more demanding context of shared usage.

Chapter 5 summarizes the key results of these studies and reflects on their implications for future vehicle technology deployment and policy.

The primary research questions presented in the body of this dissertation are briefly elaborated below.

User Preferences in Carsharing Systems: A Latent Variable Model

Research Theme 1

What service, vehicle and technology attributes are most important to carsharing users, and how do user attitudes to Plug-in Electric Vehicles vary?

When reserving a vehicle for a trip, carsharing users must trade off a number of variables.

First, there are the basic attributes of the service: can a user locate a vehicle when and where they want it and at a price they are willing to pay? Additionally, in most round-trip carsharing services users are able to select from a wide variety of vehicle types which vary in their size, brand, features and powertrain type (often including hybrids and gasoline-powered vehicles, and sometimes PEVs).

While Kriston et al. (2010) reported that users were willing to pay slightly more for vehicles with lower emissions, Le Vine et al. (2014a) did not find statistical significance on vehicle type when other service attributes were considered. Existing researchers have used regression models to identify some factors in small systems, such as that older cars are used less (de Lorimier and El-Geneidy, 2013) and a positive relationship exists between nearby rail access and usage.(Stillwater et al., 2009)

Additional anecdotal evidence from conversations with carsharing providers reveals a belief that users primarily select on two factors when a vehicle is available: walking distance to pickup a vehicle and price. For many carsharing providers, hybrid vehicles are priced cheaply, confounding the assumed price sensitivity with the willingness to use a hybrid vehicle.

This study consists of a discrete choice analysis of user preferences in a round-trip carsharing system. The data for this study is drawn from approximately 4,000 responses to an online survey conducted with a large round-trip carsharing provider. Information was gathered on user demographics and preferences in the service, and users were asked to respond to a panel of four discrete choice experiments. Additionally, users were asked to rate the vehicles they use in the service and to provide open-ended responses to questions about vehicles that they prefer and other comments about the service. The open-ended comments are analyzed using structural topic model (STM) techniques to generate indicators of carsharing user attitudes, which are then incorporated into a Hybrid Choice Model.

Research Theme 2

How do Plug-in Electric Vehicles perform in private and shared duty cycles?

Energy Consumption of Plug-in Electric Vehicles

Plug-in vehicles, while having been on the market for several years, represent a very small portion of the on-road vehicle fleet, and approximately 0.7% of new vehicle sales in 2014. (Cobb, 2015) Because of these low penetration rates, our understanding of the energy consumption of plug-in vehicles continues to evolve quickly.

Relatively little is known about the energy consumption of vehicles in carsharing services. There are a number of plausible reasons why energy consumption rates in shared vehicles could be higher or lower than those in private vehicles. On one hand, users of shared vehicles typically self-identify with support of environmental issues. And since vehicles are used more frequently, the vehicles will spend less time in energy-intensive warmup phases of driving. On the other hand, shared vehicles are typically used in urban areas with dense traffic. And since users do not pay directly for energy or maintenance, there is little incentive for users to drive in an efficient manner. Finally, since users are paying for the time they use the vehicle, drive less frequently, and may be unfamiliar with traffic patterns, they may drive more aggressively resulting in greater rates of acceleration and braking.

This study draws on two data sets of energy consumption in Plug-in Vehicles. The first is a fleet of 125 Toyota Prius Plug-In Hybrid vehicles that were loaned to government, academic and industry users for a period of one year. The second data set is drawn from a fleet of 700 BMW ActiveE electric vehicles, a portion of which were driven by private users, and a portion of which were used in the carsharing service DriveNow.

Using these data sets we analyze the energy consumption of the vehicles under a variety of conditions, and identify the variations in performance of the vehicles in shared and private usage. We also identify the importance of auxiliary energy, used for vehicle accessories, as a fraction of total energy consumption. Finally, the analysis includes a simulation to estimate the performance of PHEV vehicles under theoretical changes in vehicle design, charging frequency, and infrastructure conditions.

Deployment and Utilization of Plug-in Electric Vehicles in Round-trip Carsharing Systems

> **Research Theme 3**
>
> What is the feasibillity of greater use of PEVs in round-trip carsharing services while still meeting user demand?

As discussed in Section 1.4, BEVs are generally characterized by a shorter range than their gasoline counterparts. They also require substantial time to recharge onboard batteries, while gasoline-powered vehicles can refuel in approximately five minutes. These characteristics place additional constraints on carsharing operators: individual reservations cannot exceed the range of the vehicle, and down time must be scheduled for vehicles with a low state of charge.

Some vehicle sharing services including DriveNow and Car2Go have aggressively deployed BEVs as a means to meet California Zero-Emission Vehicle (ZEV) requirements (California Code of Regulations Title 13, Section 1962.1) without depending on private customers to adopt the new technology. However, since these services are relatively new it is not known to what extent the range constraints of PEVs will impact the ability of the services to mature. Since ZEV requirements do not stipulate any usage criteria for shared vehicles, carsharing providers could conceivably benefit from deploying more BEVs than can be efficiently used.

Data for this comes from approximately 50,000 observations from a large carsharing provider for one month in Boston, MA. Assumptions of energy consumption are informed by the real-world energy consumption values Chapter 3. This study uses an integrated approach combining Mixed Integer Programming and Simulation at two levels: how trips are assigned to vehicles, and how vehicles are assigned to Pods (common locations for round-trip shared vehicles). This combination of analytical techniques provides an upper bound to the estimated benefits of shared Electric Vehicles, as well as estimations of the feasibility under more conservative simulations that incorporate the timing of reservations.

The analysis calculates the net benefits of BEV deployment to a carsharing provider, and calculates the relative performance of a variety of heuristics to pick vehicle locations. It also extends the analysis to include the effect of user rejection of BEVs on the feasibility of deployment.

Contributions

This dissertation contributes to our knowledge in several key ways. First, it presents a new discrete choice analysis that combines traditional latent variable discrete choice modeling with unsupervised machine learning to develop latent variable indicators without the respondent burden of conventional attitudinal questions. The results of this analysis indicate that the words respondents use in open-ended questions may be a valuable indicator of attitudes that impact their choice behavior. This dissertation also develops an integrated approach to the deployment and utilization of Plug-in Electric Vehicles (PEVs) in a round-trip carsharing system. This approach allows the integrated assessment of both short-term decisions (how trips are assigned) and long-term strategies (how vehicles are placed) in order to both serve consumer travel demand and minimize gasoline consumption.

The results of this work have significant practical impact. The results of indicate that substantial portions of the existing gasoline fleet of round-trip carsharing could be converted to PEVs. Not only is the impact to service delivery small, but the deployment of BEVs would result in direct, private financial benefit to carsharing providers. These results suggest that round-trip carsharing can function as a near-term avenue for PEV deployment. However, the feasibility of this deployment depends on the willingness of users to accept PEVs when they have sufficient range for their travel needs. Chapter 2 indicates that some users may be reluctant to using PEVs even when their travel distances are short. In the near term these attitudes could be overcome with small financial incentives. Users have a positive view of hybrids, and carsharing services have acted as a conduit to hybrid technology for hundreds of thousands of drivers. Finally, results from Chapter 3 indicate that while users of carsharing vehicles have no direct financial incentive to drive efficiently, the energy consumption of shared BEVs is only marginally higher than that of privately owned BEVs, a difference which is primarily the result of accessories, not energy used to move the vehicle.

These outcomes have a number of public policy implications. Carsharing can provide a viable pathway towards greater PEV deployment and simultaneously act as avenue for consumer exposure to new vehicles types, but continued financial incentives are required for some consumers to try either PHEVs or BEVs. The benefit of introducing PEVs also varies widely depending on their placement. Current incentives for shared PEVs under the ZEV mandate do not consider utilization, only deployment, and they could be improved by more specific requirements on *how* PEVs are used to qualify for incentives.

Chapter 2

User Choices and Words: Text Responses as Indicators of Attitudes in Carsharing Rentals

2.1 Introduction

Carsharing is a business model in which users are able to rent cars for brief time periods, most commonly seen in North America and Europe. Such services differ from conventional car rentals primarily in that (1) Vehicles are distributed throughout an urban area rather than at dedicated locations, and (2) the duration of rentals is typically shorter, with pricing by the hour or minute. In selecting a vehicle to rent, users trade off a number of attributes: the price of the car, the distance they must travel to access the car and whether the vehicle is free exactly when they want it. Some carsharing services allow users to choose between a variety of cars, and many carsharing services rent advanced-technology vehicles: hybrids, plug-in hybrids and electric vehicles. Previous work has found that vehicle type is generally less important to users than the basic service attributes of price, distance, and schedule. This work examines the decision to use advanced-technology vehicles in greater detail.

The core of this work is a Multinomial Logit (MNL) Discrete Choice model based on a survey in which users of a large carsharing service were asked to choose between vehicles for their typical rental needs. Choice modeling has progressed dramatically in the past few years, and increasingly sophisticated models are now used to represent the choices users make when faced with a decision. Among these improvements is the application of a Hybrid Choice Model (HCM) which incorporates latent classes or variables that affect the choices users make. While not directly observable, we can observe indicators of latent variables based on a user's responses to attitudinal questions.

The use of indicators in choice modeling is an area of active research and is discussed in detail in Section 2.2.5. The most common approach uses the responses to Likert scale questions such as "Rate your agreement with the following statement ...." The responses to such questions are typically incorporated directly into a choice model. This work takes a different approach. Instead of including Likert scale questions about user attitudes, this study used a survey with two plain text response boxes. In order to reduce the dimensionality of the text, a Structural Topic Model was fit to each of the open ended text responses, and the results of the Structural Topic Models used as indicators of attitude in the final, integrated model.

This work presents the results of three primary research goals. First, to quantify the relative importance of basic carsharing service attributes: price, access distance, schedule and vehicle type. Second, to develop topic models of open-ended survey responses as indicators of user attitudes. And finally to show, using a latent variable model, how user attitudes about new vehicle types affect their willingness to use hybrids and Plug-in Electric Vehicles.

2.1.1 Research Questions

1. How do carsharing consumers trade off basic attributes such as distance, price, and schedule availability?

2. Does powertrain technology play a role in a user's decision to rent a vehicle through a carsharing service?

3. How many carsharing users are exposed to plug-in vehicles?

4. What user characteristics play an important role in the decision process for users?

5. How well do users recollect their own reservation behavior?

2.1.2 Overview of the Chapter

This chapter integrates three forms of modeling: a Discrete Choice Model, a Latent Variable model, and a Structural Topic Model. These three models are described briefly here and in further detail in Section 2.3.

A Discrete Choice Model Using repeated observations of which option users select when presented with a series of options, discrete choice models assume that users derive some unobservable "utility" or "usefulness" from each of the choices. The researcher then estimates coefficients that allow researchers to calculate how important certain features of a product or service are to the user.

A Latent Variable Model (MIMC) Model A latent variable model is a type of structural equation model that estimates unobservable attitudes as a function of demographic variables. While these attitudes cannot be directly observed, they are reflected in responses to attitudinal questions or other indicators.

Structural Topic Model A Structural Topic Model uses as its input a collection of open-ended text responses or documents, and a user-specified number of topics. The topic model, a form of unsupervised machine learning, produces associations of vocabulary words to topics, and assigns topic weights of each topic to each document.

In this work these three models are combined. First, basic Discrete Choice models are shown, illustrating the relative importance of vehicle powertrain and carsharing service attributes. Next, Structural Topic Models are fit to open-ended survey responses, which produce a series of topic weights. These topic weights are used as indicators in the final model presented, an simultaneous estimation of the Discrete Choice and Latent Variable models, referred to as an Integrated Choice and Latent Variable model (ICLV).

2.2 Literature Review

This research builds on three areas of work: published research on carsharing systems, the adoption of new vehicle technology, and unpublished results from a previous survey deployed by a carsharing operator. A review of literature in each of these areas follows below.

2.2.1 Published Literature on Carsharing Behavior

While the concept of carsharing has existed for decades, relatively little research has investigated the decision-making behavior when carsharing users select vehicles for use.

Several papers have investigated factors that contribute to overall levels of carsharing usage. Stillwater et al. (2009) investigate the attributes of the urban environment that influence the usage rates of carsharing vehicles, using aggregate reservation data from a single carsharing service. Barnes and Rutherford (2001) use a logit model to estimate the influence of various carsharing service attributes on the likelihood of prospective members to join carsharing. Membership fees and usage fees are found to be important, but access distance and reserve time were not found to be significant. Cervero et al. (2006) surveys City CarShare users and finds that car type was an important factor is vehicle choice for more than half of the users surveyed. Catalano et al. (2008) survey travelers in Palermo, Italy about travel preferences and model mode choice, including carsharing, as a function of cost and time. However, they do not include vehicle attributes, and the pricing levels used in the survey regarding carsharing are inconsistent (and appear far lower) than actual market pricing. As a result, we may not be able to draw meaningful conclusions about the utility of price for systems with very different pricing. The authors also are unable to separately estimate a coefficient for access time using the dataset and have not consider schedule as a decision factor.

More recently, de Lorimier and El-Geneidy (2013) use a regression approach and find that vehicle age and proximity to users are important decision factors. However, the authors use member concentration near stations use an independent variable and do not evaluate the decisions of individual users. Le Vine et al. (2014b) also model travel behavior with a binary mode choice model to use or not use carsharing, but the work focuses more strongly on trip purpose and does not include schedule covariates as in this work. While the authors ask respondents about vehicle type, they do not find vehicle type significant and do not include vehicle type in the results of the work.

Relatively little has been published regarding the role of techology in carsharing. An early study by Rutherford (2003) notes that vehicles in the Flexcar program at the time were 50% more fuel efficient than the average new vehicle sold in the U.S. More recently, a number of carsharing services have deployed large numbers of electric vehicles. According to a tweet from the Carsharing Association's conference in September, 2013 "RT @AutoShare: Half the people in San Francisco Bay Area who have driven electric cars did so thru @CityCarShare. #carsharing13."

In some carsharing services users are offered little choice between car types. At present, DriveNow in San Francisco currently uses an identical fleet of BMW ActiveE electric vehicles. Similarly, in North America Car2Go uses exclusively Smart ForTwo vehicles in gasoline and electric variants, and most regions are exclusively served by one type or the other. In systems such as these there is no diversity of price or vehicle attributes, and a reasonable assumption would be that consumers simply choose the nearest available vehicle for their trip.

In other services, however, users face a larger breadth of options. Zipcar, the largest carsharing operator in North America, uses approximately thirty different vehicles models in its fleet, and prices vary by both region and vehicle type. In such a system, users must weigh price, distance and vehicle type in their decision calculation. Zipcar also permits its members to view vehicle availability in a calendar format so that members may fit a trip between existing reservations. As a result, we anticipate that members may use schedule as a decision variable in their reservation behavior. The balance of these attributes in the decision making process is not well understood and is the focus of this research.

2.2.2 Previous Results from Industry Surveys

Zipcar has previously administered annual surveys that ask members to describe their preferences (Zipcar, 2012). Results from these survey are unpublished, but indicate that the top three factors users cite as influencing their decision to rent a particular car are proximity, time availability, and price. This survey, however, doesn't specify the relative importance of these features. We assume that users must tradeoff among these attributes in some way. A discrete choice modeling approach provides a tool with which to study these tradeoffs.

Respondents to the 2012 Zipcar survey rank environmental impact quite low in motivators for choosing a particular car. However, nearly two thirds of respondents indicated that they were either "interested" or "extremely interested" in Zipcar bringing on electric cars, and hybrids and electric cars were the two top-requested vehicles. Such responses seem to conflict with the low importance of environmental footprint users reported in selecting a vehicle and prompts a number of possible hypotheses. Perhaps users are interested in electric vehicles for some attribute besides their environmental footprint. It is also possible that respondents have the unspoken assumption that electric vehicles will be cheaper to operate than conventional vehicles, as Zipcar currently charges lower rates for hybrid vehicles. The stated-preference approach used here allows the individual utility contribution of price and fuel type to be separately explored.

2.2.3 Technology Adoption

In contrast to carsharing, there is a large body of literature that discusses the role of technology in a consumer's decision to purchase a vehicle. One of the earliest and most cited of such studies, Brownstone et al. (2000), merges revealed preference (RP) and stated preference (SP) data about consumer preferences for AFVs. The authors highlight substantial challenges in using both RP and SP methods to estimate consumer preferences. RP methods present numerical problems, such as high collinearity and limited variation among attributes in the market, while SP methods are limited by the simple fact that a user's reported selections may not reflect actual behavior.

Tanaka et al. (2014) reviews a large number of these studies, and compares attitudes in the U.S. and Japan. Al-Alawi and Bradley (2013) also perform a substantial literature review of vehicle technology adoption studies. The authors highlight a large degree of variation in the predicted adoption rate of alternative fuel vehicles (AFVs). The authors also recommend better integration of consumer choice models with policy incentives and supply side constraints.

Researchers have developed a number of insights that address public opinions about hybrid vehicles. Bolduc et al. (2008) and Glerum et al. (2012) specifically use Hybrid Choice Models, as does this work, to investigate the role of respondent attitudes in preferences for alternative fuel vehicles. Jensen et al. (2013) investigate how participant perceptions change after they are exposed to hybrids, while Sivak and Schoettle (2014) investigates the difference of attitudes between those who drive hybrids vs. those who drive conventional vehicles.

However, most studies until this point have dealt exclusively with the willingness of consumers to purchase or lease AFVs. A very small number of studies have included vehicle type within the carsharing environment. Le Vine et al. (2014b), for example, uses a stated preference method to identify willingness to use a carsharing service. However, the model does not find vehicle type to be significant.

2.2.4 Advances in Choice Modeling

The multinomial logit model (MNL) belongs to a class of models known as discrete choice models. In various forms these have been widely used for nearly fifty years (McFadden, 1974; Ben-Akiva and Lerman, 1985) and have been applied to problems in transportation, marketing, logistics and many more. The MNL model is highly effective as a tool to elicit relative preferences for product or service attributes from respondents without directly asking about them.

More recently, substantial advances have been made in discrete choice modeling. These advances have lead to a more generalized structure of discrete choice models, generally referred to as Hybrid Choice Models (HCMs), that allow the incorporation of latent variables, flexible disturbances, latent classes, and the combination of revealed and stated preference data. (Ben-Akiva et al., 2002; Walker and Ben-Akiva, 2002; Walker, 2001; Paulssen et al., 2014) Figure 2-1 illustrates the various components that may be included in Hybrid Choice Models.

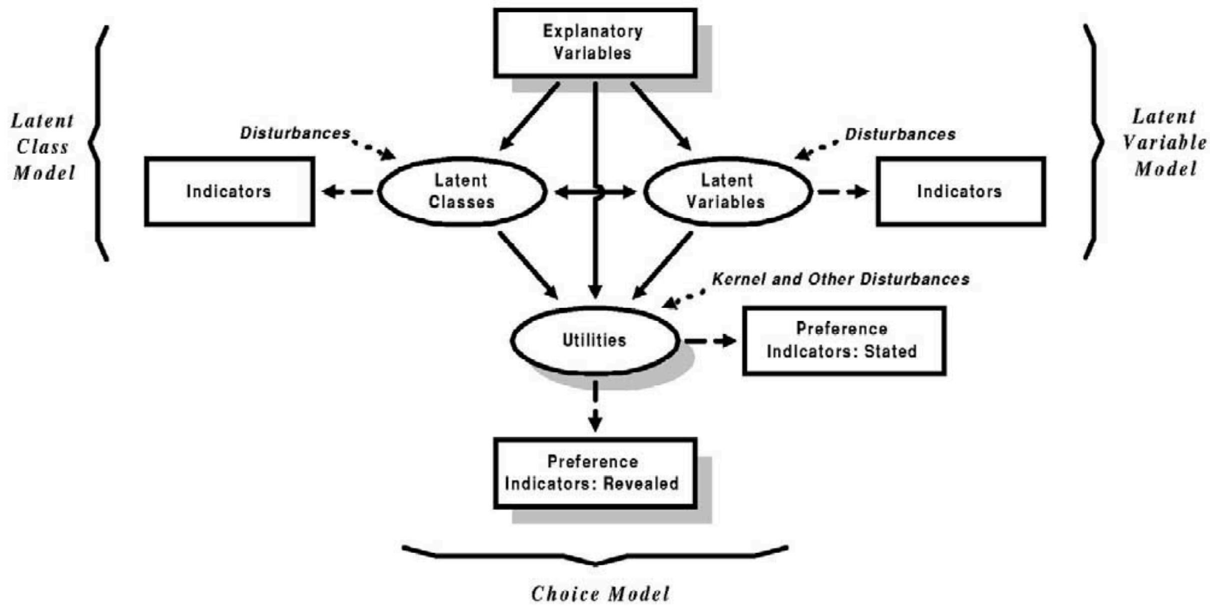

Figure 2-1: Visualization of Hybrid Choice Model components from Walker and Ben-Akiva (2002).

The work presented here builds specifically on an Integrated Choice and Latent Variable (ICLV) model, a structure which combines an MNL choice model with a MIMC (Multiple Indicators Multiple Causes) structural equation model. While relatively new, the ICLV has gained popularity in recent years. Theis (2011) applied latent variable techniques to airlines itineraries, Vij (2013) to travel modal demand, and Gaker et al. (2011) to the value of environmental-friendliness in travel preferences. Most closely related to this work is Alvarez-Daziano (2011), which attempts to incorporate latent variables into a model of consumer preferences for automotive technology using a hierarchical bayes approach. This work also develops a latent variable model of consumer preferences for automotive technology, but with two key differences. First, the application is to a carsharing environment, rather than a purchase decision, making the level of commitment much lower. Second, we here apply unsupervised machine learning techniques to develop indicators of latent variables rather than the more common approach of using a series of Likert scale responses from a questionnaire.

2.2.5 The Incorporation of Indicators

As discussed in section 2.2.4, HCMs which incorporate latent attitudes or classes are entering more common usage. A common method in these models is to incorporate Likert scale questionnaires (Likert, 1932) or Yes-No questions into a survey, and use ratings directly as indicators a respondent's attitudes. However, other techniques are increasingly being applied to discrete choice models, as shown in Figure 2-2.

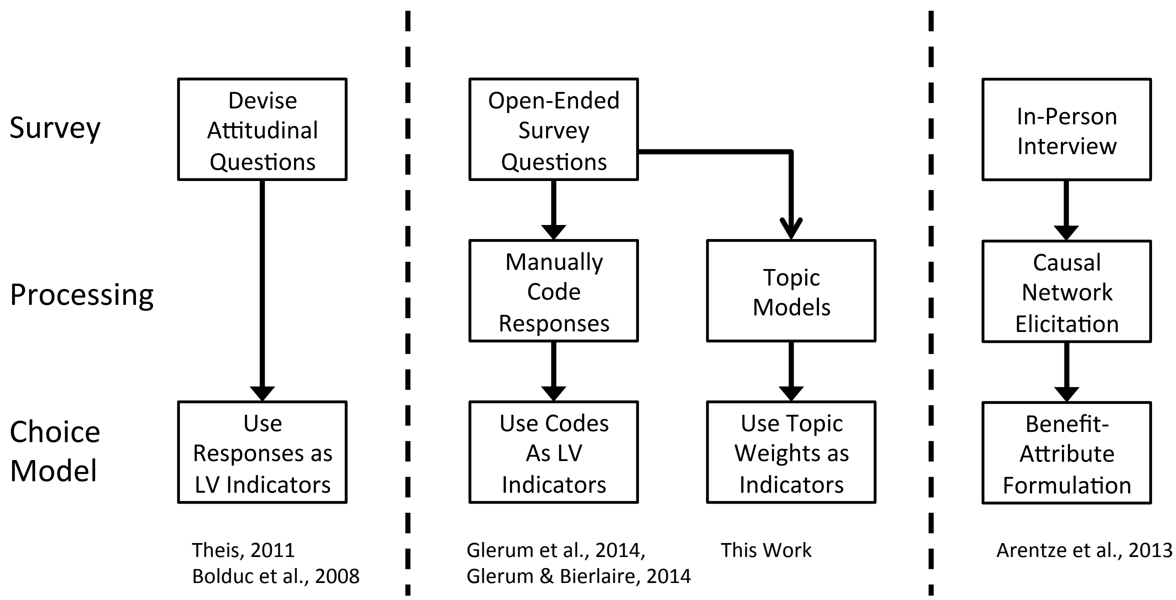

Figure 2-2: Current techniques for using indicators in choice models.

Likert Scale Responses as Indicators

Many surveys rely heavily on questions that require a respondent to rate statements on a Likert scale. Developed more than eighty years ago, Likert scale surveys are a staple in social science research. As the use of Hybrid Choice Models has grown, the majority of them continue to incorporate such responses directly as indicators of latent variables or latent classes. (Bolduc et al., 2008; Temme et al., 2008; Theis, 2011; Atasoy et al., 2010) There is continued discussion over the specific way such indicators should be included – as continuous indicators or ordinal indicators–with Daly et al. (2012) pointing to improvements in model fit with ordinal indicators.

However, researchers recognize that there are limitations to the use of these responses. Glerum and Bierlaire (2014) summarize a few of the problems: yea-saying, nay-saying, centrism, extremism, and others. The authors also cite a number of studies that note the scale used to assess agreement may itself impact the effectiveness of the rating exercise.

Likert-scale responses also introduce a number of practical difficulties. Since the number of Likert scale questions asked – typically a dozen or more – is generally much larger than the number of latent variables the rating exercise itself increases respondent burden and may limit the number of completed surveys. Another challenge is that any rating exercise incorporated into the survey instrument itself requires that the researcher develop *a priori* a set of questions that encompass all the psychological constructs to be examined. These requirements place a large burden on researchers to develop a compact, yet comprehensive survey instrument, and offer limited flexibility to incorporate mental models outside the scope of those previously envisioned by the researcher.

Open-Ended Responses

Free-form (open-ended) text responses are difficult to incorporate in models because of the vast amount of information they contain. In the past two years, researchers in the discrete choice literature have begun to incorporate more flexible response structures to incorporate latent attitudes or perceptions. One technique applied by Arentze et al. (2013) is Causal Network Elicitation Technique (CNET), which uses in-person interviews to develop a detailed representation of a respondent's thought process and incorporate decision, attribute and benefit variables into the structure of the choice model.

An alternative, applied by Glerum et al. (2014) and Glerum and Bierlaire (2014) is to use open-ended or semi-open ended responses. The researchers then apply coding techniques to the responses to quantify open-ended responses on a numerical scale.

These two techniques are promising, but the time and cost associated with data collection is high. Interviews inherently require a large time commitment by both the researcher and study participants. Coding can be done by third parties, which may limit burden to the research staff. However, coding is generally performed by multiple third parties to reduce bias and becomes costly with large numbers of responses.

For this study, confidentiality of the data set presented a challenge to use of third-party coding, as this exercise would have required multiple additional signatories to an existing non-disclosure agreement. As a result, an unsupervised machine-learning technique, topic modeling, was chosen to quantify open-ended responses and generate indicators.

2.2.6 Topic Modeling

The techniques discussed in Section 2.2.5 largely represent ways to reduce the dimensionality of natural language into a more compact form, but in a way that still captures the themes and ideas that the author of the original text expresses. If we can find a way to express the essence of the author's original intent in a more compact form, natural language content can be included usefully in models that incorporate user attitudes.

In choice models, attempts to incorporate open-ended text have been largely manual and labor intensive. However, in the past two decades there have been numerous advances in the field of machine learning. Of particular interest are topic models, which are now used to explore collections of documents and discover topics, or groups of words, from them. Blei (2012) notes that topic models conveniently provide "an algorithmic solution to managing, organizing, and annotating large archives of texts," which is precisely what is needed in the case of large-scale surveys that are common today.

The most commonly cited type of topic model, Latent Dirichlet Allocation (LDA) (Blei et al., 2003), is the root of several different types of topic models in use today. LDA models, a refinement of earlier Probabilistic Latent Semantic Indexing (Hofmann, 1999), observe a set of documents consisting of collections of words. In LDA, we observe a set of documents and the words they contain. But we cannot observe the hidden topic structure, which includes (i) the topics themselves, (ii) the topic distributions per document, and (iii) the topic assignments per document and per word. The process of LDA modeling computes the conditional distribution of these hidden variables given the observed variables, or the set of documents we begin with. LDA draws a document's mixture of topics from a Dirichlet distribution to help prevent overfitting, but imposes the assumption that topics are weakly independent. The generative process is discussed in greater detail in Blei (2012).

A modification of the LDA specification are Correlated Topic Models (Blei and Lafferty, 2006, 2007). Correlated Topic Models relax the assumption of independence of topics within a document. Rather than draw from a Dirichlet distribution, mixtures of topics are drawn from a logistic normal distribution, allowing for correlations between topics. [1]

In the past two years, researchers have further refined the CTM concept to allow the incorporation of respondent demographics and other covariates within the topic model structure. Structural Topic Models (Roberts et al., 2014b,c,a) build on the CTM model structure to incorporate covariates and allow covariates to affect both the topic *prevalence* and the topic *content*. The Structural Topic Model is the type of topic model fit to the open ended comments in this work.

It is also important to note that topic models represent documents as vectors of words counts, and the particular order in which words appear is not explicitly considered. Similarly, since topic models simply consider for relations between words, they offer a number of advantages of human coders attempting to interpret open-ended text. One advantage of this representation is that topic models are generally language-agnostic, in that they are capable of estimating hidden topic structure regardless of what language underlying words are written in. The only impact of language comes at the preprocessing stage when high

[1] Researchers have also extended the LDA model to an sLDA, or supervised LDA formulation in which the model jointly models topic allocation and a separate response variable. The authors apply these Supervised Topic Models (Blei and Mcauliffe, 2007) to model the content of movie reviews and the numerical star rating assigned to the movie.

frequency words (e.g. "and," "or," "the,") are discarded. Industry-specific words, slang, and acronyms are treated as any other word in a topic model, whereas such content might present a problem to human coders, particularly non-native speakers of the language in which the survey is written.

2.3 Analytical Approach

This section discusses the methodological approaches, the way in which they interact, and data sources. This section includes four distinct models: the choice model, the Multiple Indicators Multiple Causes model, the topic model, and the Integrated Choice and Latent Variable model. Figure 2-3 shows an outline of the process and relevant software packages.

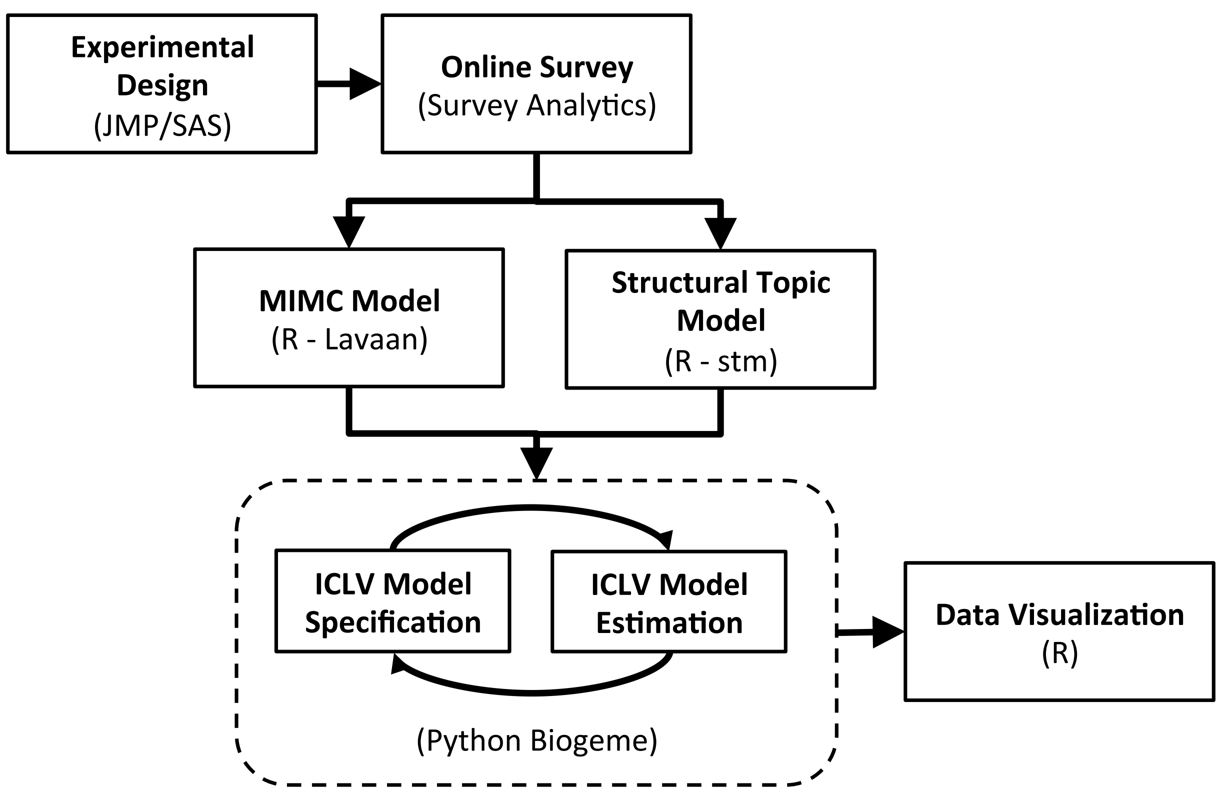

Figure 2-3: Overview of the data generation and analysis procedure used in this paper.

R and Python Biogeme (Bierlaire, 2003) are the primary software packages used for analysis. The STM package (Roberts et al., 2014b) permitted the development of topic models, and the Lavaan package (Rosseel, 2012) was used for structural equation modeling. We administered the survey using Survey Analytics (`http://www.surveyanalytics.com`).

2.3.1 Choice Model

This section briefly presents the choice model used here. Subsequent section discuss the integration of latent variables via a MIMC model, and the development of appropriate indicators using topic models. The following is a brief introduction to some of the terminology used:

Utility The unobservable "goodness" that a decision maker seeks

Decision Maker The decision maker is the person choosing amongst alternatives:

Alternatives Individual items or possibilities that the decision maker is choosing between

Choice Set All possible alternatives that a decision maker can choose between form a choice set (i.e. A, B, or C)

Attributes The aspects of an alternative that make it different from other alternatives (color, size, price, etc)

Levels The values that each attributes can take (i.e. $10/hour, red, 1.0 miles away)

Characteristics Demographic information about a decision maker

The MNL model, one of a class of random-utility models, assumes that users make decisions based on the *utility* of an alternative, which is a variable composed of both a deterministic component and a random component.

The deterministic component is a function of the attributes of that alternative and characteristics of the decision maker. The Utility U_{in} for alternative i for individual n depends on both the attributes of the alternative i and the characteristics of the decision maker n, as

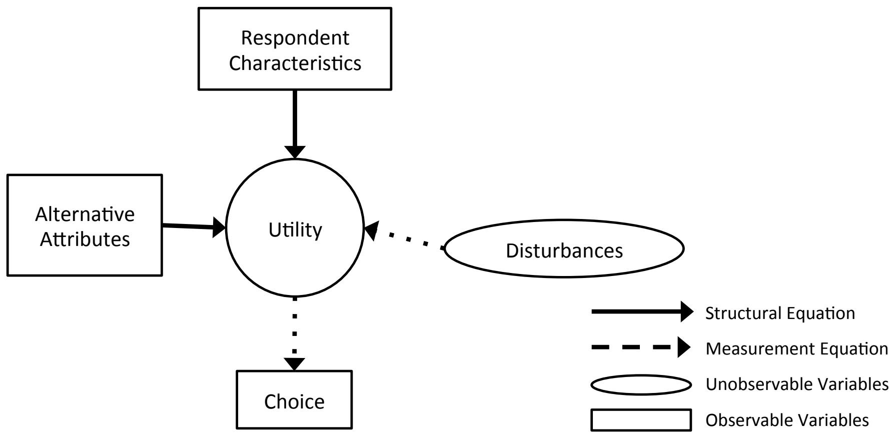

Figure 2-4: Framework of the basic MNL choice model.

shown in Equation 2.1. The error term ϵ_{in} is assumed to be i.i.d. extreme value distributed, and a vector of unknown parameters β is estimated.

Structural Equation:

$$U_{in} = X_{in}\beta + \epsilon_{in} \tag{2.1}$$

Measurement Equation:

$$Y_{in} = \begin{cases} 1 & \text{if } U_{in} \geq U_{jn}, \forall j \in C_n, j \neq i \\ 0 & \text{otherwise} \end{cases} \tag{2.2}$$

In the measurement equation 2.2, Y_{in} indicates the choice of the decision maker: zero if the respective alternative is not chosen; 1 if it is. In these equations, n denotes an individual, $n = 1, ..., N$ and i, j denote alternatives $i, j = 1, ..., J$. X_{in} is a (1 x K) vector of explanatory variables related to i and n, and β is a (K x 1) vector of unknown coefficients that we wish to estimate.

Using this format, the individual choice probability is described by:

$$P(y_{in} = 1 | X_n; \beta) = \frac{e^{X_{in}\beta}}{\sum\limits_{j \in C_n} e^{X_{in}\beta}} \tag{2.3}$$

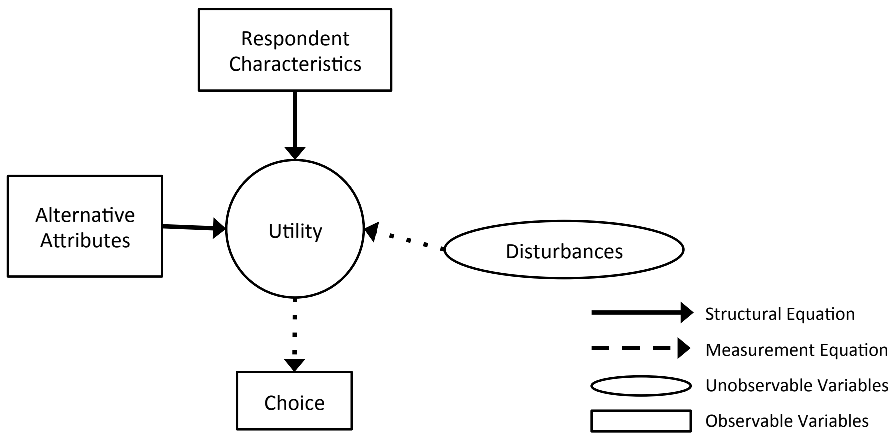

Figure 2-5: Framework of the basic MNL choice model.

2.3.2 The MIMC Model

The MIMC model (also spelled MIMIC), is a structural equation model which also incorporates both structural and measurement equations. Like utility in the choice model, we cannot observe the latent variables: only their causes and indicators, as shown in Figure 2-6. Structural equation models are widely used in social science literature, studying the interrelationship of variables by studying their covariance structure.

Structural Equation:

$$X_n^* = \Lambda \tilde{X}_n + \omega_n \tag{2.4}$$

where Λ is a matrix of unknown parameters (L x M), $\tilde{X}_n$ is an (M x 1) vector of explanatory variables causing the latent variables and ω_n is an (L x 1) vector of disturbances of individual n. ω_n is normally distributed with a zero mean and variance I and we assume

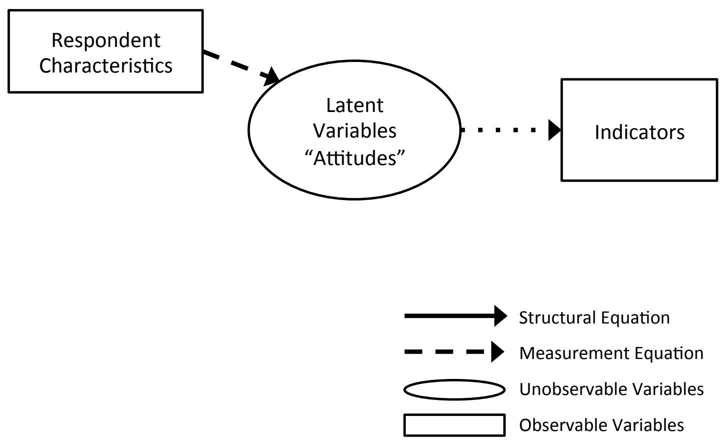

Figure 2-6: Framework of the Multiple Indicators Multiple Causes model.

all covariances are 0, or independence of the error terms.

Measurement Equation:

$$I_n = \kappa + AX_n^* + \nu_n, \nu_n \sim N(0, \Sigma_v \text{diagonal}) \tag{2.5}$$

where I_n is a (R x 1) vector of latent variables of the individual n, κ is an (R x 1) vector of intercepts, A is a matrix of unknown parameters (R x L), X_n^* is a (L x 1) vector of latent explanatory variables of the individual n, and ν_n is a (R x 1) vector of disturbances of individual n. We assume conditional independence of the error term (i.e. covariances are 0); we therefore assume all correlation between indicators is explained by the latent variable. There is thus a diagonal matrix with R variances.

Likelihood Function:

We can write the likelihood function of the model as follows:

$$f(I_n \mid X_n; \alpha, \lambda, \Sigma_v, \Sigma_\epsilon)) = \int_{X^*} f_M(I \mid X^*; \alpha, \Sigma_\nu) f_S(X^* \mid \tilde{X}; \lambda, \Sigma_\omega) dX^* \tag{2.6}$$

While early model specifications incorporated both binary and continuous indicators, the final results presented later in this chapter only use continuous indicators with the distribution as follows, from Bolduc and Alvarez-Daziano (2010).

$$f(I_r) = \frac{1}{\sigma_{\nu_r}} \Phi(\frac{I_r - \kappa_r - X^* \alpha_r}{\sigma_{\nu_r}}) \tag{2.7}$$

Where Φ is the standard normal density function.

2.3.3 The Integration of Latent Variables

We now combine the structural and measurement equations from both the choice and MIMC models, as shown in Figure 2-7. In this specification, the first structural equation is the choice model utility function, but with the latent variables integrated.

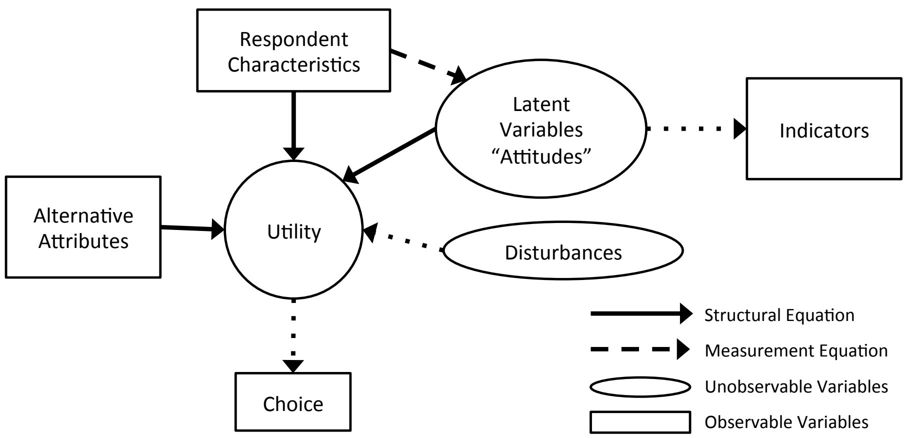

Figure 2-7: Framework of the Integrated Choice and Latent Variable model.

Structural Equations:

$$U_n = X_n \beta_1 + X_n^* \beta_2 + \epsilon_n \quad \epsilon \sim \text{i.i.d. extreme value} \tag{2.8}$$

$$X_n^* = \Lambda \tilde{X}_n + \omega_n \quad \omega \sim \text{N(0,I)} \tag{2.9}$$

Measurement Equations:

$$Y_{in} = \begin{cases} 1 & \text{if } U_{in} \geq U_{jn}, \forall j \in C_n, j \neq i \\ 0 & \text{otherwise} \end{cases} \tag{2.10}$$

$$I_n = \kappa + AX_n^* + \nu_n, \nu_n \sim N(0, \Sigma_v \text{diagonal}) \tag{2.11}$$

Likelihood Function:

For this model the likelihood function is shown in equation 2.12.

$$f(Y_n, I_n \mid X_n; \alpha, \beta, \lambda, \Sigma_\epsilon, \Sigma_\nu, \Sigma_\omega) = \int_{X^*} P(y_n \mid X_n, X^*; \beta) f_s(X^* \mid \tilde{X}; \lambda, \Sigma_\omega) f_M(I_n \mid X^*; \alpha, \Sigma_\nu) dX^* \tag{2.12}$$

Where the first component on the right-hand side of the equation is the logit formulation with variables and latent variable interactions, the second component is the structural equation of the latent variable model, and the third component is the measurement equation of the latent variable model, using continuous indicators as previously described. The variance of the structural equation is set to 1 to facilitate identification of the model.

2.3.4 Model Estimation

Choice models and ICLV models were estimated using Python Biogeme (Bierlaire, 2003). Estimation time is substantial for most models, and ICLV models were estimated on Amazon Web Services EC2 c3.8xlarge 32 core instances. ICLV models were run with only 1-2 Latent variables at a time as computational time scales exponentially with the number of estimated parameters. Estimation cost for a single model run at current Amazon EC2 prices averaged in the tens of dollars for ICLV models with one latent variable and in the hundreds of dollars for models with two latent variables.

2.4 Survey Design and Experimental Statistics

The data were collected in a survey of Zipcar users conducted online in October, 2013. 68,982 randomly-selected users were invited to participate in the survey via email. Users were instructed that they would be entered for a chance to win $50 in free driving credit for their complete response. 4673 unique respondents began the survey, and 3958 respondents completed the survey and at least a portion of the discrete choice panel, for a response rate of approximately 5.7%.

2.4.1 Data Cleaning

Additional partially incomplete responses and implausible distance/time combinations that resulted in average speeds in excess of 65mph were later removed. All zero travel distance responses were shifted to just above zero to avoid errors in discrete choice estimation using log travel distances. Duplicate responses from the same survey were also deleted, with the most complete version of the survey kept.

2.4.2 Respondent Characteristics

Respondent characteristics are summarized in Table 2.1. Users were generally young, slightly biased towards males, and approximately two-thirds did not own a vehicle. Of respondents who lived in a household with a car, the majority only had one car, with progressively fewer respondents with more cars. Few respondents (approximately 10%) had children. Respondents were most commonly from the largest metropolitan areas in which the carsharing operator is based: Boston, Chicago, New York, San Francisco, Toronto and Washington, D.C. A significant number of respondents were also classified as users in University settings, a collection of smaller areas in which the service operates.

Table 2.1: Graphical display of respondent attributes.

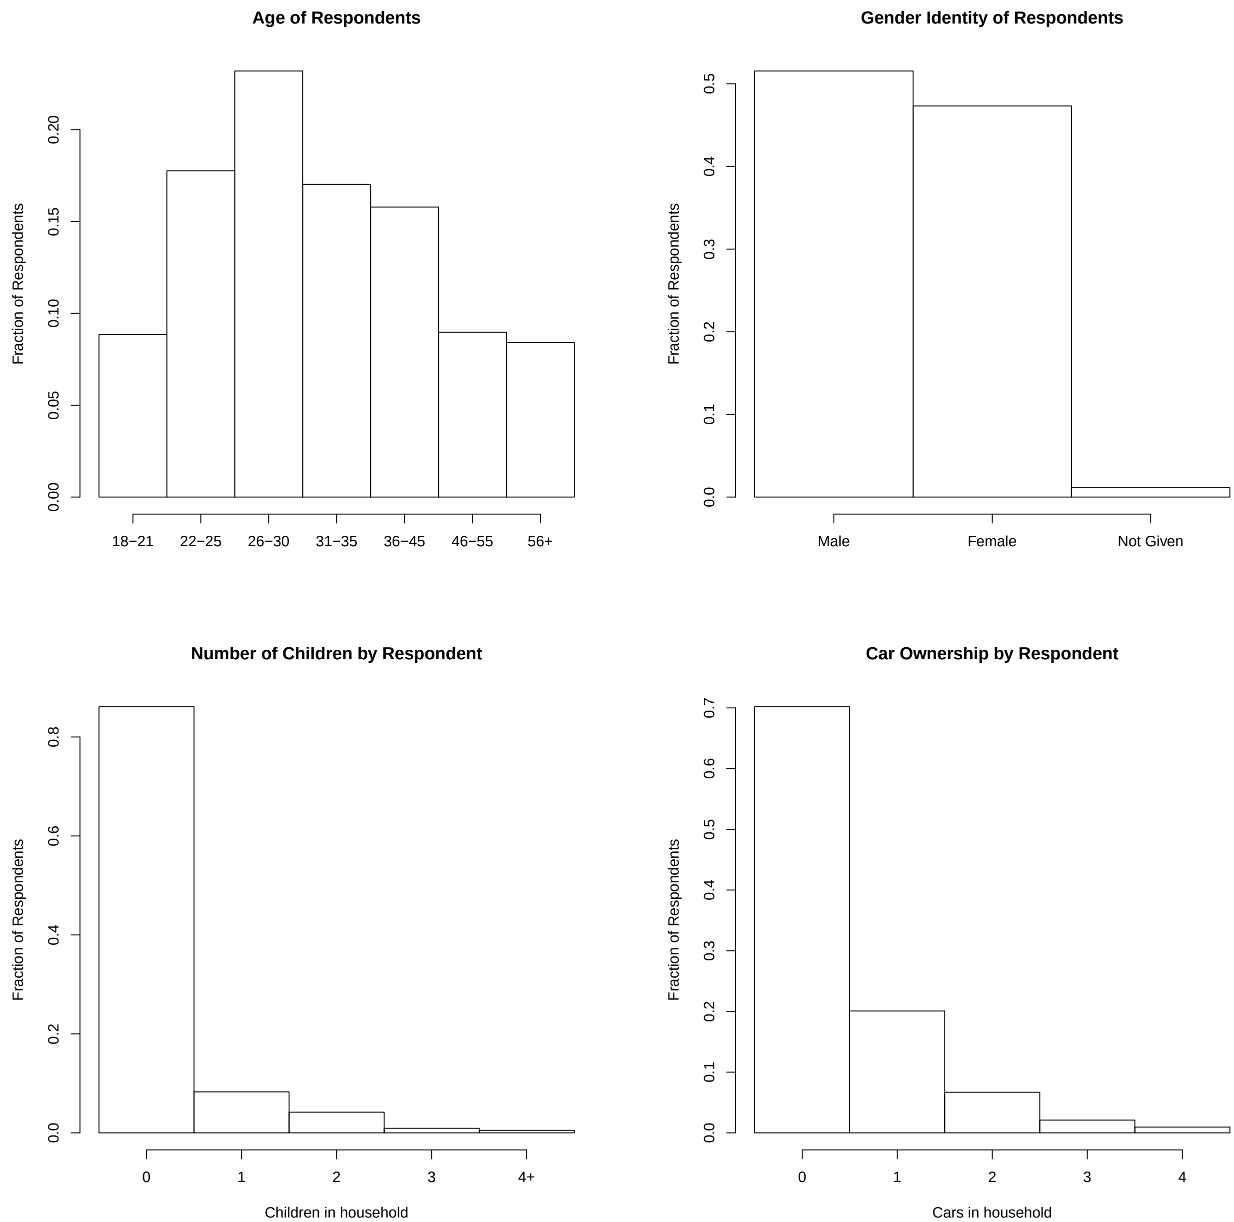

2.4.3 Experience with Alternative Fuel Vehicles

A principal goal of this work was to better understand the experience and attitudes of users towards new vehicle technology. The survey included surveys three questions about experience with alternative fuel vehicles, listed below:

1. Some Zipcars are *Hybrids* (e.g. Toyota Prius). Hybrids run on gasoline, but use batteries and an electric motor to reduce the amount of gasoline the car uses. Have you ever driven a hybrid?

2. Some Zipcars are *Plug-In hybrids* (e.g. Chevrolet Volt). Plug-In hybrids are like regular hybrids, but can be also recharged directly with electricity, to travel farther under electric power and further reduce the gasoline the cars use. Have you ever driven a Plug-In hybrid?

3. Some Zipcars are *Electric cars* (e.g. Nissan Leaf). Electric cars use no gasoline, being recharged 100% using electricity. Have you ever driven an electric car?

Each of these questions was presented in multiple choice format, with the following responses allowed (example responses for the hybrid vehicle type shown):

- Yes, I own (or previously owned) a hybrid
- Yes, I've driven a Zipcar hybrid
- Yes, I've driven a hybrid elsewhere
- No, I haven't driven a hybrid
- I'm not sure

These questions were designed to perform two functions: to gather information about user exposure to technology within and outside of carsharing, and to ensure that users had at least a minimal understanding of the definition of Hybrids, Plug-In Hybrids and Electric

Vehicles prior to being asked the discrete choice questions. Figure 2-8, Figure 2-10, and Figure 2-9 show the responses for experience with these three vehicle types and indicates that more than half of respondents (50.5%) have driven a Hybrid vehicle through Zipcar – far more than those who own or have owned a hybrid (3.3%) and those who have driven a hybrid elsewhere (13.8%). User experience levels with both PHEVs and EVs are far lower – fewer than 10% have driven either type of Plug-in Vehicle, consistent with the low penetration of these vehicle types in the marketplace and in the fleet of this carsharing provider.

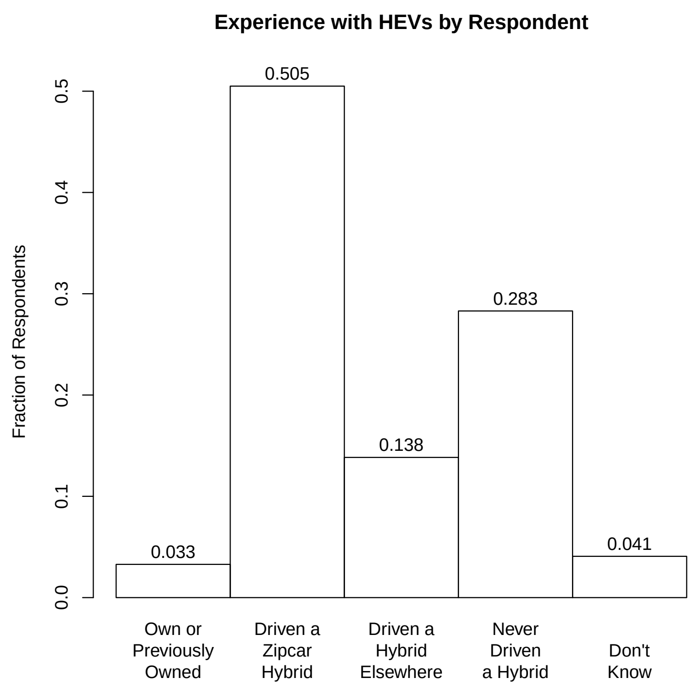

Figure 2-8: Respondent exposure to Hybrid Electric Vehicles.

As of the end of October, 2013 the carsharing provider reported membership of approximately 709,000 in the U.S. and Canada. While the size and composition of the fleet varies over time, a sample taken in the year preceding the launch of this survey suggested that approximately 892 vehicles, or approximately 10% of the fleet at the time, were hybrids.

If respondents are a representative sample of members, these results suggest that approximately 358,000 members of the carsharing service have been exposed to Hybrid vehicle

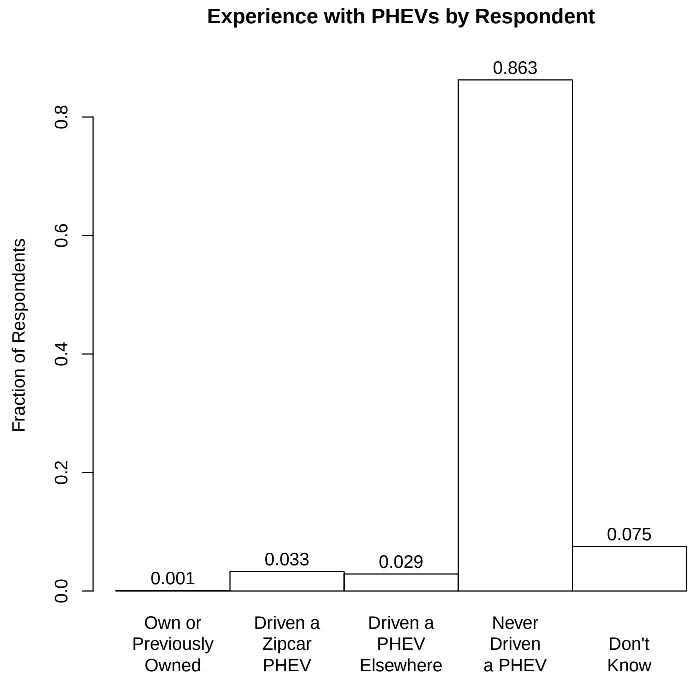

Figure 2-9: Respondent exposure to Plug-In Hybrid Electric Vehicles.

technology by their participation in carsharing through only 892 vehicles in the fleet. This represents a ratio of approximately 400 users exposed per vehicle in active service.

The large number of users driving hybrid vehicles could partially be explained by the pricing model: this carsharing service typically prices hybrid vehicles more cheaply than most other vehicles, so users looking for the cheapest vehicles will frequently drive hybrids. While users do not pay directly for gasoline for short trips. However, for longer trips users may prefer hybrids for their fuel economy.

As a reference for comparison, approximately 434,000 hybrid vehicles were sold in 2012 (EPA, 2013), and such vehicles are typically driven by only 1 or 2 members of a household.

2.4.4 User Recollection of Reservation Behavior

The survey instrument asked users three questions regarding their reservation behavior: how frequently the use Zipcar, how long their typical reservation lasts, and how far they believe

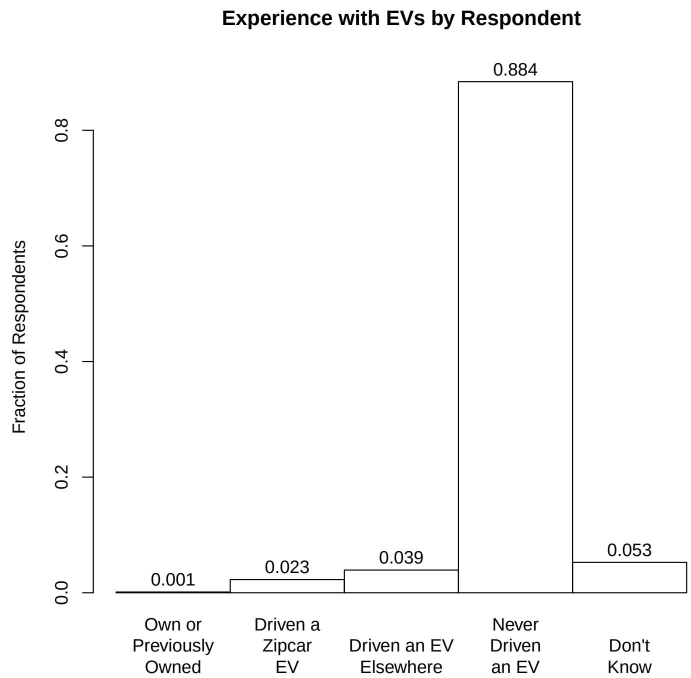

Figure 2-10: Respondent exposure to Electric Vehicles.

they drive in a typical reservation. As with questions about vehicle usage, these questions were designed to accomplish two tasks: to understand how respondent characteristics affect choice behavior and to prompt users to have a specific usage in mind when considering whether or not to reserve an electric vehicle with a limited range that might not be able to actually drive the distance the respondent believes that he or she drives. Data presented in these figures is based on approximately half the total respondent sample for which actual reservation behavior could be correlated.

In travel frequency, shown in Figure 2-11 we see that users recall their typical travel frequency fairly well. There is a large variance in the actual travel behavior, in part attributable to users who were not active for all of 2012 or who joined throughout that year.

In travel duration (time), shown in Figure 2-12 we see that users recall their typical travel duration well. A paired T-test of the actual median reservation length in hours and the reported reservation duration is not significant at the 10% level.

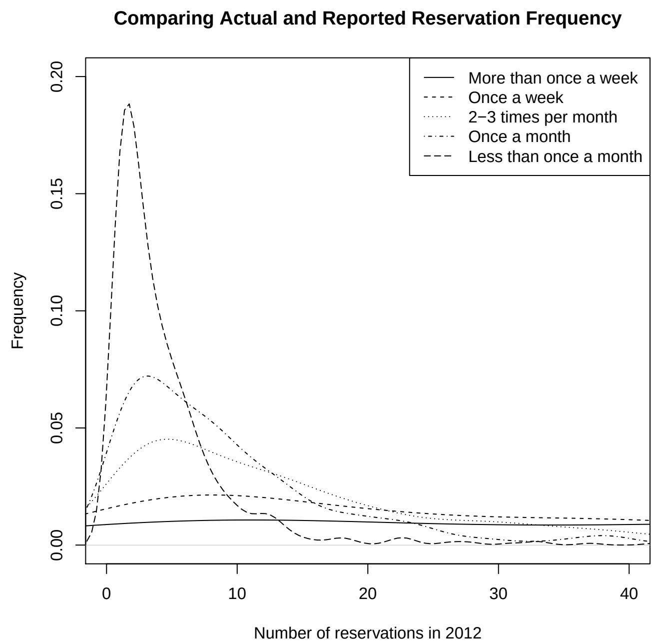

Figure 2-11: Distribution of user reservation behavior in 2012, plotted by reported reservation frequency.

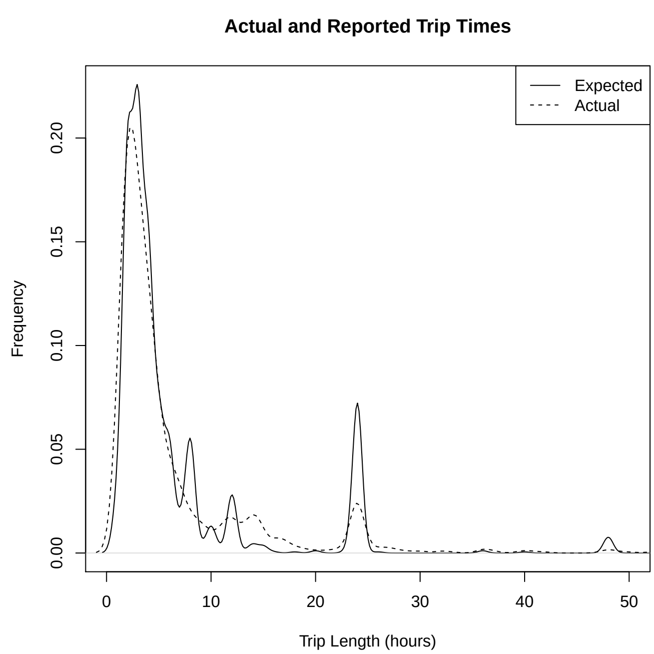

Figure 2-12: Distribution of user reported reservation duration in hours, compared to actual median reservation duration in 2012.

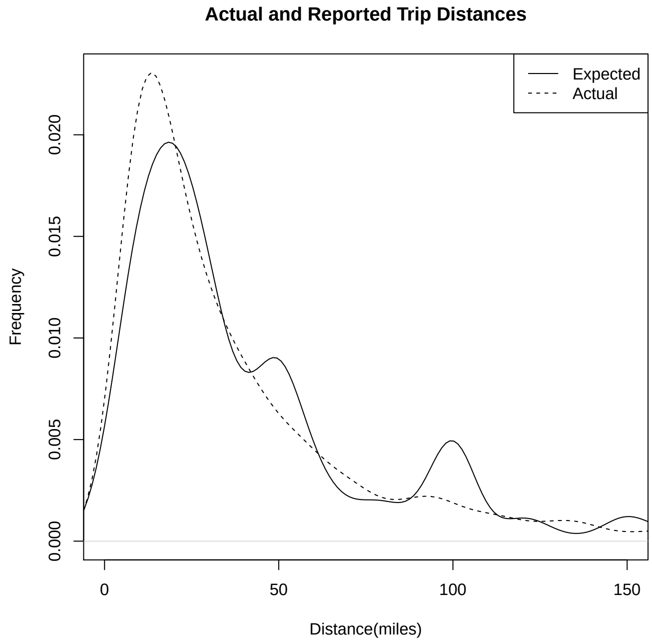

Figure 2-13: Distribution of user reported typical reservation travel distance in miles, compared to actual median reservation travel distance in 2012.

Respondents perform less well when attempting to recall the distance traveled during a typical reservation. Figure 2-13 compares actual median reservation distance from 2012 with reported typical reservation distance. A paired T-test of the actual median and the reported reservation distance shows a significant difference at the 5% level, which a difference in means of 3.1, or approximately 7% of the typical reservation travel distance.

The significant overestimation of travel distance is potentially problematic for the deployment of vehicles with a limited range, such as battery electric vehicles (BEVs) or the alternative fuel vehicles, as users may elect *not* to take such a vehicle even if it were capable of meeting the driving requirements of a specific trip.

2.4.5 Choice Experiment

As noted in Section 2.2, other researchers have used choice models of carsharing behavior. Barnes and Rutherford (2001) uses a binomial logit formulation to model probability of adoption of carsharing membership, while Le Vine et al. (2014b) uses a mixed logit formulation to identify the probability of choosing carsharing compared to other modes.

Unlike previous adoption and mode-choice formulations, the choice panels presented to respondents in this user survey asked them to pick the carsharing vehicle that would best meet their needs for a typical carsharing trip. In the survey instrument, respondents were presented which four discrete choice panels, each with four alternatives and a "none of the above" option, as shown in Figure 2-14.

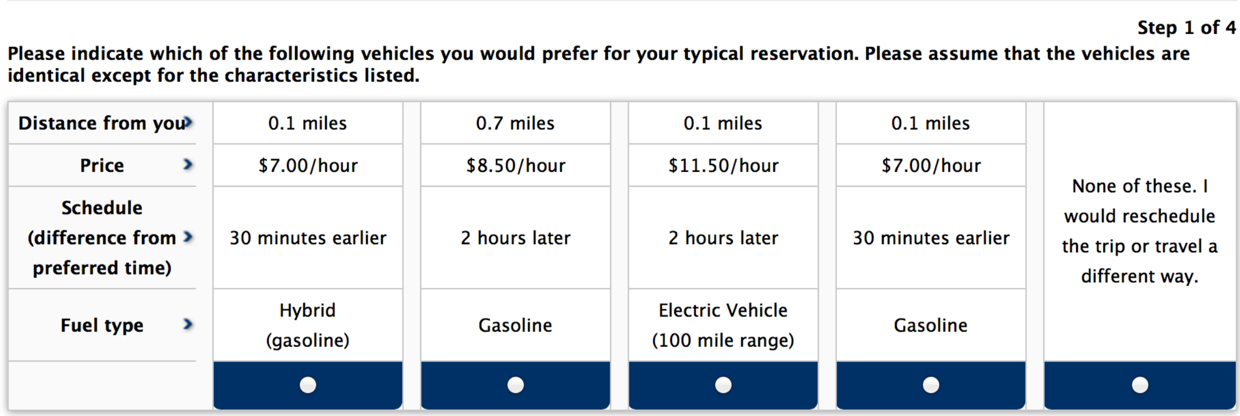

Figure 2-14: Sample discrete choice panel from survey instrument.

Based on prior work and unpublished survey data, the attributes of distance, hourly price, and schedule were chosen to be included in the discrete choice experiment. The survey also included fuel type as an attribute in order to inform our understanding of carsharing user attitudes towards alternative fuel vehicles. The attributes and levels used in the discrete choice experiment are shown in Table 2.2.

In order to avoid potential loss of valuable data due to presentation of dominated alternatives, a random experimental design was rejected. The final design was a fractional factorial design, seeded with decreasing utility with price, distance and deviation from preferred schedule. Given that no existing research permits an *a priori* hypothesis about directionality of the utility of fuel type to carsharing, users a utility seed was not used for fuel type. Additionally, the experimental design included two replicates of the "gasoline" fuel type so that conventional gasoline vehicles would makeup approximately 50% of alternatives in order to reduce overrepresentation of alternative fuel vehicles.

Table 2.2: Attributes and Levels of the discrete choice experiment.

Variable	**Levels**	**Values**
Distance from you	7	0.1 miles
		0.4 miles
		0.7 miles
		1.0 miles
		1.3 miles
		1.6 miles
		1.9 miles
Price	6	\$7.00 / hour
		\$8.50 / hour
		\$10 / hour
		\$11.50 / hour
		\$13 / hour
		\$14.50 / hour
Schedule	7	2 hours earlier
		1 hour earlier
		30 minutes earlier
		Exactly when I need it
		30 minutes later
		1 hour later
		2 hours later
Fuel Type	4	Gasoline;
		Hybrid;
		Plug-In Hybrid (30 mile EV range);
		Electric Vehicle (100 mile range)

Each user was presented with four choice experiments, each with four alternatives and a "none of the above" alternative. The fractional factorial experimental consisted of fifty individual choice surveys, for a total of 200 distinct choice panels.

2.5 Model Results

This section presents the results of modeling efforts. The section begins with several basic formulation of the choice model and provides calculations for willingness to pay for service attributes. The chapter then presents the results of the topic models used to generate indicators from open-ended survey responses, and finally presents the results from the Integrated Choice and Latent Variable Model.

2.5.1 Basic Multinomial Logit Models

Three of the four attributes in the survey presented to respondents (price, access distance, and schedule) were presented as discrete samples of continuous variables. One of the first tasks was to investigate whether the response to these attributes was linear or took on some other form. Figure 2.3 shows the results from a plot of coefficients on each variable as a series of discrete responses.

The results indicate linear response to both price and distance for the aggregate respondents. However, schedule appears non-linear, with particularly negative reactions to a two-hour delay in a reservation. As a result, schedule was left as a discrete variable in subsequent models. Table 2.4 shows the results from a based multinomial logit (MNL) model based on the alternative attributes presented to users.

2.5.2 Price and Cost Model Results

The results shown in Table 2.4 compare two basic model specifications with only basic service and vehicle attributes and no demographic covariates. In the first formulation, the vehicle price, in $/hour is modeled. In the second formulation, the price and reported reservation length are used to generate an approximate cost for the total reservation.

For both price and cost models, the sign of all coefficients is as we would expect: the coefficient on access distance, price and deviations from the desired schedule are all negative and statistically significant with very small standard errors. The only attribute with a

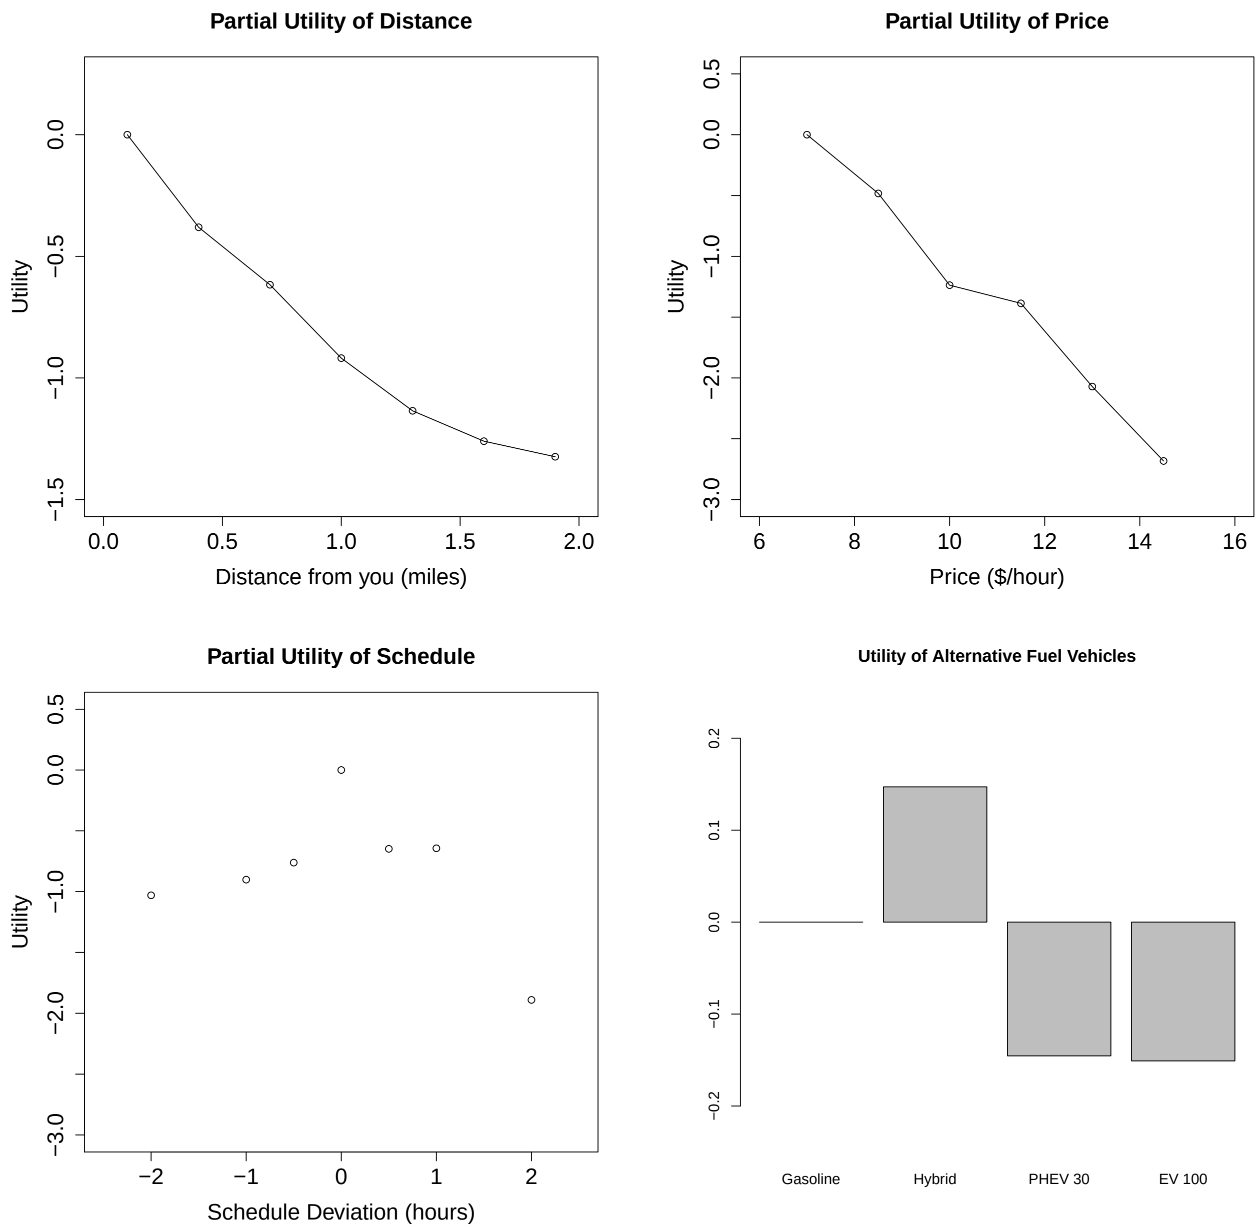

Table 2.3: Graphical display of coefficients for simple multinomial logit model of carsharing service attributes.

Table 2.4: Model 1: Basic MNL model results using Price and Cost attributes.

Attribute	Price Model			Cost Model		
	Coeff.	Std Error	t-test	Coeff.	Std Error	t-test
Price ($/hour)	-0.349	0.00625	-55.79	-	-	-
Reservation Cost* ($)	-	-	-	-0.051	0.00114	-44.86
Access Distance (mi)	-0.734	0.0205	-35.75	-0.569	0.0185	-30.74
Schedule 30min Early	-0.753	0.0323	-23.30	-0.736	0.0314	-23.44
Schedule 30min Late	-0.624	0.0386	-16.16	-0.579	0.0378	-15.32
Schedule 1hr Early	-0.872	0.0415	-20.99	-0.847	0.0409	-20.71
Schedule 1hr Late	-0.581	0.0396	-14.69	-0.469	0.0383	-12.23
Schedule 2hrs Early	-0.963	0.0424	-22.72	-0.819	0.0407	-20.13
Schedule 2hrs Late	-1.81	.0523	-34.66	-1.06	0.0404	-26.21
Hybrid	0.141	0.0289	4.88	0.160	0.0280	5.73
Plug-In Hybrid	-0.172	0.0284	-6.06	-0.140	0.0275	-5.08
Electric Vehicle	-0.159	0.0295	-5.40	-0.128	0.0288	-4.44

*Reservation Cost estimated by multiplying hourly rate by reported reservation length adjusted for daily maximum.

	Price Model	Cost Model
Sample size:	13002	13002
Init log-likelihood:	-18024.599	-18024.599
Final log-likelihood:	-14854.110	-15702.887
Likelihood ratio test for the init. model:	6340.979	4643.425
Rho bar for the init. model:	0.175	0.128

positive coefficient is the dummy variable for a hybrid vehicle, suggesting that users would prefer a hybrid vehicle to a gasoline vehicle. However, the coefficient on both Plug-In Hybrid and Electric Vehicles are negative. All coefficients on vehicle type are smaller in magnitude than those on other service attributes, suggesting that vehicle type is less important than even a 30 minute deviation in reservation time.

The carsharing service provider offers discounted daily reservation rates which, for reservations exceeding 8 hours but less than 24 hours, charges approximately 8 times the hourly rate. Figure 2-15 shows the transformation used to develop an estimate of the user's total reservation cost from the hourly price.

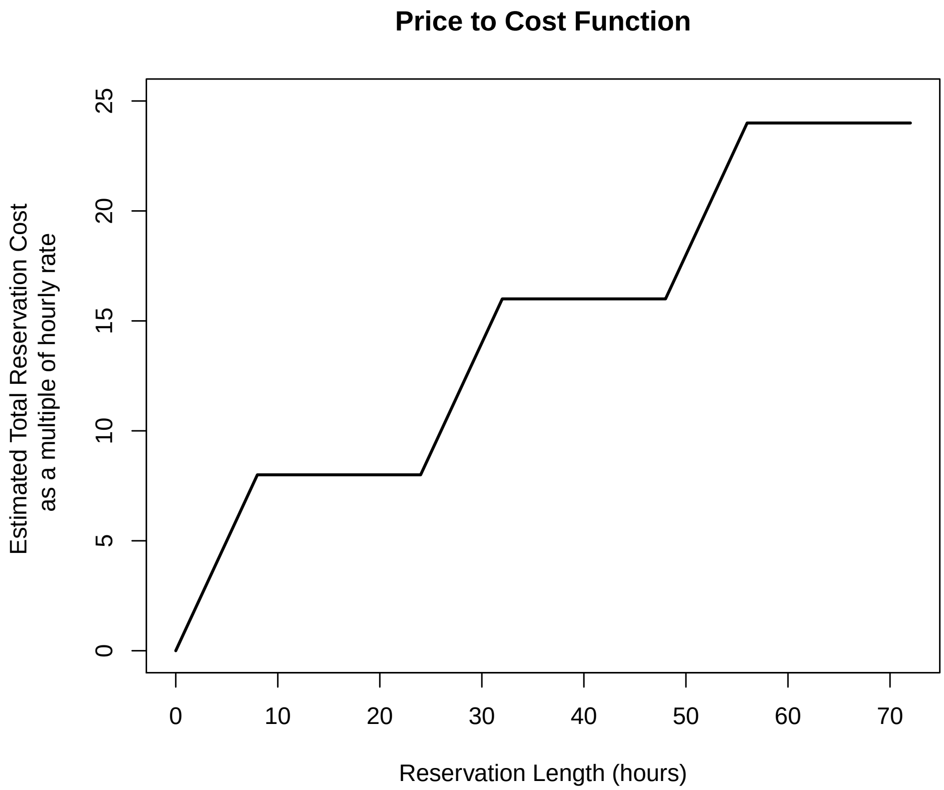

Figure 2-15: Sample discrete choice panel from survey instrument.

As we can see based on likelihood values and $\bar{\rho}^2$ values, the fit of this model is notably worse than the hourly price specification. This result suggests that users shop vehicles primarily based on the hourly rate presented to them and not based on the total estimated cost that they would face.

2.5.3 Willingness to Pay

Using the coefficients from a model, we can calculate willingness to pay for certain attributes. For instance, if β_{price} is in units of $\frac{utils}{\$/hr}$ and $\beta_{distance}$ is in units of $\frac{utils}{mile}$ then we can calculate the value of tradeoffs between these two attributes. For the price-based multinomial logit model here, we calculate willingness to pay for attributes in terms of price ($/hour), shown in Table 2.5, or in terms of cost ($ per reservation) in Table 2.6. In relative terms, the magnitude of the coefficients remain similar, although magnitude of costs becomes larger in the $ domain than in the $/hour domain.

Table 2.5: WTP Values in $/hour of reservation price.

Attribute	Coefficients	Value
Access distance	$\frac{-.734}{-.349}$	$2.10/hr per mile
Schedule	$\frac{-.581}{-.349}$	$1.66/hr for 1hr late
	$\frac{-.872}{-.349}$	$2.50/hr for 1hr early
	$\frac{-1.81}{-.349}$	$5.19/hr for 2hr late
Vehicle Type	$\frac{.141}{-.349}$	$0.40/hr for a Hybrid
	$\frac{-.172}{-.349}$	$0.49/hr for a Plug-In Hybrid
	$\frac{-.159}{-.349}$	$0.46/hr for an Electric Vehicle

Table 2.6: WTP Values in $ per total reservation cost.

Attribute	Coefficients	Value
Access distance	$\frac{-.569}{-.051}$	$11.16 per mile
Schedule	$\frac{-.469}{-.051}$	$9.19 for 1hr late
	$\frac{-.847}{-.051}$	$16.60 for 1hr early
	$\frac{-1.06}{-.051}$	$20.78 for 2hr late
Vehicle Type	$\frac{.160}{-.051}$	$3.14 for a Hybrid
	$\frac{-.140}{-.051}$	$2.75 for a Plug-In Hybrid
	$\frac{-.128}{-.051}$	$2.50 for an Electric Vehicle

2.5.4 Mixed Logit Model

This work tested a variety of Mixed Logit specifications. The results from a specification with normally distributed coefficients for price, distance and vehicle types is shown in Table 2.7. Schedule coefficients remained fixed.

Table 2.7: Model 1: Mixed logit model results.

Attribute	**Coefficient**	**Std Error**	**t-test**
Price ($/hour)	-0.587	0.0150	-39.19
Price s.d.	-0.420	0.0149	-28.23
Access Distance (miles)	-1.18	0.0451	-26.17
Access Distance s.d.	1.55	0.0588	26.40
30min Early	-0.989	0.0416	-23.74
30min Late	-0.867	0.0501	-17.29
1hr Early	-1.14	0.0530	-21.53
1hr Late	-1.05	0.0566	-18.61
2hrs Early	-1.57	0.0607	-25.80
2hrs Late	-2.68	0.0831	-32.27
Hybrid	0.125	0.0413	3.03
Hybrid s.d.	0.634	0.105	6.02
Plug-In Hybrid	-0.341	0.0471	-7.23
Plug-In Hybrid s.d.	1.02	0.0831	12.23
Electric Vehicle	-0.414	0.0520	-7.96
Electric Vehicle s.d.	1.22	0.0823	14.85

Sample size:	13002
Init log-likelihood:	-15689.874
Final log-likelihood:	-13907.453
Likelihood ratio test for the init. model:	3564.841
Rho bar for the init. model:	0.114
Number of Halton Draws	500

The mixed specification indicates that there is large variation in some coefficients, particularly in access distance to the vehicle and in attitudes toward alternative fuel vehicle types. This result is particularly revealing, since the fixed coefficients in basic model specifications suggested that vehicle powertrain types had relatively small practical significance. These themes will be explored further in subsequent model specifications.

2.5.5 Interaction Model

Tables 2.8, 2.9 and 2.10 show the results from a model that includes interaction effects with distance and other travel modes, as well as an interaction between schedule and advance planning, and a piecewise linear specification of the utility of distance interacted with fuel type.

2.5.6 Schedule Flexibility and Planning

Figure 2-16 shows the disutility of schedule shifts interacted with the reported time in advance the respondent plans travel from Model 4. As expected, 2 hour delays are universally disliked by respondents, but this interaction highlights several additional details.

Respondents who plan far in advance appear to be generally more flexible about shifts in schedule – particularly those who plan more than a month in advance. One possible explanation is that users who plan travel this far in advance have the capability to shift their trip purpose (e.g. appointment) to coincide with the vehicle availability, whereas users who reserve only a few days or hours in advance may not have this option.

Additionally, the interaction highlights an asymmetry for users who book less than an hour in advance and for users who plan several hours in advance. Both groups show a greater disutility for brief schedule shifts earlier than schedule shifts later. A possible explanation for users who plan less than an hour in advance is that such shifts would result in an abbreviated reservation, not just a change in plans.

2.5.7 Utility of Vehicle Type by Travel Distance

One major objective of this work is to understand how users of carsharing react to different fuel types. Using the specification in Table 2.10, we can investigate the utility of vehicles

		Vehicle Availability						
		2 hours early	1 hour early	30 min early	Exactly when I prefer	30 min late	1 hour late	2 hours late
Typically Reserve...	...less than an hour in advance	-0.85	-0.89	-0.68	0	-0.44	-0.42	-1.33
	...several hours in advance	-0.88	-0.88	-0.77	0	-0.62	-0.58	-1.72
	...more than a day in advance	-1.11	-0.99	-0.78	0	-0.70	-0.62	-1.90
	...more than a week in advance	-0.78	-0.53	-0.79	0	-0.58	-0.60	-2.01
	...more than a month in advance	-0.64	0.01	-0.73	0	-0.49	-0.17	-1.60

Figure 2-16: Utility of different schedule shifts interacted with advanced planning.

with different fuel types to Zipcar members who travel different distances during a typical reservation. The results of this analysis are shown in Figure 2-17, derived from the coefficients of Model 4 described in the previous section. In this model vehicle powertrain type (gasoline, hybrid, plug-in hybrid and electric vehicle) was interacted with the travel distance in miles, as reported by the respondent. A piecewise specification was used, with breaks at 20, 40, 60, 80 and 100 miles.

In this analysis, the utility of a gasoline vehicle has been normalized to zero and the utilities of other fuel types are relative to a gasoline vehicle. This analysis shows a positive utility for hybrid vehicles up to approximately 60 miles, with approximately zero utility for larger reservations.

The coefficient on Plug-in hybrids and Electric Vehicles, however, never becomes strongly positive, no matter how short the travel distance. Utility becomes progressively more neg-

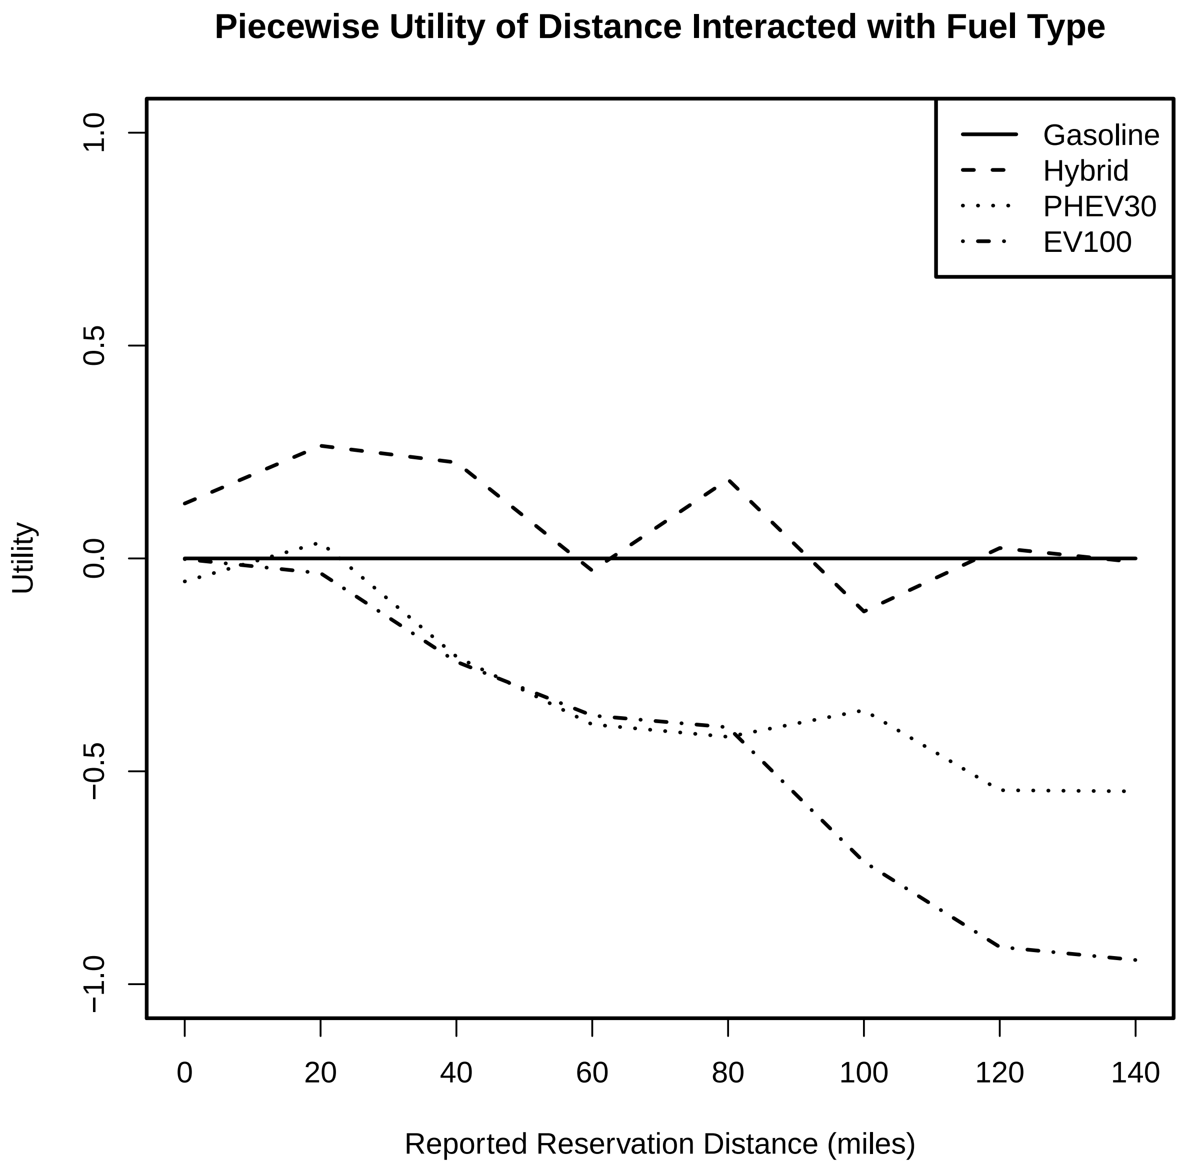

Figure 2-17: Utility of different fuel types by reported typical reservation distance.

ative with larger travel distances, and utility of PHEVs and EVs remains approximately consistent from very short travel distances all the way until reported travel distances of approximately 80 miles. Beyond this point, the utility of electric vehicles falls off sharply, most likely attributable to users becoming concerned about the range limitations of the vehicles.

Table 2.8: Model 4: Price and Distance Interactions.

Attribute	**Coefficient**	**Std Error**	**t-test**
Price & Gender: Male	-0.327	0.00741	-44.11
Price & Gender: Female	-0.374	0.00795	-46.99
Access Distance	-0.686	0.058	-11.81
Access Distance & Kids	0.174	0.0718	2.42
Access Distance & Train	0.081	0.0383	2.12
Access Distance & Bus	0.0716	0.037	1.93
Access Distance & Subway	-0.0472	0.0465	-1.02
Access Distance & Atlanta	-0.162	0.167	-0.97
Access Distance & Austin	0.0538	0.221	0.24
Access Distance & Baltimore	0.00871	0.136	0.06
Access Distance & Boston	-0.186	0.06	-3.1
Access Distance & Chicago	-0.345	0.0729	-4.74
Access Distance & Denver	0.232	0.31	0.75
Access Distance & LosAngeles	0.0752	0.116	0.65
Access Distance & Miami	-1.04	0.451	-2.3
Access Distance & Milwaukee	-0.414	0.321	-1.29
Access Distance & Minneapolis	-0.298	0.298	-1
Access Distance & Philadelphia	-0.0733	0.105	-0.7
Access Distance & Pittsburgh	0.0268	0.162	0.17
Access Distance & Portland	-0.233	0.117	-1.99
Access Distance & Providence	-0.0686	0.177	-0.39
Access Distance & San Diego	-0.1	0.334	-0.3
Access Distance & San Francisco	-0.0669	0.0677	-0.99
Access Distance & Seattle	-0.249	0.104	-2.39
Access Distance & Toronto	0.0435	0.0712	0.61
Access Distance & Universities	0.003	0.0742	0.04
Access Distance & Vancouver	0.21	0.119	1.76
Access Distance & Washington, DC	-0.336	0.0709	-4.74

Sample size:	13002
Init log-likelihood:	-18024.599
Final log-likelihood:	-14735.635
Likelihood ratio test for the init. model:	6577.929
Rho bar for the init. model:	0.178

Table 2.9: Model 4: Schedule and planning interactions.

Attribute	**Coefficient**	**Std Error**	**t-test**
Schedule: 2hrs Early & Plan_under1hr	-0.852	0.122	-6.96
Schedule: 2hrs Early & Plan_hours	-0.882	0.0725	-12.17
Schedule: 2hrs Early & Plan_day	-1.11	0.0631	-17.66
Schedule: 2hrs Early & Plan_week	-0.778	0.119	-6.51
Schedule: 2hrs Early & Plan_month	-0.635	0.412	-1.54
Schedule: 1hr Early & Plan_under1hr	-0.89	0.133	-6.69
Schedule: 1hr Early & Plan_hours	-0.875	0.0733	-11.94
Schedule: 1hr Early & Plan_day	-0.992	0.0624	-15.89
Schedule: 1hr Early & Plan_week	-0.53	0.115	-4.61
Schedule: 1hr Early & Plan_month	0.0112	0.409	0.03
Schedule: 30min Early & Plan_under1hr	-0.675	0.1	-6.73
Schedule: 30min Early & Plan_hours	-0.766	0.0573	-13.38
Schedule: 30min Early & Plan_day	-0.777	0.0473	-16.43
Schedule: 30min Early & Plan_week	-0.791	0.101	-7.85
Schedule: 30min Early & Plan_month	-0.725	0.374	-1.94
Schedule: 30min Late & Plan_under1hr	-0.435	0.113	-3.87
Schedule: 30min Late & Plan_hours	-0.619	0.0684	-9.04
Schedule: 30min Late & Plan_day	-0.695	0.057	-12.2
Schedule: 30min Late & Plan_week	-0.581	0.115	-5.03
Schedule: 30min Late & Plan_month	-0.487	0.433	-1.13
Schedule: 1hr Late & Plan_under1hr	-0.422	0.118	-3.57
Schedule: 1hr Late & Plan_hours	-0.577	0.0679	-8.49
Schedule: 1hr Late & Plan_day	-0.621	0.0569	-10.92
Schedule: 1hr Late & Plan_week	-0.604	0.12	-5.05
Schedule: 1hr Late & Plan_month	-0.168	0.415	-0.41
Schedule: 2hrs Late & Plan_under1hr	-1.33	0.119	-11.21
Schedule: 2hrs Late & Plan_hours	-1.72	0.0754	-22.84
Schedule: 2hrs Late & Plan_day	-1.9	0.0663	-28.69
Schedule: 2hrs Late & Plan_week	-2.01	0.116	-17.43
Schedule: 2hrs Late & Plan_month	-1.6	0.518	-3.1

Table 2.10: Model 4: Vehicle Powertrain type and travel distance piecewise interactions.

Attribute	**Coefficient**	**Std Error**	**t-test**
Hybrid	0.129	0.118	1.1
Plug-In Hybrid	-0.054	0.113	-0.48
Electric Vehicle	0.0878	0.115	0.76
EV & Travel Distance $\leq$20 mi	-0.00165	0.00746	-0.22
HEV & Travel Distance $\leq$20 mi	0.00678	0.00757	0.9
PHEV & Travel Distance $\leq$20 mi	0.00458	0.00723	0.63
BEV & Travel Distance 20-40 mi	-0.0104	0.00607	-1.71
HEV & Travel Distance 20-40 mi	-0.00196	0.006	-0.33
PHEV & Travel Distance 20-40 mi	-0.0134	0.00581	-2.31
BEV & Travel Distance 40-60 mi	-0.00632	0.00845	-0.75
HEV & Travel Distance 40-60 mi	-0.0127	0.00842	-1.5
PHEV & Travel Distance 40-60 mi	-0.00798	0.00815	-0.98
BEV & Travel Distance 60-80 mi	-0.00138	0.0104	-0.13
HEV & Travel Distance 60-80 mi	0.0107	0.0101	1.07
PHEV & Travel Distance 60-80 mi	-0.00145	0.00999	-0.14
BEV & Travel Distance 80-100 mi	-0.0158	0.0101	-1.56
HEV & Travel Distance 80-100 mi	-0.0155	0.00948	-1.64
PHEV & Travel Distance 80-100 mi	0.00312	0.00955	0.33
BEV & Travel Distance 100-120 mi	-0.01	0.00957	-1.05
HEV & Travel Distance 100-120 mi	0.00745	0.0083	0.9
PHEV & Travel Distance 100-120 mi	-0.00939	0.00837	-1.12
BEV & Travel Distance 120+ mi	-0.00153	0.00153	-1
HEV & Travel Distance 120+ mi	-0.00165	0.00132	-1.25
PHEV & Travel Distance 120+ mi	-0.000132	0.00123	-0.11

2.6 Vehicle Ratings

The second-to-last section of the survey asked users to rate vehicles offered in the service by clicking on star icons from one to five stars in half-star increments. These are interpreted as a 1-10 rating scale for 21 different vehicles included in the survey. To allow for variations in overall levels of satisfaction, these "raw" star ratings (S_{in}) were converted into relative star ratings, taken as the difference for a specific vehicle relative to the mean of all star ratings a particular user assigned to all vehicles. Thus for vehicle i and person j the relative rating R_{ij} is defined as:

$$R_{ij} = S_{ij} - \frac{\sum_{i=1}^{n} S_i}{n}$$

Figure 2-18 shows the distributions of relative star ratings R_i for a variety of conventional and advanced powertrain vehicles from all users. The vertical dashed line in each graph represents the mean value of R_i for the vehicle. A complete list of star ratings and their distributions in shown in Appendix B.

A number of effects can be observed by investigating these relative star ratings. Most conventional vehicles had R_i values centered near zero, with customers having a relatively neutral attitude towards them. Since the distribution of R_i depends on the number of responses, narrower distributions tend to reflect vehicles which were rated more frequently. Most conventional, non-luxury vehicles, such as the Honda Civic shown in Figure 2-18a, are distributed with little skew around a mean near zero. As would be expected, most users rated luxury vehicles higher than conventional vehicles, as evidenced by Audi A3 R_i distribution skewed to the right and mean value above zero as shown in Figure 2-18b. Customer star ratings for hybrids, such as the Toyota Prius shown in Figure 2-18c were generally similar to conventional vehicles, suggesting that customers either have few reservations about them, or they may not know that they are hybrids. PEVs, on the other hand, such as the Honda Fit EV shown in Figure 2-18d, were generally rated lower than conventional vehicles, suggesting that customers may have concerns about using them even for short durations.

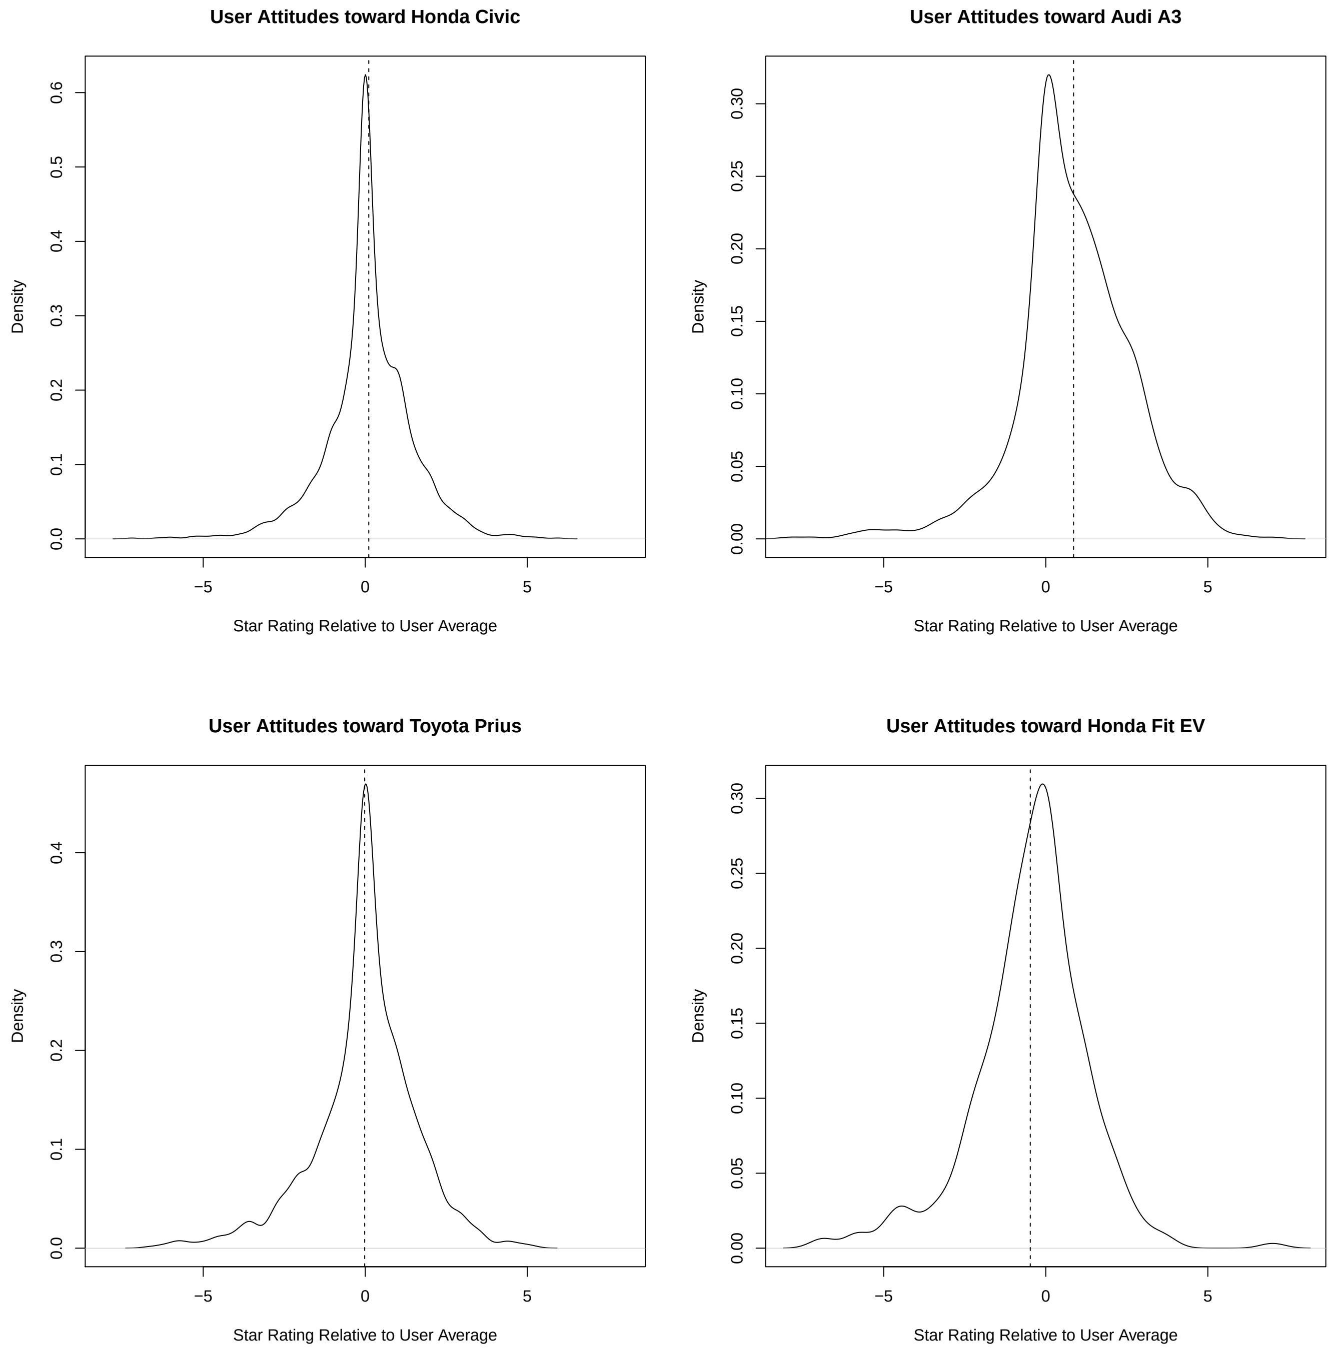

Figure 2-18: Relative star ratings R_i for a variety of conventional and advanced vehicles.

2.7 Open-Ended Comments and Topic Models

The last two questions of the survey instrument were open-ended questions that form the basis for the topic models presented here. These comments were intended to allow survey respondents to communicate a wide variety of needs or concerns to the service provider, and allowed for the observation of a wide variety of respondent sentiments.

In contrast to previous surveys in which users were asked to rank a small number of vehicle types that they would like to see more, users were allowed to enter any text they wished in response to the question "What kinds of cars would you like to see more of in our fleet, and why?" Of the 4133 respondents, 2159 or approximately 52% provided a valid response to the question. Figure 2-19 shows a word cloud of the most common responses to the open-ended question. While the words "hybrid," "electric" and "Prius" all appear prominently, a number of themes seem to be expressed: specific brands, cars with certain powertrain technology, cars with certain utilitarian attributes and a number of service-related comments.

The last question of the survey was intended as a "Catch all" question to allow respondents to express any specific concerns, comments or requests. It asked respondents "Please let us know if you have any other comments or suggestions about our service." Of the 4133 respondents, 1229 or approximately 30% provided a valid response to the questions. Comment length ranged from a single word to several paragraphs of text. Figure 2-20 shows a word cloud of common words used in response to open-ended question about the service. Comments were quite diverse, and fewer words appeared frequently. Customers expressed a wide variety of sentiments about the service in their area, concerns about fees or convenience of the service. Some customers also used general expressions of happiness or dissatisfaction about the service provider – the words "love" and "hate" both appeared with relative frequency.

While more respondents answered the question "What kinds of cars would you like to see more of in our fleet, and why?" the responses were often brief and focused largely on the vehicles themselves. After stemming and the removal of common words these comments

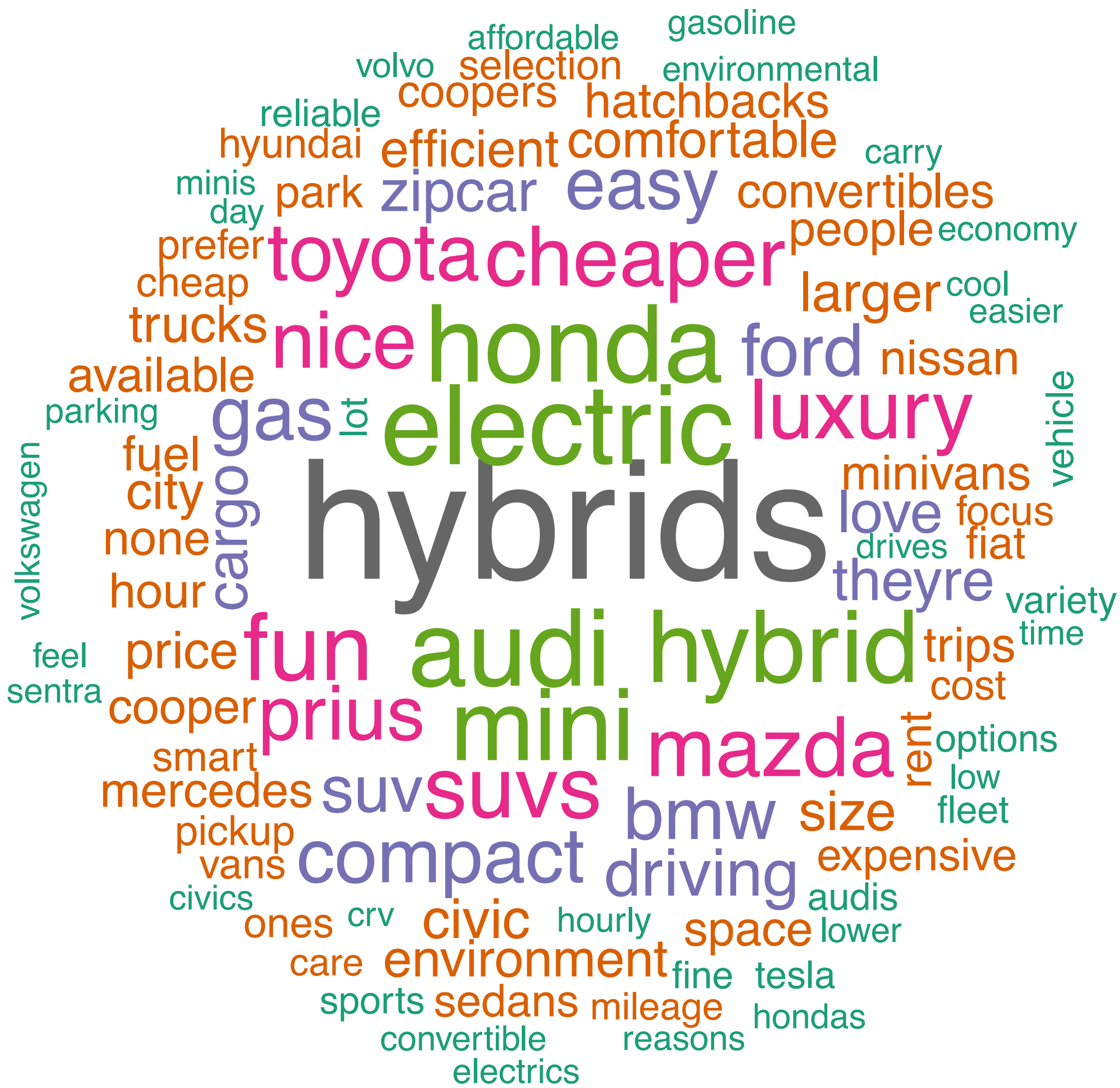

Figure 2-19: Word cloud of all vehicle-related comments.

formed a 698-word dictionary used in generating "Vehicle Comments" topic models.

In contrast, while fewer respondents choose to answer the question "Please let us know if you have any other comments or suggestions about our service" the responses were generally longer and more diverse in their content, with a 1112-word dictionary used for modeling after stemming.

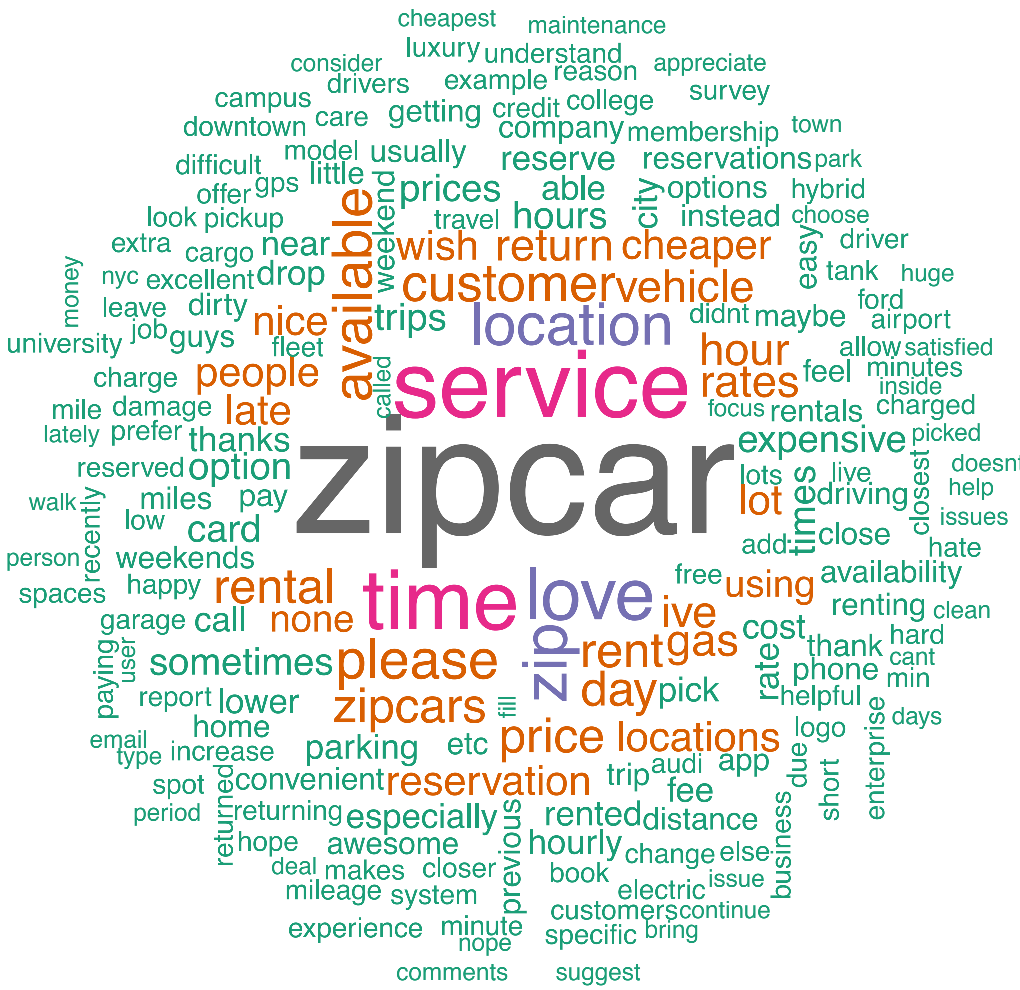

Figure 2-20: Word cloud of all service-related comments.

These documents were analyzed indepently using the package STM package in R. (Roberts et al., 2014b) Multiple topic models were analyzed with and without covariates. In all cases the words were stemmed using common words in the English language. Extremely common words from this specific set of documents such as “car” and “Zipcar” were also removed before modeling topics.

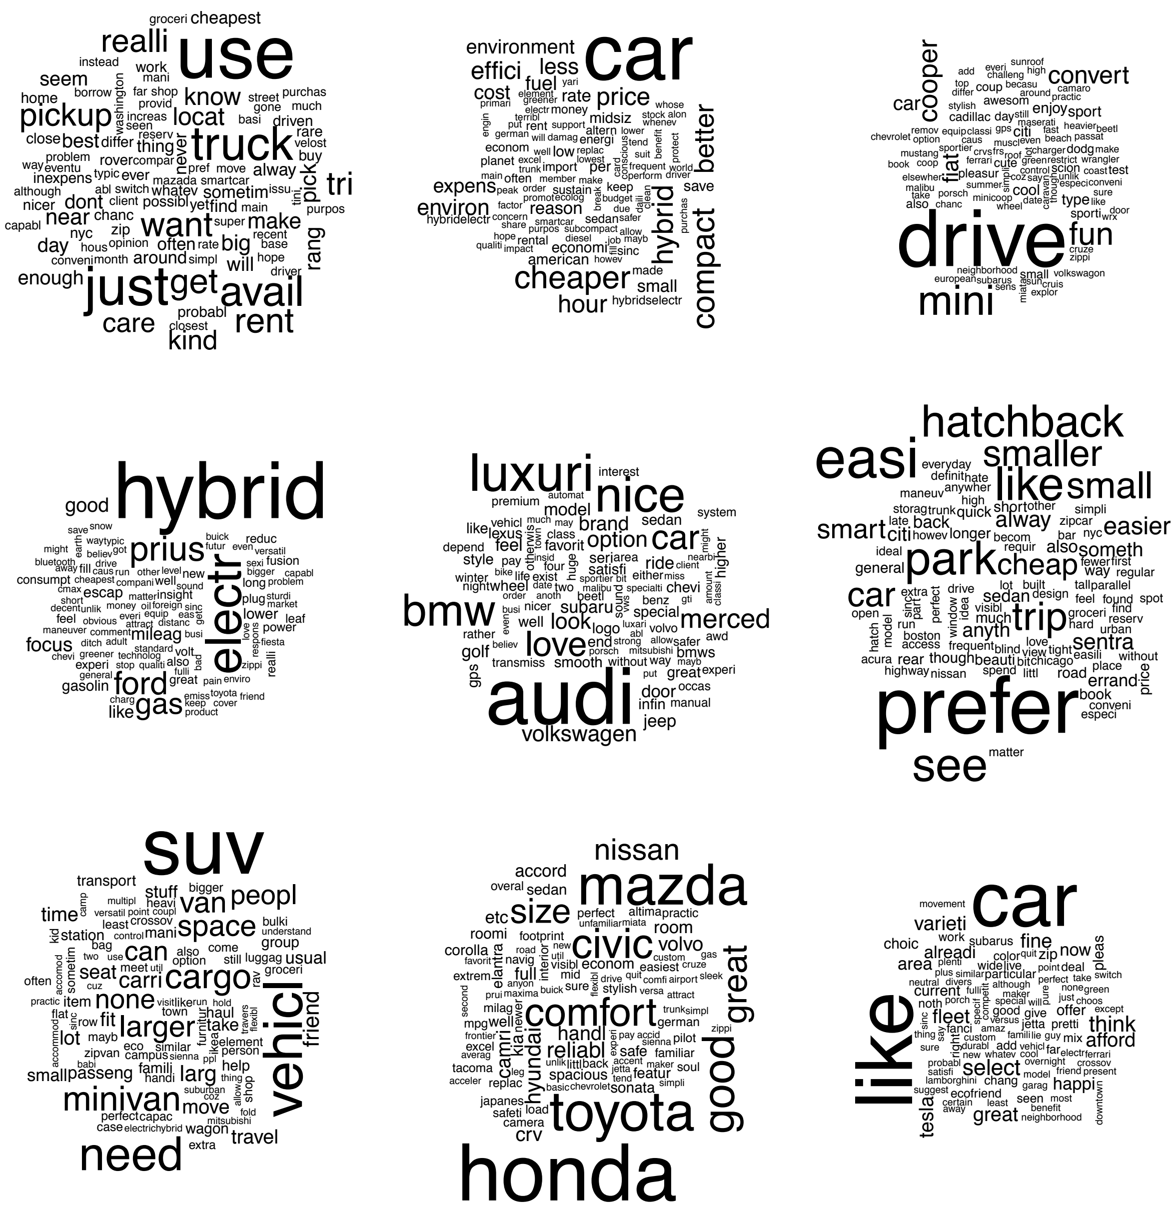

Figure 2-21: Word cloud of individual topics modeled for vehicle-related comments.

In the final models chosen, which are presented in the following section, models with nine topics, and without covariates were chosen to be incorporated in the Intergrated Choice and Latent Variable model presented later. The topic models chosen are described further in the following section.

Figure 2-22: Word cloud of individual topics modeled for service-related comments.

2.7.1 Exploring Topic Models

There are two ways to evaluate topics that have been fit with topic models. The first of these is to explore the words that have been associated with the topics. The STM package offers a number of ways to do this, described in Roberts et al. (2014a), that include inspecting common words, words that are exclusive to a topic, or both.

Figures 2-21 presents individual word clouds of the full distribution of words associated with vehicle comment topics 1-9, starting without Vehicle Topic 1 at the top left and Vehicle Topic 9 at the bottom right. Similarly, Figure 2-22 presents the individual word distributions associated with service topic models 1-9, with Service Topic 1 at the top left and Service Topic 9 at the bottom right. Tables 2.11 and 2.12 show four ways of subsetting complete topic-word distributions to present topical content in a more manageable fashion.

Topic 1 Top Words	
Highest Prob	one, use, just, zipcar, truck, avail, rent
FREX	truck, locat, one, care, kind, never, tri
Lift	hous, buy, never, mazada, pref, close, closest
Score	one, truck, use, just, pickup, avail, zipcar
Topic 2 Top Words	
Highest Prob	car, cheaper, compact, hybrid, better, hour, price
FREX	hour, cheaper, effici, fuel, per, better, environ
Lift	energi, hybridelectr, daili, world, per, damag, fuel
Score	cheaper, car, hour, effici, hybrid, fuel, environ
Topic 3 Top Words	
Highest Prob	drive, mini, fun, cooper, convert, fiat, car
FREX	fun, convert, cooper, mini, fiat, dodg, sport
Lift	frs, challeng, minicoop, convert, fun, cooper, mustang
Score	drive, mini, fun, cooper, convert, fiat, sport
Topic 4 Top Words	
Highest Prob	hybrid, electr, prius, ford, gas, focus, good
FREX	ford, focus, prius, plug, consumpt, gasolin, electr
Lift	comment, consumpt, market, focus, ford, emiss, bluetooth
Score	hybrid, electr, ford, prius, gas, focus, mileag
Topic 5 Top Words	
Highest Prob	audi, luxuri, nice, bmw, love, car, merced
FREX	bmw, audi, luxuri, merced, volkswagen, seri, nice
Lift	automat, seri, night, lexus, transmiss, manual, specialti
Score	audi, bmw, luxuri, nice, merced, volkswagen, love
Topic 6 Top Words	
Highest Prob	prefer, easi, park, like, see, hatchback, trip
FREX	park, prefer, trip, cheap, smaller, smart, easi
Lift	hatch, acura, open, access, blind, part, window
Score	prefer, park, easi, hatchback, see, trip, smaller
Topic 7 Top Words	
Highest Prob	suv, vehicl, need, cargo, minivan, larger, van
FREX	suv, cargo, minivan, van, larg, carri, need
Lift	visit, campus, camp, kid, accomod, babi, ppl
Score	suv, need, minivan, cargo, van, vehicl, larger
Topic 8 Top Words	
Highest Prob	honda, mazda, toyota, good, civic, comfort, size
FREX	honda, mazda, civic, camri, accord, reliabl, roomi
Lift	accord, honda, prui, basic, camri, camera, milag
Score	honda, mazda, civic, toyota, comfort, good, size
Topic 9 Top Words	
Highest Prob	like, car, select, think, varieti, fleet, tesla
FREX	select, think, alreadi, varieti, fine, fleet, current
Lift	alreadi, mix, select, current, divers, porch, fine
Score	like, car, select, think, alreadi, fine, fleet

Table 2.11: Words and scoring methods for 9 topic model of vehicle-related comments.

Three broad themes were identified in the topics presented in the final model. First were topics that seemed to discuss vehicle, technology – primarily topic 2 and topic 4. Second were drivers that focused on fun and luxury – primarily topic 3 and topic 5. Topic 7 seemed to focus on utility of the vehicle such as carrying capacity or size. The remaining topics 1, 6, 8 and 9 seemed to indicate a preference for convenience of the service.

Service topics were more difficult to identify. Topics 1, 5, 6, 7 and 8 indicated some level of focus on the convenience of the service, while remaining topics seemed to focus on happiness or complaints about a very specific issue: customer service concerns or happiness about a positive service recovery, or comments about the condition in which a vehicle was left by a previous customer. Some respondents used language that indicated price sensitivity. While this work does not include price sensitivity as a latent variable, it is a likely candidate for further investigation.

2.7.2 Examining Comments and Topics for coherence.

Another method that Roberts et al. (2014a) propose to investigate modeled topics is by exploring the documents (comments) that are highly associated with each topic. While the authors generally use normalized topic proportions to do this, we here present representative comments with high unnormalized values associated with each topic. As a result, representative comments include those that score highly on multiple topics.

Topic 2 and Topic 4 were both identified as Topics associated with preference for vehicle technology, and to confirm the accuracy of this assessment, individual comments were examined. Figure 2-23 shows a sample of representative comments that score highly on η_4 and Figure 2-24 shows a sample of comments that score highly on η_2.

This examination seems to confirm that the Topic models have done a relatively good job of identifying comments related to vehicle technology preferences, as nearly all the comments with large weights on Topic 2 and 4 refer to specific technology types. There also appears to be some overlap with price sensitivity in Topic 2. While price sensitivity is not separately analyzed in this work, it appears to be a good candidate for future study.

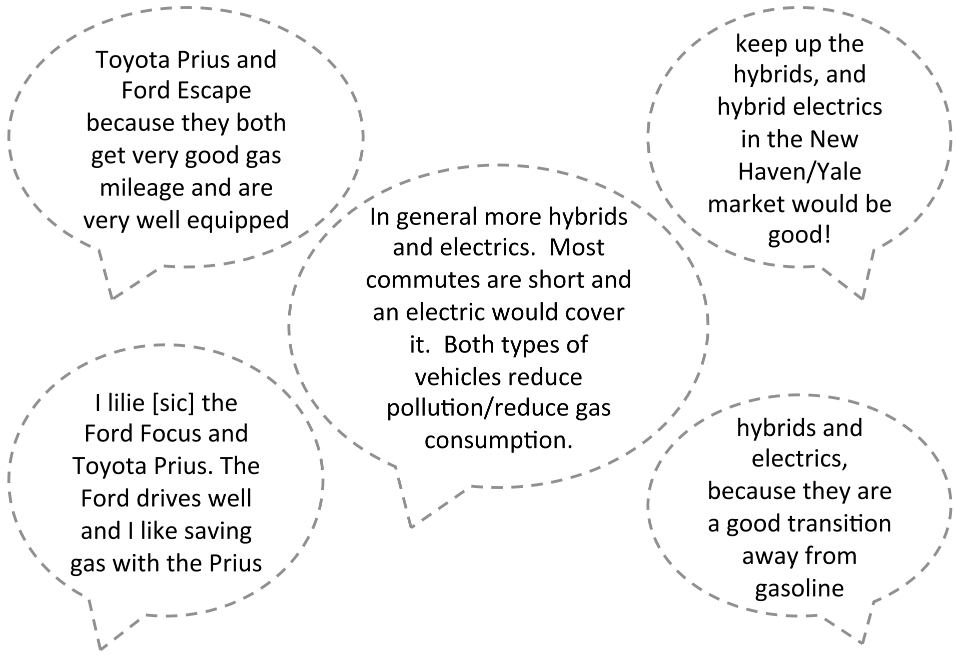

Figure 2-23: Representative vehicle-related comments with a high weight for Topic 4.

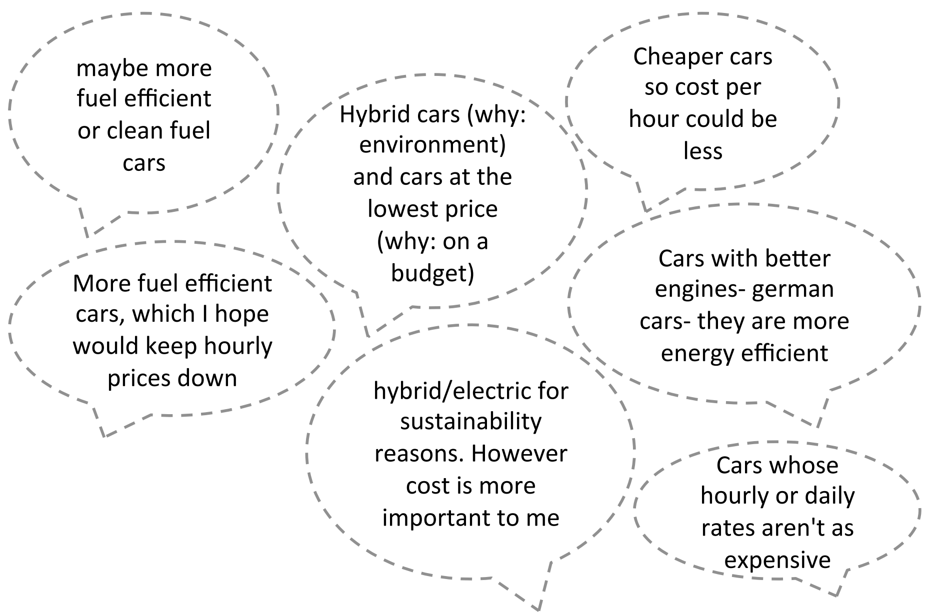

Figure 2-24: Representative vehicle-related comments with a high weight for Topic 2.

Topic 1 Top Words	
Highest Prob	locat, zipcar, one, option, pick, differ, return
FREX	pick, locat, drop, differ, near, option, home
Lift	columbia, drop, pick, plaza, someplac, syracus, nashvill
Score	locat, drop, pick, differ, near, return, option
Topic 2 Top Words	
Highest Prob	servic, great, custom, thank, zipcar, call, guy
FREX	excel, servic, comment, great, happi, nope, custom
Lift	gotten, nope, surpris, win, agent, excel, glad
Score	servic, custom, great, thank, call, excel, guy
Topic 3 Top Words	
Highest Prob	car, gas, card, wish, also, dirti, seem
FREX	dirti, tank, gas, clean, fill, card, cleaner
Lift	declin, refil, sticker, wipe, bldg, cleaner, fluid
Score	gas, dirti, card, tank, clean, car, fill
Topic 4 Top Words	
Highest Prob	hour, price, rent, rental, rate, day, cheaper
FREX	rate, hour, rental, competit, cheaper, daili, per
Lift	effect, flat, competit, daili, except, overnight, packag
Score	hour, rate, price, day, rental, cheaper, lower
Topic 5 Top Words	
Highest Prob	use, car, zip, avail, take, need, expens
FREX	take, use, zip, expens, closest, less, much
Lift	hybridelectr, blvd, chicago, dislik, econom, fanci, liabil
Score	use, zip, car, expens, avail, trip, take
Topic 6 Top Words	
Highest Prob	vehicl, just, realli, need, drive, one, want
FREX	survey, answer, vehicl, want, larg, just, realli
Lift	alot, complic, item, scenario, terribl, answer, bonus
Score	vehicl, survey, answer, just, want, drive, know
Topic 7 Top Words	
Highest Prob	time, reserv, return, late, get, charg, report
FREX	minut, report, credit, penalti, gps, someon, time
Lift	buck, extens, grace, improv, increment, indic, quarter
Score	time, report, return, minut, late, previous, min
Topic 8 Top Words	
Highest Prob	zipcar, love, keep, pleas, work, get, park
FREX	love, hope, space, univers, work, campus, build
Lift	urban, build, hope, leas, minivan, suffici, behind
Score	love, work, park, keep, good, live, garag
Topic 9 Top Words	
Highest Prob	make, like, see, better, avail, phone, reserv
FREX	app, civic, hybrid, honda, model, make, key
Lift	adapt, app, bluetooth, energi, interfac, power, usb
Score	make, app, model, hybrid, audio, usb, ford

Table 2.12: Words and scoring methods for 9 topic model of service-related comments.

2.8 Integrated Choice and Latent Variable Model

This section presents the results of two variants of Intergrated Choice and Latent Variable models. Due to computational constraints, models were run with no more than two latent variables at a time. A graphical diagram of the ICLV model is shown in Figure 2-25.

2.8.1 ICLV Model: Technology Affinity and Powertrain Type

In the integrated model, Technology Affinity is interacted with the vehicle powertrain type. As with earlier model specifications, vehicle powertrain type is also integrated with reported travel distance. The Latent Variable "Technology Affinity" is interacted with the intercept of all non-gasoline vehicle types. As a result, varying levels of Technology Affinity manifest themselves as a vertical shift upwards or downwards in the utility of hybrid, Plug-in Hybrid and Electric Vehicles relative to gasoline vehicles.

The results from this model are shown in Tables 2.13, 2.14 and 2.15. Table 2.13 shows results of the choice model, Table 2.14 shows the results of the latent variable structural model, and Table 2.13 shows the results of the latent variable measurement model.

The coefficients of the choice model maintain their sign and the magnitude and are relatively stable from MNL models. All coefficients on service attributes of price, Access Distance and Schedule are highly significant, and result in similar values of willingness to pay for Access Distance and Schedule.

For vehicle types, coefficients on Hybrid and Plug-In Hybrid vehicles are similar to those from MNL models. The coefficient on Electric Vehicles is not statistically significant in this model. However, the utility of both types of plug-in vehicles decreases with travel distance. The decrease in utility is greatest with Electric Vehicles, consistent with the limited range of such vehicles, suggesting that users believe that traveling greater distances with electric vehicles will introduce difficulties of some sort – perhaps compromising their travel plans or obligating them to locate charging facilities.

However, as with MNL models, the utility of Plug-in Hybrid Vehicles also decreases with travel distance, despite the fact that such vehicles are not subject to hard range limits and

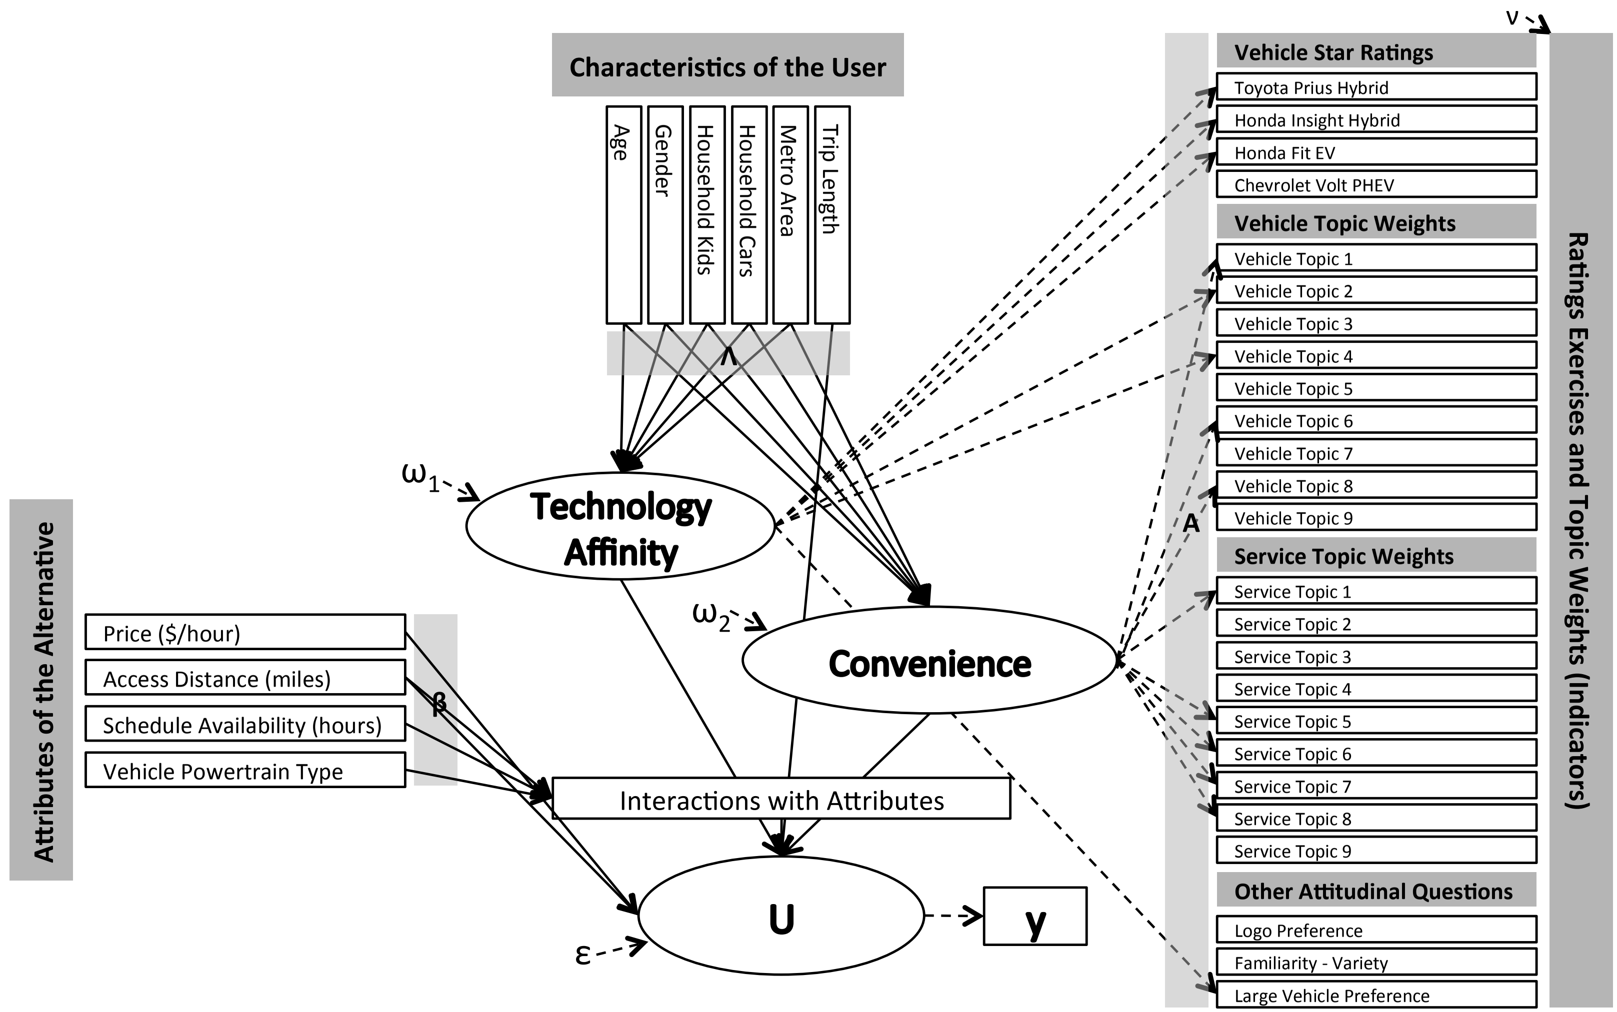

Figure 2-25: Representation of ICLV Model.

can continue to operate as conventional hybrid vehicles once their battery is depleted.

Table 2.13: ICLV Model: Choice Model Results.

Attribute	**Coefficient**	**Std Error**	**t-test**
Price ($/hour)	-0.354	0.00638	-55.44
Access Distance (miles)	-0.748	0.0208	-35.93
Schedule 30min Early	-0.768	0.0327	-23.48
Schedule 30min Late	-0.631	0.0391	-16.15
Schedule 1hr Early	-0.883	0.042	-21.03
Schedule 1hr Late	-0.594	0.0401	-14.8
Schedule 2hrs Early	-0.985	0.043	-22.9
Schedule 2hrs Late	-1.85	0.0532	-34.71
Hybrid	0.139	0.0539	2.58
PHEV30	-0.135	0.0537	-2.51
EV100	-0.00981	0.0555	-0.18
Interactions:			
Hybrid and Travel Distance (mi)	-0.00199	0.000552	-3.61
PHEV and Travel Distance (mi)	-0.00295	0.00057	-5.18
EV and Travel Distance (mi)	-0.00545	0.000671	-8.12
Technology Affinity and Hybrid/PHEV/EV	0.902	0.156	5.78

Sample size:	13002
Init log-likelihood:	-47084.564
Final log-likelihood:	-30187.268
Likelihood ratio test for the init. model:	33794.593
Rho bar for the init. model:	0.358
Number of Halton Draws	200

There is a strong interaction effect between the Technology Affinity latent variable and all non-gasoline vehicles. As expected, the interaction is positive, with users exhibiting a higher level of Technology Affinity deriving higher utility from non-gasoline vehicles. In this specification the latent variable is interacted exclusively with the intercept of non-gasoline vehicles: the distance interaction term is not further interacted with the latent variable.

In the structural equation of the latent variable, a number of characteristics are significant. Respondents from Boston, Portland, University campuses and Seattle have a significantly higher Technology Affinity, whereas respondents from Toronto have a lower Technology Affinity with weak statistical significance. Since this study did not include information on income or education of respondents, the attitudes reflected in the users in these metro areas may in fact be the product of systematic differences in income or education of the users.

Both younger and older respondents tended to exhibit a lower Technology Affinity, with the lowest levels occurring at the 18-21 year age group. This result is contrary to common beliefs about younger people typically being highly represented within early adopter groups. One possible explanation for such behavior would be that younger users, who typically have less driving and car-buying experience, are less knowledgeable about the vehicle types presented in this survey and therefore base their decisions on incorrect assumptions about the capabilities of the vehicles.

Table 2.14: ICLV Model: Latent Variable Structural Model Results for Technology Affinity Latent Variable.

Characteristic	**Coefficient**	**Std Error**	**t-test**
Metro Area: Boston	0.117	0.0407	2.88
Metro Area: Chicago	0.0192	0.0429	0.45
Metro Area: Philadelphia	-0.00224	0.0671	-0.03
Metro Area: Portland	0.101	0.0705	1.43
Metro Area: Universities	0.0956	0.0472	2.03
Metro Area: San Francisco	0.0543	0.0414	1.31
Metro Area: Seattle	0.0903	0.0613	1.47
Metro Area: Toronto	-0.0618	0.0475	-1.3
Metro Area: Washington DC	0.0445	0.0444	1
Gender: Female	0.0518	0.0243	2.13
Age: 18-21	-0.112	0.0464	-2.42
Age: 22-25	-0.067	0.0326	-2.06
Age: 46-55	0.0369	0.0397	0.93
Age: 56+	-0.0478	0.0415	-1.15
No Children Household	0.0808	0.0358	2.26
No Car Owned	0.0173	0.0253	0.69
Standard Deviation	0.466	0.0702	6.65

Figure 2-26 shows the distribution of the Technology Affinity latent variable for all users,

Table 2.15: ICLV Model: Latent Variable Measurement Model Results for Technology Affinity Latent Variable.

Indicator	Coefficient	Std Error	t-test
Technology Affinity → Honda Insight Rating	1*	-	-
Technology Affinity → Honda Fit EV Rating	0.783	0.312	2.51
Technology Affinity → Toyota Prius Rating	0.763	0.163	4.68
Technology Affinity → LargeVehFrequency	0.205	0.053	3.86
Technology Affinity → Vehicle Eta2	1.63	0.251	6.5
Technology Affinity → Vehicle Eta4	1.69	0.261	6.45

*Factor loading on Honda Insight Rating fixed to one for model estimation purposes

taking into account all associated demographic characteristics of each individual. Similarly, Figure 2-27 shows the distribution in Technology Affinity for men (solid line) and women (dashed line) separately, incorporating all other demographics as before. The effect of gender on Technology Affinity is relatively weak, but women throughout the population have higher levels of Technology Affinity than their male counterparts.

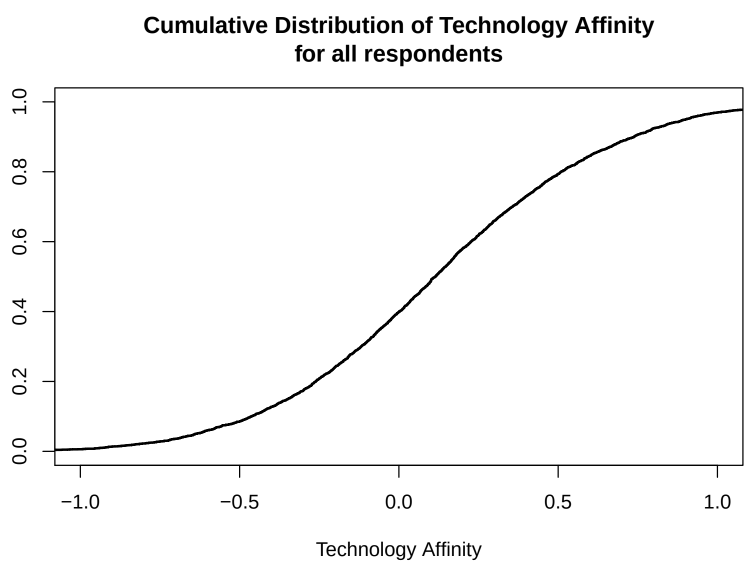

Figure 2-26: CDF of Technology Affinity Latent Variable.

Figure 2-28, Figure 2-29 and Figure 2-30 present the results of interactions of the Technology Affinity latent variable with vehicle type and reported travel distance. Figure 2-28

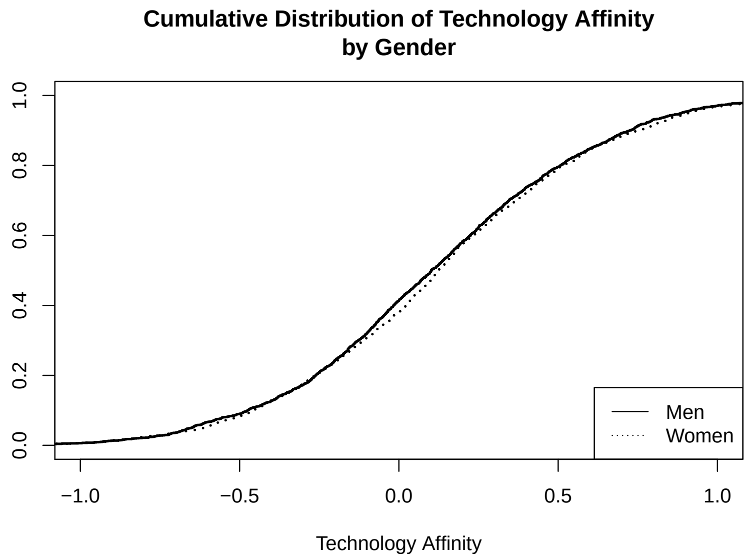

Figure 2-27: CDF of Technology Affinity Latent Variable by Gender.

presents the results for users with a mean value of Technology Affinity, Figure 2-29 presents results for users with a value of Technology Affinity one standard deviation above average, while Figure 2-30 presents results for users with a value of Technology Affinity one standard deviation below average.

For users with a mean value of Technology Affinity (Figure 2-28), hybrids are weakly preferred to gasoline vehicles for most travel distances, and gasoline vehicles are weakly preferred to hybrids for longer travel distances. While this could be perceived as a misunderstanding of the capabilities of hybrids, it may also indicate that hybrid vehicles generally have other characteristics disliked by users traveling long distances, despite the note in the discrete choice survey which instructed users to assume that vehicles are identical except for powertrain type.

For respondents with a mean value of Technology Affinity, respondents are essentially indifferent to Plug-in Vehicle types (Electric Vehicles and PHEVs) for very short travel distances, but interaction effect of reported travel distance with PEVs means that utility of these vehicles rapidly falls below that of a gasoline vehicle, even for moderate travel distance well within the range of a 100 mile Electric Vehicle.

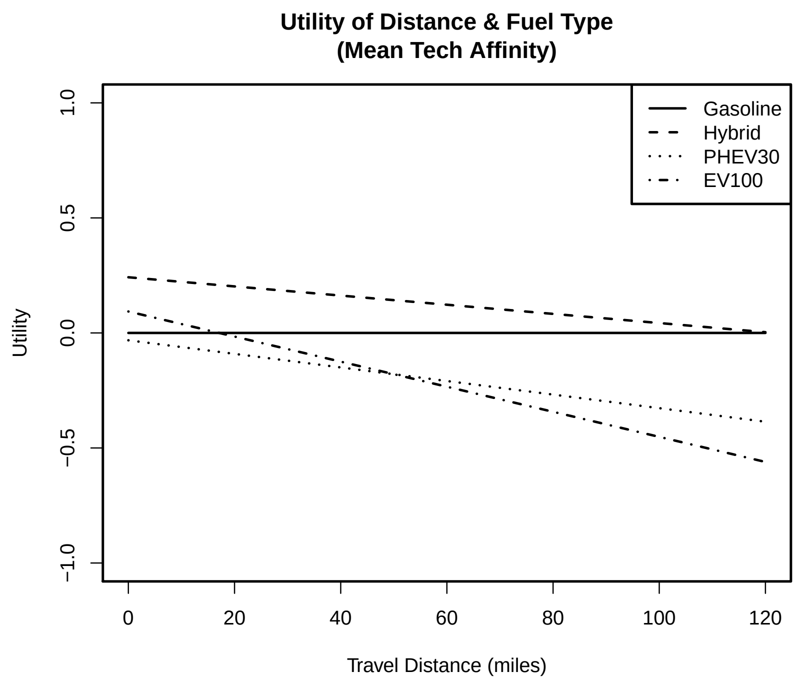

Figure 2-28: Preference for Vehicle Technology for Mean Technology Affinity.

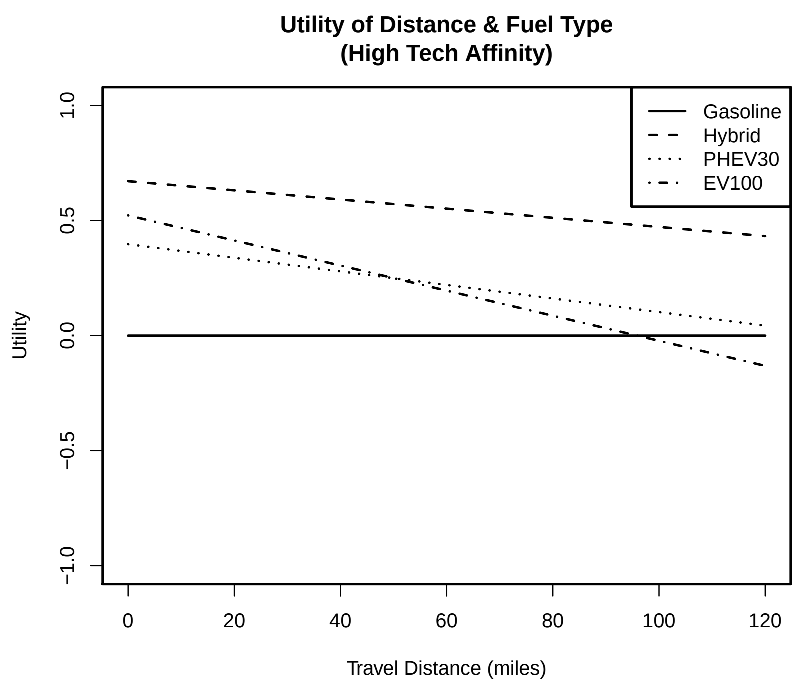

Figure 2-29: Preference for Vehicle Technology for High Technology Affinity.

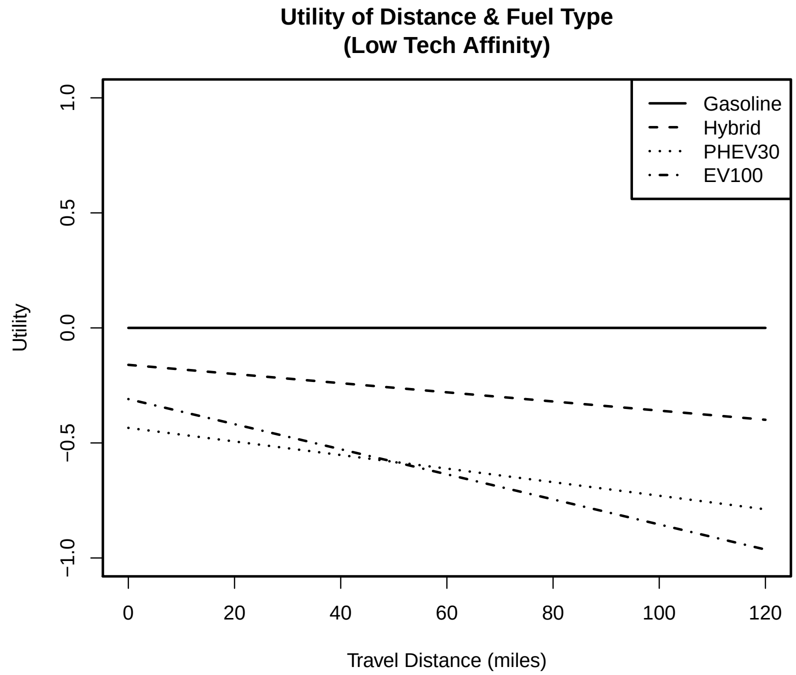

Figure 2-30: Preference for Vehicle Technology for Low Technology Affinity.

For users with a value of Technology Affinity one standard deviation above the mean, shown in Figure 2-29 users have a much higher utility for non-gasoline vehicles. Hybrids are preferred to gasoline vehicles for all but the very longest trips, and Plug-in Vehicles are preferred for all trips up to 100 miles. Electric Vehicles are only not preferred to gasoline vehicles once the reported travel distance exceeds 100 miles, the specified range of the Electric Vehicle in the survey. This result suggests that users with a high Technology Affinity are relatively well informed about the capabilities of PEVs, and take a rational approach towards vehicles with a limited range.

However, users with a low Technology Affinity, one standard deviation below the mean, are shown in Figure 2-30. These users strongly prefer gasoline vehicles under all circumstances. If we calculate approximate willingness to pay values for such users, the drop in utility of an Electric Vehicle for a trip with a reported distance of 100 miles would be approximately equivalent to an increase in price of \$3/hour for a user with a low Technology Affinity.

In Table 2.15, we can see the results of the factor loadings of the measurement model for Technology Affinity. We can see that the factor loadings on relative star ratings for all advanced vehicle types are positive, as are the factor loadings on Vehicle Topics 2 and 4. Of particular interest is that the largest factor loadings are found on vehicle topic weights, suggesting that topic content of open-ended questions provides valuable insight into the attitudes of users.

2.9 Summary

The analysis performed in this chapter has presented insights in a number of areas. First, we identified new information about perceived and actual usage behavior, and used new choice model formulations to improve our understanding of the role that schedule and vehicle technology play in a user's decision over which vehicle to rent. We also demonstrated a way to integrate large numbers of open-ended comments in a choice model using machine learning techniques. We also identified the role that carsharing plays in exposing users to advanced vehicle technology.

Takeaway 1

Carsharing serves as a conduit to expose large numbers of users to advanced vehicle technology.

In this survey, members belonged to a carsharing service where only 892, or 10% of the vehicles are hybrids. Yet more than 50.5% of respondents had driven a hybrid. If the sample representative, this suggests that the ratio of users exposed to vehicles is nearly 400:1. 400 users exposed for a single car is an enormous increase over the 1-2 drivers that typically drive a vehicle purchased purchased by private users in a household.

Consumer adoption literature commonly notes that "trialability," or the ability to interact risk-free with an innovation, is typically an important component of adoption. In this sense, carsharing is a *de facto* test drive for some users, offering them the chance to drive dozens of different types of vehicles without ever visiting a dealer.

However, the"stickiness" of this exposure remains unknown. While most respondents do not own a car, many will buy one in the future, and a clear extension of this work would be to study vehicle purchasing behavior of former carsharing users. However, finding a control group or valid counterfactual would be significant challenges to such work.

Takeaway 2

Getting a car when and where they need it at the cheapest price possible is still most important to users.

This survey included a large number of observations, and the choice formulation as a user survey, rather than adoption or mode choice survey, allowed the estimation of coefficients on vehicle type and schedule that were previously not observed. In virtually all models estimated the coefficients on price, distance and schedule were larger than those on fuel type.

Takeaway 3

An "average" carsharing user is still cautious about renting a plug-in vehicle, but attitudes vary widely.

For most users, hybrid vehicles are mildly preferred to gasoline vehicles. But for electric vehicles and plug-in hybrids, user preferences are strongly dependent on travel distance.

Decreased willingness to use electric vehicles for longer trips is easily attributable to the inability to complete travel without recharging. However, preferences against plug-in hybrids for these use cases are more difficult to explain. One explanation would be that respondents simply believe that the capability of PHEVs is diminished after the battery is exhausted, while in reality PHEVs (except Range-Extended EVs) perform as well as conventional hybrids once their battery is exhausted.

User overestimation of travel distance is potentially problematic for range-limited vehicles such as EVs. At the margin, this behavior could lead to consumers eschewing vehicles that would be perfectly suitable for their travel needs, and underutilization of EVs.

The incorporation of a latent variable highlights large heterogeneity in user preferences for vehicle technology. While at the mean level of technology affinity users only weakly prefer plug-in vehicles at very short travel distances, there is enormous variation in technology affinity, and a substantial portion of users prefer plug-in vehicles to gasoline vehicles, and all else equal, would pay slightly more for them.

Surprisingly, young users have lower technology affinity than older users. These attitudes run counter to conventional wisdom that younger users are more willing to adopt new technology, and perhaps suggest that education is necessary to ensure that users fully understand the capabilities of plug-in vehicles.

Takeaway 4

What users say can be an important indicator of their attitude, and can be used in choice modeling.

This work also integrated Structural Topic Models and Hybrid Choice Models, using topics weights generated from Structural Topic Models of open-ended survey responses as indicators of latent variables. The topic models used produced models that resulted in semantically coherent topics. When used as indicators of the Technology Affinity variable the factor loadings on topic weights associated with vehicle technology were higher than conventional star ratings of advanced technology vehicles such as the Toyota Prius, Honda

Insight and Honda Fit EV.

Topic weights associated with vehicle-specific comments proved more valuable than service-related comments as indicators of latent variables. While this is partially explained by the greater response rate, another plausible explanation is that users simply attempted to describe quite a large and diverse set of topics.

The combination of topic models with discrete choice models demands further investigation. One clear direction that could be pursued is the simultaneous estimation of topic models and choice models, which could be expected to be more efficient than the sequential estimation process used here.

Chapter 3

Energy Consumption of Plug-in Electric Vehicles

3.1 Introduction

Growing concern over the long-term availability of petroleum and the environmental impact of its combustion products has led to the development of various alternative fuel vehicle technologies. Plug In Vehicles (PEVs), or vehicles which source some of their energy from grid electricity, are increasingly viewed as a way to reduce petroleum consumption and emissions from the transportation sector while maintaining the levels of individual mobility that humans have become accustomed to in many parts of the world.

PEVs are broadly divided into two types: Electric Vehicles (EVs) which use exclusively electric power stored in a battery, and Plug-In Hybrid Vehicles (PHEVs) which use both electricity and gasoline for propulsion. In both EVs and PHEVs reductions in emissions are dependent on drawing electricity from a clean electrical grid. Like conventional vehicles, the emissions of PEVs depend on driver behavior, usage, and ambient conditions which impact the efficiency of the vehicle operation. PHEVs, which operate on two energy sources, introduce the additional complexity that emissions profile and energy consumption can be different when operating on gasoline or electricity, and therefore depend on the portion of the driving distance and time that the vehicle operates on each fuel. This in turn depends on when and where the user recharges the vehicle.

This chapter presents the results of the analysis of energy consumption in two separate vehicle studies. The first is a study of 125 Toyota Prius PHEV prototype vehicles deployed in the U.S. from 2011 - 2012 as part of an evaluation of real-world vehicle performance. The second is a study of BMW ActiveE vehicles in operation from June, 2012 through early 2015, deployed to both private users and as part of the carsharing system DriveNow.

3.2 Related Work

The study of the performance of new vehicle technology is an area of active research. Despite availability for several years, plug-in vehicles make up a relatively small portion of the new vehicle fleet, and therefore many studies of the energy consumption of such vehicles are based on small samples and/or early vehicle adopters.

Early studies of PHEV usage and energy consumption, such as the impact of battery size and the grid impact of recharging, have been largely based on analysis of known mobility patterns, surveys, and retrofitted hybrid vehicles (Denholm and Short, 2006; Hadley and Tsvetkova, 2009). Various efforts have attempted to develop more realistic assessments of how PHEVs will perform in the real world. Vehicle-level simulation has been used to model the effects of design attributes and control strategies (Gonder and Simpson, 2006; Vyas et al., 2009), while survey data and, more recently, GPS-based datalogging are used to characterize driving patterns. (Vyas et al., 2009; Lin and Greene, 2011; Khan and Kockelman, 2012; Gonder et al., 2007; Williams et al., 2011) The validity of these approaches requires an assumption that driving behavior will be the same for PHEVs as for conventional vehicles. Charging behavior is an area of even greater uncertainty. Due to a lack of real-world data, charging behavior in existing work has been largely assumption-driven (Khan and Kockelman, 2012) or based on small samples. Axsen and Kurani (2009) surveyed respondents about possible charging behavior, based on availability and perceived importance. Davies and Kurani (2010) reported results from a study of 40 vehicles for a one-week period during which the authors identified a mean of one daily charge, including two participants that did not recharge at all. Williams et al. (2011) noted the paucity of real-world information on recharging behavior, and presented the results of one prototype PHEV vehicle rotated among twelve households over one year to gather more information on real-world charging behavior. Using small samples to predict fleet-wide impact generates substantial uncertainty. (Gonder et al., 2007)

Understanding real-world energy consumption of both gasoline and plug-in vehicles is challenging. A large number of factors, including driving style, vehicle condition, traffic conditions, ambient temperature, vehicle load all may affect the energy consumption of a vehicle. Researchers at MIT Berry (2010) has identified that driver aggressiveness alone can cause as much as a fifty percent variation in fuel consumption in a given vehicle on a given route. In the United States, CAFE regulations assume that gasoline vehicles will fall 20% short of their test-cycle economy under "real world" conditions. PEVs are estimated to further deviate from their test-cycle energy consumption, and the on-road "gap" between real-world and test-cycle energy consumption for these vehicles is estimated at 30%. (EPA, 2012) Researchers have recently observed that test-cycle energy consumption is decreasingly representative of actual energy consumption. (Mock et al., 2013)

The EV Project, managed by the Idaho National Laboratories (INL, 2015) is currently one of the largest repositories of electric vehicle operational data, and has produced a number of papers which have identified differences in EV energy consumption by region, differences in driving and charging behavior for PEVs from a large fleet of early Chevrolet Volt and Nissan Leaf vehicles. (Smart and Schey, 2012; Smart et al., 2013; Project, 2013; Smart et al., 2014; Smart, 2014) However, the work does not consider the theoretical effect of changes to vehicles, nor are the reports able to differentiate between accessory and drive energy or between shared and private vehicles.

Recent research has identified the role of climate and grid emissions on the range and emissions profile of electric vehicles. Yuksel and Michalek (2015) highlight the reduced range of electric vehicles in areas with colder climates, while Nichols et al. (2015) identify potential detrimental effects of grid emissions of criteria pollutants associated with recharging electric vehicles from coal-fired electric grids.

This work expands this body of knowledge by analyzing the driving and charging behavior observed during a large, long-term, geographically diverse deployment of 125 prototype PHEVs and 700 BEVs in the United States. A range of vehicle performance measures are investigated, including gasoline consumption in charge-sustaining (CS) and charge-depleting (CD) modes, electricity consumption in CD mode, effective electricity consumption per electrified km, real-world effective electric range and utility factors and petroleum displacement factors at a fleet and vehicle level. The fleet impact of vehicle and behavioral changes on the consumption of gasoline and electricity is simulated. We show variations in measured energy consumption from auxiliary sources, and compare the energy consumption of shared and private vehicles of the same type.

3.3 PHEV Operating Concepts

PHEVs may operate in three modes: charge-sustaining (CS), all-electric (EV), or blended. In CS mode, the batterys state of charge (SoC) fluctuates within a limited range like that of a regular HEV, but exhibits no long-term trend, so the vehicle is considered to use only gasoline. When using only electricity a PHEV is operating in EV mode, and when using both gasoline and electricity a PHEV is operating in blended mode. Both EV mode and blended mode are charge-depleting (CD) modes, because SoC trends downward over time.

PHEVs are commonly designated PHEV-x, where x is some measure of electric range. However, there is ambiguity as to the exact meaning of this term. Kurani et al. (2010) identify at least three published interpretations of the PHEV-x designation: x = equivalent miles of gasoline displaced by electricity (the interpretation preferred in this work), x = distance before the engine first turns on, or x = distance that the vehicle travels in CD mode. (Kurani et al., 2010) Two other important PHEV-related concepts are utility factor (UF) and petroleum displacement factor (PDF). UF, as defined by SAE Standard J2841, is the ratio of the distance the vehicle travels in CD mode to the total distance traveled, and

PDF is the ratio of distance attributable to the non-petroleum fuel (i.e. grid electricity) to total distance traveled. (SAE, 2010) PDF depends on vehicle design, driving patterns, and charging behavior: (Gonder and Simpson, 2006; Dowds et al., 2010; Samaras and Meisterling, 2008)

These two concepts UF and PDF are functionally equivalent for vehicles that do not use blended mode. However, for vehicles that do use blended mode, UF will overestimate fuel displacement because a portion of the tractive force during CD mode is petroleum-derived. This work reports estimates for both UF and PDF. To calculate PDF when blended mode is used, electrified distance is defined as the amount by which the distance traveled in CD mode exceeds the distance that can be explained by the amount of gasoline consumed in CD mode. The latter is taken to be the CD mode gasoline usage (GasolineCD) divided by the CS mode fuel consumption rate (FCcs):

There are potential problems imputing gasoline distance in CD mode based on fuel consumption in CS mode, since systematic differences in the driving patterns that characterize each mode may exist. These risks are partially mitigated by calculating the CS mode fuel consumption separately for each vehicle in the trial, and using that vehicles specific CS mode fuel consumption to impute its gasoline miles in CD mode. However, bias may remain if, for instance, trips that occur soon after a charge have different speed, acceleration, or accessory load profiles than those occurring later. It should be possible to adjust for these differences using more disaggregated data from each mode, but such adjustments are not considered in this analysis.

3.4 Case Study: Toyota Prius PHEV Trial and Data

The test fleet studied here consisted of 125 pre-production Toyota Prius PHEVs deployed in the U.S. from approximately April, 2011 to April, 2012. The deployment was part of a global fleet evaluation and learning program to test the real-world usage of the vehicles against intended and expected usage and performance. The evaluation also provided a platform for assessing the potential merits of changing the availability and accessibility of level 1 (110V) or

2 (220V) at-home charging, work place charging, and other vehicle-grid interactions. A final objective was to disseminate information on the driving patterns, charging habits, and other factors related to the real-world operation of PHEVs. Vehicle Specifications The powertrain configuration of the vehicles was a prototype, adapted from the 2010 Toyota Prius. It was not representative of current or future production PHEVs from the manufacturer. Key specifications for the vehicle included:

- 5.4 kWh Li-ion battery, estimated 21 km range (EPA test cycle)
- Permanent magnet synchronous motor (maximum 60 kW / 207 Nm)
- 1.8-liter internal combustion engine (maximum 73 kW / 142 Nm)
- Maximum combined output 100 kW

3.4.1 Vehicle Operation & Charging

The PHEVs in this trial operate in three modes: All-electric, blended, and charge-sustaining. The selection of operating mode is dictated by battery SoC and power demands from driver inputs and accessory loads. The vehicle operates in EV mode with sufficient SoC and low-moderate loads consistent with typical operation and accessory loads. If power requirements exceed battery limits as a result of throttle input, HVAC settings, or vehicle speed exceeding 100 km/h, the vehicle enters blended mode. Additionally, as the battery nears a state of discharge, the vehicle may enter blended mode in order to reduce the current flow from the battery to extend battery life. When SoC is low, the PHEV operates like a conventional HEV, and uses a limited portion of the battery capacity for regenerative braking and supplemental torque. Each vehicle was supplied with a compact, readily portable level 1 charger using a standard 15 A household outlet, though it is unknown whether participants carried the chargers with the vehicles or used them in a fixed location. Other permanent and semi-permanent level 1 and level 2 charging facilities were available in some locations. The vehicles were not compatible with level 3 charging.

Participants for the trial in the U.S. were a variety of corporate, governmental, and educational partner institutions, selected to meet the following primary objectives according to Future Fuels and Environmental Strategy Manager Jaycie Chitwood:

- A range of operational use cases and conditions including personal, business, and demonstration fleets.

- Geographic distribution that provide wide coverage, but clustered to minimize the cost of support organizations.

- Some preference was given to partner organizations with an existing relationship with the manufacturer, and to those with known outreach capabilities.

All vehicles were equipped with proprietary data loggers collecting more than 100 channels of data with 1-second resolution. The data were cached, periodically transmitted via cellular modem to the data logger manufacturer, and then transferred to the vehicle manufacturer. The data available to researchers included time, speed, location, temperature, battery SoC, operating mode (CS/CD), and information on HVAC and regenerative braking use. Researchers at the National Renewable Energy Laboratory (NREL) aggregated the second-by-second data into 59,287 individual trips.

3.4.2 Gasoline and Electricity Consumption of Vehicles in the Trial

The performance of the PHEVs in this trial was characterized according to several common figures of merit. When operating in CS mode, the gasoline consumption was similar to that of the standard Prius HEV. In CD mode, the gasoline consumption was approximately halved, as grid electricity provided a portion of the vehicles power demands. The use of blended mode complicates the attribution of distance traveled to either gasoline or electricity, so a simple method for assigning CD mode travel to each fuel was established, as described in Section 3.3.

For each vehicle in the trial, average fuel consumption in CS mode was calculated by dividing total gasoline consumption by total distance traveled in CS mode. Nearly all of the vehicles returned average CS mode fuel consumption between 4.0 and 6.0 l/100 km (39 59 mpg), as shown in Figure 3-1. The overall average was 4.90 l/100 km (48.0 mpg), which is close to the 4.70 l/100 km (50 mpg) EPA rating of the standard 2010 Prius HEV. When operating in charge-depleting mode, the average rate of gasoline consumption was approximately halved, to an average of 2.48 l/100 km (94.6 mpg). Most of the vehicles returned rates between 1.0 and 4.0 l/100 km. The variance in gasoline consumption was higher in CD mode than in CS mode, reflecting the wider range of variables that may influence gasoline consumption in CD mode.

The effective electricity consumption rate per electric kilometer was calculated as the distance attributed to grid electricity divided by the total electricity consumed in CD mode (Figure 3-2). The large majority of vehicles used between 150 and 350 Wh for each additional kilometer of travel beyond that explained by their gasoline usage. The overall mean across all vehicles in the trial was 218 Wh/km, which is similar to the EPAs electricity consumption ratings for the Nissan Leaf EV and Chevrolet Volt PHEV (213 and 225 Wh/km, respectively). For 3 kWh of working capacity in CD mode, the distribution of electricity consumption rates implies a distribution of electric range values; on average, the vehicles in this trial gained 13.8 km of electrified travel from a 3 kWh charge.

Operation in blended mode means that vehicles generally drive further than their effective electric range in order to fully deplete their CD battery capacity. Although the vehicles used an average of 218 Wh for each electric km, the actual rate of electricity consumption in CD mode averaged 107 Wh/km. Thus, even in CD mode, gasoline was providing approximately half the vehicles energy requirements on average, and it took an average of 28 km before a fully charged vehicle exited CD mode. This pattern will tend to reduce the PDF, since the vehicle must be driven further in order to fully exploit the batterys stored electricity. Put differently, the finding that gasoline still accounts for half the distance traveled in CD mode effectively puts an upper limit of 50% on the PDF, although it could be higher if driving

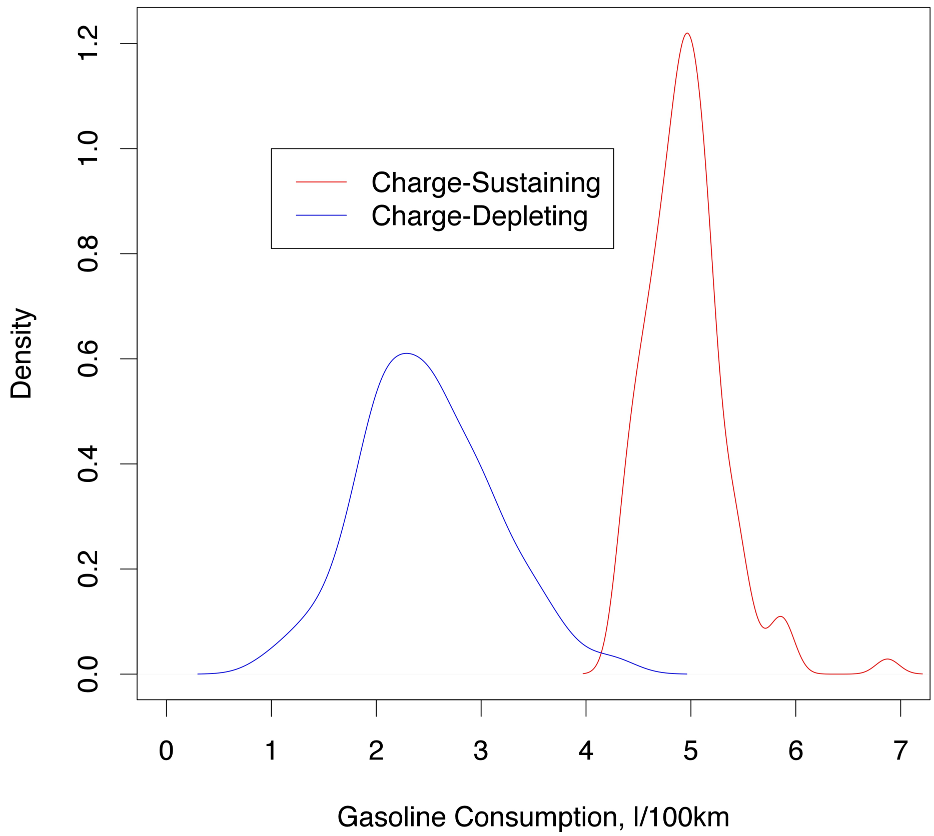

Figure 3-1: Density plots of fuel consumption in Charge-Depleting and Charge-Sustaining modes.

habits were adjusted to increase pure EV mode usage.

The PDF was calculated as the ratio of the distance attributed to grid electricity to the total distance traveled. The distribution of PDF values is shown Figure 3-3, and the overall average in this trial was found to be 13.7%. Also shown is the distribution for UFs across the vehicles in this trial, which averaged 28.1% overall. The PDFs and UFs observed in this trial are lower than predicted by the SAE J2841 standard for UF, a discrepancy that is further explored below.

There was a very wide spread in the values of these metrics across different vehicles. The

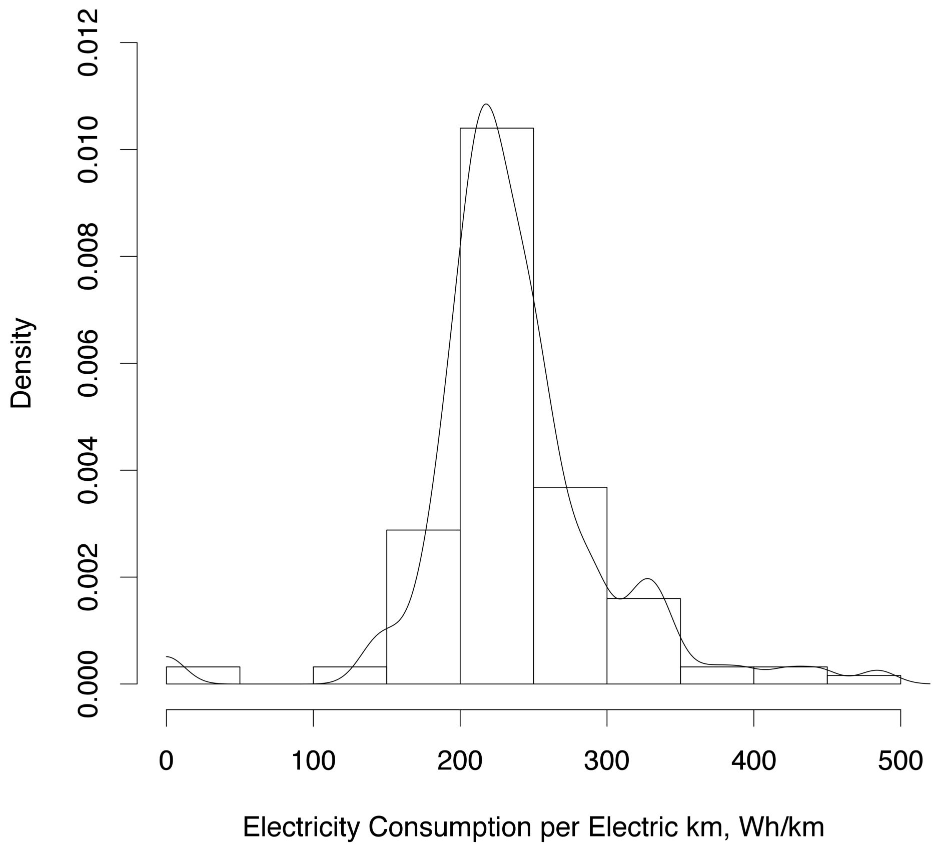

Figure 3-2: Distributions of electricity consumption when operating in electric mode.

highest PDF was 59%, showing that the right combination of driving patterns and charging habits can make very effective use of even a small battery to displace gasoline. On the other hand, five of the 125 vehicles in the study had PDFs of less than 1%, indicating that they derived almost none of their energy usage from grid electricity. Another 16 vehicles had PDFs between 1% and 5%, typically coincident with infrequent charging.

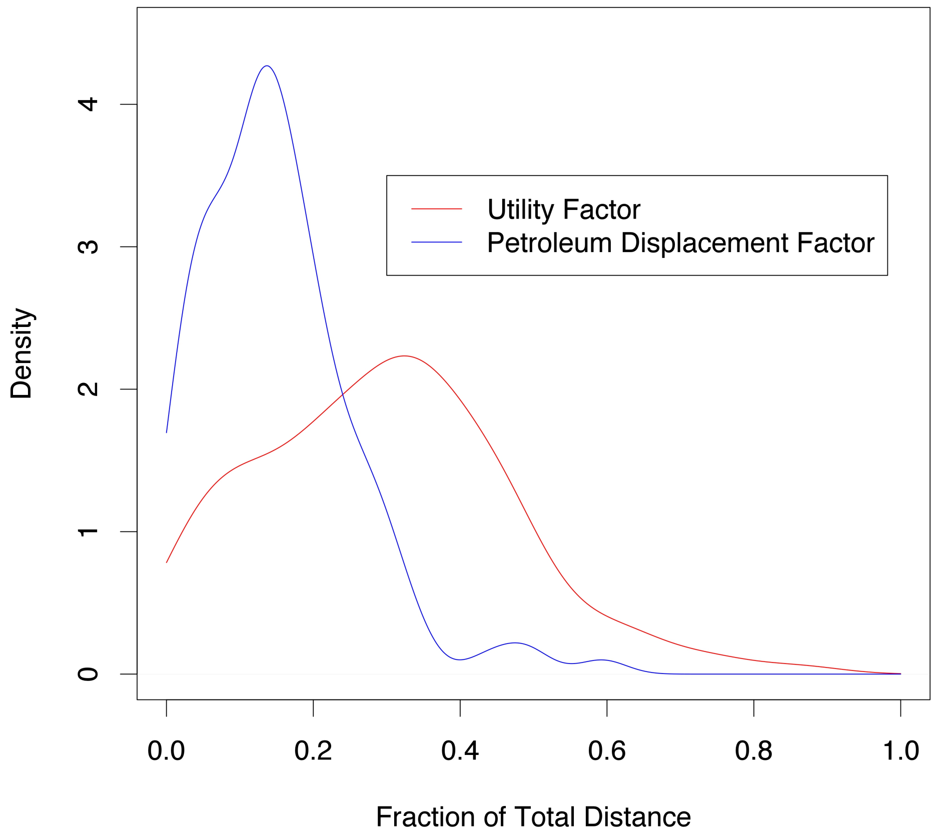

Figure 3-3: Density plot of the Utility Factor and Petroleum Displacement Factor.

3.4.3 Analysis of Trips, VMT and Charging Behavior

Trips and Daily Distance

To help assess the generalizability of the trial data, the distribution of daily travel distances from the trial was compared with that observed in the 2001 National Household Transportation Survey (NHTS). The distribution of daily distances observed in this trial is positively skewed, with more low-mileage days (Figure 3-4). The distribution of individual trip lengths in this trial was also positively skewed when compared with NHTS 2001.

The daily distance distribution from the trial was used to calculate the fleet utility factor

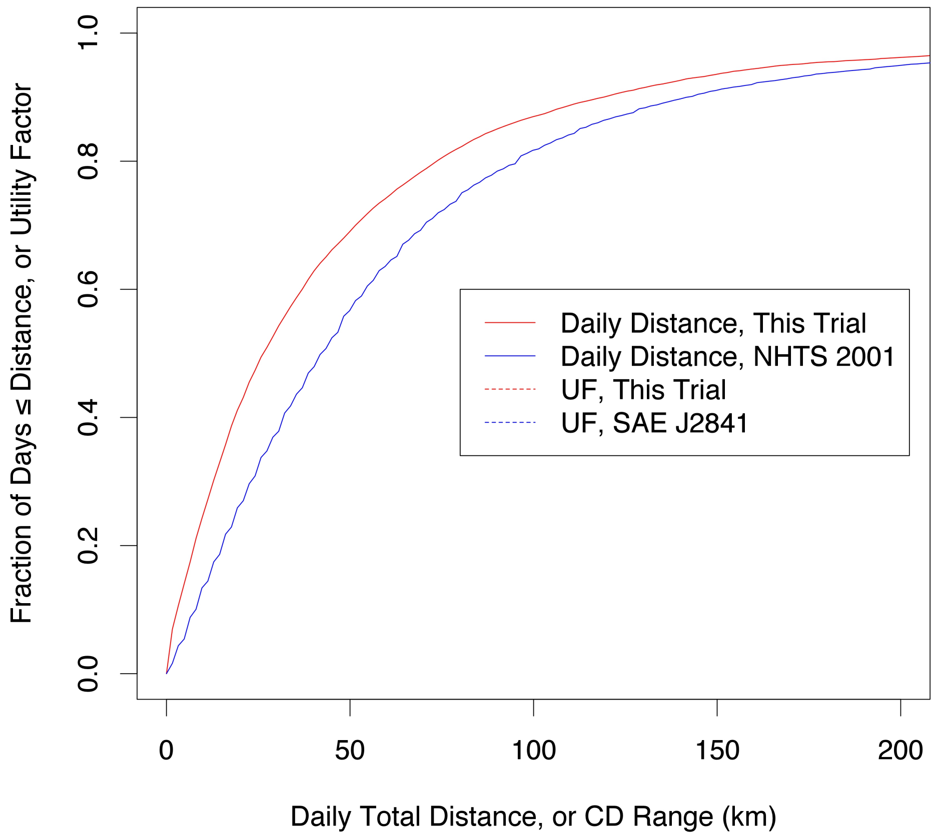

Figure 3-4: Cumulative distribution of daily travel distance and corresponding Utility Factor curves.

(FUF) using the method described in SAE standard J2841. (12) The resulting FUF curve is shown in FIGURE 4 along with the FUF curve from SAE J2841 (which is based on NHTS 2001). The FUF from this trial predicts that a fleet of PHEVs with a 28 km CD range would drive 42% of its overall distance in CD mode, assuming each vehicle is fully charged each night and only at night. In contrast, SAE J2841 predicts that such a fleet would cover 36% of its overall distance in CD mode. The larger UF predicted for this trial is a result of the bias toward shorter daily total distances in this trial when compared with NHTS 2001.

As discussed in Section 4, the 3 kWh of charge-depleting battery capacity is estimated to provide an equivalent electric range of 13.8 km. If the vehicles control strategy strictly preferred electric operation when in CD mode (i.e. if the vehicles did not have a blended mode), then the predicted PDF and the FUF for this 13.8 km range would be 20% (based on SAE J2841) or 23% (based on the daily distance distribution from this trial).

The actual UF averaged across all vehicles in this trial was 28.1%, and the average PDF was 13.7%. Both of these values are lower than the respective predictions based on SAE J2841 and on the theoretical FUF based on the trip distribution in this work. This appears to be due both to between-days variation in distance traveled, and to vehicles being charged less than once per day. (Lin and Greene, 2011)

Aggregate Charging Behavior

The trial reported here provides important insights into the charging habits of actual drivers who used PHEVs on a regular basis. Because the action of plugging the vehicle in was not actually logged, a charge was deemed to have occurred whenever the SoC between the end of one trip and the beginning of the next trip increased by at least 5 percentage points. The rationale for this condition is to avoid misclassifying as a charging event mere drifts in SoC resulting from changes in temperature or other conditions that can influence the measured SoC of the battery.

Figure 3-5 shows the distribution of charging start times across all charging events identified in the trial. Charging was assumed to start immediately at the end of the trip preceding the charging period, a reasonable assumption since the vehicles were not equipped with smart chargers. The most common times for the start of a charge were between 2:00 PM and 4:00 PM, and the initialization of new charging events fell off abruptly after 6:00 PM. This is somewhat surprising but may be due to the way in which the cars were deployed in this trial. If the cars were used mainly in corporate fleets, then the afternoon peak in charging may reflect them being plugged in after the last work-related trip rather than after an evening commute.

Approximately half of the charging events observed in this study involved nearly a full

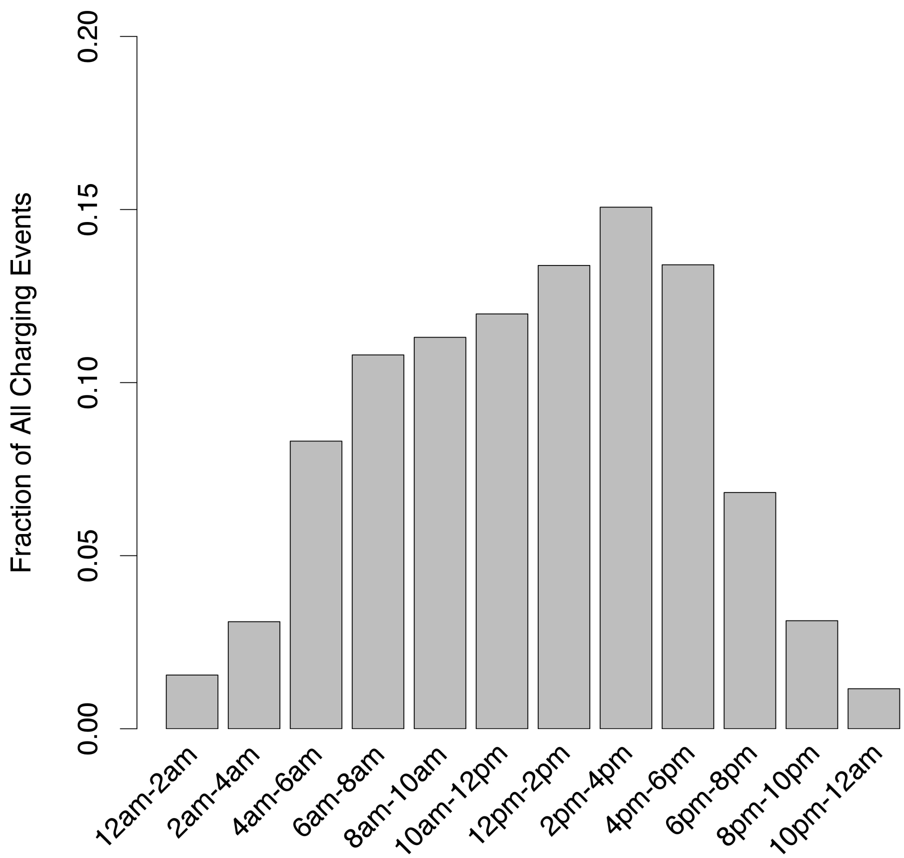

Figure 3-5: Distribution of charging start times.

charge (i.e. a charge of more than 2.5 kWh for the 3 kWh battery pack), indicating that the battery was fully discharged before charging and was allowed to charge fully before being unplugged. The other half of the charging events were uniformly distributed between 0 and 2.5 kWh.

More than 40% of the days on which a vehicle was driven in the trial saw no charging, while a similar number saw one charge. Just 10% saw two charges, and smaller numbers saw more. Thus nearly half the time, the PHEVs were not being charged even once on days when they were driven, even though the overall average rate was 0.75 charges per day per vehicle.

Differences in charging behavior were apparent between different vehicles. For example, one vehicle in the study was never charged, one charged an average of 1.7 times per day over the course of the trial, while the others were broadly dispersed between 0 and 1.25 charges per day. These observations tend to discredit the simplistic assumption that each vehicle charges once and only once each day, which is commonly employed in PHEV analyses. Not only was observed charging behavior different than commonly assumed, it also varied substantially between vehicles, suggesting differences in drivers preferences, incentives, or abilities to plug in.

3.4.4 Simulation Analysis

The gasoline and electricity consumption in this trial result from specific charging and driving behaviors applied to one vehicle design. This section reports on simulations of the electricity and gasoline consumption when small, one factor at a time (OFAT) changes are made to the vehicle design or charging logic, assuming that all trips remained the same. These scenarios are described below, and results are summarized in the subsequent section.

Scenarios

Each scenario used original trip cycles with new CD and CS percentages calculated based on the parameters of the scenario, such as battery size and new charging schedules. In all cases vehicle-level average values of CD and CS gasoline and electricity consumption were used.

Charged Once Daily

To simulate once-daily charging, SoC for each vehicle was reset to 100% at the start of each new calendar day. Battery capacity was unchanged. This assumption is present in existing UF definitions and embodies the common-sense assumption that PHEVs will be charged primarily overnight at home. Opportunistic Charging The Opportunistic Charging case simulates a scenario in which PHEV users charge their vehicles whenever they will be parked longer than a given time threshold. This assumes ubiquitous availability of 110V charging facilities. The study was conducted at nine time thresholds from 0 8 hours. Zero hours represents a limiting case in which the vehicle is plugged in immediately whenever switched off.

Fast Charging

The Fast Charging scenario simulates the same vehicle fleet, charged at exactly the same times and for the same duration but charged at higher rates. In theory this should primarily impact short charge cycles that are not capacity-limited. Alternative Battery Capacity Alternative Battery Capacity cases apply different battery pack sizes to the same vehicle, which permits a higher SoC and longer CD driving distances. This scenario considers the impact of battery capacity only through the mechanism of extending CD range. A larger pack size will also increase vehicle weight and decrease overall efficiency, but this effect is not considered.

Strictly Prefer EV mode

The Strictly Prefer EV Mode scenario models the effect of an alternative vehicle control strategy in which blended mode is not used and vehicles operate strictly on electric power until the battery pack is exhausted, then reverting to CS operation.

Scenario Outcomes

The results of these scenarios are summarized in 3.1 and compared against an actual UF of 28.1% and PDF of 13.7% With daily charging, the PDF improves to 18.2%, similar to the 19.1% achieved under opportunistic charging with an 8-hour threshold. With more aggressively opportunistic charging, improvement continues until a maximum of 28.3% is reached at zero hours, simulating performance if each vehicle is charged at every stop.

Table 3.1: Summarized results for Charging and Vehicle Design Scenarios

Scenario	kWh	Fuel (L)	Liters Saved[1]	PDF	UF
Actual Fleet Performance	21804	31328	4908	0.137	0.281
Charged Once Daily	30947	29534	6703	0.186	0.382
Opportunistic Charging (Charge if parked N hours)					
N = 8	31784	29351	6885	0.191	0.392
N = 7	32694	29148	7088	0.196	0.404
N = 6	33524	28960	7277	0.201	0.416
N = 5	34659	28707	7530	0.208	0.430
N = 4	36403	28315	7922	0.219	0.452
N = 3	38580	27839	8398	0.232	0.480
N = 2	41403	27234	9003	0.249	0.514
N = 1	44580	26559	9678	0.268	0.553
N = 0	47154	26012	10224	0.283	0.584
Fast Charging					
1 kWh / hour[2]	22451	31346	4891	0.137	0.280
2 kWh / hour	23055	31216	5020	0.140	0.288
4 kWh / hour	23310	31161	5075	0.142	0.291
8 kWh / hour	23399	31142	5094	0.142	0.292
Alternative Battery Capacity					
0 kWh	0	36237	0	0.000	0.000
1.5 kWh	13357	33330	2907	0.081	0.166
3 kWh**	22451	31346	4891	0.137	0.280
6 kWh	33534	28928	7309	0.204	0.420
12 kWh	44490	26571	9666	0.270	0.556
18 kWh	49991	25403	10833	0.303	0.624
24 kWh	53396	24684	11552	0.323	0.666
30 kWh	55851	24168	12068	0.337	0.696
Strictly Prefer EV mode	25644	30687	5550	0.155	0.155

As expected, increasing battery capacity would increase petroleum displacement, but with diminishing marginal returns. Quadrupling CD capacity would have increased PDF to approximately 27%, but increasing by a factor of ten would only have increased this to 33.7% (assuming that driving and charging patterns remained unchanged).

Other scenarios yielded only minimal changes. Deployment of rapid charging alone, even at rates up to 8x as fast as current 110V charging, left the PDF substantially unchanged at 14.2%, offsetting only 203 additional liters of gasoline. Eliminating blended mode via the

[1] Liters versus simulation of HEV
[2] Approximates actual performance

Strictly Prefer EV Mode scenario yielded a PDF of 15.5%, saving an additional 642 liters of gasoline.

This paper summarizes findings from the largest fully instrumented, real-world trial of prototype PHEV vehicles thus far. The results of this analysis highlight the important role blended mode plays in the operation of this type of PHEV and consequently the importance in distinguishing between UF and PDF. While this fleet returned an in-use UF of 28% the PDF was just half, at 14%, due to the use of gasoline in blended mode.

Scenario analysis indicates that deployment of fast charging with a small battery capacity brings little benefit, but ubiquitous use of conventional 110V chargers more than doubles UF and PDF. Increasing battery capacity decreases gasoline consumption, but most benefits are realized by increasing effective electric range to 55 km. Charging the vehicles in this study every time they are stopped for 3 hours or more would increase PDF to 23%, the same level expected from increasing electric range to 41km. When charging, more than half of all charge events result in a charge delivered of at least 2.5 kWh. On days that they were used, 40% of vehicles were charged only once and 40% were not charged at all.

3.4.5 Toyota Prius PHEV Case Study Discussion

Under the right conditions, PHEVs with a relatively small battery can achieve significant reductions in petroleum usage reductions of up to 60% below a comparable HEV were observed in this work. However, average petroleum displacement was 14% (and was below 5% for one-sixth of vehicles). For the driving and charging patterns observed in this work, increasing CD range and increasing charging frequency can raise PDF by similar amounts. Policymakers should bear in mind that petroleum displacement does not scale linearly with battery size or charging frequency, and that interactions between battery size and charging patterns may be significant in the real world. This suggests that if petroleum displacement is the goal, then PHEV policies should carefully weigh the costs of enabling more frequent charging against the costs of subsidizing larger batteries. Discrete choice analysis revealed that charging probability was highest when stopping for more than 3 hours, suggesting that

Level 1 or 2 charging infrastructure development may be most effective if it targets locations where stops at least this length are common.

3.5 Case Study: The BMW ActiveE trial

The BMW ActiveE was the second of two field trails as part of the development of its production electric vehicle, the i3. The first trial, the MINI E, was a three year study of approximately 600 converted MINI vehicles in the United States, Mexico, Brazil, Germany, the United Kingdom, France, China, Hong Kong, Japan and South Africa. The second EV trial, the ActiveE is the subject of the work presented here. The ActiveE is also an EV conversion but based on the BMW 1 Series coupe with approximately 1000 converted vehicles. This study analyzes the energy consumption of approximately 700 BMW ActiveE vehicles in the United States, used by both private customers and carsharing users. This study updates earlier work by Rodgers et al. (2014) with additional metrics for shared and private vehicle energy consumption.

3.5.1 The DriveNow fleet

DriveNow is a one-way carsharing service which operates in Germany, Austria, the United Kingdom and the United States. Users pickup vehicles and drop off vehicles in different locations around a metropolitan area, and are billed per minute for their usage. In the United States at present, the DriveNow service operates in the San Francisco area. While the DriveNow service uses a variety of different vehicles worldwide, the service in San Francisco uses exclusively the ActiveE. The fleet currently consists of 150 ActiveE vehicles, 70 of which were initially deployed in 2012, and an additional 80 vehicles which were added during the month of June, 2014.

3.5.2 The BMW ActiveE and Data

This study is based on the energy consumption of a fleet of 700 BMW ActiveE electric vehicles operated in private and commercial service for approximately 2.5 calendar years from mid-2012 to early 2015. Aggregated vehicle data is transferred from each vehicle to a central server periodically. Data transferred includes vehicle use, charging behavior, energy consumption figures, and distance traveled, along with diagnostic information. The data is aggregated on-board by electronic control units (e.g. the electric motor control unit, the high voltage battery management unit) and is transmitted via a built-in GSM card. Periodic transmission of the data offers several advantages: it reduces the need for constant connectivity, and allows for interruptions in cellular service. However, under conditions of extended cellular service disruption (especially in rural areas, or in parking garages) certain details cannot be transmitted before they are overwritten with updated values. In general, an average data quality and completeness of 80-95% can be assumed.

The dataset consists of 199,504 unique drive-charge observations from the United States once scrubbed of missing data and duplicate observations. Drive-charge cycles were recorded each time the vehicle was charged, so individual observations include one or more drive cycles and a single charge cycle. Drive cycles of 0 and 1 km were removed from the analysis due to large rounding error and incomplete data.

Figure 3-6 presents a histogram of travel distances for drive-charge observations in the data set. Nearly all travel distance are under 160km (100 miles), in keeping with the electric driving range of the vehicle. Among private vehicle observations, shorter drive-charge cycles are the most common, while drive-charge cycles for shared vehicles are somewhat longer. In private vehicles, 44% of observation distances are under 40km, while only 38% of shared vehicle observation distances are under 40km.

There are a number of possible explanations for differences in drive-charge observations in shared and private vehicles. One possible explanation is simply selection bias: shared vehicle drivers might use other travel modes for shorter trips and only use shared vehicles for longer trips, or they might "save up" errands and trip chain when they do rent a shared vehicle. Another explanation might have less to do with travel behavior and more to do with charging behavior: shared vehicle users might forget to plug-in vehicles. Alternatively, parking facilities for shared vehicles might not have adequate recharging facilities and vehicles might be used for multiple rentals between recharges. Since we only observe a single timestamp per observation, we are unable to determine which of these possibilities are most likely.

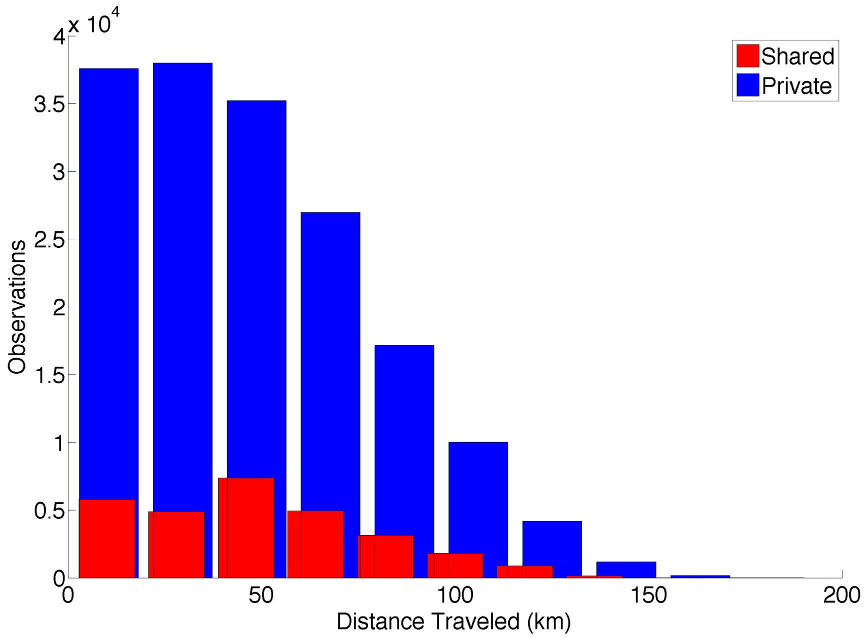

Figure 3-6: Histogram of travel distances in shared and private drive-charge cycles from the dataset. A greater fraction of trips are moderate distances (50-60km) in shared than private vehicles.

Data was gathered on the energy consumption, time, and distance of each observation, and energy data was split into three categories: drive energy, regenerative energy, and accessory energy. The definitions of these recorded values are explained below.

Drive The energy used by the motor to propel the vehicle. When this value is positive, energy current flow is measured as Drive Energy.

Regenerative When current flow to the motor is negative, energy from the motor is stored as regenerative energy and measured separately.

Auxilliary Energy consumed by accessories such as heating, air conditioning, lights, entertainment system and other onboard equipment.

Total Total energy is the sum of drive energy and auxiliary energy, less regenerative energy (multiplied by an efficiency factor).

Regenerative braking is not 100% efficient, and a portion of the kinetic energy of the vehicle is lost during the conversion back to electric energy and storage in the vehicle battery. Since regenerative energy measurements were taken at the motor, these measurements do not include the losses associated with the regenerative braking system, so measurements were scaled by an estimated regenerative braking efficiency factor of 0.5. This process is explained in greater detail in Rodgers et al. (2014).

This analysis particularly focuses on trip-to-trip variation, sources of energy consumption, and differences between privately owned and shared vehicles. These are discussed in greater detail in the following sections.

3.5.3 Sources of Energy Consumption

Unlike the PHEV in the Section 3.4, the ActiveE is a battery electric vehicle (BEV) and all energy to move the vehicle, heat and cool the vehicle, and power accessories comes from the battery pack, which is charged from grid electricity. Heating and air conditioning represent the largest sources of power consumption among auxiliary devices, each consuming several kilowatts at full power.

While some aspects of energy consumption are a product of vehicle design, there are a large number of factors under the control of the vehicle operator that impact energy consumption, as shown in Figure 3-7. In this Figure, purple boxes represent driver decisions and controls, red boxes represent areas of energy loss, and green boxes are forms of energy storage.

Figure 3-8 shows a distribution of the split of energy consumption between drive energy and auxiliary energy in each drive-charge observation. In virtually all drives, at least 10-15% of energy was used for accessories in the vehicle, and the average auxiliary energy

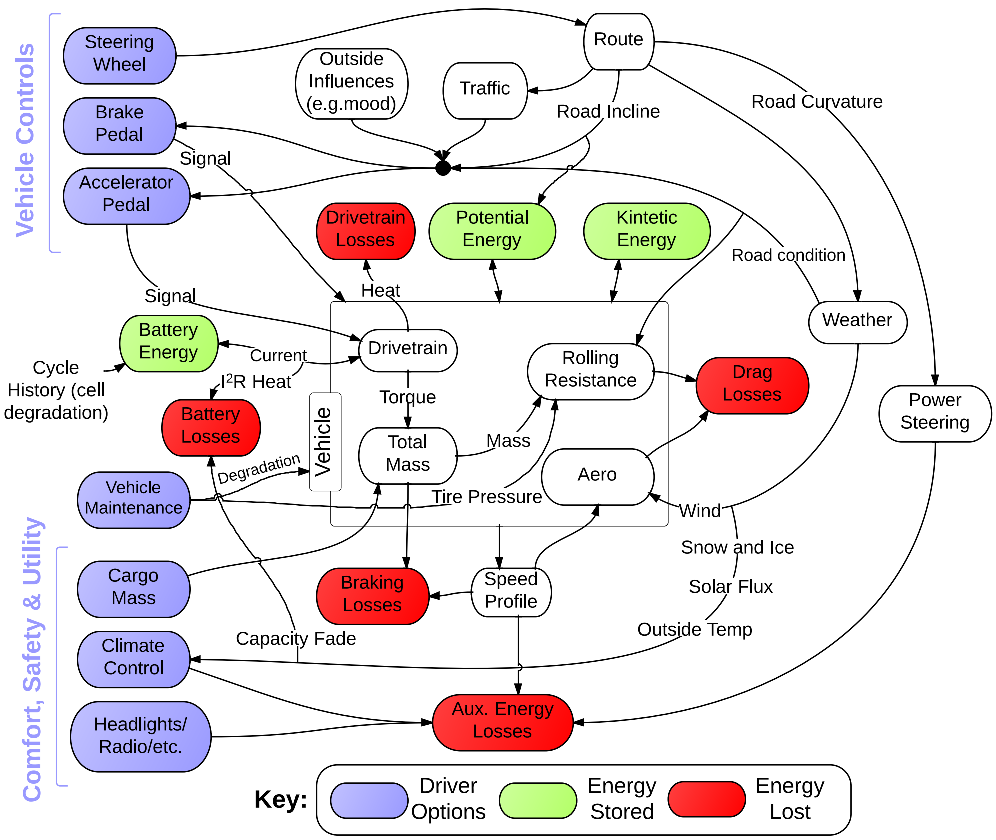

Figure 3-7: Sources of energy consumption and losses in Plug-in Vehicles from Rodgers et al. (2014).

consumption was 22.7%. In a few observations auxiliary energy consumption accounted for more than half the energy consumption.

3.5.4 Variability of Energy Consumption

In referring back to Figure 3-7 many sources of energy consumption would plausibly be somewhat consistent from observation to observation: a vehicle in the same place, driven by the same person, and in the same climate would seem to minimize a number of sources of variability: driving style, climate controls, and vehicle settings would be similar over the short term. However, Figure 3-9 and Figure 3-10 indicate that this belief is not entirely true.

Figure 3-9 shows a distribution in the *difference* in energy consumption between trip n and trip $n+1$. Observations where the energy consumption differed by 25% or more between

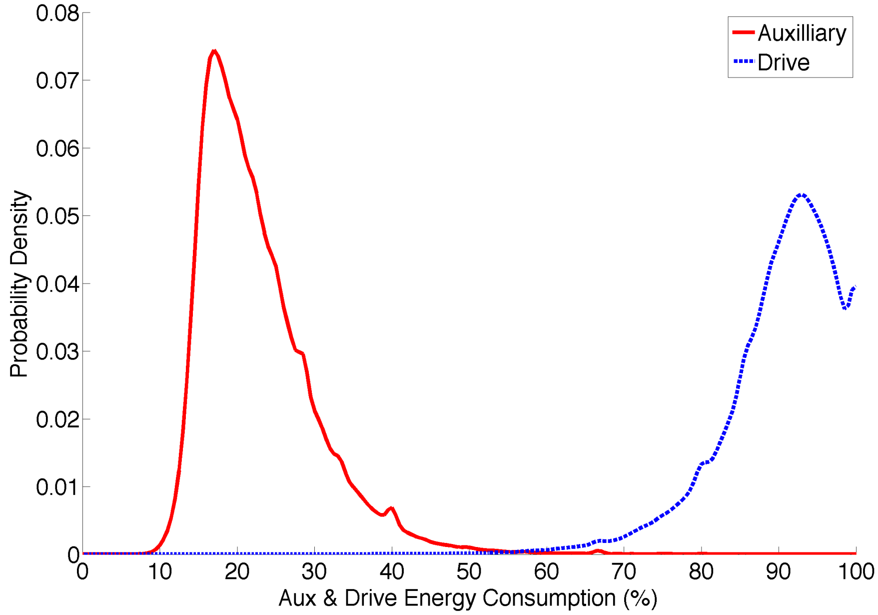

Figure 3-8: Fraction of energy used for Drive (propulsion of the vehicle) and Auxilliary (accessories & climate).

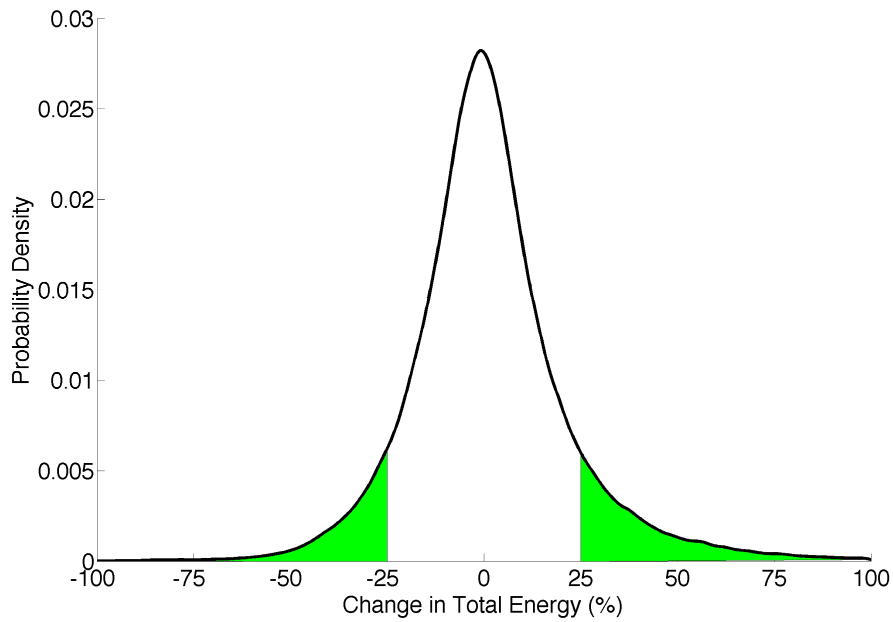

Figure 3-9: Difference in energy consumption rate from trip n to trip $n + 1$ for total energy consumption.

successive drive-charge events, indicated by the green shaded areas, accounted for more than 17% of trips. The high variability in energy consumption from one drive-charge to the next presents a number of challenges. Most fundamentally, variability in energy consumption combined with a fixed battery capacity means that the range of the vehicle varies. As a result, customers may have a difficult time planning their driving in a way that they will be able to accomplish their travel needs without needing to recharge. Similarly, variability in energy consumption presents a challenge to automotive manufacturers who wish to present a realistic distance to empty (DTE) estimate to vehicle operators.

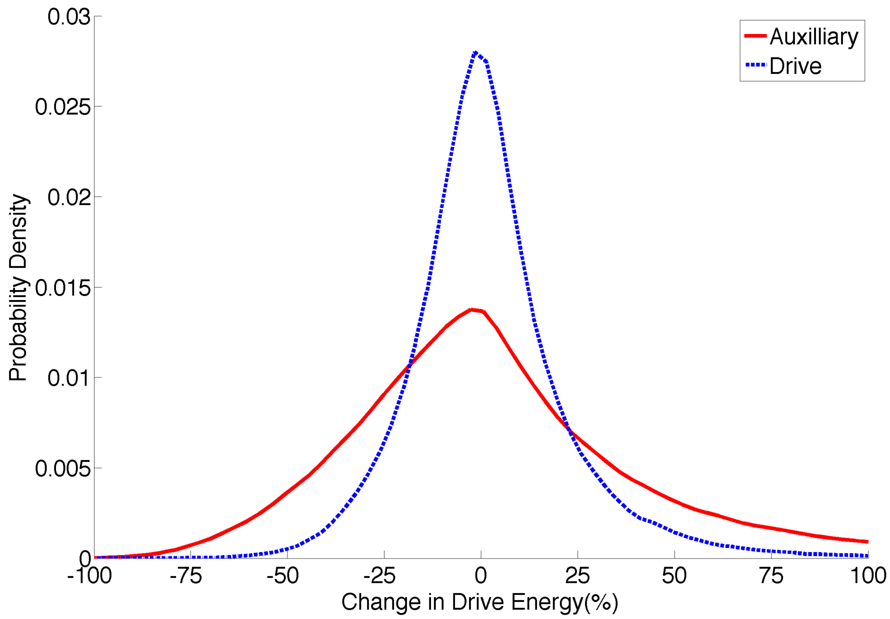

Figure 3-10: Difference in energy consumption rate from trip n to trip $n + 1$ for Drive (propulsion of the vehicle) and Auxilliary (accessories & climate).

As in Figure 3-9, Figure 3-10 shows the variation in energy consumption from trip n to trip $n + 1$, but separated into Drive and Auxiliary energy. Figure 3-10 indicates that variability between successive observations is substantially higher for Auxiliary loads than for the energy used to propel the vehicle, and models of total energy consumption are more highly correlated with accessory energy consumption than drive energy.

3.5.5 Energy Consumption in the Private and Shared Fleet

This study includes vehicles that were used by private owners and those which were used by carsharing users for individual trips. We analyzed the energy consumption of these vehicles separately to understand whether shared vehicles consumed systematically more or less energy than their private counterparts.

Figure 3-11, Figure 3-13 and Figure 3-12 show the differences in energy consumption between shared and private vehicles. Energy consumption for Drive energy, Auxiliary energy and Total energy are all higher for shared vehicles than for private vehicles. Mean comfort energy consumption for shared vehicles is approximately 6.8% higher for shared vehicles, while mean drive energy consumption is approximately 5.2% higher for shared vehicles.

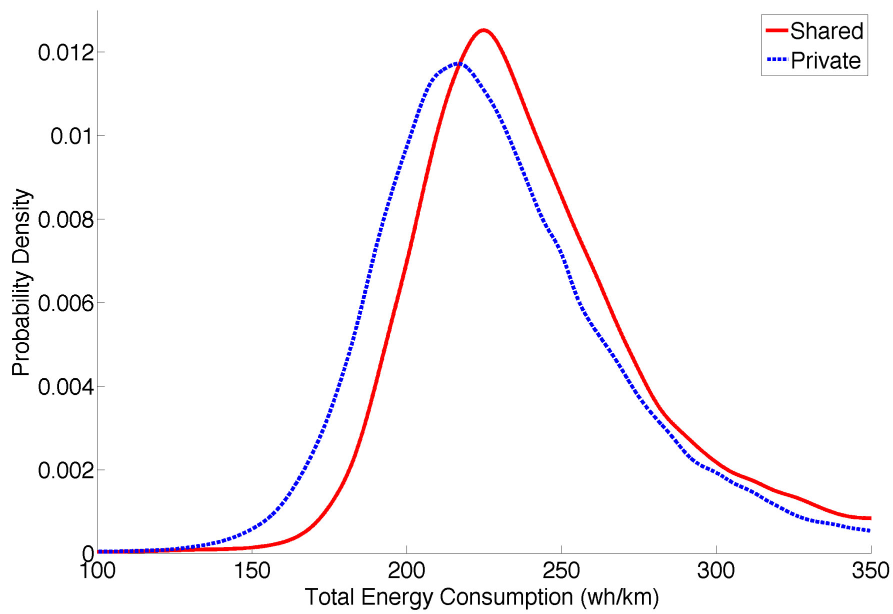

Figure 3-11: Energy consumption (Wh/km) is higher for shared vehicles.

There are a number of plausible explanations for the higher drive energy consumption in shared vehicles. The first could be simple unfamiliarity with the vehicle controls. The ActiveE incorporates regenerative braking when the gas pedal is released, while the brake pedal activates the friction brakes. Users who are unfamiliar with this configuration may use the friction brakes unnecessarily, while vehicle owners learn to use the regenerative braking

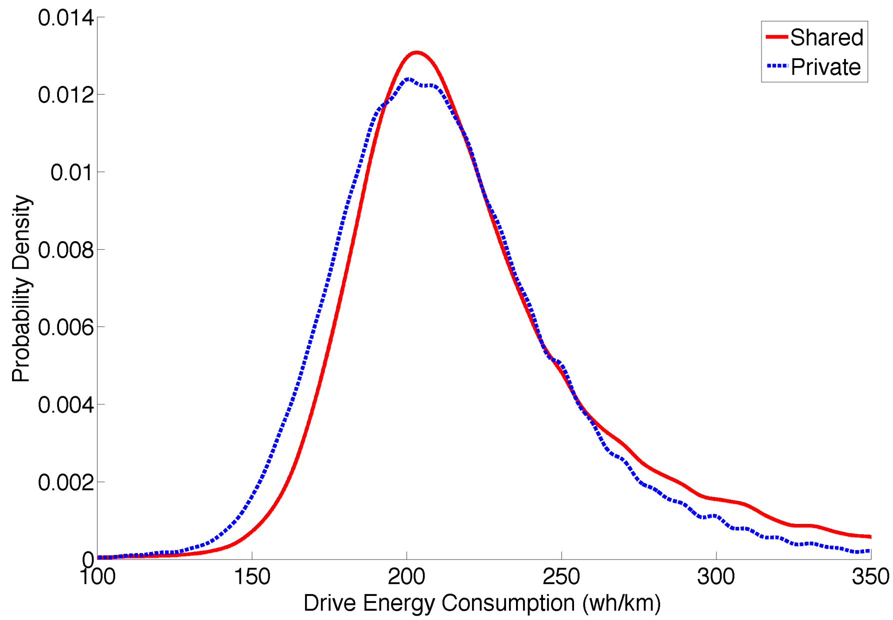

Figure 3-12: Average drive energy consumption is 5.2% higher for shared vehicles than for private vehicles.

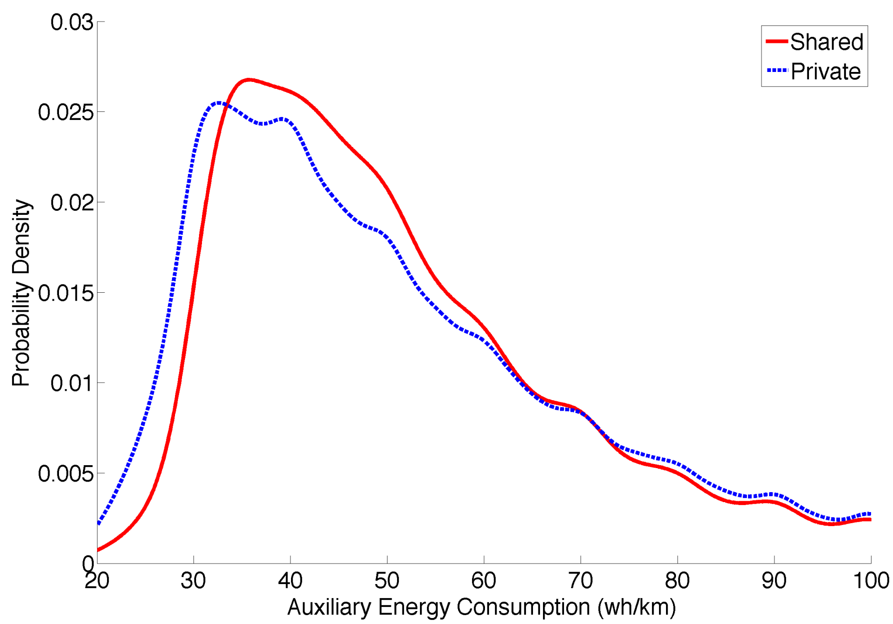

Figure 3-13: Average auxiliary energy consumption is 6.8% higher for shared vehicles than for private vehicles.

more effectively. An alternative explanation has to do with the type of driving that shared vehicles perform: they are located in primarily dense, urban environments. While we do not have sufficient information on the drive cycles to be able to perform a comprehensive energy analysis, driving in heavier traffic would generally lead to higher energy consumption. A final plausible explanation is that users may simply suffer from "rental car syndrome" in which they drive the car aggressively since they do not pay directly for energy use or for wear and tear to the vehicle.

Similarly, there are a number of plausible explanations for higher energy Auxiliary (Comfort) energy consumption in shared vehicles. If comfort loads are approximately constant, then lower average speeds in the shared vehicle service would lead to higher comfort energy consumption per unit distance, so differences in energy consumption may be a consequence of different driving environments. Alternatively, as with drive energy, vehicle accessories can be used in a more or less efficient manner. The ActiveE includes a number of energy-saving modes which reduce fan speeds, air conditioning and heating loads. If shared vehicle drivers are unfamiliar with these modes, they would be unable to operate vehicle accessories as efficiently as owners.

3.6 BMW ActiveE Case Study Discussion

This analysis highlights a number of interesting details from approximately 200,000 BEV drive cycles. As with the Toyota Prius PHEV discussed in Section 3.4, energy consumption per unit distance varied widely – by nearly a factor of 3 from best-case to worst-case observations. High variability between successive observations of the same vehicle introduces challenges for BEV producers and operators alike: vehicle manufacturers will need to develop more sophisticated algorithms to predict future energy consumption and range, while BEV owners will need to become more attuned to the capabilities of their vehicles to avoid exhausting the battery pack. Work by Pichelmann et al. (2013) suggests that consumer are likely doing this over the lifetime of their vehicle.

User adaptation to the capabilities of PEVs may be more challenging for users in the shared vehicle environment who drive less frequently and may drive different vehicle types for each trip. Data from these observations indicate that energy consumption is higher for shared vehicles. With additional data on the length, speed and driving environment it would be possible to determine whether this higher energy consumption is a function of user behavior or simply of the environment in which the vehicles are used.

Chapter 4

Deployment and Utilization of Plug-in Electric Vehicles in Round-trip Carsharing Systems

4.1 Introduction

The development of Plug-In Vehicles (PEVs) holds considerable premise to mitigate the environmental footprint of the automobile sector by displacing gasoline consumption to electricity consumption. However, the adoption of PEVs, including Battery Electric Vehicles (BEVs, or EVs) that are powered exclusively by electricity, and Plug-in Hybrid Vehicles (PHEVs) that are powered alternatively by electricity or fuel, has been historically slow in the private market. (EPA, 2014) In contrast, the use of PEVs in carsharing dates back as far as the 1970s (Bendixson and Richards, 1976) and is expected to increase in the next 5 to 10 years as growth of carsharing businesses, technological developments and policy incentives for adoption are expected. (Shaheen and Chan, 2015) At the same time, the adoption of PEVs involves typically higher acquisition costs than comparable gasoline vehicles, and their usage is constrained by technological factors such as range limitations and the need for sufficient recharging time between trips. There is thus a strong need to understand how PEVs can be deployed and utilized in a carsharing system in order to improve the sustainability of automobile systems while satisfying technological constraints, carsharing business objectives and consumer demand.

This work develops a comprehensive approach to PEV deployment and vehicle utilization in round-trip carsharing systems. PEV deployment refers to the introduction of PEVs in carsharing fleets in addition or replacement of gasoline vehicles. Vehicle utilization refers to the assignment of trips to vehicles as consumer demand realizes, i.e., as trip reservations are made. The former is a strategic decision that typically ranges several months or years, while the latter is a tactical decision that is made on a daily or hourly basis by carsharing operators. This work provides decision-making support to optimize and assess the deployment of PEVs in a carsharing fleet of vehicles, while anticipating how PEVs and gasoline vehicles can be utilized as consumer demand materializes. Before presenting its contributions in Section 4.1.3, we detail and motivate the scope of this work (Section 4.1.1) and review the related literature (Section 4.1.2).

4.1.1 Scope and Motivation

This work analyzes the potential gasoline consumption reduction associated with the near-term deployment and utilization of PEVs in round-trip carsharing systems. Carsharing is a business in which users can rent cars for short time periods (as little as one hour in 30 minute increments), providing access to vehicles without the financing, maintenance, parking, and registration costs typically associated with car ownership in a metropolitan area. In a round-trip carsharing model, cars are grouped in clusters called "pods", and users rent a vehicle and return it to the same pod upon completion of the rental. This is the most prevalent system in the United States, although some operators also offer one-way reservations.

The use of PEVs in round-trip carsharing systems eases a number of the challenges associated with PEV adoption in the private market. First, the high vehicle utilization in carsharing systems eases the upfront price premium currently associated with PEVs. Second, carsharing providers are both responsible for the costs of investing in charging infrastructure *and* able to capture the benefits associated with the reduction in gasoline consumption resulting from the utilization of PEVs. Thus the deployment of PEVs in carsharing fleets alleviates the chicken-and-egg problem associated with PEV deployment in private car fleets, in which station operators may be reluctant to install infrastructure due to low utilization and vehicle buyers may be reluctant to invest in PEVs due to the small number of charging stations. (Kley et al., 2011; Dijk et al., 2013; Green et al., 2014; Shaheen and Chan, 2015) Finally, round-trip carsharing models ensure that PEVs will always return to a charging station, thus enabling recharging of any idle vehicle between successive trips through the installation of a single charger.

At the same time, there are also inherent technical challenges associated with the use of electric vehicles. The most obvious of these challenges is their limited range, which imposes a hard constraint on the distance that can be traveled with an EV. While Green et al. (2014) note that carsharing organizations are suited to dense areas with generally lower daily distances traveled, it is not clear that shared vehicles will actually travel shorter distances as users may elect to use shared vehicles for longer trips for which other public transit modes are unsuitable or unavailable. In turn, the utilization, hence the deployment, of EVs may be constrained by range availability. Note that, although PHEVs do not face strict range constraints, the gasoline savings that can be captured from PHEV deployment is also limited by their electric ranges, which are in general significantly smaller than those of comparable EVs. Additionally, PEVs require sufficient charging time between trips. On the fastest available chargers, EVs can recharge in as little as 15-20 minutes, but full recharging on more common 240V chargers takes several hours. The time needed to recharge EVs inherently conflicts with the business model of carsharing services, which earn revenue only when vehicles are in use. This conflict presents carsharing operators with a choice between investing in rapid recharging infrastructure and enforcing longer downtime between trips.

This discussion underscores the interdependencies between strategic decisions related to PEV deployment and tactical decisions regarding vehicle utilization. On the one hand, vehicle utilization on any given day of operations depends on the vehicle fleet available on that day. Conversely, PEV deployment depends on the way vehicles can be utilized to operate trips, and on the resulting gasoline savings and other business implications. These interdependencies are not accounted for in the current legislation aimed to incentivize PEV deployment. Most notably, the Zero Emissions Vehicle (ZEV) mandate (California Code of Regulations (CCR) 1962.1 - 1962.3) in California and seven collaborating states requires vehicle manufacturers to produce and sell an increasing percentage of hydrogen or electric vehicles and are granted bonus credits for introducing ZEVs in carsharing or mobility services. (Shaheen et al., 2002) But these credits are not tied to usage or demand patterns in carsharing systems, which may result in over-deployment of PEVs in some locations and under-deployment of PEVs in others. As a result, there is a need to understand how PEVs can be utilized in carsharing systems and how it informs PEV deployment in carsharing fleets.

This work focuses on Electric Vehicles, and this choice is motivated by three reasons. First, carsharing enables drivers to choose their car for a specific trip, whose distance is at least approximately known. In other words, any driver can select an EV for a given, short trip, and a gasoline vehicle for another, longer trip. This contrasts with the private car market, where a given vehicle needs to serve most of the travel needs for a single driver, whose trip distances may exhibit significant variability. As a result, the potential for EV adoption is significantly higher in carsharing systems than in the private market. A second reason for the focus on EVs is that, unlike PHEVs, their in-use efficiency does not depend on recharging behavior. PHEVs offer substantial fuel savings in theory, but in practice if they are not recharged they can still be powered by gasoline, eliminating any benefit from their introduction. The impact of PHEVs is thus sensitive to the wide variability in recharging behavior (Zoepf et al., 2013). Third, policy incentives (including the ZEV program) favor the use of EVs to PHEVs. Despite the focus on EVs in this work, the frameworks and models developed here can be easily extended to PHEVs.

4.1.2 Related Research

An extensive literature has investigated the deployment of PEVs in private car markets and their potential for mitigating greenhouse gas emissions in the transportation sector. On the demand side, several studies have characterized and quantified demand for privately owned electric vehicles and plug-in hybrid vehicles, using a combination of surveys, trials, and statistical models. (Golob and Gould, 1998; Glerum et al., 2012; Tamor et al., 2015) On the supply side, demand for PEVs has implications on the design of PEVs, including range and charging capabilities (Pearre et al., 2011; Lin, 2014) and the development and deployment of public charging infrastructure. (Hoke et al., 2011; Chen et al., 2013; He et al., 2013; Dong et al., 2014)

At the operational level, research aims to quantify the impact of PEV deployment on energy consumption. Several studies have analyzed the feasibility and potential benefits of introducing electric vehicles into private fleets, using national survey-based trip distance information and, recently, more detailed multi-day GPS data. (Lin and Greene, 2011; Khan and Kockelman, 2012; Jing Dong, 2014; Chang, 2015) Realized reductions in gasoline consumption depend, however, on several operational factors, including the heterogeneity of consumer charging behavior (Vyas et al., 2009; Dong and Lin, 2012; Zoepf et al., 2013), user preferences regarding vehicle technologies (Le Vine et al., 2014b; Zoepf and Keith, 2015) and vehicle routing (Schneider et al., 2014). In turn, an active body of work aims to better understand how PEVs are operated and to develop decision-making models to improve such operations.

Moving to the large and rapidly growing body of literature on shared vehicle systems, one of the major operational challenges involves the operations of vehicles in one-way carsharing systems. A number of studies have developed simulation and optimization techniques to balance and reposition vehicles in response to consumer demand (Marouf et al., 2014; Kaspi et al., 2014; Repoux et al., 2015; Boyac et al., 2015). At a more macroscopic level, the fact that the benefits of PEV deployment may be larger and easier to capture in carsharing systems than in private car markets has motivated recent research on PEV deployment in carsharing systems. Spieser et al. (2014) estimated the number of autonomous, shared, electric vehicles necessary to serve a population in Singapore. Such systems of shared and autonomous vehicles have been found to result in more vehicle miles traveled, fewer cars, and lower emissions (Fagnant and Kockelman, 2014).

However, little research has studied the near-term possibility of converting portions of existing fleets of carsharing systems from gasoline to electric power. The quantification of the benefits (gasoline savings) and the costs (operational constraints and larger acquisition costs) of PEVs in carsharing fleets is rendered difficult by the interactions between PEV deployment at the strategic level, and PEV utilization, at the tactical level. This is the focus of this work.

4.1.3 Approach and Contributions

This work formulates and solves the problems of PEV deployment and vehicle utilization in a round-trip carsharing system. The vehicle utilization problem takes as inputs a set of trips to be completed in each pod in a given period of time (e.g., one month) and the number of gasoline vehicles and PEVs available in each pod. It involves assigning trips to vehicles with the objective of minimizing gasoline consumption in each pod, subject to service constraints and technological constraints (e.g., range limitations, charge requirements). We refer to this problem as "Vehicle-Trip Assignment (VTA)". The PEV deployment problem takes as inputs a set of trips to be completed in all pods in a given period of time. It involves determining the number of gasoline vehicles and the number of PEVs to allocate to each pod, with the objective of minimizing overall gasoline consumption. We refer to this problem as "Vehicle-Pod Allocation" (VPA).

For each of these two problems, we develop an optimization approach and heuristic-based simulations (see Table 4.1). For the VTA problem, our optimization approach provides a computationally efficient way to quantify the potential gasoline savings associated with deploying a PEV on any given pod, assuming an ideal assignment of trips to vehicles. In contrast, our heuristic simulates the assignment of trips to vehicles by order of reservation.

This heuristic prioritizes the assignment of trips to EVs, under feasibility constraints, but may ultimately result in a sub-optimal assignment given *ex post* information. In turn, the VTA optimization provides an upper bound of the gasoline savings that can be captured with a given PEV (e.g., assuming that the carsharing operator could forecast demand with perfect accuracy for the entire period considered and assign trips to vehicles in an optimal way) while the VTA simulation provides a more realistic estimate of these gasoline savings that can be achieved in a carsharing system where reservations are revealed dynamically. In combination, VTA optimization and simulation bound our expectations of the benefit that could be achieved by PEV utilization: a best-case estimate via optimization, and conservative estimate via simulation.

For the VPA problem, we develop a simple optimization procedure that uses the VTA results, i.e., the benefits associated with the utilization of PEVs in any pod, in order to optimize the allocation of vehicles to pods in a way that minimizes gasoline consumption. We then develop alternative heuristics that represent the kind of "reasonably smart" PEV deployment strategies one might apply in practice, in the absence of an advanced VTA model such as the ones developed in this work. These heuristics are, instead, based on aggregate metrics and summary statistics (e.g., number of trips per pod, distribution of traveled distance per pod, etc.). The objectives of these VPA heuristics is twofold. First, they provide simple guidelines to inform PEV deployment in carsharing systems. Second, the comparison the outcomes under the VPA heuristics and the VPA optimization quantifies the benefits associated with integrating tactical decisions regarding PEV utilization into strategic decisions regarding PEV deployment.

The contributions of this work fall into four categories:

- *Developing the first optimization approach of PEV utilization in a round-trip carsharing system.* First, we develop a baseline Integer Program that minimizes fleet size, assuming an all-gasoline fleet. Then, we introduce Mixed Integer Programs for VTA optimization under a mixed fleet comprising gasoline vehicles and EVs.

- *Developing realistic simulation models of PEV utilization in a round-trip carsharing*

Table 4.1: Summary of optimization and simulation strategies for the VTA and VPA problems

		Vehicle-Trip Assignment (VTA)	
		Optimized	Simulated
Vehicle-Pod Allocation (VPA)	Optimized	Best case scenario: (i) perfect *ex post* knowledge about list of trips *and* reservations; (ii) optimal tactical VTA to minimize gasoline consumption for a given fleet; (iii) optimal strategic VPA to yield highest savings for any PEV "budget".	Informed VPA scenario: (i) perfect *ex post* knowledge about list of trips; (ii) realistic tactical VTA as a carsharing operator might behave, given order of reservations; (iii) optimal strategic VPA to yield highest savings for any PEV "budget".
	Simulated	Informed VTA scenario: (i) perfect *ex post* knowledge about list of trips *and* reservations; (ii) optimal tactical VTA to minimize gasoline consumption for a given fleet; (iii) heuristic strategic VPA to maximize anticipated gasoline savings using aggregate metrics.	Baseline scenario: (i) no data requirements; (ii) realistic tactical VTA as a carsharing operator might behave; (iii) heuristic strategic VPA to maximize anticipated gasoline savings using aggregate metrics.

system, using actual reservation data. We develop a simple heuristic that assigns trips to vehicles by order of reservation. The heuristic prioritizes the assignment of trips to PEVs, under feasibility constraints. We provide an extension of this heuristic that accounts for the heterogeneity of users' technological preferences.

- *Applying the models to a major carsharing system in the United States to characterize vehicle utilization and quantify gasoline savings under various fleeting scenarios.* We use data from Zipcar, a major carsharing service provider, related to operations during January, 2012 in Boston, MA. We quantify the gasoline savings (or the petroleum displaced by electricity) resulting from transferring vehicle miles traveled (VMT) from gasoline vehicles to PEVs. We show that (i) the deployment of PEVs can yield very significant gasoline savings; (ii) gasoline savings increase with the number of PEVs deployed, but marginal returns are non-increasing; (iii) gasoline savings are very sensitive to vehicle ranges; and (iv) gasoline savings are not very sensitive to charging rates.

- *Integrating tactical models of PEV utilization into strategic models of PEV deployment.* We compare informed PEV deployment, using the VTA results, to heuristic PEV deployment based on aggregate metrics. We show that educated PEV deployment and utilization can improve the operator's profitability. We show that the value of information regarding how PEVs can be utilized at the tactical level can be very significant when it comes to making decisions regarding PEV deployment. In turn, the comprehensive approach to PEV deployment and utilization developed in this work informs major carsharing operators' decisions on PEV deployment and public policy on PEV adoption incentivization.

The remainder of this work is organized as follows. Section 4.2 presents Zipcar's operations data and the types of gasoline and plug-in vehicles considered in this work. Section 4.3 formulates our modeling approach to VTA optimization, and shows implementation results for the Zipcar carsharing system in Boston. Section 4.4 presents our approach to VTA simulation and implements it in the Zipcar Boston system. We also compare the results of the VTA optimization and the VTA simulation. Section 4.5 discusses the implications for system-wide deployment of PEVs in carsharing systems by comparing the VPA optimization approach to heuristics. Section 4.6 concludes and discusses extensions of the work.

4.2 Data and Parametrization

This project is based on trip-level data from Zipcar, currently the largest carsharing provider in North America. Zipcar operates primarily round-trip carsharing. It has also recently launched one-way services in some markets, but these are left out of scope of this work. The service is available in most larger cities in the United States, and nearly all vehicles in the service are powered by gasoline, although the operator has experimented with a number of plug-in hybrid and electric vehicles over the past 2-3 years. In this work, we focus on the company's operations in the city of Boston, MA. Zipcar vehicles are parked in public and private garages throughout the city, and grouped in "pods" which may comprise as few as a single car or more than a dozen cars in a single location.

The data used for this analysis consist of 49,518 trips taken during the month of January 2012 in Boston across 330 pods. The number of reservations at individual pods ranged from a minimum of 10 to a maximum of 777 during the month considered. Limiting the scope of the problem to a single month of data ensures computational tractability of the VTA models while allowing for daily and weekly variation in demand. For each trip we observe the unique pod ID where the trip originated and ended, the time the reservation was made, the start time and end time of the reservation, and the total distance traveled.

We note that this data represents an observation of censored demand, and that there may have been additional users who wished to access a vehicle at these locations who were unable to do so because no vehicles were available. Other users with travel flexibility may have shifted their reservations in time by a few hours or a few days, or shifted their demand in space to nearby pods. Both of these issues represent spillover in demand. Additional representations of demand have not been explicitly modeled.

Figure 4-1 shows the cumulative distribution function (cdf) of the trips in the dataset considered (in red), compared to the cdf of the 2009 National Household Travel Survey (NHTS) trip distances (in blue) and daily travel totals (in black). It reveals that reservations are generally longer in distance than single trips in private vehicles, but generally less than distances in a household travel day. This makes sense as shared vehicles in a round-trip carsharing system are typically used for more than a single trip, but generally for less than a day. Note, nonetheless, that the trip distance distribution exhibits a long tail, reflective of the fact that some users reserve vehicles for a full day or multiple travel days. In turn, the trip data considered in this work seem representative of car usage in the United States.

The distribution of reservation lengths in hours, shown in Figure 4-2, reflects that reservations are generally relatively brief, with trips of 2-4 hours in duration being the most common. However, some users reserve vehicles for much longer periods. Peaks are observed at 24, 48, and 72 hours, reflective of a pricing model in which users are charged a daily

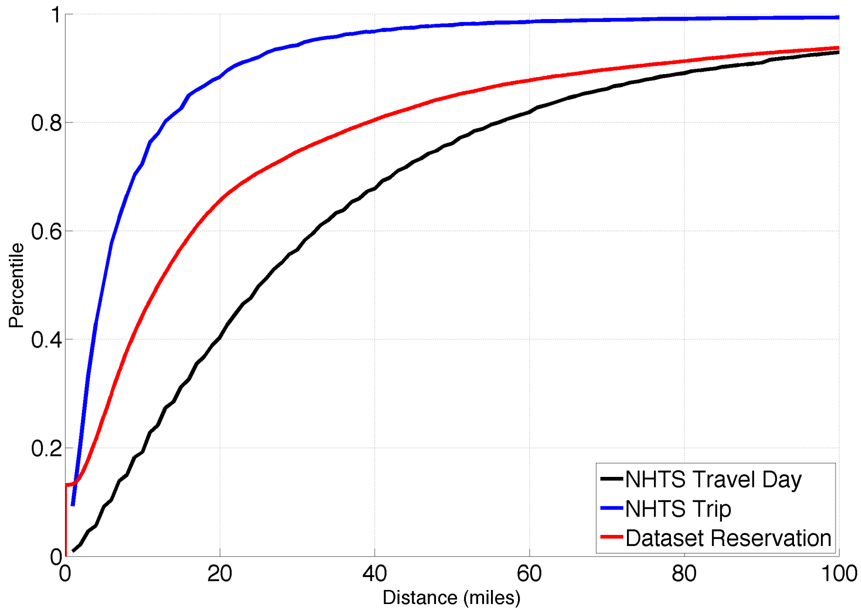

Figure 4-1: Distribution of trip distances and 2009 NHTS data

maximum rate and therefore incur no additional cost for keeping a vehicle for the remainder of a 24 hour day.

The distribution of lead times in hours, shown in Figure 4-3. Users most commonly reserve a vehicle within 24 hours of their reservation, but a significant portion also plan several days or more in advance. The X axis of the figure is truncated for legibility as a small portion of users reserve a vehicle several weeks or months in advance.

We now turn to the presentation of the vehicle technologies considered in this work. First, we characterize gasoline vehicles by their rate of gasoline consumption, denoted by α. We consider a value of $\alpha = 20$ miles per gallon. When it comes to PEVs, their performance is determined by the following factors: (i) the size of its battery (denoted by B, in kWh), (ii) its rate of energy consumption (denoted by δ, in kWh/mile), and (iii) its rate of battery charging (denoted γ, in kW). The range of the vehicle, i.e., how far it can travel on a single charge, is denoted by R and determined by the following relationship: $R = \frac{B}{\delta}$ (in miles). Similarly, the speed of battery recharging is frequently specified in terms of a "C-rate", determined as: $C = \frac{\gamma}{B}$ (per hour). A rate of $1C$ corresponds to a vehicle ability to fully

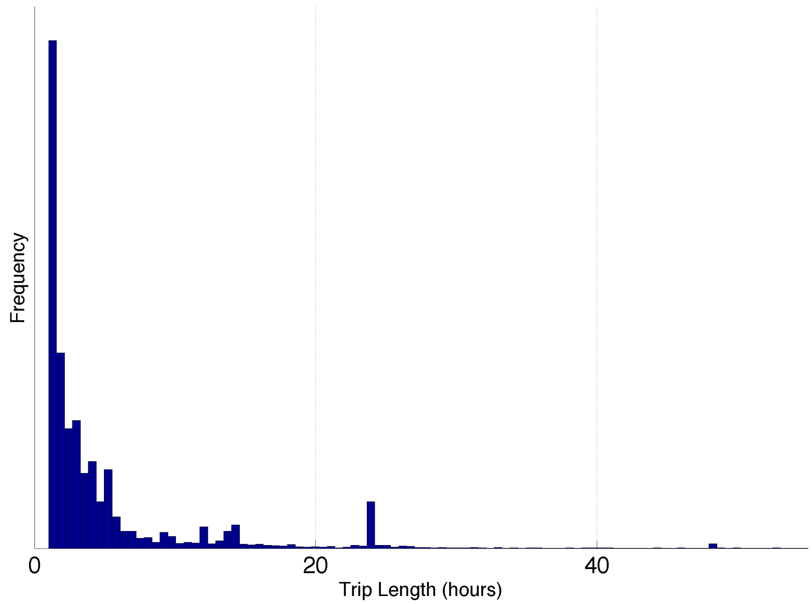

Figure 4-2: Histogram of trip duration in hours. Most last only a few hours, but some last a day or more.

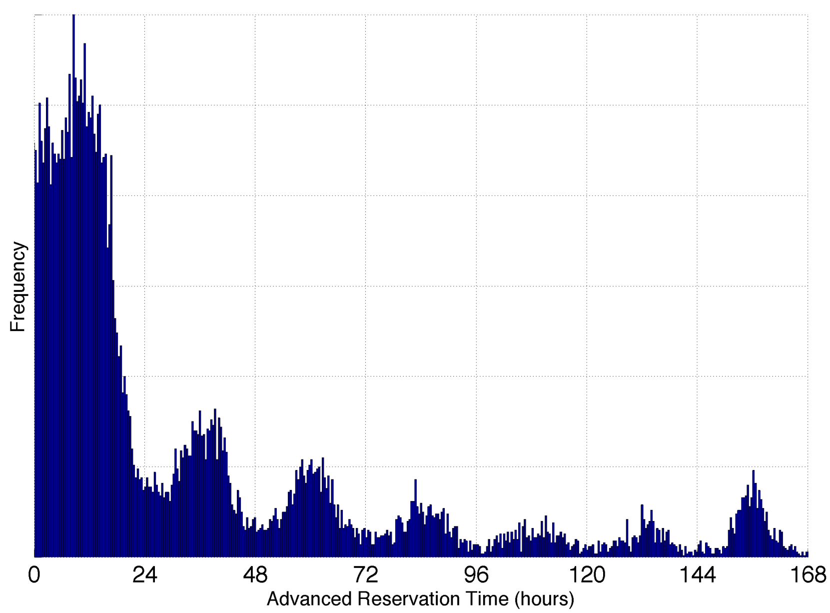

Figure 4-3: Histogram of the lead times by reservation. Most are under 24 hours, although some users reserve days in advance.

recharge its battery in one hour, a rate of $2C$ corresponds to a full recharge in 30 minutes, etc.

We assume that the rates of gasoline and electricity consumption are identical for all trips. While in practice vehicle energy consumption varies based on driving conditions and speeds, trip data does not allow us to distinguish driving profiles. The assumption of homogeneous energy consumption captures the gasoline savings associated with the transfer of vehicle miles traveled (VMT) from gasoline vehicles to electric vehicles. Moreover, we also assume that all gasoline vehicles in the fleet are identical (i.e., have the same rate of gasoline consumption α) and that all EVs in the fleet are identical (i.e., have the same electricity charging rate γ and the same rate of electricity consumption δ). We vary the various parameters to perform sensitivity analyses.

Without loss of generality, we fix the battery size, and vary the energy consumption rate δ and the recharging rate γ. We consider a battery size equal to $B = 24$ kWh. We vary the values of δ from a minimum of 0.05kWh/mi to a maximum of 1kWh/mi, and the values of γ from a minimum of 1kW to a maximum of 96kW. These values correspond to electric ranges spanning from 24 miles to 480 miles, and to C-rates spanning from $C/24$ (i.e., a battery recharged in a full day) to $4C$ (i.e., a battery recharged in 15 minutes). These values are intended to span the range of vehicle capabilities on the market today or in the near future.

The parameters used in this study are shown in Table 4.2. In the remaining sections we characterize vehicles by referring to their range R and their recharging rate γ. As described above, there is a one-to-one mapping between these values and the rate of energy consumption δ and the C-rate. Where multiple values of the range R or γ are not shown, the default values of $\gamma = 8$ kW and $\delta = 0.2$ kWh/mi are used. This corresponds to a range of 120 miles and a C-rate of $C/3$ (a battery charged in 3 hours).

Table 4.2: Summary of parameters used in this work

Vehicle	Parameter	Symbol	Values Used
Gasoline	Fuel Consumption (gal/mi)	α	0.05
EV	Battery Size (kWh)	B	$\underline{24}$
	Energy Consumption (kWh/mi)	δ	0.05, 0.1, $\underline{0.2}$, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0
	Charging Rate (kW)	γ	1, 4, $\underline{8}$, 12, 16, 24, 32, 48, 96
	C-Rate	C	$C/24$, $C/6$, $\underline{C/3}$, $C/2$, $2C/3$, C, $4C/3$, $\overline{2C}$, $4C$
	Electric Range (miles)	R	480, 240, $\underline{120}$, 80, 60, 48, 40, 34.3, 30, 26.7, 24

4.3 Potential Benefits from PEV Deployment and Utilization: A VTA Optimization Approach

In this section, we formulate and implement Mixed Integer Programs to optimize PEV utilization, i.e., the vehicle-trip assignment (VTA) in each pod, in order to minimize gasoline consumption, subject to trip feasibility and fleet availability constraints. Their results will inform, in turn, PEV deployment in a carsharing fleet, i.e., the vehicle-pod allocation (VPA). As we focus on round-trip carsharing systems, for which both the observation of passenger demand and the fleeting decisions are made at the level of each individual pod, the VTA problem can be solved for each pod independently. This decoupling in space reduces the computational requirements of the models considerably. In contrast, the VPA problem optimizes the deployment of vehicles across all pods.

Specifically, our approach follows a three-step structure:

Fleet Size Optimization (FSO). We determine the minimum number of vehicles required to serve the set of trips for the pod considered. In this stage, we assume an all-gasoline fleet.

Vehicle-Trip Assignment (VTA). We introduce PEVs in replacement of gasoline vehicles, while keeping the fleet size constant. We then optimize the assignment of individual trips to available vehicles to minimize gasoline consumption, subject to demand constraints, fleet availability constraints, electric range constraints, and charging dynamics considering a fleet comprising gasoline vehicles and EVs.

Vehicle-Pod Allocation (VPA). We optimize the deployment of PEVs in the carsharing system to minimize overall gasoline consumption (across all pods).

Note that all inputs and outputs of the FSO and VTA models are defined for each individual pod. For notational ease, we omit these dependencies in the presentation of these models. We reintroduce them in the formulation of the VPA model.

4.3.1 Fleet Size Optimization

The Fleet Size Optimization (FSO) model determines the minimal number of vehicles necessary to serve a given set of trips on a given pod. We develop a baseline vehicle-trip assignment model in the case of a fleet comprising only gasoline vehicles. We denote by n the number of trips in the pod considered, and we introduce the following parameters:

$$D_i = \text{the distance of trip } i = 1, ..., n$$
$$t_i^S / t_i^E = \text{the start/end time of trip } i = 1, ..., n$$

The assignment of individual trips to vehicles is determined through *connections* between trips. A connection between two trips i and j is defined by these two trips being traveled by the same vehicle and j being the immediate successor of i. We introduce below a parameter τ that characterizes feasible connections, and a variable z that characterizes the connections that are assigned by the model.

$$\tau_{ij} = \begin{cases} 1 & \text{if a connection between trip } i \text{ and trip } j \text{ is feasible, } i.e., \text{ if: } t_j^S \geq t_i^E \\ 0 & \text{otherwise} \end{cases}$$
$$z_{ij} = \begin{cases} 1 & \text{if trip } i \text{ and trip } j \text{ connect} \\ 0 & \text{otherwise} \end{cases}$$

Any vehicle-trip assignment involves partitioning the set of trips into subsets that are served by the same vehicle. Figure 4-4 provides a schematic representation of such a partition using a Gantt diagram. For instance, Vehicle 1 makes 3 trips during the period considered, Vehicle 2 makes 4 trips, etc. Based on this representation of the vehicle-trip assignment, we express below the relationship between the number of trips, the number of vehicles used, and the connectivity of the vehicle-trip assignment.

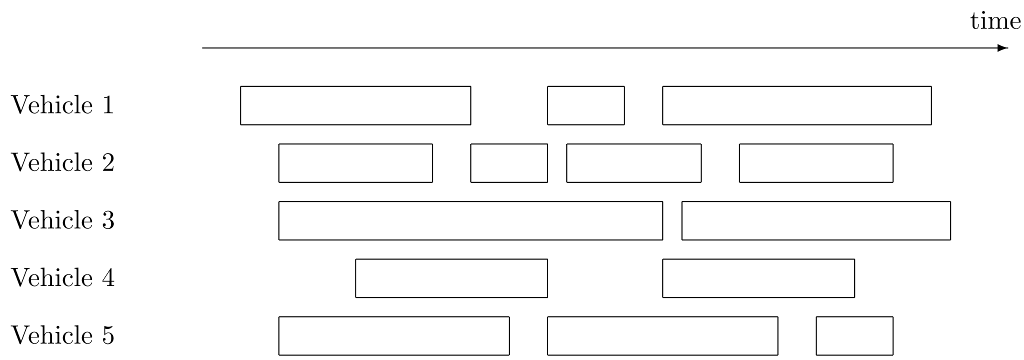

Figure 4-4: A schematic representation of the vehicle-trip assignment using a Gantt diagram

Minimizing fleet size is equivalent to maximizing the number of connections between trips (since the total number of trips n is given). The resulting problem is formulated as follows:

$$\begin{aligned}
\max \quad & \sum_{i=1}^{n}\sum_{j=1}^{n} z_{ij} && && (4.1)\\
\text{st} \quad & \sum_{i=1} z_{ij} \leq 1 && \forall j = 1, ..., n && (4.2)\\
& \sum_{j=1} z_{ij} \leq 1 && \forall i = 1, ..., n && (4.3)\\
& z_{ij} \leq \tau_{ij} && \forall i, j = 1, ..., n && (4.4)\\
& z \text{ binary} && &&
\end{aligned}$$

Constraints (4.2) and (4.3) ensure that each trip has at most one predecessor and one successor. Constraint (4.4) ensures that connections are only made when feasible in time (i.e. trip j start after trip i ends) to prevent simultaneous trip assignments to the same vehicle. We denote by z^* the optimal connectivity matrix and by $N^* = n - \sum_{i=1}^{n}\sum_{j=1}^{n} z^*_{ij}$

the minimal number of vehicles in the pod considered.

4.3.2 Vehicle-Trip Assignment

The Vehicle-Trip Assignment (VTA) model takes as inputs the list of trips per pod, as the FSO model, but it also varies the composition of the fleet. We fix the number of gasoline vehicles and the number of PEVs, and assign trips to vehicles to minimize gasoline consumption. The feasibility of such assignment depends on the fleet mix, i.e., the number of gasoline vehicles and PEVs, and on the characteristics of the PEVs (e.g., battery size, charging rate, distance range). We assume that vehicles are not recharged *during* a trip—a reasonable assumption for carsharing systems.

In addition to the parameters introduced in Section 4.3.1, we denote by N^G and N^{EV} the number of gasoline vehicles and EVs, respectively. Vehicles are characterized by the parameters introduced in Section 4.2. We introduce the following variables:

$$z_{ij}^{EV} = \begin{cases} 1 & \text{if trip } i \text{ and trip } j \text{ connect with an EV} \\ 0 & \text{otherwise} \end{cases}$$

$$z_{ij}^{G} = \begin{cases} 1 & \text{if trip } i \text{ and trip } j \text{ connect with a gasoline vehicle} \\ 0 & \text{otherwise} \end{cases}$$

$$r_{i}^{EV} = \begin{cases} 1 & \text{if trip } i \text{ is served by an electric vehicle} \\ 0 & \text{otherwise} \end{cases}$$

$$c_i^S/c_i^E = \text{state of charge at the start/end of trip } i \text{ (in kWh)}$$

The resulting problem is formulated as follows, where $M >> 1$ designates a very large scalar:

$$
\begin{aligned}
\min \quad & \sum_{i=1}^{N} \alpha D_i (1 - r_i^{EV}) && && (4.5)\\
\text{st} \quad & \sum_{i=1}^{n} \left(z_{ij}^{EV} + z_{ij}^{G} \right) \leq 1 && \forall j = 1, ..., N && (4.6)\\
& \sum_{j=1}^{n} \left(z_{ij}^{EV} + z_{ij}^{G} \right) \leq 1 && \forall i = 1, ..., N && (4.7)\\
& z_{ij}^{EV} + z_{ij}^{G} \leq \tau_{ij} && \forall i, j = 1, ..., N && (4.8)\\
& \sum_{j=1}^{n} z_{ij}^{EV} \leq r_i^{EV} \leq 1 - \sum_{j=1}^{n} z_{ij}^{G} && \forall i = 1, ..., N && (4.9)\\
& \sum_{i=1}^{n} z_{ij}^{EV} \leq r_j^{EV} \leq 1 - \sum_{i=1}^{n} z_{ij}^{G} && \forall j = 1, ..., N && (4.10)\\
& c_i^S \leq B r_i^{EV} && \forall i = 1, ..., N && (4.11)\\
& c_i^E \leq c_i^S - \delta D_i r_i^{EV} && \forall i = 1, ..., N && (4.12)\\
& c_j^S - c_i^E - M(1 - z_{ij}^{EV}) \leq \gamma (t_j^S - t_i^E) && \forall i, j = 1, ..., N && (4.13)\\
& \sum_{i=1}^{N} (1 - r_i^{EV}) - \sum_{i=1}^{N} \sum_{j=1}^{N} z_{ij}^{G} \leq N^G && && (4.14)\\
& \sum_{i=1}^{N} r_i^{EV} - \sum_{i=1}^{N} \sum_{j=1}^{N} z_{ij}^{EV} \leq N^{EV} && && (4.15)\\
& c^S, c^E \geq 0 && &&\\
& z^G, z^{EV}, r \text{ binary} && &&
\end{aligned}
$$

Equation (4.5) formulates the model's objective of minimizing gasoline consumption, which is proportional to the total distance traveled by gasoline vehicles. Constraints (4.6) and (4.7) ensure that each trip has at most one predecessor and one successor. Constraint (4.8) ensures the feasibility of each connection. Constraints (4.9) and (4.10) ensure the consistency of variables z^G, z^{EV}, and r^{EV} by forcing r_i^{EV} and r_j^{EV} to be both equal to 1 (resp. to 0) if $z_{ij}^{EV} = 1$ (resp. if $z_{ij}^G = 1$). Constraint (4.11) ensures that the state of charge is lower than the battery size of each EV. Constraint (4.12) quantifies the battery consumption during each trip made with an EV, and Constraint (4.13) quantifies the charging dynamics when the EV is idle (The large scalar M is used to apply this constraints only for EVs, not for gasoline vehicles). Finally, Constraints (4.15) and (4.14) ensure that the number of gasoline vehicles and EVs do not exceed the number of available vehicles N^G and N^{EV}. In these constraints, we express the number of gasoline vehicles and EVs used as the difference between the number of trips made with gasoline vehicles and EVs (i.e., $\sum_{i=1}^{N}(1 - r_i^{EV})$ and $\sum_{i=1}^{N} r_i^{EV}$, respectively), and the number of connections made by gasoline vehicles and EVs (i.e., $\sum_{i=1}^{N}\sum_{j=1}^{N} z_{ij}^G$ and $\sum_{i=1}^{N}\sum_{j=1}^{N} z_{ij}^{EV}$, respectively), as described in Section 4.3.1.

4.3.3 Vehicle-Pod Allocation

We now describe our algorithm to optimize the deployment of PEVs in the carsharing system, i.e., the Vehicle-Pod Allocation (VPA), for a given "budget" of PEVs. For a given total number of PEVs available (denoted by N_{all}^{EV}), we aim to find the number of PEVs per pod that will yield the largest gasoline savings, as estimated by the VTA optimization procedure developed in this section. In the remainder of this work, we refer to this VPA procedure by "VPAopt/VTAopt".

Algorithm 1 details our VPAopt/VTAopt approach to EV deployment. We first solve the FSO for all pods, which provides the value of the minimal number of vehicles required in each pod p to serve the demand. We now denote it by N_p^*. Then, we solve the VTA with an increasingly large number of EVs in *replacement* of gasoline vehicles, i.e., we keep the fleet size constant. This choice is motivated by the objectives of minimizing changes in the carsharing system's operations, and of isolating the benefits associated with the introduction of EVs. Under the assumption of a constant fleet size, the more EVs are deployed, the greater the potential reduction in gasoline consumption, but the more constrained the vehicle-trip assignment problem. For each pod p, and for each number of electric vehicles N_p^{EV} (and thus $N_p^G = N_p^* - N_p^{EV}$ gasoline vehicles), we update a boolean variable (denoted by $feas$) that characterizes whether the problem is feasible, or not, and we denote by $\Phi_p^*\left(N^{EV}\right)$ the minimal gasoline consumption. The algorithm stops when the problem becomes unfeasible, i.e., when the balance of electric vehicles and gasoline vehicles is such that all trips can no longer be served because of the limitations in range and charge of EVs. Note that this problem is clearly monotonic: if the problem is unfeasible with N^{EV} electric vehicles and $N^* - N^{EV}$ gasoline vehicles, then it will also be unfeasible with $N^{EV} + 1$ electric vehicles and $N^* - N^{EV} - 1$ gasoline vehicles. We denote by $\overline{N}_p^{EV}$ the maximal number of EVs that can be deployed in pod p in replacement of gasoline vehicles.

Then, in a second time, we solve the problem of deploying the N_{all}^{EV} EVs in all pods to minimize overall gasoline consumption, which is formulated as follows:

$$\begin{aligned} \min \quad & \sum_{p=1}^{P} \Phi_p^* \left(N_p^{EV}\right) \\ \text{st} \quad & \sum_{p=1}^{P} N_p^{EV} \leq N_{\text{all}}^{EV} \\ & N_p^{EV} \leq \overline{N}_p^{EV}, \forall p \end{aligned} \tag{4.16}$$

This problem can be easily solved by sorting the marginal gasoline savings, i.e., $\Phi_p^* \left(N_p^{EV}\right) - \Phi_p^* \left(N_p^{EV} - 1\right)$, for all pods p and for all values of $N_p^{EV} \leq \overline{N}_p^{EV}$, and by then selecting its first N_{all}^{EV} components (see Algorithm 1). Note that each additional EV can be either allocated to a pod that previously only had gasoline vehicles, or to a pod that already had EVs, depending on the gasoline consumption reductions resulting from such deployment. In other words, the deployment of EVs involves deciding to which pods the EVs should be allocated *and* how many EVs should be allocated per pod.

4.3.4 VTA Optimization Results

We apply the Vehicle-Trip Assignment models presented above to the carsharing system detailed in Section 2. Figure 4-5 shows the VTA results for a given pod and for a given week of operations using Gantt diagrams (see Figure 4-4). Figure 4-5a considers a fleet of

Algorithm 1 VPA optimization based on VTA optimization (VPAopt/VTAopt)

Inputs: number of EVs to deploy N_{all}^{EV}, list of trips per pod, vehicle characteristics (α, B, γ, δ)
for each pod p **do**
  Solve Fleet Size Optimization (FSO)$\longrightarrow N_p^*$
  Initialization: $N_p^{EV} \leftarrow 0$, $\overline{N}_p^{EV} \leftarrow -1$, $feas \leftarrow TRUE$
    while $feas$ **do**
      $\overline{N}_p^{EV} \leftarrow \overline{N}_p^{EV} + 1$
      Store gasoline consumption with $N_p^{EV} \longrightarrow \Phi_p^*\left(N^{EV}\right)$
      $N_p^{EV} \leftarrow N_p^{EV} + 1$, $N^G \leftarrow N_p^* - N_p^{EV}$
      Solve Vehicle-Trip Assignment (VTA) $\longrightarrow feas$
    end while
end for
Sort marginal gasoline savings: $\Phi_p^*\left(N_p^{EV}\right) - \Phi_p^*\left(N_p^{EV} - 1\right)$, $\forall p$, $\forall N_p^{EV} \leq \overline{N}_p^{EV}$
Deploy the N_{all}^{EV} EVs to the pods that yield the largest gasoline savings.

10 gasoline vehicles, while Figure 4-5b considers a fleet comprising 8 gasoline vehicles (in blue) and 2 EVs (in green). Note that the VTA model decouples the set of trips into a subset a trips made with PEVs and another subset of trips made with gasoline vehicles. As expected, EVs are used to operate a large number of short trips (shorter than their range of 120 miles), while gasoline vehicles serve fewer trips, but all the longer trips. Moreover, the usage of each vehicle varies, as some trips follow each other in short periods of times, while the period of idleness between other trips may be longer. In some periods of time (e.g., Day 2), all trips can be served by EVs. In some other periods of time, in contrast, EVs are used to serve the set of trips that yield the largest gasoline savings (under feasibility constraints), and gasoline vehicles are used to serve the remaining demand.

In the remainder of this section, we show the results of the VPA and VTA models together, i.e., for any total number of PEVs in the fleet (N_{all}^{EV}), we jointly optimize their allocation to pods (based on the VPAopt/VTAopt strategy shown in Algorithm 1) and the utilization of vehicles in the fleet to serve demand. In order to illustrate how such deployment and utilization of PEVs is performed, Table 4.3 shows the distribution of distance served with EVs for increasing numbers of EVs (which we represent as a percentage of the total fleet). Note, first, that EVs are never used to serve trips that are longer than their range

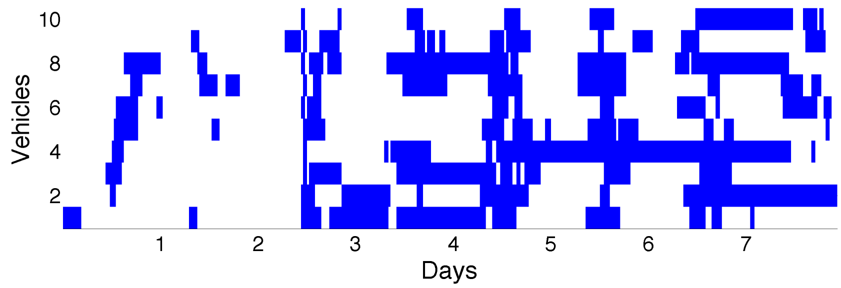

(a) 10 gasoline vehicles

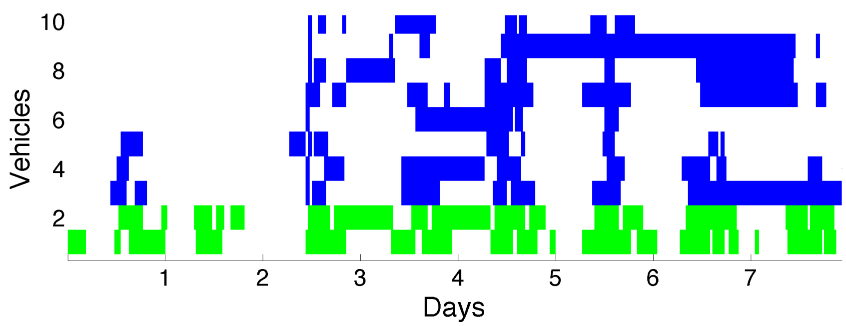

(b) 8 gasoline vehicles, 2 EVs

Figure 4-5: Vehicle-trip assignment with and without EVs

limit of 120 miles, consistently with their technological constraints. Second, the VTA aims to use EVs to serve primarily the *longest* of the trips that meet range limitations. This is because the gasoline savings are more significant if EVs are allocated to longer trips than to shorter trips. However, EVs cannot serve only the longest of the trips under the range limit of 120 miles, either because these trips are simultaneous, or because the charging time between these trips is not sufficient. In turn, the optimal VTA involves using some EVs to serve shorter trips, as well. Finally, up to 40% of the fleet can be converted to EVs while serving all observed demand. In this case, only 60-70% of the trips under the range limit of 120 miles are, in fact, served by EVs. Even though some other short trips could be served by additional EVs if these were introduced in the fleet, the removal of additional gasoline vehicles (under assumption of a constant fleet size) would not enable to serve all the demand (either because of a large number of above-range trips, or because of insufficient charging times).

Table 4.3: Proportion of trips served by EVs as a function of distance and fleet composition

	Distance (in miles)					
%EVs	0-20	20-40	40-60	60-80	80-100	100-120
5%	9%	14%	18%	19%	22%	25%
10%	18%	24%	31%	34%	35%	38%
15%	26%	33%	41%	45%	46%	49%
20%	33%	42%	50%	54%	53%	64%
25%	39%	49%	57%	61%	61%	54%
30%	45%	55%	63%	66%	66%	68%
35%	50%	60%	67%	69%	69%	70%
40%	54%	63%	69%	70%	70%	70%

Table 4.4 shows the distributions of charge at the end of each trip for various levels of deployment of EVs. We can see that the state of charge is clearly a binding constraint in the optimization: the majority of trips (at all deployment levels) end with a very low state of charge. This result illustrates one of the challenges of adopting exclusively an optimization-based approach. For real-world applications, it is unlikely that a carsharing operator would wish to leave so little margin for error when assigning trips to vehicles, as users will have imperfect information about their travel plans, and small variations in energy consumption could lead to stranded users.

We now turn to the allocation of vehicles to pods at the aggregate level. Figure 4-6 and Table 4.5 show the number of EVs per pod, for an increasing number of vehicles converted to EVs. As expected, the first PEVs that are deployed are allocated to pods that do not have any PEVs, in replacement of gasoline vehicles. Specifically, the optimal deployment strategy involves allocating only one EV per pod for the first 15 EVs. Then, for larger numbers of PEVs, the optimal strategy involves deploying some PEVs to pods that had all-gasoline fleets beforehand, and some other PEVs to pods whose fleet already comprised one or several PEVs. The more PEVs are deployed, the more the balance shifts from the former case to the latter case, and the last 50 EVs are predominantly allocated to pods that already had mixed fleets. These results suggest that the deployment of PEVs should balance allocating many PEVs to the busiest locations (a “deep” strategy), on the one hand, and allocating

Table 4.4: Distribution of charge levels at the end of trips made by EVs

	N_{all}^{EV}							
C_i^E (kWh)	5%	10%	15%	20%	25%	30%	35%	40%
0-1	46%	48%	50%	52%	53%	55%	56%	58%
1-2	3%	3%	2%	2%	2%	2%	2%	2%
2-3	3%	2%	2%	2%	2%	2%	2%	2%
3-4	2%	3%	3%	2%	2%	2%	2%	2%
4-5	2%	3%	3%	2%	2%	2%	2%	2%
5-6	1%	3%	3%	2%	2%	2%	2%	2%
6-7	2%	3%	3%	2%	2%	2%	2%	2%
7-8	2%	3%	3%	2%	2%	2%	2%	2%
8-9	2%	3%	3%	2%	2%	2%	2%	2%
9-10	1%	3%	3%	2%	2%	2%	2%	2%
10-11	2%	3%	3%	2%	2%	2%	2%	2%
11-12	2%	3%	3%	2%	2%	2%	2%	2%
12-13	1%	3%	3%	2%	2%	2%	2%	2%
13-14	1%	3%	3%	2%	2%	2%	2%	2%
14-15	2%	3%	3%	2%	2%	2%	2%	2%
15-16	2%	3%	3%	2%	2%	2%	2%	2%
16-17	2%	3%	3%	2%	2%	2%	2%	2%
17-18	2%	3%	3%	2%	2%	2%	2%	2%
18-19	2%	3%	3%	2%	2%	2%	2%	2%
19-20	2%	3%	3%	2%	2%	2%	2%	2%
20-21	3%	3%	3%	2%	2%	2%	2%	2%
21-22	3%	3%	3%	2%	2%	2%	2%	2%
22-23	3%	3%	3%	2%	2%	2%	2%	2%
23-24	3%	3%	3%	2%	2%	2%	2%	2%

PEVs to as many locations as possible (a "wide" strategy), on the other hand, based on demand and PEV utilization patterns in the various pods. We further describe the benefits of integrating PEV utilization patterns into PEV deployment strategies in Section 4.5.

Finally, we present the gasoline savings resulting from the deployment and utilization of PEVs in the carsharing fleet. Figure 4-8 and Figure 4-7 show the reduction in gasoline consumption achieved by converting increasing numbers of gasoline vehicles to EVs for different values of the energy consumption rate δ and of battery charging rate γ. Several observations are noteworthy. First, the relationship between the number of PEVs deployed and the resulting gasoline savings is increasing and concave. In other words, deploying PEVs reduces gasoline consumption, but the rate of return associated with PEV deployment is

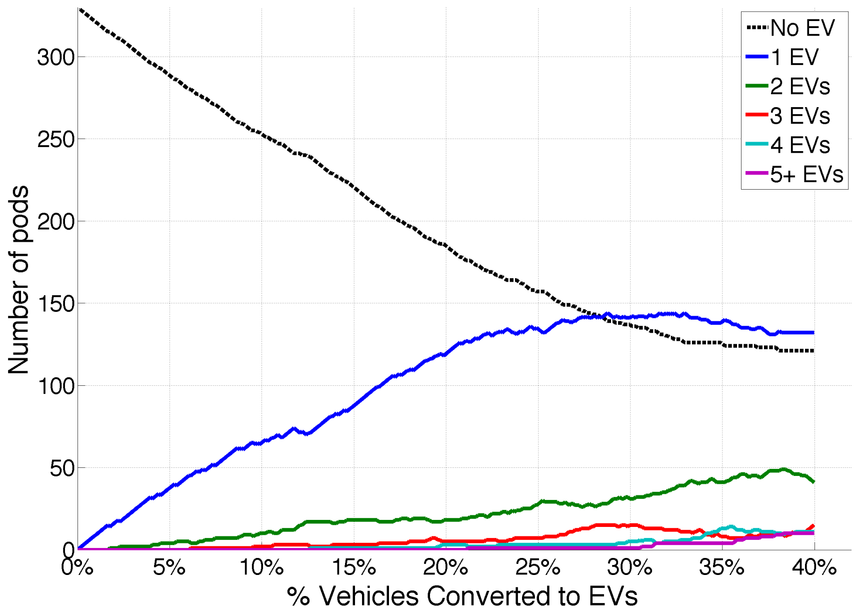

Figure 4-6: Number of EVs per Pod.

Table 4.5: Number of pods with 0, 1, 2, 3, 4 or more EVs for various levels of PEV deployment

%EVs	0	1	2	3	4	5+
5%	288	38	4	0	0	0
10%	253	65	10	2	0	0
15%	220	88	18	3	1	0
20%	185	119	18	5	3	0
25%	157	134	28	7	3	1
30%	137	141	31	15	5	1
35%	125	139	41	8	13	4
40%	121	132	41	15	11	10

non-increasing. This is expected because as more PEVs are deployed, a larger share of vehicle miles traveled gets converted from gasoline to electricity, and because PEVs are allocated to pods by decreasing order of marginal gasoline savings (see Section 4.3.3). Second, improving EV capabilities (e.g., increasing range or charging rate) results in larger numbers of EVs that can be deployed while serving all the demand (i.e., fewer unfeasibility scenarios) as well as higher gasoline savings. Third, system performance seems to be very sensitive to the range of the PEVs. As the range increases, the gasoline savings become significantly larger and more EVs can be deployed. In contrast, system performance does not seem to vary significantly with the battery charging rate. Performance gains seem to be achieved from increasing PEV capabilities from Level 1 charging ($\gamma = 1$ kW) to Level 2 charging ($\gamma = 4$ kW or $\gamma = 8$ kW), but very little gains (if any) seem to be achieved through further increases in charging capabilities, both in terms of the number of EVs that can be deployed and of the resulting gasoline savings.

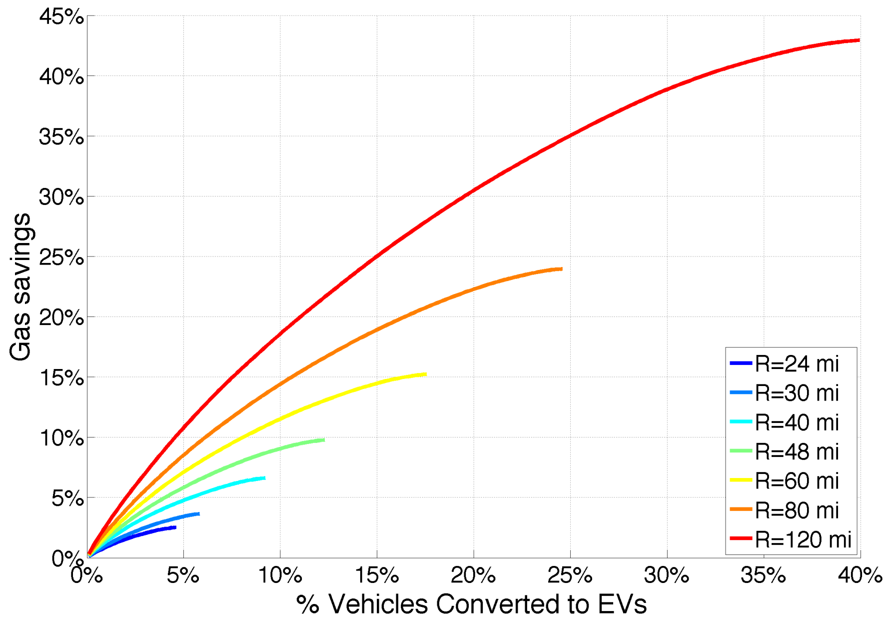

Figure 4-7: Sensitivity to δ ($\gamma = 8$ kW)

In sum, the Vehicle-Trip Assignment (VTA) model decouples the set of trips into a subset of trips made by PEVs and another made by gasoline vehicles to minimize gasoline consumption in any pod, under feasibility and fleet availability constraints, and the Vehicle-Pod Allocation (VPA) model optimizes the deployment of PEVs across all pods to minimize overall gasoline consumption. Results suggest that gasoline savings can be very significant, and are more sensitive to vehicle range than to vehicle charging capabilities. In the next section, we compare this VTA optimization approach to a VTA simulation approach (see Table 4.1).

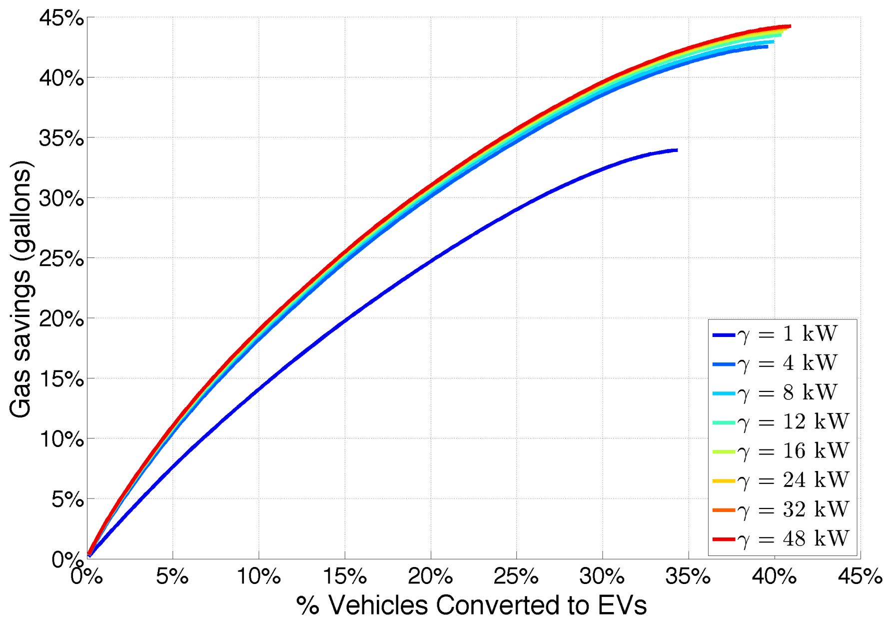

Figure 4-8: Sensitivity to γ ($R = 120$ miles)

4.4 Impact of Reservation Timing on Realized PEV Benefits: A VTA Simulation Approach

In this section, we develop and implement our simulation-based approach to the VTA problem. The approach considers a given vehicle-trip assignment heuristic and quantifies the realized benefits associated with PEV utilization. It differs from the VTA optimization approach developed in the previous section in three major ways. First, while the VTA optimization determined an ideal assignment of trips to vehicles given the full list of trips in a given period of time, the VTA simulation considers the timing of reservations and assigns trips to vehicles by order of reservations. Second, while the VTA optimization imposed that no trip may be dropped as a constraint, the order of reservations may be such that some trips may not be served in practice, PEV deployment thus creates a trade-off between reductions in gasoline consumption, on the one hand and higher likelihood of trips being dropped, on the other. The VTA simulation approach quantifies this trade-off. Third, while the VTA optimization assumed a centralized assignment of trips to vehicles, carsharing users may have preferences over vehicle types. The VTA simulation approach enables to quantify the impact of users' preferences and technology choices on vehicle utilization.

4.4.1 VTA Simulation Presentation

As the VTA optimization, the VTA simulation takes as inputs the list of trips in each pod and the composition of the fleet, i.e., the number of vehicles (N^{EV} and N^{G}) and their characteristics (α, B, R and γ). For each pod, we consider trip reservations in the order that they were received by the carsharing operator, and aim to assign them to one of the available vehicles in the pod. The heuristic prioritizes the assignment of trips to one of the EVs. For each EV, we check if (i) the EV is available for the trip considered (i.e., the trip does not overlap in time with any other trip served by the EV) and (ii) the EV has sufficient range and charge to complete the trip (i.e., the state of charge of the vehicle will not drop below zero).

We assign the trip to the first EV that meets these conditions. Note that these feasibility checks must be re-evaluated for the entire trip chain served by the EV considered each time a trip is added, as the addition of any trip may affect the future state of charge of the vehicle. For instance, as illustrated in Figure 4-9, consider an EV that is currently assigned to the feasible trip chain of reservation (A) from 12-4pm and reservation (B) from 8-10pm, and assume that a new reservation is received from 5-7pm. When deciding to assign this trip to this EV, we must consider not only its state of charge between 5pm and 7pm, but also the impact of the new trip on the future state of charge for previously assigned trips in the chain, e.g. reservation (B). If the trip cannot be assigned to any EV, then the algorithm assigns the trip to the first gasoline vehicle that is available to serve the trip (i.e., a vehicle such that the trip does not overlap in time with any other trips it serves). If no vehicle is available to serve the trip, then the trip is dropped (i.e., not served).

Note that this VTA simulation procedure relies on a simple greedy heuristic to assign trips to vehicles. While it is beyond the scope of this work, a large number of alternative heuristic schemes could be envisioned, e.g., (i) the assignment of trips to the EV with the

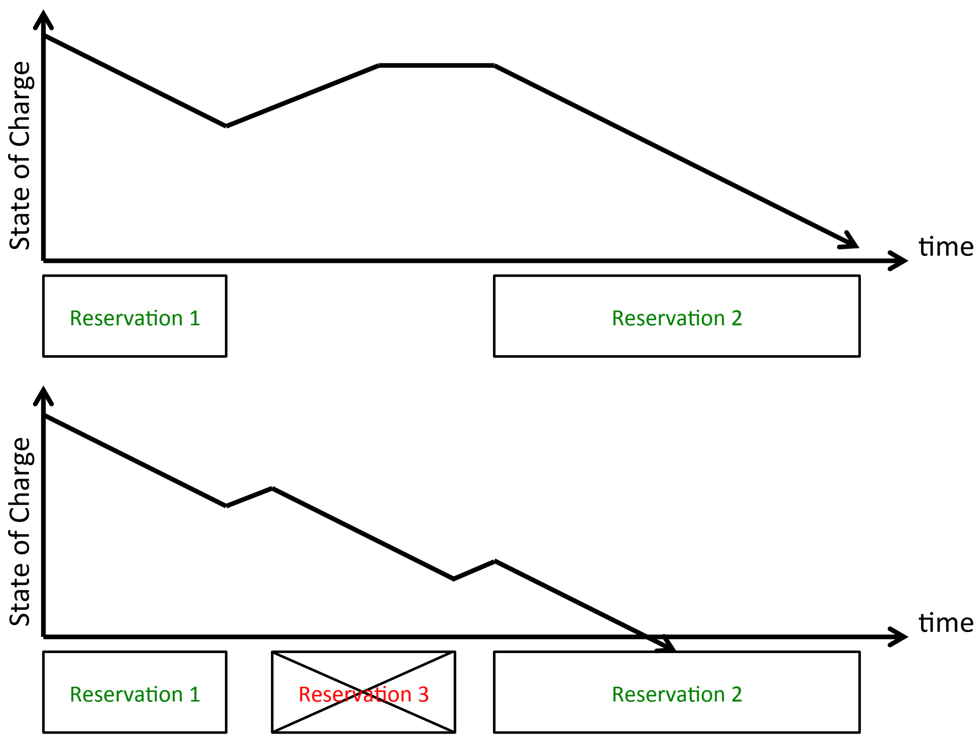

Figure 4-9: State of Charge is simulated for each potential trip chain. Trips reserved later are rejected if they introduce infeasibility anywhere along the chain.

lowest state of charge that still meets the needs of the given trip, or (ii) the re-optimization of the vehicle-trip assignment each time a new reservation is received. Nonetheless, this simulation procedure relies on a first-come first-assigned process that is relatively close to the assignment dynamics in place in practice and that is consistent with the business model of carsharing operators (e.g., no reservation gets re-assigned to a different vehicle after it has been made). It is therefore relatively similar to what could be implemented, if the carsharing operator prioritized the assignment of trips to EVs over gasoline vehicles. As the VTA optimization, this simulation procedure assumes that users know their travel distances accurately and that no vehicle gets recharged during a trip. Finally, for consistency with the optimization results, we perform all tests with N_p^* vehicles in each pod p, i.e., we do not vary the total number of vehicles in any pod.

We perform three sets of simulations:

All Gasoline. We consider a fleet comprising exclusively gasoline vehicles and no EV. The number of gasoline vehicles in each pod p is equal to N_p^* obtained from the FSO model (see Section 4.3.1). This aims to quantify the inefficiency of the simulation process itself relative to optimization and to establish a baseline level of fuel consumption and dropped trips for an all-gasoline fleet.

Mixed Fleet. We consider a mixed fleet comprising N_p^G gasoline vehicles and N_p^{EV} EVs in pod p, with $N_p^G + N_p^{EV} = N_p^*$. In this scenario we assume that users are always willing to accept an EV if it is available and has sufficient charge to travel as far as their planned trip distance. The comparison of the Mixed Fleet scenario to the All Gasoline scenario quantifies the impact of replacing gasoline vehicles by EVs on gasoline consumption and on the system's ability to serve demand.

EV Rejection. We also consider a mixed fleet comprising N_p^G gasoline vehicles and N_p^{EV} EVs in pod p, but we introduce heterogeneity in renter attitudes towards EVs, consistently with the findings of Zoepf and Keith (2015). This heterogeneity materializes in some users accepting to use an EV and some other users rejecting to use an EV due to

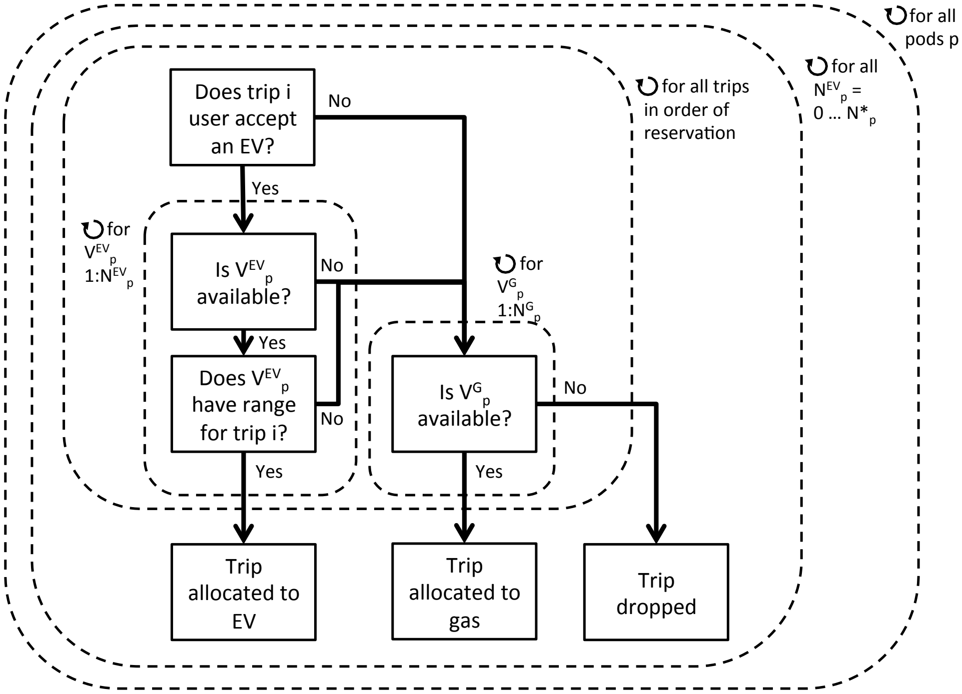

Figure 4-10: Flow chart of the VTA simulation procedure

unknown reasons (e.g., technological preferences, uncertain travel distance, overestimated travel distance, etc.). We represent these attitudes by a random binary variable. At the time of vehicle assignment, each renter has a probability P_{accept} of accepting to use an EV (if he/she gets assigned to an EV), and a probability $1 - P_{accept}$ of rejecting to use an EV. This aims to quantify the impact of users' technology choices on gasoline consumption and on the system's ability to serve demand. A summary of our VTA Simulation approach is provided in Figure 4-10. Note that the vehicle-trip assignment depends not only on the number of users that accept to use EVs, but also on *which* users accept to use EVs. In turn, for any given value of P_{accept}, we run 20 entire VTA simulations where each user' acceptance to use an EV is determined by Monte Carlo sampling.

For each simulation, we store both the gasoline consumption, denoted by $\Phi_p^{\text{sim}}\left(N_p^{EV}\right)$, and the set of dropped trips, denoted by $\mathcal{T}_p^{\text{sim}}(N_p^{EV})$, as a function of the number of EVs allocated to pod p, i.e., N_p^{EV}. Note that it is necessary to store the *set* of dropped trips (instead of simply the number of dropped trips) as the financial impact of dropped trips depends on *which* trips are dropped (since carsharing revenues are a function of the duration of the trips served). We then use the results of the VTA simulation to inform the Vehicle-Pod Allocation (VPA), i.e., the deployment of EVs in the different pods. We describe this procedure in Algorithm 2, and we refer to it by "VPAopt/VTAsim" in the remainder of this work. This algorithm is similar to Algorithm 1, with the exception that the VTA simulation involves both gasoline savings and dropped trips. We solve the associated trade-off by fixing the maximum proportion of dropped trips per pod due to the replacement of gasoline vehicles by EVs, which we denote by ε. We then determine the largest number of EVs that can be feasibly allocated to pod p, denoted by $\overline{N}_p^{EV}$, beyond which the proportion of trips that are dropped due to the replacement of gasoline vehicles by EVs exceeds ε. In other words, $\overline{N}_p^{EV} = \max\left\{N_p^{EV} = 1, ..., N_p^* \middle| \frac{\left|\mathcal{T}_p^{\text{sim}}(N_p^{EV})\right| - \left|\mathcal{T}_p^{\text{sim}}(0)\right|}{n_p - \left|\mathcal{T}_p^{\text{sim}}(0)\right|} \leq \varepsilon\right\}$, where n_p designates the number of trips in pod p. We then allocate EVs to the pods to maximize the overall gasoline savings, under feasibility constraints.

In the remainder of this section, we provide results of the VTA simulation in the All Gasoline case, in the Mixed Fleet case, and in the EV Rejection case.

4.4.2 VTA Simulation Results: All Gasoline

First, we apply the VTA simulation procedure to establish a baseline in the case of an all-gasoline fleet in each pod. As the VTA simulation assigns trips to vehicles in order of reservations, we can expect that a certain number of trips will go unallocated due to the sub-optimality of the VTA heuristic considered in this work. These dropped trips result in a reduction in gasoline consumption that is unrelated to the relative performance of EVs vs. gasoline vehicles.

Results for the all gasoline simulations are shown in Table 4.6 Of the 49,518 trips in the dataset, 242 trips are unallocated in the all gasoline simulation (fewer than 0.5% in aggregate), with the worst-case pod dropping 6 trips (2%) as a result of simulation inefficiency.

Algorithm 2 VPA optimization based on VTA simulation (VPAopt/VTAsim)

Inputs: number of EVs to deploy N_{all}^{EV}, list of trips per pod, vehicle characteristics (α, B, γ, δ), dropped trips threshold ε

for each pod p **do**

  Simulate Trip Allocation for $N_p^{EV} = 0, N_p^G = N_p^* \longrightarrow \mathcal{T}_p^{\text{sim}}(0), \Phi_p^{\text{sim}}(0)$

  Initialization: $N_p^{EV} \leftarrow 0$, $\overline{N}_p^{EV} \leftarrow -1$, $feas \leftarrow TRUE$

  while $feas$ **do**

    $\overline{N}_p^{EV} \leftarrow \overline{N}_p^{EV} + 1$

    Store gasoline consumption and dropped trips with $N_p^{EV} \longrightarrow \Phi_p^{\text{sim}}\left(N^{EV}\right)$, $\mathcal{T}_p^{\text{sim}}\left(N^{EV}\right)$

    $N_p^{EV} \leftarrow N_p^{EV} + 1$, $N^G \leftarrow N_p^* - N_p^{EV}$

    Simulate Vehicle-Trip Assignment

    if $\frac{\left|\mathcal{T}_p^{\text{sim}}(N_p^{EV})\right| - \left|\mathcal{T}_p^{\text{sim}}(0)\right|}{n_p - \left|\mathcal{T}_p^{\text{sim}}(0)\right|} > \varepsilon$ **then** $feas \leftarrow FALSE$

    end if

  end while

end for

Sort marginal gasoline savings: $\Phi_p^{\text{sim}}\left(N_p^{EV}\right) - \Phi_p^{\text{sim}}\left(N_p^{EV} - 1\right)$, $\forall p$, $\forall N_p^{EV} \leq \overline{N}_p^{EV}$

Deploy the N_{all}^{EV} EVs to the pods that yield the largest gasoline savings.

This corresponds to small reductions in overall gasoline consumption (lower than 0.5% in aggregate). In only 10 of 330 pods were more than 3 trips unallocated in the all gasoline VTA simulation. These unallocated ("dropped") trips were not strongly associated with large or small pods.

Table 4.6: VTA simulation results in the case of an all-gasoline fleet

# Dropped Trips	Occurrences	Largest N^*	Smallest N^*
0	196	16	1
1	73	10	2
2	31	7	2
3	20	10	2
4	6	7	2
5	1	7	7
6	3	5	4

Because the number of dropped trips is very small and consequent reduction in gasoline consumption is minimal, nearly all the gasoline savings shown in the remainder of this section can be predominantly attributed to the utilization of EVs instead of gasoline vehicles.

4.4.3 VTA Simulation Results: Mixed Fleet

We now turn to the case of a fleet comprising both gasoline vehicles and EVs. Similarly to Section 4.3.4, we first show the results of the vehicle-trip assignment at the level of a given pod, and we then look at aggregate results at the system-wide level. One difference, however, with the results from Section 4.3.4, is that any reduction in gasoline consumption may result from a shift of Vehicle Miles Traveled (VMT) from gasoline vehicles to electric vehicles or from some trips being dropped (hence, revenue being lost) because of the inefficiency of the vehicle-trip assignment heuristic considered and the specific order of reservations. This contrasts with the VTA optimization, under which any gasoline savings resulted from the displacement of petroleum by electricity.

Figure 4-11 shows the vehicle-trip assignment for the same week and the same pod as those shown in Figure 4-5, and for various fleet mixes. First, if two EVs of 120-mile range are allocated to the pod considered, no trips are dropped, and the two EVs are used to operate a large number of short trips (Figure 4-11a). This is consistent with the optimization results shown in Figure 4-5b. In the case where 5 120-mile range EVs are deployed to the pod considered, all reservations are still served, but more trips are served by EVs. However, if the range of the five EVs is reduced to 24 miles, then the ten vehicles in the pod (i.e., the 5 EVs and the 5 gasoline vehicles) no longer have the capability to serve the existing demand, and 29 trips get unassigned to vehicles. This is because the demand that cannot be served by EVs cannot be fully served by the 5 gasoline vehicles. In an extreme case with 9 24-mile EVs, the utilization of gasoline vehicles is very low, but 159 trips, or nearly a third of the trips are dropped.

We present the results of VPA optimization based on the VTA simulation results (see Table 4.1), as described in Algorithm 2. Figure 4-12 and Figure 4-13 show the allocation of EVs per pod using the VTA Simulation and VPA Optimization with a feasibility threshold of $\varepsilon = 1\%$ and $\varepsilon = 5\%$. Note, first, that the VPA strategy obtained with $\varepsilon = 1\%$ is relatively

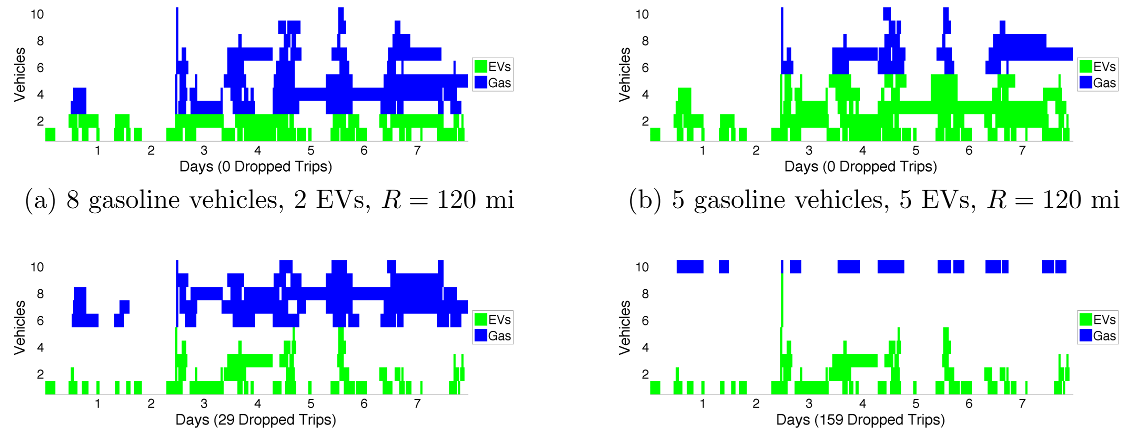

(a) 8 gasoline vehicles, 2 EVs, $R = 120$ mi

(b) 5 gasoline vehicles, 5 EVs, $R = 120$ mi

(c) 5 gasoline vehicles, 5 EVs, $R = 24$ mi

(d) 1 gasoline vehicle, 9 EVs, $R = 24$ mi

Figure 4-11: Vehicle-Trip Assignment in different fleeting scenarios

similar to that based on the VTA optimization results (see Figure 4-6). In this case, $\sim 47\%$ of vehicles converted to EVs. However, with $\varepsilon = 5\%$ progresses much further, with nearly all pods assigned at least 1 EV, and more than 85% of the fleet converted. Beyond 50% EVs, note that most new EVs added are deployed to pods in which an EV is already present, reflected by the negative slope of the "EV = 1" curve and the increasing number of pods with two EVs or more.

Figures 4-14 and 4-15 show the system-wide gasoline savings and the system-wide number of dropped trips, respectively, obtained from this procedure as a function of the number of EVs deployed in the system, for values of the feasibility threshold ε spanning from 1% to 10% (i.e., we vary the number of permitted dropped trips per pod from 1% to 10%). We also plot the corresponding gasoline savings obtained from the VTA optimization (Section 4.3). As expected, increases in ε lead to increases in the number of dropped trips, but also to larger numbers of EVs deployed and to larger gasoline savings. In the case of $\varepsilon = 1\%$ (i.e., no more than 1% of dropped trips per pod), the gasoline savings are below those of the VTA optimization. In other words, the order of reservation is such that our heuristic results in lower gasoline savings than could be achieved if the vehicle-trip assignment was optimized using *ex post* information. In contrast, for the largest values of ε, we can convert nearly the entire fleet of vehicles to EVs up to a point where nearly all trips are either served by

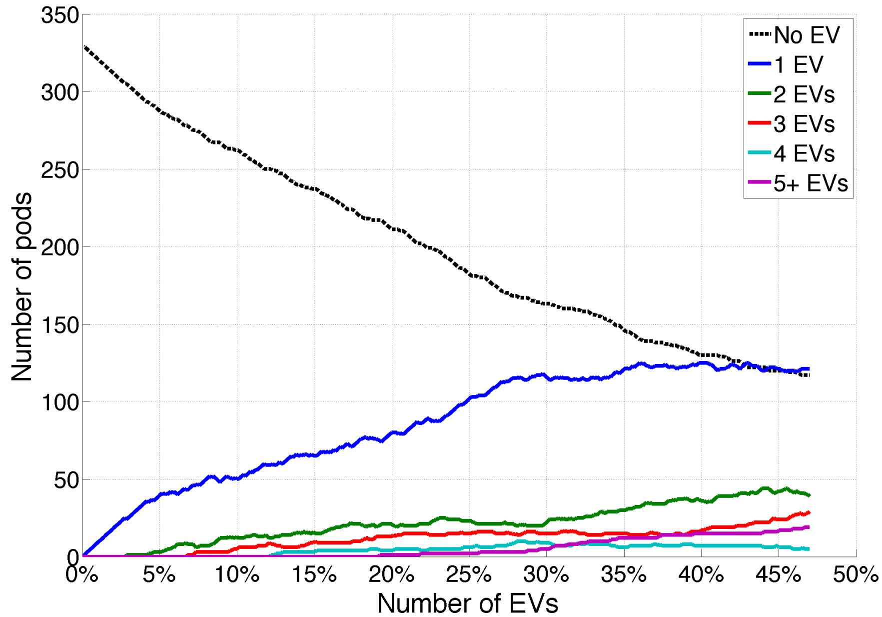

Figure 4-12: Number of EVs assigned per pod with $\varepsilon = 1\%$.

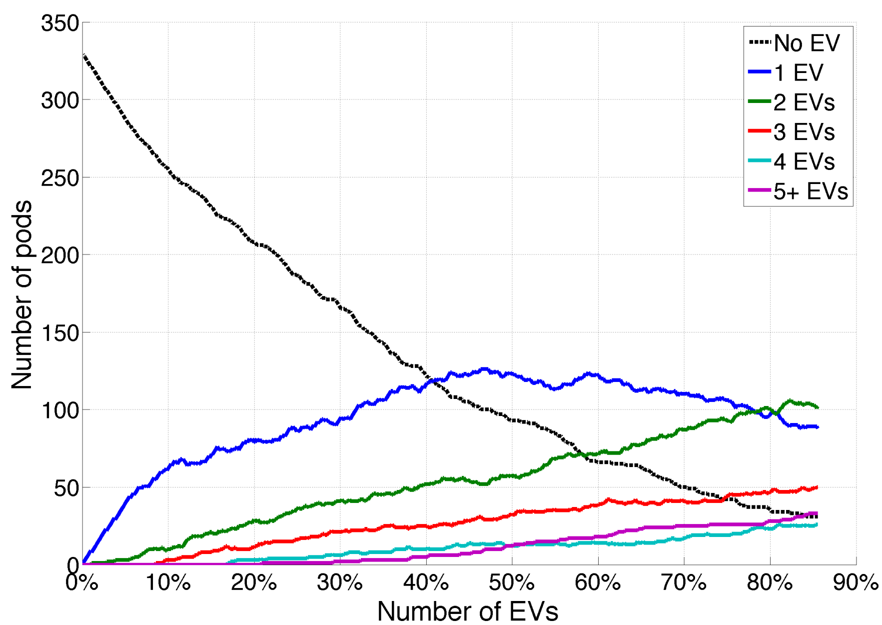

Figure 4-13: Number of EVs assigned per pod with $\varepsilon = 5\%$.

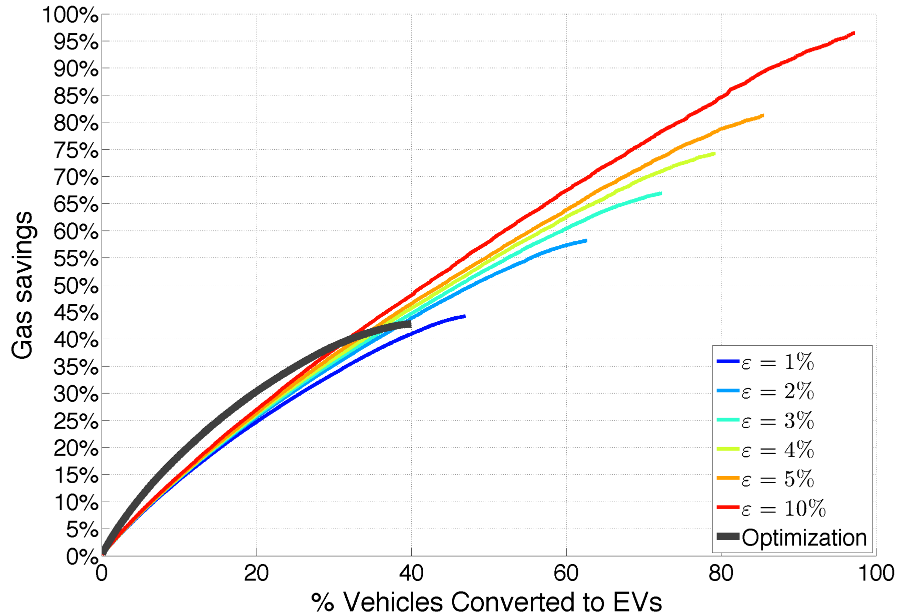

Figure 4-14: Aggregate gasoline savings for different feasibility thresholds ε

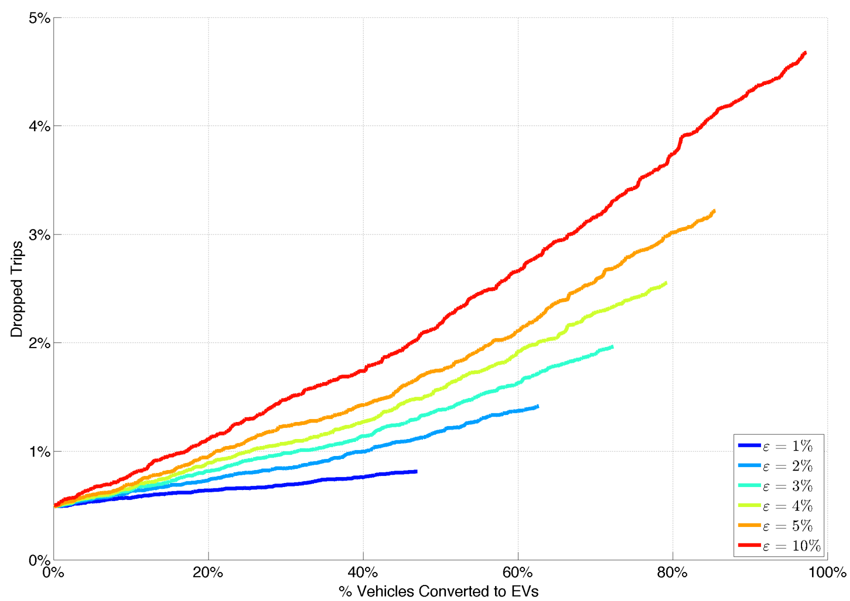

Figure 4-15: Aggregate number of dropped trips for different feasibility thresholds ε

an electric vehicle or dropped, so gasoline consumption is reduced virtually to zero. On the other hand, this comes at the cost of a larger number of dropped trips (up to 5% of all trips) and, thus, lost revenue. Note, also, that the number of dropped trips increases super-linearly as increasing fractions of the fleet get converted to EVs. Significant numbers of EVs can be deployed in the system without forcing many trips to be dropped, while the most aggressive deployment strategies lead to higher numbers of dropped trips. In combination, the concavity of the gasoline savings and the convexity of the number of dropped trips as a function of the number of EVs deployed suggest that significant benefits can be captured through the allocation of EVs. Specifically, for all levels of permitted dropped trips up to 5%, the conversion of 40% to 85% of the fleet to EVs (with a 120-mile range and a 8 kW charging rate), would induce 40% to 80% reductions in aggregate gasoline consumption, without dropping more than 1% to 4% of trips systemwide.

In summary, the VTA simulation introduces a new trade-off: the deployment of EVs increases gasoline savings, but it also creates situations where all rentals cannot be served, resulting in dropped trips and lost revenue. When dropped trips are held to a maximum of 1% per pod, the gasoline savings and ability to deploy EVs are both slightly reduced from the levels achieved in VTA optimization. With larger permitted levels of dropped trips, the simulation strategy permits greater levels of EV deployment, but potentially at higher costs of lost revenue.

4.4.4 VTA Simulation Results: EV Rejection

The results shown thus far relied on the assumption that all trips get assigned in priority to EVs, subject to feasibility constraints. However, the utilization of electric vehicles in carsharing systems is subject to some uncertainty regarding users' technological preferences and, in turn, their willingness to accept PEVs. (Le Vine et al., 2014b; Zoepf and Keith, 2015) Large numbers of users who are unwilling to accept EVs would have a negative impact on the benefits of EV deployment and utilization: at the limit, if all users rejected EVs, electric vehicles would be completely unused and the remaining portion of the fleet that is powered by gasoline may not be able to serve all the demand. This section investigates the impact that random unwillingness of users to accept EVs may have on EV utilization and the resulting gasoline savings. As described earlier, we model user decisions by a binary random variable, such that each user accepts (resp. rejects) to use an EV with probability P_{accept} (resp. $1 - P_{accept}$). In the case of rejection, the trip is assigned to a gasoline vehicle, if available, and dropped otherwise. We run 20 instances of the VTA simulation, and we perform Monte Carlo sampling on users' EV acceptance. We consider values of P_{accept} from 0.50 to 1 in increments of 0.05, and we store the gasoline savings and the set dropped trips in each scenario.

Note that, because of the stochastic nature of refusal of electric vehicles, we must redefine our "feasibility" criteria since the number of dropped trips and gasoline savings will vary from trial to trial based on *which* trips are assigned to EVs, or not (even though the probability of acceptance remains unchanged). For the purposes of simulating user rejection of EVs, we define an allocation of EVs to a pod "feasible" if no more than $\varepsilon = 5\%$ of the trips per pod are dropped due to the replacement of gasoline vehicles by EVs in 90% of trials.

The results of the EV rejection simulations are shown in Figure 4-16 for the different values of P_{accept}. As expected, reductions in the probability of acceptance of EVs result in a higher likelihood of unfeasibility in each trial, and, in turn, in lower numbers of EVs that can be feasibly deployed in the system and in lower resulting gasoline savings. At very high levels of user acceptance of EVs (e.g., $P_{accept} = 90 - 95\%$), the feasibility of EV deployment is very high, and more than half of the fleet could be converted to EVs, yielding reductions in gasoline consumption of up to 40-50%. However, at lower levels of acceptance, the ability to deploy electric vehicles as well as the resulting gasoline savings are reduced significantly. For instance, a 65% probability of acceptance reduces the feasible deployment of EVs to approximately 22% of the fleet and the gasoline savings to 20% of total consumption. By the time the probability of acceptance falls to 50%, only 15% of the fleet could be converted to EVs, and gasoline savings fall to 12%.

In summary, user acceptance of electric vehicles is a critical part of the utilization, hence

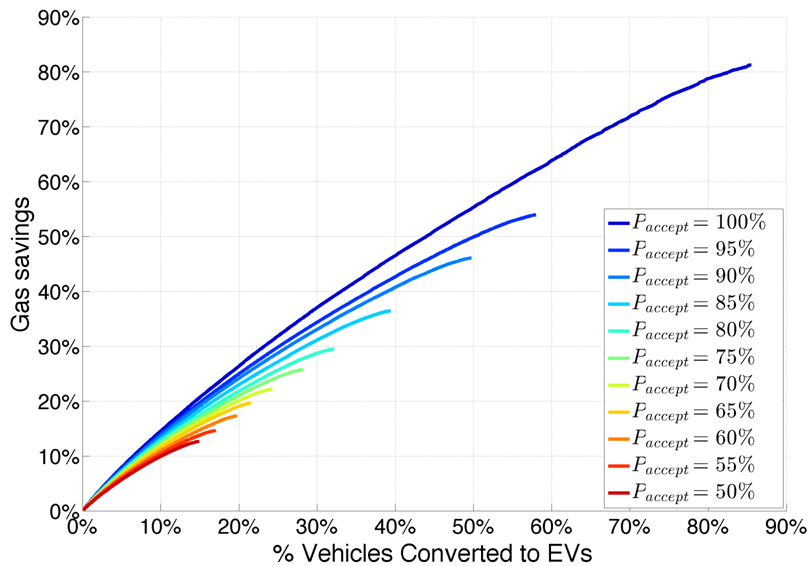

Figure 4-16: Gasoline savings for various levels of user acceptance of EVs.

the deployment, of EVs into carsharing services. If users are unwilling to use the electric vehicles that are deployed, then EV utilization will be lower and thus gasoline consumption will be higher than optimal. This suggests that carsharing providers may consider education or pricing strategies that emphasize the environmental benefits or incentivize the usage of electric vehicles to unlock the performance benefits associated with EV deployment and utilization.

4.5 Implications for System-wide PEV Deployment

As we have seen in the previous sections, the deployment of electric vehicles in replacement of gasoline vehicles involves a trade-off between shifts in vehicle miles traveled (VMT) from gasoline vehicles to electric vehicles and a higher likelihood of unserved demand (this materializes in the form of unfeasibility in the VTA optimization or higher numbers of dropped trips in the VTA simulation). In addition, the acquisition costs of EVs are typically higher than those of their gasoline counterparts. The quantification of these trade-offs is therefore necessary to inform Vehicle-Pod Allocation (VPA) decisions, i.e., (i) how many EVs to deploy and (ii) in which pods to allocate them.

In this section, we use the VTA results (Sections 4.3 and 4.4) to assess and compare several VPA strategies. The objective is twofold. First, we develop VPA heuristics that represent EV deployment strategies that a carsharing operator could easily implement, even in the absence of an advanced model such as the ones developed in this work. The comparisons of several heuristics informs which ones are likely to be the most effective. Second, comparisons between the VPA optimization developed in this work and these VPA heuristics quantify the benefits associated with integrating tactical vehicle utilization decisions (VTA) into strategic EV deployment (VPA).

4.5.1 VPA Strategies

We consider three classes of strategies: (i) VPA optimization, (ii) VPA heuristics, and (iii) random VPA.

VPA optimization involves deploying EVs in the carsharing fleet as described in Algorithm 1 (VPAopt/VTAopt, based on VTA optimization results) and in Algorithm 2 (VPAopt/VTAsim, based on VTA simulation results). For the VPAopt/VTAsim strategy, we use of feasibility threshold equal to $\varepsilon = 1\%$, i.e., we do not allow for any more than 1% of dropped trips per pod. They result in 364 EVs deployed in the system (out of 911 vehicles) for VTAopt/VTAopt and 428 for VTAopt/VTAsim.

In contrast, random VPA involves a completely random allocation of vehicles to pods without relying on any information about pod size, distance traveled, etc. For each total number of EVs N_{all}^{EV}, we sample without replacement the gasoline vehicles that are converted to EVs. We perform 1,000 Monte Carlo samples of VPA, and evaluate the impacts of such EV deployment on vehicle utilization for each sample.

The VPA heuristics lie in-between. They are simple, readily implementable strategies for EV deployment that rely on aggregate pod-level statistics (and not on the VTA results provided in this work). We base these heuristics on metrics that are correlated with the benefits of EV utilization. Two such metrics we have found effective are: (i) the number of vehicles per pod p (i.e., N_p^*), and (ii) the total distance of trips in pod p under EV range (i.e., $\sum_{i=1,\dots,n_p|D_i \leq R} D_i$, which we denote by $D_R(p)$). For each pod p, the heuristics considered involve replacing up to 50% of the fleet (i.e., we replace $\left\lfloor \frac{N_p^*}{2} \right\rfloor$ gasoline vehicles by EVs, for each pod p). This results in 387 EVs deployed in the system. For each of the two proxys of EV deployment performance, we sort pods from the most attractive one (i.e., the one with the largest number of vehicles, the one with the largest distance under range) to the least attractive one, and we consider a "wide" deployment strategy that allocates EVs one pod at a time, in order, until up to 50% of vehicles in all pods are replaced by EVs, and a "deep" deployment strategy that replaces up to 50% of vehicles in each pod, one at a time and in order. In sum, we consider a total of four heuristics:

Large Pods, Wide (LPW) For each pod p in decreasing order of N_p^*, replace one gasoline vehicle by an EV, if the resulting number of EVs in pod p is lower than 50% of its fleet. Continue until up to 50% of the fleet of each pod is replaced.

Large Pods, Deep (LPD) For each pod p in decreasing order of N_p^*, replace up to 50% of the gasoline vehicles of pod p by EVs.

Distance under Range, Wide (DRW) For each pod p in decreasing order of $D_R(p)$, replace one gasoline vehicle by an EV, if the resulting number of EVs in pod p is lower than 50% of its fleet. Continue until up to 50% of the fleet of each pod is replaced.

Distance under Range, Deep (DRD) For each pod p in decreasing order of $D_R(p)$, replace up to 50% of the gasoline vehicles of pod p by EVs.

Note that all the heuristics are equivalent when all the 387 EVs are deployed (they will be allocated in a way that fills up to 50% of the fleet of each pod). However, they differ by the *order* through which the EVs are allocated to the different pods.

4.5.2 Assessment of VPA strategies

We assess the different VPA strategies presented in the previous section using the VTA optimization and VTA simulation results. Under VTA optimization, we deploy only the vehicles that do not result in unfeasibility cases. In other words, if a VPA strategy results in more EVs allocated to a given pod than the maximal number of EVs that can be feasibly deployed to this pod, then we "accept" only the maximal number of EVs, and store the resulting gasoline savings. Under VTA simulation, we do not limit the number of EVs deployed to any pod, and store the gasoline savings, the set of dropped trips and associated revenue hours. We use these results to quantify the net private benefits of EV deployment and utilization for the carsharing operator, i.e., the improvement in the operator's profitability, as compared to the all-gasoline scenario. It is equal to the value of the reduction in gasoline consumption, minus the added vehicle acquisition costs and the lost revenue due to dropped trips. The value of gasoline savings is equal to the difference between the price of saved gasoline and the price of electricity required to travel the same distance with EVs. We denote by P_g the price of gasoline (in \$/gallon) and by P_e the price of electricity (in \$/kWh). The impact of EV deployment on acquisition costs is modeled by a monthly premium paid for each EV deployed, denoted by P_{EV} (in \$/month). This assessment captures the difference in the monthly cost of leasing an EV and that of leasing a comparable gasoline vehicle. Finally, the lost revenue associated with dropped trips is equal to the total duration of dropped trips times the hourly price of rental, denoted by p_h (in \$/hour). The net benefit of EV deployment and utilization is a function of the number of EVs per pod (i.e., of $N_1^{EV}, ..., N_P^{EV}$) and of the VTA strategy (e.g., the VTA optimization, the VTA simulation developed in this work), which we refer to by a superscript S. We denote the net benefit by $\Delta\pi^S\left(N_1^{EV}, ..., N_P^{EV}\right)$, and express it in Equation (4.5.2) below, where $\Phi_p^S\left(N_p^{EV}\right)$ and $\mathcal{T}_p^S\left(N_p^{EV}\right)$ denote the gasoline savings and the set of dropped trips, respectively, in pod p under the VTA strategy S.

$$\Delta\pi^S\left(N_1^{EV}, ..., N_P^{EV}\right) = \sum_{p=1}^{P}\left\{\Phi_p^S\left(N_p^{EV}\right)\left(P_g - \frac{\delta}{\alpha}P_e\right) - N_p^{EV}P_{EV} - \sum_{i \in \mathcal{T}_p^S\left(N_p^{EV}\right)}(t_i^E - t_i^S)p_h\right\}$$

This model does not capture several complexities that affect the costs and benefits of EV deployment and utilization for a carsharing operator. First, the benefits of reductions in gasoline consumption depend on the variations in the prices of fuel and electricity. Second, modeling the difference in acquisition costs between EVs and gasoline vehicles by a fixed monthly premium does not account for the differences in appearance or equipment between EVs and gasoline vehicles, the uncertainty associated with EV adoption subsidies from manufacturers and regulators, the vehicle acquisition strategies by carsharing operators, etc. Third, the assumption of revenues being proportional to rental durations, while consistent with the typical carsharing business model, does not account for the daily rates of rentals for the longest reservations. While these considerations are beyond the scope of this work, we address them through sensitivity analyses. We fix the price of electricity P_e to \$0.1/kWH and vary the price of gasoline P_g from \$2/gal to \$5/gal. These values reflect the range of gasoline prices observed nationwide recently. We vary the monthly premium associated with EV acquisition from \$50/month to \$200/month. Finally, we set the rental price to \$10/hour.

Figure 4-17 and Figure 4-18 shows the net benefits of the different VPA strategies under VTA optimization (Figure 4-17) and VTA simulation (Figure 4-18) obtained with a price of gasoline equal to \$3/gal and an EV premium equal to \$100/month. The observations are threefold. First, the heuristics perform significantly better than random VPA, as they result in higher gasoline savings and fewer dropped trips, thus higher net benefits. This suggests that filling the largest pods first or the pods with the longest distance under the EV range, seem to be effective guidelines to inform EV deployment by carsharing operators. Second, all heuristics do not perform equally well for all values of N_{all}^{EV}. The first EVs that are deployed seem to yield the largest marginal benefits under the "wide" strategies (one vehicle per pod at a time), as compared to the "deep" strategies (one pod at a time). In contrast, the last EVs that are deployed yield the largest marginal benefits under the "deep" strategies. This is consistent with the results shown in Figures 4-6. Third, the VPA optimization strategies seem to perform better than *any* of the heuristics, i.e., they result in larger gasoline savings and/or fewer dropped trips, hence in higher net benefits. This suggests that the deployment of EVs should balance deploying EVs in several pods and allocating several EVs to the busiest pods, as a function of the demand patterns and vehicle capabilities. The integration of tactical vehicle utilization into strategic EV deployment proposed in this work provides decision-making support for such decisions and can thus result in significant system-wide benefits.

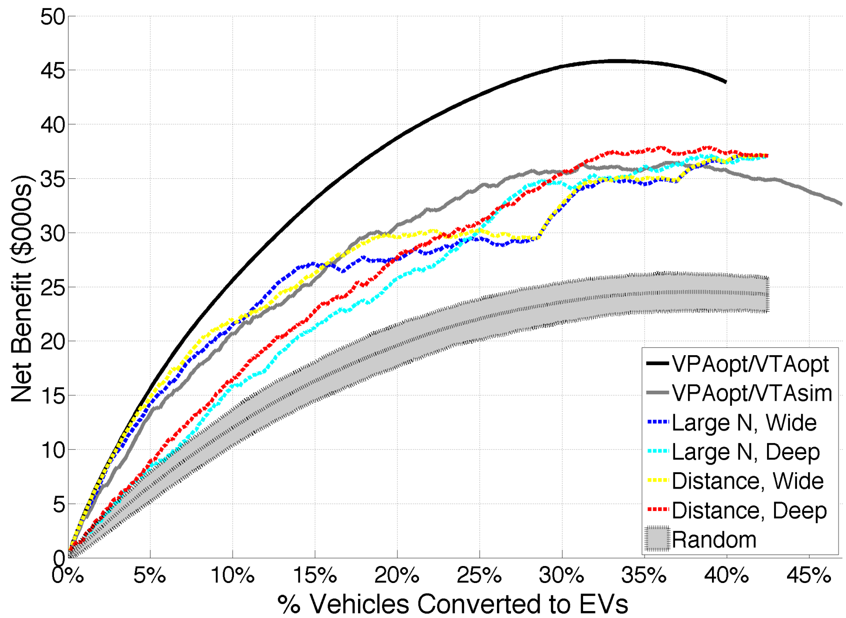

Figure 4-17: Net benefits of EV deployment and utilization under various VPA strategies for VTA Optimization results.

Finally, Figure 4-19 and Figure 4-20 quantify the magnitude of the net benefits under the VPAopt/VTAopt deployment strategy, evaluated under the VTA optimization (Figure 4-19) and the VTA simulation (Figure 4-20), for different values of the price of gasoline P_g (represented by different colors) and of the EV premium P_{EV} (represented by solid, dashed, or dotted lines). Note that the values of P_g and P_{EV} impact the profit-maximizing number of EVs deployed in the system, as well as the magnitude of the benefits resulting from EV deployment and utilization. In a large number of scenarios, the deployment of EVs can result in significant improvements in carsharing operators' profitability, as compared to the scenario where all vehicles are powered by gasoline. Under realistic assumptions (e.g., P_g = \$3/gal and P_{EV} = \$100/mo), we estimate that 20% to 40% of the fleet can be converted to EVs, yielding a net benefit of the order of \$20,000 in the month of January 2012 in Boston, MA,

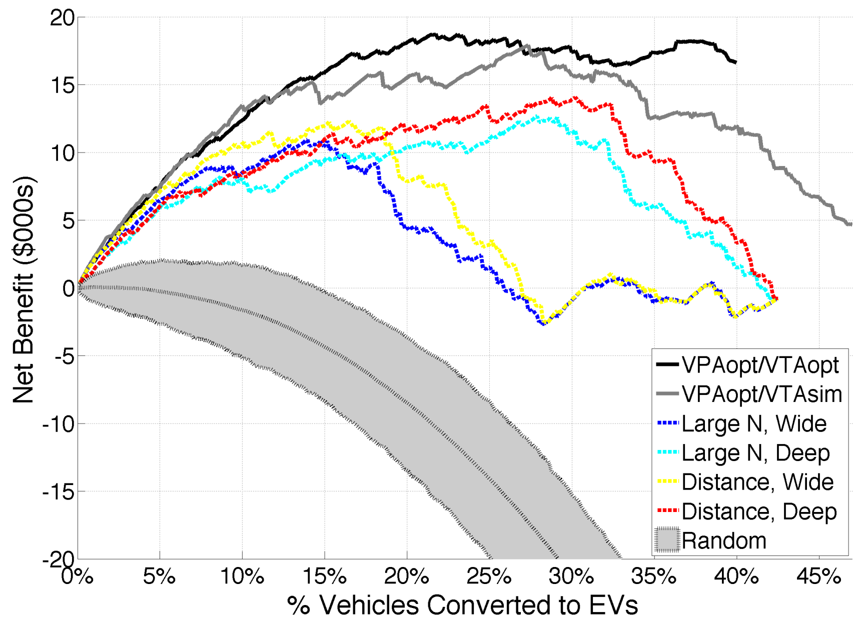

Figure 4-18: Net benefits of EV deployment and utilization under various VPA strategies for VTA Simulation results.

even under the simple heuristic-based VTA simulation considered in this work. Given the uncertainty in parameters, maximizing this benefit is subject to variability, but we observe that the net benefit is are relatively "flat" around its optimal value. This suggests that even if a suboptimal number of EVs is deployed (e.g., because of errors in the estimation of the price of gasoline or the EV premium), this may not affect the operator's profitability significantly. Finally, the net benefits are significantly larger under VTA optimization than under VTA simulation. In the case discussed above (P_g = \$3/gal and P_{EV} = \$100/mo), they are estimated at over \$40,000 for the month of operations under VTA optimization (vs. around \$20,000 under VTA simulation). This suggests that an important avenue for future research involves the design of VTA policies for carsharing operators to optimize the dynamic assignment of trips to gasoline vehicles and EVs as reservations are made, and ultimately capture more of the potential benefits associated with EV utilization.

In summary, the analysis from this section has shown that: (i) the deployment and utilization of EVs can result in significant improvements in carsharing operators' profitability,

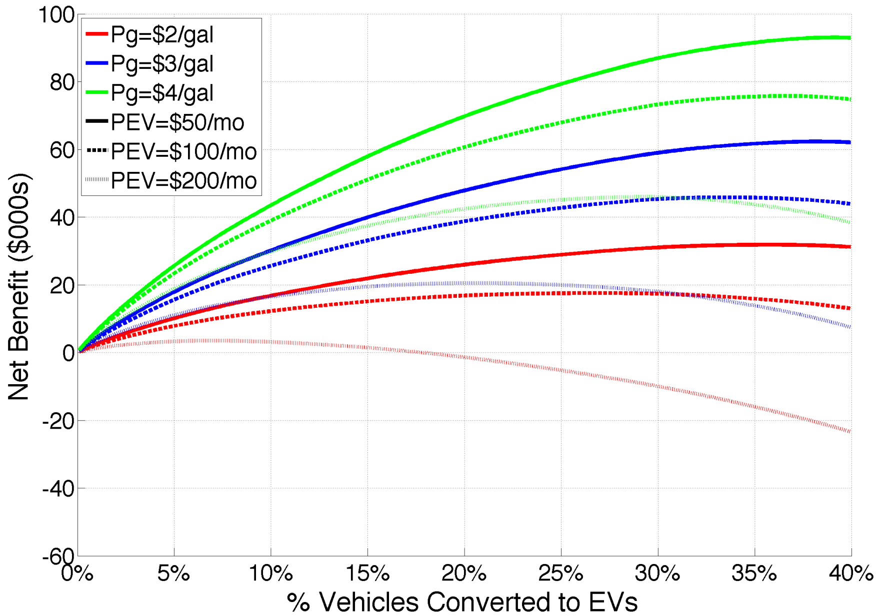

Figure 4-19: Net benefits of EV deployment (under the VPAopt/VTAopt strategy) for VTA Optimization results.

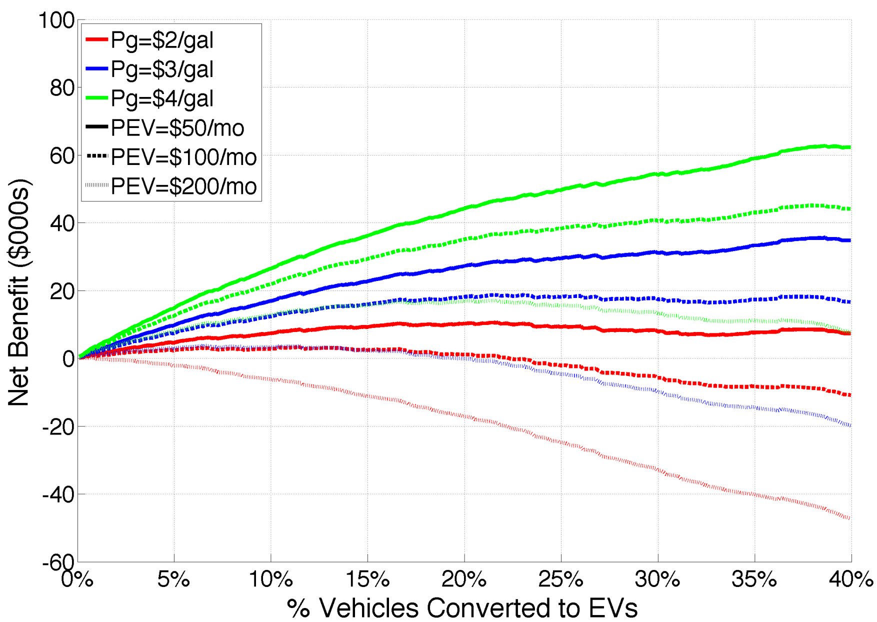

Figure 4-20: Net benefits of EV deployment (under the VPAopt/VTAopt strategy) for VTA Simulation results.

(ii) some easily implementable VPA and VTA heuristics can capture some of these benefits, (iii) the integration of VTA utilization patterns into the design of more advanced VPA strategies (Algorithms 1 and 2) can result in added benefits, and (iv) the gap between VTA optimization and VTA simulation can be large, thus creating the scope for potential improvements in VTA policies.

Note, finally, that the analysis presented in this section has focused on the carsharing operator's perspective and modeled exclusively the private costs of gasoline consumption, of electricity consumption, and of lost demand. The results therefore suggest that carsharing operators may have significant incentives to deploy EVs in their systems even in the absence of any incentivizing policy at current market prices of electric vehicles. Nonetheless, the social costs of gasoline and electricity consumption as well as the social cost of lost demand may include external components that are not borne by the carsharing operator itself. In turn, the welfare-maximizing number of EVs may differ from the optimal number of EVs from the carsharing operator's perspective. If this is the case, incentivizing policies may be considered to spur EV adoption (e.g., subsidies to reduce the EV acquisition premium, adoption mandates, etc.). A full analysis of this policy question is beyond the scope of this work, but represents an important avenue for future research. The approach developed in this work provides the methodological foundation to tackle these strategic issues, while integrating how EVs and gasoline vehicles can be utilized at the tactical level in carsharing systems.

4.6 Findings

This chapter developed an integrated approach to the deployment and utilization of Plug-in Electric Vehicles (PEVs) in a round-trip carsharing system. This approach first analyzes how gasoline vehicles and PEVs can be utilized at the tactical level to serve consumer demand in order to reduce gasoline consumption, subject to vehicle availability constraints and to technological constraints associated with PEV utilization. We referred to this problem as the Vehicle-Trip Assignment (VTA) problem and developed an optimization model and a simulation algorithm to solve it. The approach, then, uses VTA results to inform the deployment of PEVs in a carsharing system at the strategic level to capture the largest gasoline savings and profitability improvements, given how vehicles can then be tactically utilized. We referred to this problem as the Vehicle-Pod Allocation (VPA) problem.

The application of the modeling approach to the Zipcar system in Boston, MA, has yielded four major results. First, and foremost, the deployment and utilization of PEVs can yield very significant system-wide benefits that materialize in the form of reduced gasoline consumption and limited impacts on the system's ability to serve demand, which manifest themselves as improvements in the carsharing operator's profitability while simultaneously reducing the environmental footprint of the system. Second, PEV benefits seem much more sensitive to the range of PEVs than to their charging capabilities. Third, the allocation of vehicles to pods (VPA) has significant impacts of system-wide performance. This work provides decision-making support to optimize vehicle specifications and placement. Finally, the carsharing operator's policies to assign trips to vehicles has an important effect on the ultimate benefits of PEV deployment and utilization.

The major practical implication of this work is that the incentives for PEV deployment in carsharing fleets are significant. From the perspective of a carsharing operator whose fleet comprises gasoline vehicles exclusively, the deployment of PEVs to replace some portion of the gasoline fleet would result in gasoline savings likely to outweigh the constraints on vehicle utilization that it creates. From a policy perspective, these results suggest that carsharing systems can represent an important avenue for PEV deployment Since deployment at limited levels has been shown to be profitable, incentives for PEV deployment in carsharing systems would reinforce the private incentives demonstrated in this work.

The benefits associated with the deployment and the utilization of PEVs have been found to depend not only on the number of PEVs deployed, but also on *how* they are deployed (e.g., to which pods they are allocated) and on *how* they are utilized (e.g., how trips are assigned to them). Based on these interdependencies, tighter integration between tactical PEV utilization and strategic PEV deployment (from an operator's standpoint) or

PEV adoption incentivizes (from a policy standpoint) is recommended. At a policy level, this finding suggests that incentives for PEVs should take their usage into consideration. Current ZEV policies, which provide bonus ZEV credits for shared PEVs only in ZEV states with no consideration of usage, may create an incentive to under-deploy Plug-in Vehicles in some locations and over-deploy them in others. The integrated approach developed in this work provides a modeling architecture that can systematically address such problems of PEV deployment and utilization to promote efficient and sustainable carsharing systems.

Chapter 5

Conclusions and Future Work

> **Takeaway 1**
>
> Integration of Plug-in Vehicles into existing carsharing services offers the potential for economic and environmental benefits.

Carsharing services facility a transition from the "swiss army knife" mentality of a vehicle purchaser, to the "right tool for the job" mentality approach of carsharing. Rather than using one large vehicle for all purposes, carsharing users can select small, efficient cars for basic travel, but choose vehicles with more passenger or cargo space for specific trips where these capabilities are needed. The deployment of some EVs in carsharing services enhances this flexibility in vehicle assignment, allowing EVs to be used for shorter trips and gasoline vehicles for longer trips beyond the range of EVs, or spaced too closely to allow time for recharging.

The results shown in this thesis validate the intuition expressed by Kley et al. (2011), Dijk et al. (2013) and Green et al. (2014) that carsharing and electric vehicles are natural complements. By analyzing actual usage data and plausible assumptions of charging rates and range, we find that fractions of the round-trip carsharing fleet could be switched from gasoline to electric vehicles, simultaneously lowering private operating costs *and* reducing gasoline consumption. However, private benefits are maximized when a only a portion of the fleet (20-40% depending on simulation parameters) of the fleet is converted. Further deployment of EVs results in costs and service loss that outweigh the benefits of gasoline savings.

Actual EV deployment strategies employed by carsharing services today differ from this recommendation. Only a small fraction of round-trip carsharing vehicles in the United States are PEVs. On the other hand, one-way carsharing services in a number of cities worldwide are converting their entire fleets, or very large fractions of them, to PEVs. While the conclusions from this work are not generally applicable to one-way carsharing systems due to their different demand and operational characteristics, it is unclear whether such a strategy is an efficient use of PEVs given that such systems also must balance the range and recharging requirements, and in some cases are unable to ensure that a vehicle will be parked near a charging facility, or may need to be driven out of the way to charging facilities.

Takeaway 2

Effectiveness of PEVs in reducing fuel consumption depends strongly on how they are used and their range, and relatively less on the availability of rapid recharging.

More broadly, these results highlight that not all PEVs are equal: how and where electric vehicles are deployed is critical. Current policy incentives fail to capture the nuances of *how* PEVs are used, and focus simply on getting more of them on the road. Even within the carsharing environment, large variations between performance in vehicle locations can drastically affect the benefit from the deployment of a single PEV. Evidence from the simulation in Chapter 4 indicates that random deployment of electric vehicles is generally a money-losing proposition. A carsharing provider that deploys EVs without sufficient strategic planning may see disappointing results, but a few simple heuristics can greatly improve environmental and financial outcomes.

From a technical perspective, the feasibility of deploying EVs is strongly dependent on their range. However, sensitivity to recharging rate is far lower: the Level 2 chargers commonly in use today, which charge at rates of approximately 4-8kW, are sufficient to realize the the vast majority of benefits. Results from Chapter 4 suggest that increasing charging rates beyond Level 2 standards only minimally increased the feasible deployment of EVs in this study.

Takeaway 3

Carsharing services expose large numbers of users to hybrids, but few users have any experience with PEVs.

Responses to the carsharing user survey indicate that a vast majority of users have experience with a hybrid vehicle: nearly 400 users have been exposed to hybrid technology for each hybrid in service. While this work does not address the question of how exposure through carsharing alters attitudes over time, the large exposure rate of hybrid technology and positive attitudes to the technology suggest a relationship between exposure and attitudes. Whether exposure to PHEVs and EVs will lead to the same level of comfort remains to be seen.

Takeaway 4

User attitudes to Plug-in Electric Vehicles vary widely, even for short-term usage.

However, as shown in Chapter 2, the importance of vehicle type to carsharing users is small relative to other service attributes. In the near term, carsharing operators could overcome reluctance to use PEVs with incentives that compensate in other ways. One easy way to make PEVs more attractive would be with lower prices (a strategy some carsharing operators have adopted with hybrids). Alternatively, the vehicles could be made more attractive with non-monetary incentives: increased schedule or location flexibility, for instance. If preferences towards PEVs are changed with greater exposure, the need for direct incentives will be reduced over time, and barriers to their wider deployment will fall away naturally.

Takeaway 5

As PEVs grow in numbers, users will need to become more aware of their own travel behavior.

Feasible deployment of EVs depends on the willingness of users to accept them for trips when an EV has sufficient range. However, it is clearly *not* desirable for users to select a car that is unsuitable for their trip – particularly if the vehicle has insufficient range for their travel. In order to properly choose a vehicle that suits their trip, users of range-limited vehicles will inevitably need to become more aware of their travel behavior if they are to make an informed decision about whether or not an EV will meet their needs. Paradoxically, if carsharing users drive less frequently, they may generally become less familiar with driving distances and evidence from Chapter 2 suggests that user distance estimates may be biased high. Carsharing providers can help users with travel planning tools and contingency plans for stranded users, but inevitably some of this responsibility will rest with users.

5.1 Future Work

While Chapter 2 identifies heterogeneity in preferences for vehicle types, there are additional details to be investigated. The stated preference (survey-based) work presented will ideally be complemented by revealed preference studies that analyze actual user rentals, but there are a number of challenges to such work. First, some carsharing models offer little or no choice between vehicle types, so travel model choice between services is confounded with the choice of vehicle technology. For carsharing services that offer a mix of vehicle types, many incorporate systematic differences in pricing between gasoline and Plug-in Vehicles, confounding pricing and technology effects. One potential service that escapes these problems are peer-to-peer carsharing services, which allow vehicle owners to set their own price. However, in peer-to-peer business models the choice decision is no longer one-sided, and both parties must mutually agree on a transaction, introducing additional the need to evaluate the choices of both the owner and renter.

The integration of open-ended survey responses into choice modeling efforts has shown promise as a way to identify user attitudes while simultaneously minimizing respondent burden. In this thesis the choice model and latent variable model have been integrated, but topic modeling has been performed separately. In the longer term these techniques could be combined, simultaneously estimating topic models and choice responses, increasing the efficiency of the models, although potentially with the consequence of even greater estimation times.

The results of Chapter 4 are intended to provide a range of estimates, from best-case optimization results, to more conservative results based on a simple set of heuristics. However, these estimates can be improved in a number of key ways. First, the simulation of trip allocation (VTASim) could be improved while maintaining realistic restrictions on the use of future information. For instance, (i) re-optimizing the deployment of vehicles after each new reservation; (ii) the integration of users' behaviors, e.g., users' decisions to recharge PEVs upon completion of a rental or not, the heterogeneity of users' preferences regarding vehicle technologies; and (iii) incorporating heterogeneity of driving behavior that can affect the rates of gasoline or electricity consumption.

5.2 PEVs and Other Mobility Models

Personal mobility models and vehicle technologies are both evolving quickly. Round trip carsharing has now existed for more than a decade and is commonplace in many larger cities. One-way carsharing (both free floating and station-based) are increasingly common. All of these models now compete with Transportation Network Companies (e.g. Uber and Lyft) which offer conveniently-scheduled replacements for traditional taxi services in many of the same markets. The increasingly competitive field of urban transportation means that the types of trips for which carsharing is used may change as travelers select and thus identify the best options for their trip type.

In systems where the traveler is not driving (TNCs, or in the very long term, autonomous vehicles) passengers will probably not know or care which energy source is powering the vehicle. However, the same fundamental tradeoffs of range and recharging time vs. the need to keep expensive assets (vehicles) earning revenue will remain important even as electric vehicles are incorporated into more highly centralized or automated transportation systems.

However, the "chicken and egg solution" of PEV deployment in carsharing does not necessarily hold for peer-based systems or TNCs. Since vehicle purchase decisions are decentralized, vehicle owners may be unlikely to purchase electric vehicles with limited range as these vehicles may limit their revenue earning potential. Such systems may need to provide additional incentives to vehicle owners such as subsidized charging facilities or passenger trip length information if they wish to support greater use of PEVs.

Bibliography

Al-Alawi, B. M. and Bradley, T. H. (2013). Review of hybrid, plug-in hybrid, and electric vehicle market modeling Studies. *Renewable and Sustainable Energy Reviews*, 21:190–203.

Alvarez-Daziano, R. (2011). Incorporating pro-environmental preferences toward green automobile technologies through a Bayesian Hybrid Choice Model. *Transportmetrica*.

Arentze, T., Dellaert, B., and Chorus, C. (2013). Incorporating Mental Representations in Discrete Choice Models of Travel Behaviour. Netspar Discussion Papers 11/2013-075, Network for Studies on Pensions, Aging and Retirement.

Atasoy, B., Glerum, A., Hurtubia, R., and Bierlaire, M. (2010). Demand for public transport services: Integrating qualitative and quantitative methods. *10th Swiss Transport Research Conference, Ascona*.

Axsen, J. and Kurani, K. (2009). Early U.S. Market for Plug-In Hybrid Electric Vehicles. *Transportation Research Record: Journal of the Transportation Research Board*, 2139:64–72.

Bandivadekar, A. P. (2008). *Evaluating the Impact of Advanced Vehicle and Fuel Technologies in the U.S. Light-Duty Vehicle Fleet*. PhD Dissertation, Massachusetts Institute of Technology, Technology, Management and Policy.

Barnes, J. A. and Rutherford, G. S. (2001). Stated Preference Survey Design and Analysis for Pre-Implementation Evaluation of the Seattle Car Share Program.

Ben-Akiva, M. and Lerman, S. R. (1985). *Discrete Choice Analysis: Theory and Application to Travel Demand*.

Ben-Akiva, M., Mcfadden, D., Train, K., Walker, J., Bhat, C., Bierlaire, M., Bolduc, D., Boersch-Supan, A., Brownstone, D., Bunch, D. S., Daly, A., Palma, A. D., Gopinath, D., Karlstrom, A., and Munizaga, M. A. (2002). Hybrid Choice Models: Progress and Challenges. *Marketing Letters*, 13(3):163–175.

Bendixson, T. and Richards, M. G. (1976). Witkar: Amsterdam's self-drive hire city car. *Transportation*, 5(1):63–72.

Berry, I. (2010). *The Effects of Driving Style and Vehicle Performance on the Real-World Fuel Consumption of U.S. Light-Duty Vehicles*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA.

Bierlaire, M. (2003). BIOGEME: A free package for the estimation of discrete choice models. Ascona, Switzerland.

Blei, D. M. (2012). Probabilistic Topic Models. *Commun. ACM*, 55(4):77–84.

Blei, D. M. and Lafferty, J. D. (2006). *Correlated Topic Models*.

Blei, D. M. and Lafferty, J. D. (2007). A correlated topic model of science. *The Annals of Applied Statistics*, pages 17–35.

Blei, D. M. and Mcauliffe, J. D. (2007). Supervised Topic Models. In Platt, J. C., Koller, D., Singer, Y., and Roweis, S. T., editors, *Advances in Neural Information Processing Systems 20*, pages 121–128. Curran Associates, Inc.

Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent Dirichlet Allocation. *J. Mach. Learn. Res.*, 3:993–1022.

Bolduc, D. and Alvarez-Daziano, R. (2010). On estimation of Hybrid Choice Models. Technical report.

Bolduc, D., Boucher, N., and Alvarez-Daziano, R. (2008). Hybrid Choice Modeling of New Technologies for Car Choice in Canada. *Transportation Research Record: Journal of the Transportation Research Board*, 2082(1):63–71.

Boyac, B., Zografos, K. G., and Geroliminis, N. (2015). An optimization framework for the development of efficient one-way car-sharing systems. *European Journal of Operational Research*, 240(3):718–733.

Brownstone, D., Bunch, D. S., and Train, K. (2000). Joint mixed logit models of stated and revealed preferences for alternative-fuel vehicles. *Transportation Research Part B: Methodological*, 34(5):315–338.

Cardell, N. S. and Dunbar, F. C. (1980). Measuring the societal impacts of automobile downsizing. *Transportation Research Part A: General*, 14(5-6):423–434.

Catalano, M., Lo Casto, B., and Migliore, M. (2008). Car sharing demand estimation and urban transport demand modelling using stated preference techniques. pages 33–50.

Cervero, R., Golub, A., Nee, B., University of California, B., and Institute of Urban & Regional Development (2006). San Francisco City CarShare: longer-term travel-demand and car ownership impacts. Technical report, University of California at Berkeley, Institute of Urban and Regional Development, Berkeley, CA.

Chang, M. (2015). *Performance Targets for Electric Vehicle Batteries*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA.

Chen, T., Kockelman, K., and Khan, M. (2013). Locating Electric Vehicle Charging Stations. *Transportation Research Record: Journal of the Transportation Research Board*, 2385(1):28–36.

Chon, D. M. and Heywood, J. (2000). Performance Scaling of Spark-Ignition Engines: Correlation and Historical Analysis of Production Engine Data. Technical Report SAE Paper No. 2000-01-0565, SAE Interational.

Cobb, J. (2015). December 2014 PEV Dashboard.

Daly, A., Hess, S., Patruni, B., Potoglou, D., and Rohr, C. (2012). Using ordered attitudinal indicators in a latent variable choice model: a study of the impact of security on rail travel behaviour. *Transportation*, 39(2):267–297.

Davies, J. and Kurani, K. S. (2010). Recharging Behavior of Households' Plug-In Hybrid Electric Vehicles. *Transportation Research Record: Journal of the Transportation Research Board*, 2191(1):75–83.

Davis, S. C., Diegel, S. W., and Laboratory, O. R. N. (2012). *Transportation energy data book*. Oak Ridge National Laboratory, Oak Ridge, TN :, ed. 31. edition.

de Lorimier, A. and El-Geneidy, A. M. (2013). Understanding the Factors Affecting Vehicle Usage and Availability in Carsharing Networks: A Case Study of Communauto Carsharing System from Montreal, Canada. *International Journal of Sustainable Transportation*, 7(1):35–51.

Denholm, P. and Short, W. (2006). An Evaluation of Utility System Impacts and Benefits of Optimally Dispatched. *National Renewable Energy Laboratory, Tech. Rep. NREL/TP-620-40293*.

Dijk, M., Orsato, R. J., and Kemp, R. (2013). The emergence of an electric mobility trajectory. *Energy Policy*, 52:135–145.

Dong, J. and Lin, Z. (2012). Within-day recharge of plug-in hybrid electric vehicles: energy impact of public charging infrastructure. *Transportation Research Part D: Transport and Environment*, 17(5):405–412.

Dong, J., Liu, C., and Lin, Z. (2014). Charging Infrastructure Planning for Promoting Battery Electric Vehicles: An Activity-based Approach using Multiday Travel Data. *Transportation Research Part C: Emerging Technologies*, 38(1):44–55.

Dowds, J., Hines, P., Farmer, C., Watts, R., and Letendre, S. (2010). Plug-in hybrid electric vehicle research project: Phase II report. *UVM Transportation Res. Center, Tech. Rep*, 20.

EIA (2010). Annual Energy Review, 2009. Technical Report DOE/EIA-0384(2009), Energy Information Administration.

EPA (2010). Light-Duty Automotive Technology, Carbon Dioxide Emissions, and Fuel Economy Trends: 1975 Through 2010. Technical Report EPA-420-R-10-023, Office ofTransportation and Air Quality, U.S. Environmental Protection Agency.

EPA, U. (2012). Regulatory Impact Analysis: Final Rulemaking for 2017-2025 Light-Duty Vehicle Greenhouse Gas Emission Standards and Corporate Average Fuel Economy Standards. Regulatory Impact Analysis EPA-420-R-12-016, U.S. Environmental Protection Agency.

EPA, U. (2013). Light-Duty Automotive Technology, Carbon Dioxide Emissions, and Fuel Economy Trends: 1975 Through 2013 | OTAQ | US EPA. Technical Report EPA-420-R-13-011. This report, published in October 2013, provides data on the fuel economy, carbon dioxide (CO2) emissions, and technology trends of new light-duty vehicles (cars, minivans, sport utility vehicles, and pickup trucks) for model years 1975 through 2013 in the United States.

EPA, U. (2014). Light-Duty Automotive Technology, Carbon Dioxide Emissions, and Fuel Economy Trends: 1975 Through 2014 | OTAQ | US EPA. Technical Report EPA-420-R-14-023a. This report, published in October 2014, provides data on the fuel economy, carbon dioxide (CO2) emissions, and technology trends of new light-duty vehicles (cars, minivans, sport utility vehicles, and pickup trucks) for model years 1975 through 2014 in the United States.

Fagnant, D. and Kockelman, K. (2014). The Travel and Environmental Implications of Shared Autonomous Vehicles, using Agent-based Model Scenarios. *Transportation Research Part C: Emerging Technologies*, 40:1–13.

Gaker, D., Vautin, D., Vij, A., and Walker, J. L. (2011). The power and value of green in promoting sustainable transport behavior. *Environmental Research Letters*, 6(3):034010.

Glerum, A., Atasoy, B., and Bierlaire, M. (2014). Using semi-open questions to integrate perceptions in choice models. *Journal of Choice Modelling*, 10:11–33.

Glerum, A. and Bierlaire, M. (2014). Capturing response behavior heterogeneity in the quantification of semi-open questions: a mixed discrete-continuous approach. Technical Report Report TRANSP-OR 140305, Transport and Mobility Laboratory: Ecole Polytechnique Fe de rale de Lausanne.

Glerum, A., Stankovikj, L., Thmans, M., and Bierlaire, M. (2012). Forecasting the demand for electric vehicles: accounting for attitudes and perceptions. Technical Report TRANSP-OR 120217, Transport and Mobility Laboratory, Ecole Polytechnique Fdrale de Lausanne.

Golob, T. and Gould, J. (1998). Projecting Use of Electric Vehicles from Household Vehicle Trials. *Transportation Research Part B: Methodological*, 32(7):441–454.

Gonder, J., Markel, T., Thornton, M., and Simpson, A. (2007). Using global positioning system travel data to assess real-world energy use of plug-in hybrid electric vehicles. *Transportation Research Record: Journal of the Transportation Research Board*, 2017(1):26–32.

Gonder, J. and Simpson, A. (2006). Measuring and Reporting Fuel Economy of Plug-In Hybrid Electric Vehicles. Technical report, National Renewable Energy Laboratory (NREL), Golden, CO.

Green, E. H., Skerlos, S. J., and Winebrake, J. J. (2014). Increasing electric vehicle policy efficiency and effectiveness by reducing mainstream market bias. *Energy Policy*, 65:562–566.

Hadley, S. W. and Tsvetkova, A. A. (2009). Potential impacts of plug-in hybrid electric vehicles on regional power generation. *The Electricity Journal*, 22(10):56–68.

He, F., Wu, D., Yin, Y., and Guan, Y. (2013). Optimal Deployment of Public Charging Stations for Plug-in Hybrid Electric Vehicles. *Transportation Research Part B: Methodological*, 47(1):83–101.

Hofmann, T. (1999). Probabilistic latent semantic indexing. In *Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*, pages 50–57. ACM.

Hoke, A., Brissette, A., Maksimovic, D., Pratt, A., and Smith, K. (2011). Electric vehicle charge optimization including effects of lithium-ion battery degradation. In *Vehicle Power and Propulsion Conference (VPPC), 2011 IEEE*, pages 1–8. IEEE.

INL (2015). Idaho National Lab, Advanced Vehicle Testing.

Jaffe, A. B., Newell, R. G., and Stavins, R. N. (2000). Technology Change and the Environment. *NBER Working Paper Series*, (7970).

Jensen, A. F., Cherchi, E., and Mabit, S. L. (2013). On the stability of preferences and attitudes before and after experiencing an electric vehicle. *Transportation Research Part D: Transport and Environment*, 25:24–32.

Jing Dong, Z. L. (2014). Stochastic Modeling of Battery Electric Vehicle Driver Behavior. *Transportation Research Record: Journal of the Transportation Research Board*, 2454(-1):61–67.

Kaspi, M., Raviv, T., and Tzur, M. (2014). Parking Reservation Policies in One-way Vehicle Sharing Systems. *Transportation Research Part B: Methodological*, 62:35–50.

Keith, D. R. (2012). *Essays on the Dynamics of Alternative Fuel Vehicle Adoption: Insights from the Market for Hybrid-Electric Vehicles in the United States*. PhD thesis, Massachusetts Institute of Technology.

Khan, M. and Kockelman, K. M. (2012). Predicting the market potential of plug-in electric vehicles using multiday GPS data. *Energy Policy*, 46:225–233.

Kley, F., Lerch, C., and Dallinger, D. (2011). New business models for electric carsA holistic approach. *Energy Policy*, 39(6):3392–3403.

Kriston, A., Szabo, T., and Inzelt, G. (2010). The marriage of car sharing and hydrogen economy: A possible solution to the main problems of urban living. *International Journal of Hydrogen Energy*, 35.

Kromer, M. A. and Heywood, J. B. (2007). Electric Powertrains: Opportunities and Challenges. In *Pub. No. LFEE 2007-02 RP, Sloan Automotive Laboratory, Massachusetts Institute of Technology*.

Kurani, K. S., Axsen, J., Caperello, N., Davies-Shawhyde, J., Dempster, P., Kempster, M., Nesbitt, K. A., and Stillwater, T. (2010). *Plug-In Hybrid Electric Vehicle (PHEV) Demonstration and Consumer Education, Outreach, and Market Research Program: Volumes I and II, University of California at Davis, November*. Davis.

Le Vine, S., Adamou, O., and Polak, J. (2014a). Predicting new forms of activity/mobility patterns enabled by shared-mobility services through a needs-based stated-response method: Case study of grocery shopping. *Transport Policy*, 32:60–68.

Le Vine, S., Adamou, O., and Polak, J. (2014b). Predicting new forms of activity/mobility patterns enabled by shared-mobility services through a needs-based stated-response method: Case study of grocery shopping. *Transport Policy*, 32:60–68.

Levine, S. (2012). Car Rental 2.0: Car Club Innovations and Why They Matter. Report, RAC Foundation.

Likert, R. (1932). A technique for the measurement of attitudes. *Archives of Psychology*, 22 140:55.

Lin, Z. (2014). Optimizing and Diversifying Electric Vehicle Driving Range for U.S. Drivers. *Transportation Science*, 48(4):635–650.

Lin, Z. and Greene, D. (2011). Assessing Energy Impact of Plug-In Hybrid Electric Vehicles. *Transportation Research Record: Journal of the Transportation Research Board*, 2252:99–106.

Marouf, M., Pollard, E., and Nashashibi, F. (2014). Automatic parallel parking and platooning to redistribute electric vehicles in a car-sharing application. In *2014 IEEE Intelligent Vehicles Symposium Proceedings*, pages 486–491.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In *Frontiers in Econometrics*, pages 105–142. Academic Press, New York.

Mock, P., Tietge, U., Franco, V., German, J., Bandivadekar, A., Lambrecht, U., Khlwein, J., and Riemersma, I. (2013). From laboratory to road: From Laboratory to Road: A Comparison of Official and Real- World Fuel Consumption and CO2 Values for Cars in Europe and the United States. *communications*, 49(30):847129–102.

Nichols, B. G., Kockelman, K. M., and Reiter, M. (2015). Air quality impacts of electric vehicle adoption in Texas. *Transportation Research Part D: Transport and Environment*, 34:208–218.

Paulssen, M., Temme, D., Vij, A., and Walker, J. L. (2014). Values, attitudes and travel behavior: a hierarchical latent variable mixed logit model of travel mode choice. *Transportation*, 41(4):873–888.

Pearre, N., Kempton, W., Guensler, R., and Elango, V. (2011). Electric Vehicles: How Much Range is Required for a Day's Driving? *Transportation Research Part C: Emerging Technologies*, 19(6):1171–1184.

Pichelmann, S., Franke, T., and Krems, J. F. (2013). The timeframe of adaptation to electric vehicle range. In *Human-Computer Interaction. Applications and Services*, pages 612–620. Springer.

Project, E. (2013). The EV Project: Q2 2013 Report. Technical report, Electric Transportation Engineering Corporation.

Repoux, M., Boyaci, B., and Geroliminis, N. (2015). Simulation and optimization of one-way car-sharing systems with variant relocation policies.

Roberts, M. E., Stewart, B. M., and Airoldi, E. M. (2014a). Structural Topic Models. Technical report.

Roberts, M. E., Stewart, B. M., and Tingley, D. (2014b). stm: R Package for Structural Topic Models. Technical report.

Roberts, M. E., Stewart, B. M., Tingley, D., Lucas, C., Leder-Luis, J., Gadarian, S. K., Albertson, B., and Rand, D. G. (2014c). Structural Topic Models for Open-Ended Survey Responses. *American Journal of Political Science*, pages n/a–n/a.

Rodgers, L., Zoepf, S., and Prenninger, J. (2014). Analysing the energy consumption of the BMW ActiveE field trial vehicles with application to distance to empty algorithms. *Transportation Research Procedia*, 4(2014):42–54.

Rogers, C. (2015). Resale Prices Tumble on Electric Cars. *Wall Street Journal*.

Rogers, E. M. (2003). *Diffusion of Innovations*. Free Press, New York, NY, 5 edition.

Rosseel, Y. (2012). lavaan: An R Package for Structural Equation Modeling. *Journal of Statistical Software*, 48(2):1–36.

Rutherford, G. S. (2003). Flexcar Seattle Carsharing Program Evaluation.

SAE (2010). Utility Factor Definitions for Plug-In Hybrid Electric Vehicles Using Travel Survey Data. SAE Technical Paper SAE J2841-2010, Society of Automotive Engineers, Warrendale, PA.

Samaras, C. and Meisterling, K. (2008). Life cycle assessment of greenhouse gas emissions from plug-in hybrid vehicles: implications for policy. *Environmental science & technology*, 42(9):3170–3176.

Schneider, M., Stenger, A., and Goeke, D. (2014). The Electric Vehicle-Routing Problem with Time Windows and Recharging Stations. *Transportation Science*, 48(4):500–520.

Shaheen, S., Wright, J., and Sperling, D. (2002). California's Zero-Emission Vehicle Mandate: Linking Clean-Fuel Cars, Carsharing, and Station Car Strategies. *Transportation Research Record: Journal of the Transportation Research Board*, 1791:113–120.

Shaheen, S. A. and Chan, N. D. (2015). Evolution of E-Mobility in Carsharing Business Models. In Beeton, D. and Meyer, G., editors, *Electric Vehicle Business Models*, Lecture Notes in Mobility, pages 169–178. Springer International Publishing.

Shaheen, S. A. and Cohen, A. P. (2014). INNOVATIVE MOBILITY CARSHARING OUTLOOK: CARSHARING MARKET OVERVIEW, ANALYSIS, AND TRENDS. Technical report, Transportation Sustainability Research Center - University of California Berkeley.

Sivak, M. and Schoettle, B. (2014). What do current owners of hybrids and non-hybrids think about hybrids? Technical Report UMTRI-2014-25, UMTRI.

Smart, J., Bradley, T., and Salisbury, S. (2014). Actual Versus Estimated Utility Factor of a Large Set of Privately Owned Chevrolet Volts. SAE Technical Paper 2014-01-1803, SAE International, Warrendale, PA.

Smart, J., Powell, W., and Schey, S. (2013). Extended Range Electric Vehicle Driving and Charging Behavior Observed Early in the EV Project. Technical report, SAE Technical Paper.

Smart, J. and Schey, S. (2012). Battery electric vehicle driving and charging behavior observed early in the EV project. Technical report, SAE Technical Paper.

Smart, J. G. (2014). EV Project in Q4 2013.

Spieser, K., Treleaven, K., Zhang, R., Frazzoli, E., Morton, D., and Pavone, M. (2014). Toward a Systematic Approach to the Design and Evaluation of Automated Mobility-on-Demand Systems: A Case Study in Singapore. In Meyer, G. and Beiker, S., editors, *Road Vehicle Automation*, Lecture Notes in Mobility, pages 229–245. Springer International Publishing.

Stillwater, T., Mokhtarian, P., and Shaheen, S. (2009). Carsharing and the Built Environment. *Transportation Research Record: Journal of the Transportation Research Board*, 2110(1):27–34.

Tamor, M., Moraal, P., Reprogle, B., and Milacic, M. (2015). Rapid Estimation of Electric Vehicle Acceptance using a General Description of Driving Patterns. *Transportation Research Part C: Emerging Technologies*, 51:136–148.

Tanaka, M., Ida, T., Murakami, K., and Friedman, L. (2014). Consumers willingness to pay for alternative fuel vehicles: A comparative discrete choice analysis between the US and Japan. *Transportation Research Part A: Policy and Practice*, 70:194–209.

Temme, D., Paulssen, M., and Dannewald, T. (2008). Incorporating Latent Variables into Discrete Choice Models A Simultaneous Estimation Approach Using SEM Software. *BuR - Business Research*, 1(2):220–237.

Theis, G. W. (2011). *Incorporating Attitudes in Airline Itinerary Choice: Modeling the Impact of Elapsed Time*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA.

Vij, A. (2013). Incorporating the Influence of Latent Modal Preferences in Travel Demand Models. *University of California Transportation*.

Vyas, A., Santini, D., and Johnson, L. (2009). Potential of Plug-In Hybrid Electric Vehicles to Reduce Petroleum Use: Issues Involved in Developing Reliable Estimates. *World Transit Research*.

Walker, J. and Ben-Akiva, M. (2002). Generalized random utility model. *Mathematical Social Sciences*, 43(3):303–343.

Walker, J. L. (2001). *Extended discrete choice models : integrated framework, flexible error structures, and latent variables*. Thesis, Massachusetts Institute of Technology. Thesis (Ph.D.)–Massachusetts Institute of Technology, Dept. of Civil and Environmental Engineering, 2001.

Williams, B., Martin, E., Lipman, T., and Kammen, D. (2011). Plug-in-hybrid vehicle use, energy consumption, and greenhouse emissions: An analysis of household vehicle placements in northern California. *Energies*, 4(3):435–457.

Yuksel, T. and Michalek, J. J. (2015). Effects of Regional Temperature on Electric Vehicle Efficiency, Range, and Emissions in the United States.

Zipcar (2012). Zipcar 2012 Annual Vehicle Survey.

Zoepf, S. and Heywood, J. B. (2012). Characterizations of Deployment Rates in Automotive Technology. Technical Report 2012-01-1057, SAE.

Zoepf, S. and Keith, D. (2015). User Decision-Making and Technology Choices in the US Carsharing Market (Submitted for Review). Technical report.

Zoepf, S., MacKenzie, D., Keith, D., and Chernicoff, W. (2013). Charging choices and fuel displacement in a large-scale plug-in hybrid electric vehicle demonstration. *Journal of the Transportation Research Board*, (13-0533).

Appendix A

Appendix A: Zipcar Survey Questions

1. What is your age?
   - 18-21
   - 22-25
   - 26-30
   - 31-35
   - 36-45
   - 46-55
   - 56+
2. What is your gender identity?
   - Male
   - Female
   - Prefer not to answer
3. In what region do you primarily use Zipcar services? If you are not based in one of our larger markets, please choose the "Universities" option.
4. How many people live in your household, including you?

- 1
- 2
- 3
- 4
- 5+

5. How many children under the age of 18 live in your household?

   - 0
   - 1
   - 2
   - 3
   - 4+

6. Besides Zipcar, what other modes of transportation do you use regularly?

   - Bicycle
   - Bus
   - Subway
   - Train
   - Walk
   - Other

7. How many personal cars are owned or leased by your household?

   - None, Zipcar all the way!
   - One
   - Two
   - Three

- Four or more

8. Please list the make and models of the cars owned and leased in your household. (i.e. 2010 Toyota Camry Hybrid)

9. Please select the primary reason your household has one or more cars.

    - We don't make car payments and don't have to pay for parking, so the costs aren't that high
    - We just like having a personal car or truck
    - We use Zipcar when we want a nicer car than our personal cars or trucks
    - One or more family member has to commute to work or school regularly
    - We use Zipcar only for special cars for a specific task, like if we need a van or pickup
    - There aren't enough Zipcars in my area that I can rely on the service
    - We have one or more kids and it's more convenient
    - Other

10. How often do you rent a Zipcar?

    - More than once a week
    - Once a week
    - 2-3 times per month
    - Once a month
    - Less than once a month

11. Some Zipcars are Hybrids (e.g. Toyota Prius). Hybrids run on gasoline, but use batteries and an electric motor to reduce the amount of gasoline the car uses. Have you ever driven a hybrid?

- Yes, I own (or previously owned) a hybrid
- Yes, Ive driven a Zipcar hybrid
- Yes, Ive driven a hybrid elsewhere
- No, I havent driven a hybrid
- Im not sure

12. Some Zipcars are Plug-In hybrids (e.g. Chevrolet Volt). Plug-In hybrids are like regular hybrids, but can be also recharged directly with electricity, to travel farther under electric power and further reduce the gasoline the cars use. Have you ever driven a Plug-In hybrid?
    - Yes, I own (or previously owned) a plug-in hybrid
    - Yes, Ive driven a Zipcar plug-in hybrid
    - Yes, Ive driven a plug-in hybrid elsewhere
    - No, I havent driven a plug-in hybrid
    - Im not sure
13. Some Zipcars are Electric cars (e.g. Nissan Leaf). Electric cars use no gasoline, being recharged 100% using electricity. Have you ever driven an electric car?
    - Yes, I own (or previously owned) an electric car
    - Yes, Ive driven a Zipcar electric car
    - Yes, Ive driven an electric car elsewhere
    - No, I havent driven an electric car
    - Im not sure
14. How far ahead do you typically make your reservation for a Zipcar?
    - More than a month before my trip

- More than a week before my trip
- More than a day before my trip
- Several hours before my trip
- Up to an hour before my trip

15. Which of the following best describes your preferences when selecting a Zipcar?
    - I always take the same brand and model Zipcar if possible
    - I prefer the same brand and model Zipcar
    - No Preference
    - I prefer to try different brands and models of cars
    - I always try to take something new

16. Please rate your agreement with the following statement: "I like Zipcars that have logos and other Zipcar branding on the car."
    - Strongly agree
    - Agree
    - Neither agree nor disagree
    - Disagree
    - Strongly disagree

17. In a typical reservation, how many hours do you keep your Zipcar?

18. In a typical Zipcar reservation, how many miles do you drive your Zipcar?

19. "In the following 4 questions, we ask you to select which vehicle you would reserve for your typical Zipcar trip given a range of vehicle options.

    Please select the vehicle that best suits your needs. Assume that all gasoline cars are filled and electric cars are fully charged when you take them, and that the vehicles are otherwise identical except for the differences shown"

[Conjoint Analysis (Discrete Choice)]

20. Zipsters sometimes need vehicles with large passenger and cargo capacity (such as SUVs, Zipvans, and pickup trucks) for certain trips. How often do you take trips that require larger vehicles?

    - Always
    - Often
    - Sometimes
    - Rarely
    - Never

21. Which of the following apply to you? (Check all that apply)

    - I always prefer a larger vehicle
    - I take a larger vehicle when many people are traveling with me
    - I take a larger vehicle when I need to carry bulky cargo
    - I take a larger vehicle only when nothing else is available
    - Other

22. In a trip which requires a larger vehicle, which of the following would you prefer?

    - Minivan (e.g. Mazda5 or Toyota Sienna)
    - Pickup Truck (e.g. Toyota Tacoma)
    - Compact SUV (e.g. Honda CRV)
    - Zipvan (Cargo Van e.g. Ford Econoline)

23. Please rate your level of satisfaction with the following brands and models of Zipcars. Please only rate the vehicles which you have used through Zipcar.

- Audi A3
- BMW 3 Series
- Chevrolet Volt
- FIAT 500
- Ford E-150 (Zipvan)
- Ford Escape
- Ford Focus
- Honda Civic
- Honda CR-V
- Honda Fit EV
- Honda Insight Hybrid
- Hyundai Veloster
- Mazda 3
- Mercedes C250/C300
- MINI Cooper
- Nissan Frontier
- Nissan Sentra
- Toyota Prius Hybrid
- Toyota Sienna
- Toyota Tacoma
- Volkswagen Golf

24. What kinds of cars would you like to see more of in our fleet, and why?

25. Please let us know if you have any other comments or suggestions about our service.

Appendix B

Appendix B: Relative Star Ratings of Vehicles

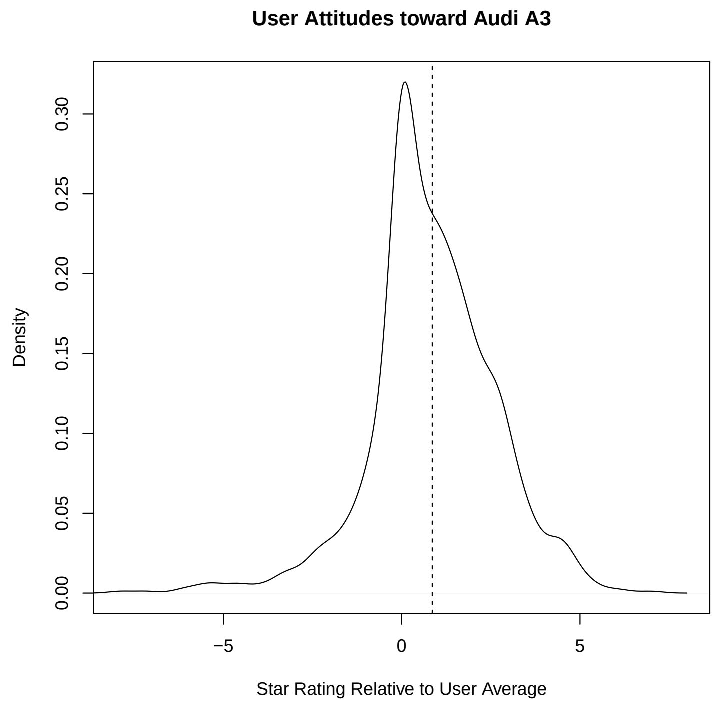
User Attitudes toward Audi A3
Density
0.00
0.05
0.10
0.15
0.20
0.25
0.30
−5
0
5
Star Rating Relative to User Average

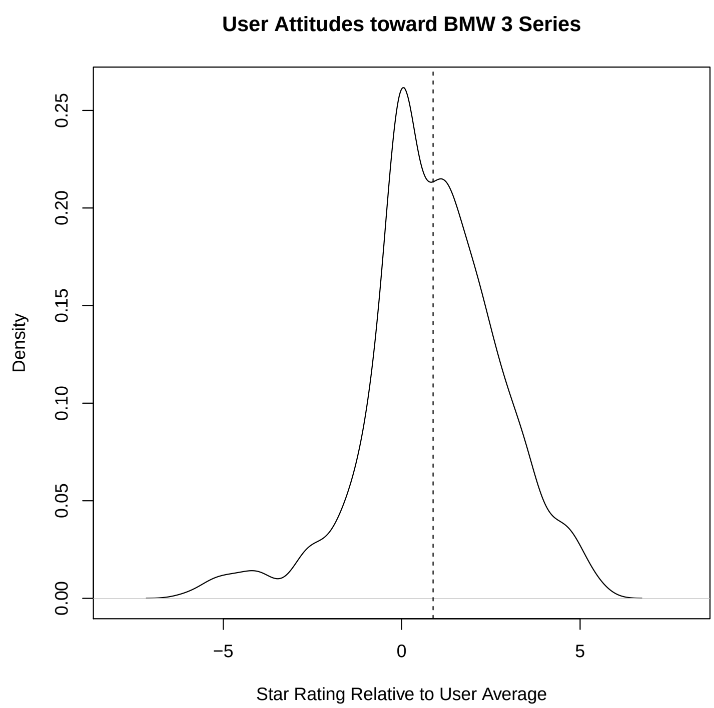
User Attitudes toward BMW 3 Series
Density
0.00
0.05
0.10
0.15
0.20
0.25
−5
0
5
Star Rating Relative to User Average

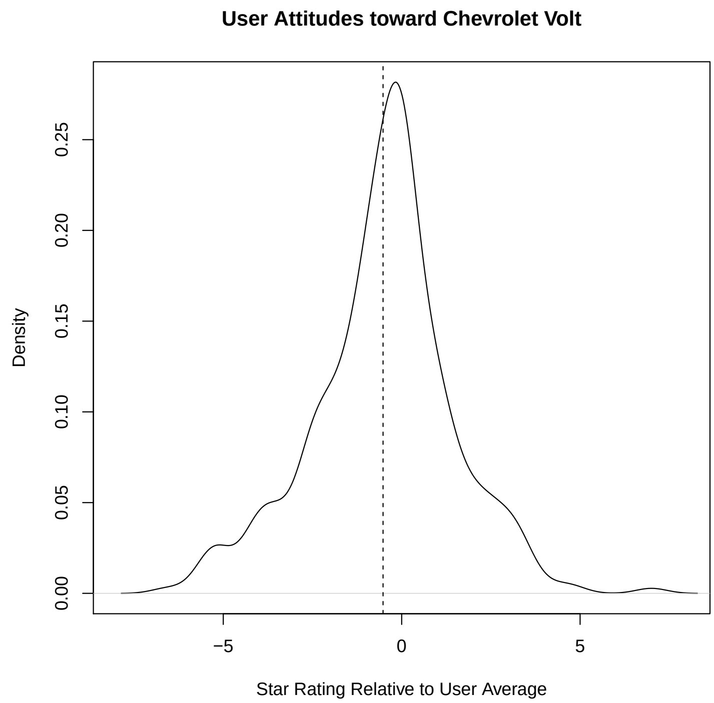
User Attitudes toward Chevrolet Volt
Density
0.00
0.05
0.10
0.15
0.20
0.25
−5
0
5
Star Rating Relative to User Average

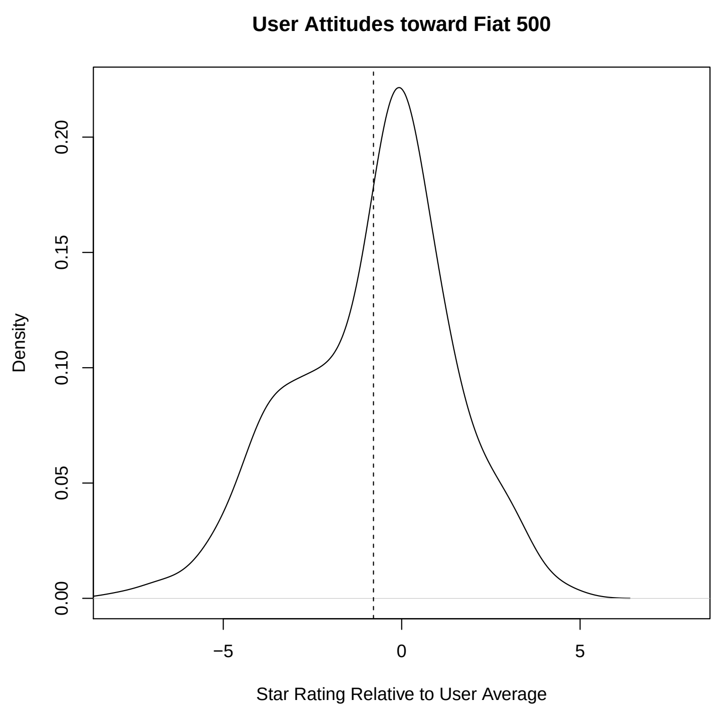
User Attitudes toward Fiat 500
Density
0.00
0.05
0.10
0.15
0.20
−5
0
5
Star Rating Relative to User Average

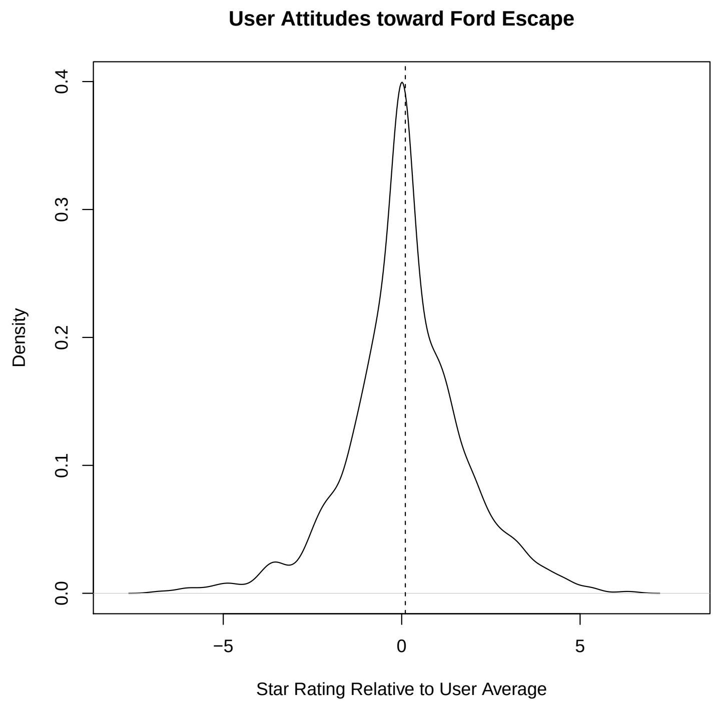
User Attitudes toward Ford Escape
Density
0.0
0.1
0.2
0.3
0.4
−5
0
5
Star Rating Relative to User Average

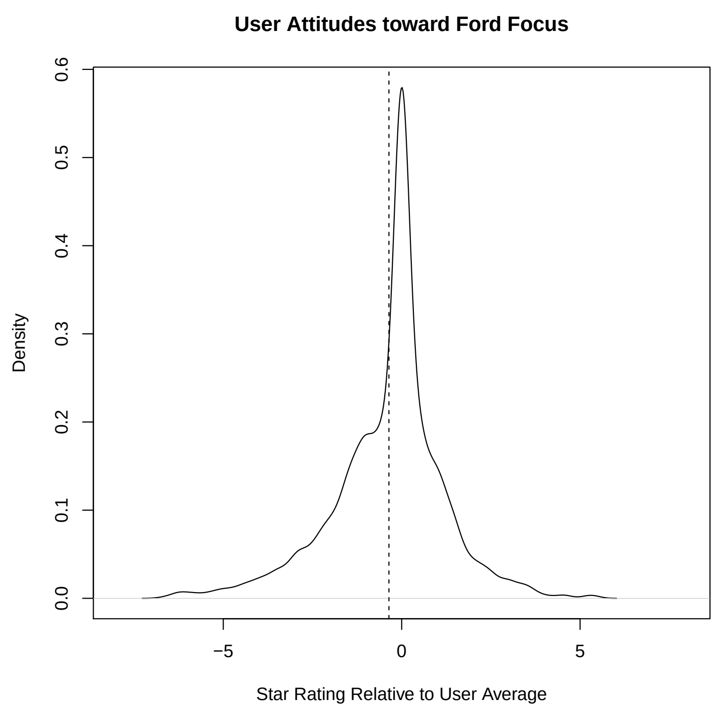
User Attitudes toward Ford Focus
Density
0.0
0.1
0.2
0.3
0.4
0.5
0.6
−5
0
5
Star Rating Relative to User Average

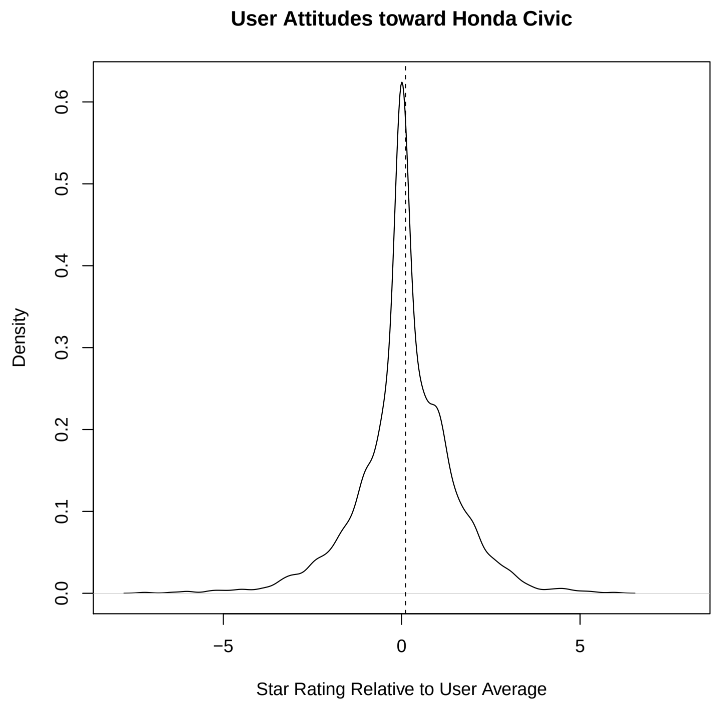
User Attitudes toward Honda Civic
Density
0.0
0.1
0.2
0.3
0.4
0.5
0.6
−5
0
5
Star Rating Relative to User Average

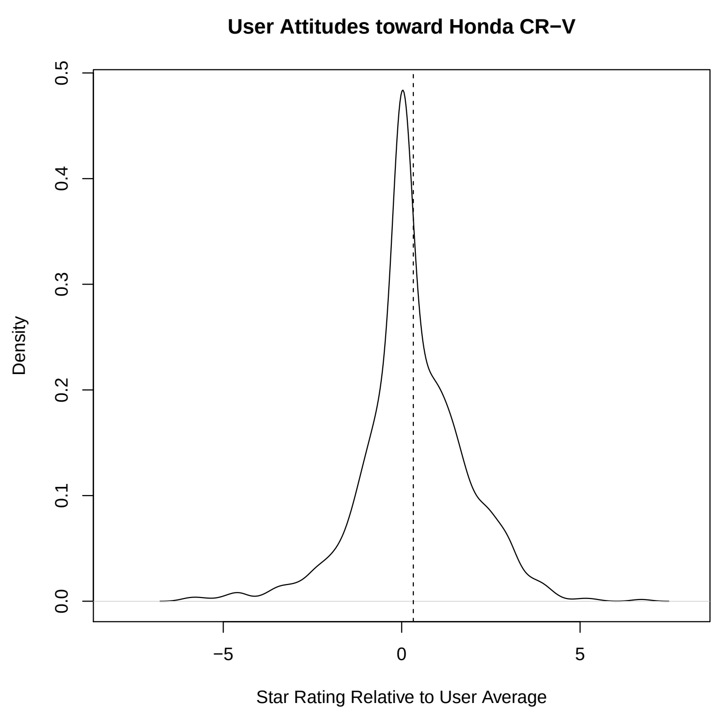
User Attitudes toward Honda CR−V
Density
0.0
0.1
0.2
0.3
0.4
0.5
−5
0
5
Star Rating Relative to User Average

User Attitudes toward Honda Fit EV

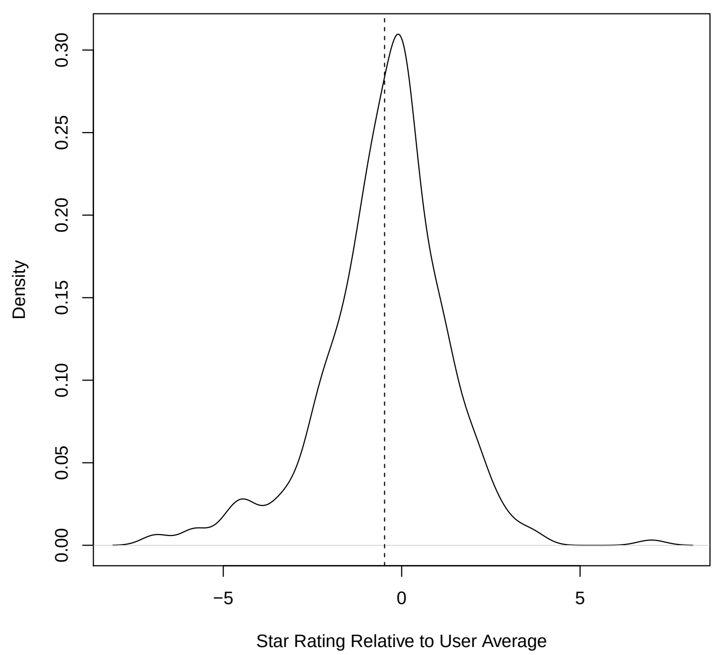

User Attitudes toward Honda Insight Hybrid

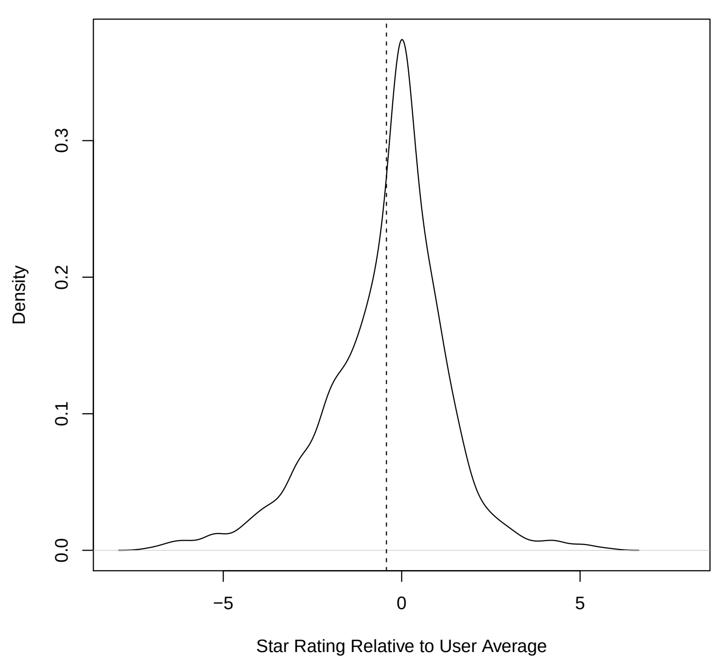

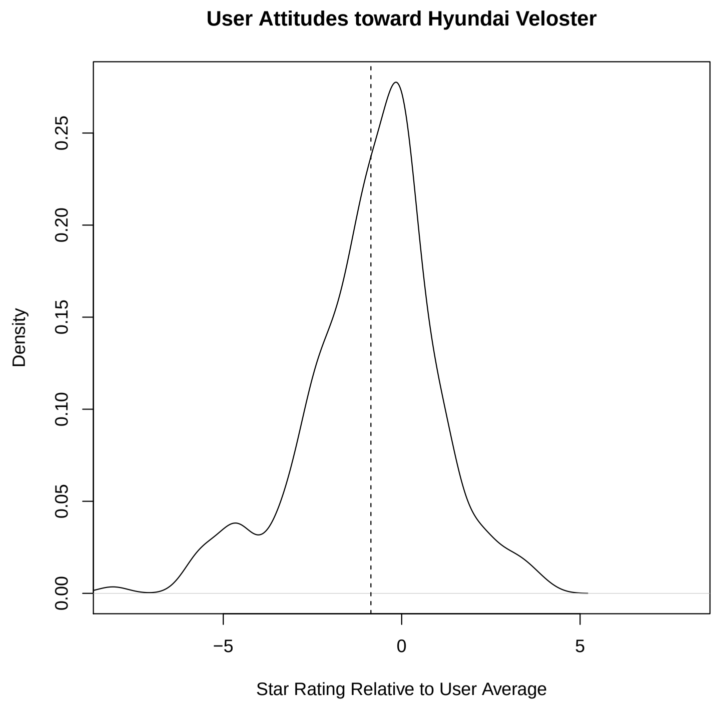
User Attitudes toward Hyundai Veloster
Density
0.00
0.05
0.10
0.15
0.20
0.25
−5
0
5
Star Rating Relative to User Average

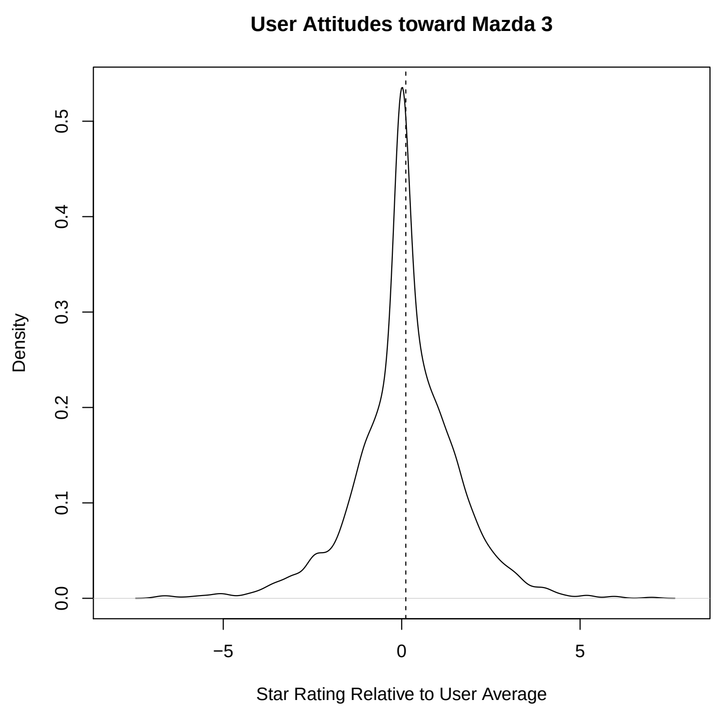
User Attitudes toward Mazda 3
Density
0.0
0.1
0.2
0.3
0.4
0.5
−5
0
5
Star Rating Relative to User Average

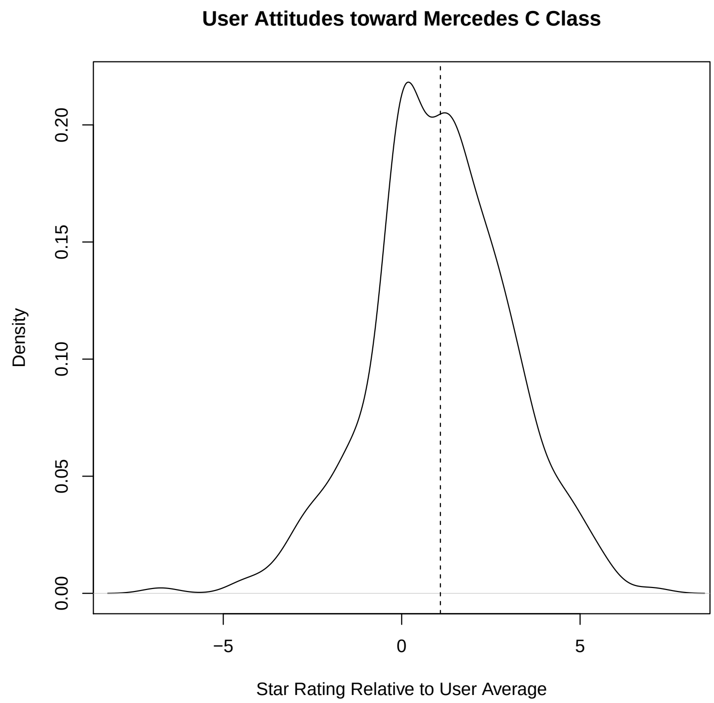
User Attitudes toward Mercedes C Class
Density
0.00
0.05
0.10
0.15
0.20
−5
0
5
Star Rating Relative to User Average

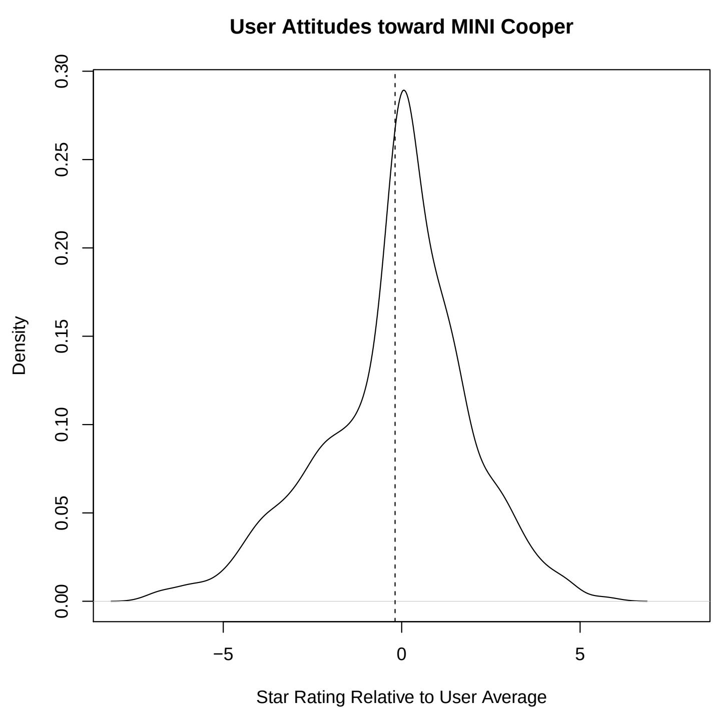
User Attitudes toward MINI Cooper
Density
0.00
0.05
0.10
0.15
0.20
0.25
0.30
−5
0
5
Star Rating Relative to User Average

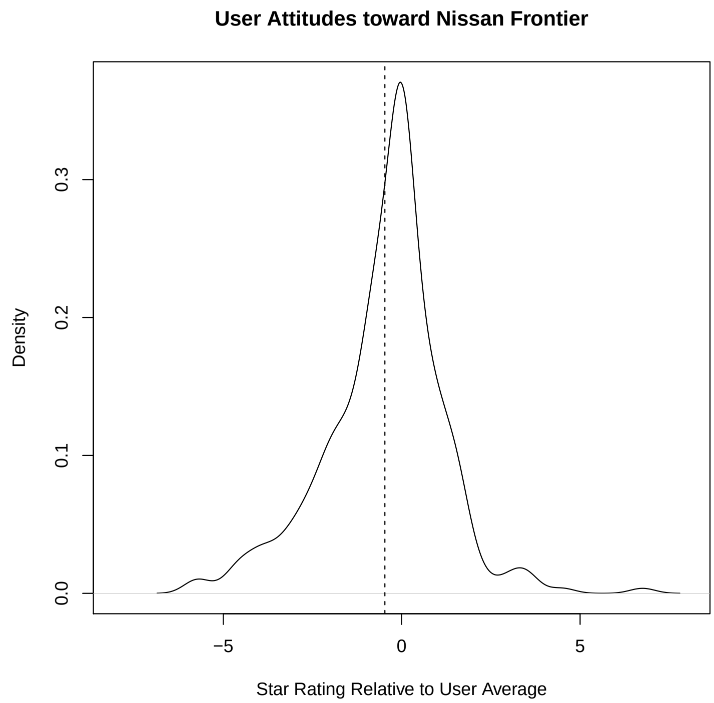
User Attitudes toward Nissan Frontier
Density
0.0
0.1
0.2
0.3
−5
0
5
Star Rating Relative to User Average

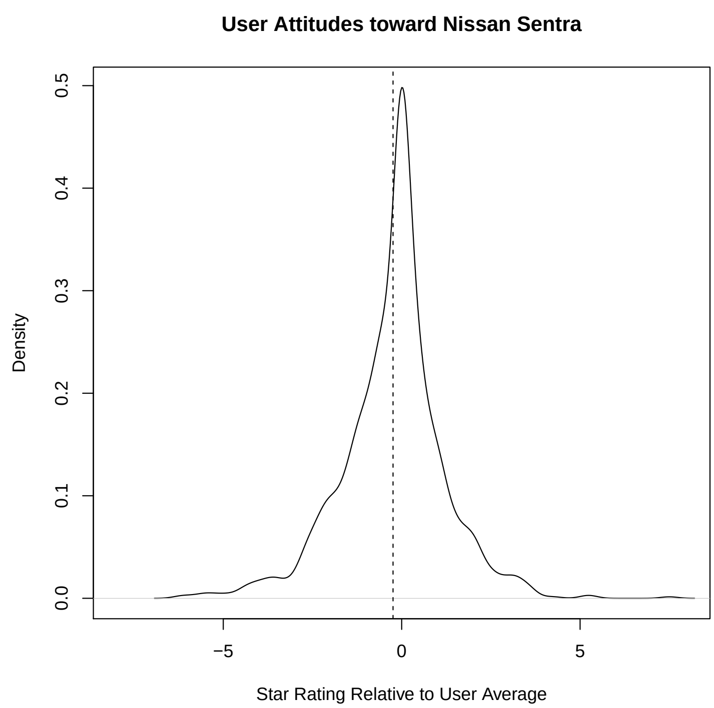
User Attitudes toward Nissan Sentra
Density
0.0
0.1
0.2
0.3
0.4
0.5
−5
0
5
Star Rating Relative to User Average

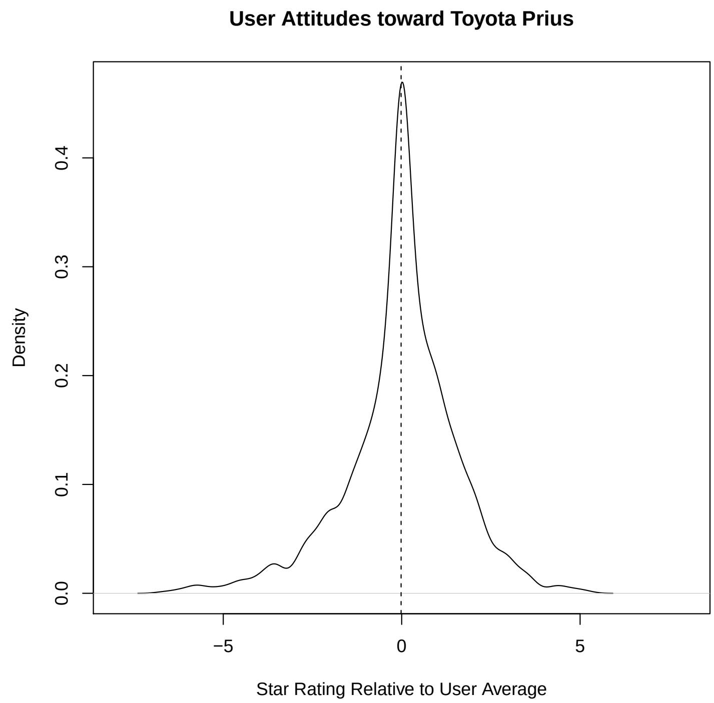
User Attitudes toward Toyota Prius
Density
0.0
0.1
0.2
0.3
0.4
−5
0
5
Star Rating Relative to User Average

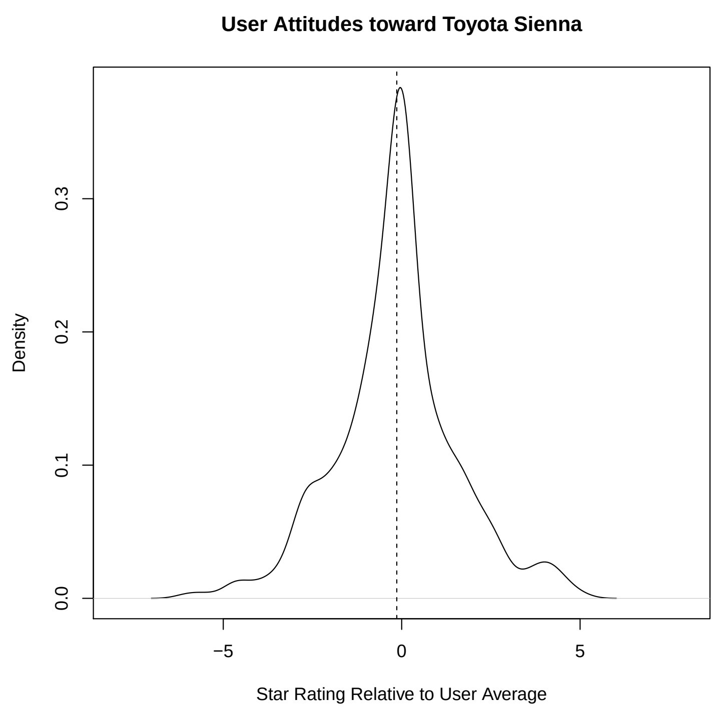
User Attitudes toward Toyota Sienna
Density
0.0
0.1
0.2
0.3
−5
0
5
Star Rating Relative to User Average

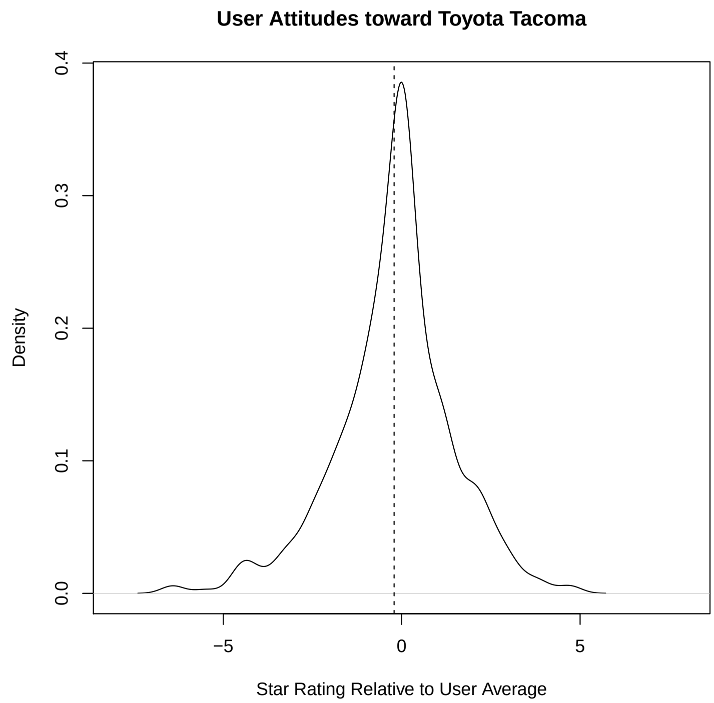
User Attitudes toward Toyota Tacoma
Density
0.0
0.1
0.2
0.3
0.4
−5
0
5
Star Rating Relative to User Average

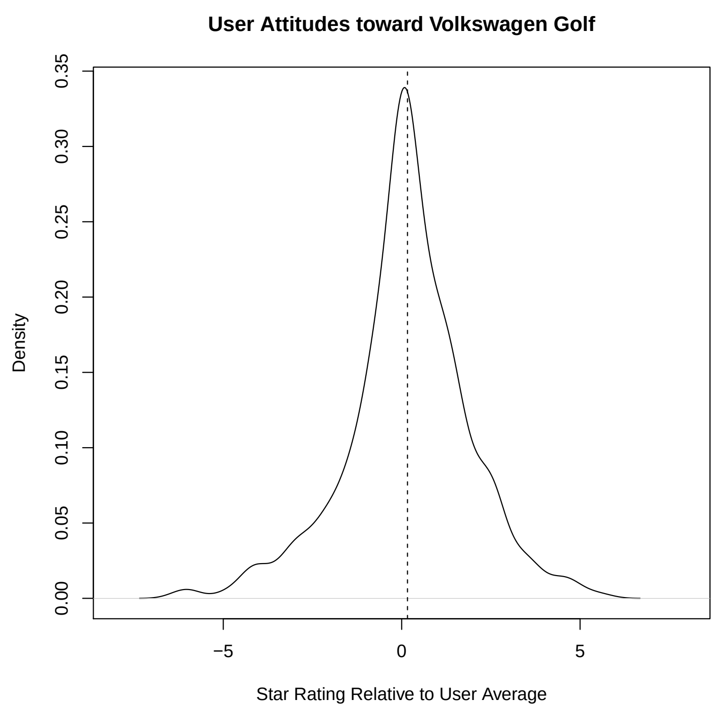
User Attitudes toward Volkswagen Golf
Density
0.00
0.05
0.10
0.15
0.20
0.25
0.30
0.35
−5
0
5
Star Rating Relative to User Average

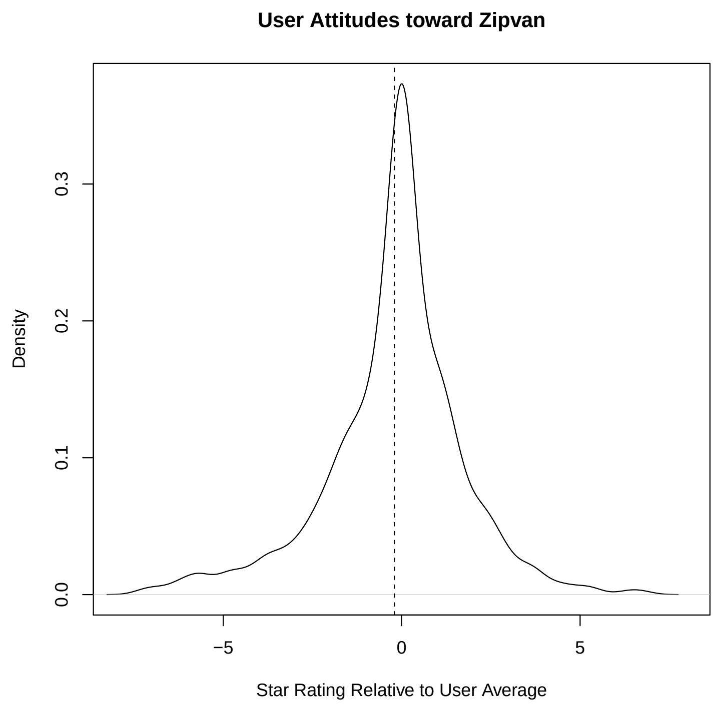
User Attitudes toward Zipvan
Density
0.0
0.1
0.2
0.3
−5
0
5
Star Rating Relative to User Average