The Dissertation Committee for Chioke Bem Harris
certifies that this is the approved version of the following dissertation:

An Assessment of the System Costs and Operational Benefits of Vehicle-to-Grid Schemes

Committee:

Michael E. Webber, Supervisor

J. Eric Bickel

Dongmei Chen

Ross Baldick

Kara Kockelman

An Assessment of the System Costs and Operational Benefits of Vehicle-to-Grid Schemes

by

Chioke Bem Harris, B.S.E., M.S.E.

DISSERTATION

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

DOCTOR OF PHILOSOPHY

The University of Texas at Austin

December 2013

Acknowledgments

I would like to thank my advisor, Dr. Michael E. Webber, for providing professional development opportunities, financial support, feedback, and advice instrumental to my development as a graduate student. His commitment to teaching, mentorship, and research is extraordinary, and I cannot overstate the value of his support.

Several of my colleagues in the Webber Energy Group provided valuable feedback on this work throughout its development. In particular, I appreciate the insights of Aaron Townsend, Stuart Cohen, Jared Garrison, Robert Fares, and Joshua Rhodes. The support provided by the administrative staff, research associates, and the rest of the undergraduate and graduate students in the group was also invaluable.

Data used in this dissertation were provided by the Electric Reliability Council of Texas (ERCOT), the National Renewable Energy Laboratory (NREL), and the US Department of Transportation. I would like to thank the Vehicles and Fuels research staff at NREL for their efforts to make available to the academic research community a wide array of GPS-based travel survey data collected in the US. I also appreciate the assistance of Michael Legatt and Jane Cates at ERCOT to enable me to obtain some of the electricity market data required for this work.

Much of the analysis in this dissertation was executed on the high performance computing (HPC) cluster maintained by the Department of Mechanical Engineering. I appreciate the department's continued investment in maintaining the HPC, as it was an immeasurably valuable resource for this work, and is an extraordinary asset to all of the students in the department. I also appreciate the timely technical support

provided by Cengiz Vural in maintaining both the HPC and our research group's high performance computer.

During a significant portion of my doctoral studies, I was supported by the National Science Foundation Graduate Research Fellowship Program (NSF GRFP). This NSF GRFP award (grant number DGE-1110007) made it possible for me to pursue this research topic. I also appreciate the additional financial support I received from the Arthur H. Hausman Endowed Graduate Fellowship in Engineering, part of the Thrust 2000 Graduate Fellowship program in the Cockrell School of Engineering at The University of Texas at Austin.

Finally, I would like to thank my parents for their support throughout my academic career. Without their support and sacrifice, I would not be the scholar I am today.

An Assessment of the System Costs and Operational Benefits of Vehicle-to-Grid Schemes

Chioke Bem Harris, Ph.D.
The University of Texas at Austin, 2013

Supervisor: Michael E. Webber

With the emerging nationwide availability of plug-in electric vehicles (PEVs) at prices attainable for many consumers, electric utilities, system operators, and researchers have been investigating the impact of this new source of electricity demand. The presence of PEVs on the electric grid might offer benefits equivalent to dedicated utility-scale energy storage systems by leveraging vehicles' grid-connected energy storage through vehicle-to-grid (V2G) enabled infrastructure. Existing research, however, has not effectively examined the interactions between PEVs and the electric grid in a V2G system. To address these shortcomings in the literature, longitudinal vehicle travel data are first used to identify patterns in vehicle use. This analysis showed that vehicle use patterns are distinctly different between weekends and weekdays, seasonal interactions between vehicle charging, electric load, and wind generation might be important, and that vehicle charging might increase already high peak summer electric load in Texas. Subsequent simulations of PEV charging were performed, which revealed that unscheduled charging would increase summer peak load in Texas by approximately 1%, and that uncertainty that arises from unscheduled charging would require only limited increases in frequency regulation procurements.

To assess the market potential for the implementation of a V2G system that provides frequency regulation ancillary services, and might be able to provide financial incentives to participating PEV owners, a two-stage stochastic programming formulation of a V2G system operator was created. In addition to assessing the market potential for a V2G system, the model was also designed to determine the effect of the market power of the V2G system operator on prices for frequency regulation, the effect of uncertainty in real-time vehicle availability and state-of-charge on the aggregator's ability to provide regulation services, and the effect of different vehicle characteristics on revenues. Results from this model showed that the V2G system operator could generate revenue from participation in the frequency regulation market in Texas, even when subject to the uncertainty in real-time vehicle use. The model also showed that the V2G system operator would have a significant impact on prices, and thus as the number of PEVs participating in a V2G program in a given region increased, per-vehicle revenues, and thus compensation provided to vehicle owners, would decline dramatically. From these estimated payments to PEV owners, the decision to participate in a V2G program was analyzed. The balance between the estimated payments to PEV owners for participating in a V2G program and the increased probability of being left with a depleted battery as a result of V2G operations indicate that an owner of a range-limited battery electric vehicle (BEV) would probably not be a viable candidate for joining a V2G program, while a plug-in hybrid electric vehicle (PHEV) owner might find a V2G program worthwhile. Even for a PHEV owner, however, compensation for participating in a V2G program will provide limited incentive to join.

Table of Contents

List of Acronyms

ACE	area control error
BEV	battery electric vehicle
ERCOT	Electric Reliability Council of Texas
EVSE	electric vehicle supply equipment
GPS	Global Positioning System
HEV	hybrid electric vehicle
ICE	internal combustion engine
LDV	light duty vehicle
LP	linear programming
MIP	mixed-integer programming
NHTS	National Household Travel Survey
NSRS	non-spinning reserve
PEV	plug-in electric vehicle
PHEV	plug-in hybrid electric vehicle
PSRC	Puget Sound Regional Council
RRS	responsive reserve
SP	stochastic programming
V2G	vehicle-to-grid

List of Tables

List of Figures

Chapter 1

Introduction

1.1 Motivation

In the past few decades, there has been increasing interest in the potential uses of energy storage on the grid. Research has revealed a wide range of configurations and applications for grid-scale energy storage. An application of significant interest is the use of energy storage to provide backup power to support renewable generators, increasing their effective capacity factors [1]. Some forms of renewable generation, especially wind and solar resources, are highly variable, and are often not well-matched to times of peak electricity demand [2]. The past decade has seen significant investment in renewable electricity generation across the United States [3]. As the percentage of total capacity constituted by renewable generators increases, it is expected that reserve capacity must increase to accommodate increases in the magnitude of potential undersupply and forecast errors [4, 5]. The variability of electricity demand and renewable generation, as well as the temporal mismatch between them, are currently moderated with ancillary services.

Aeroderivative gas turbines commonly provide ancillary services, but energy storage could provide these services with lower emissions [6–8, 13], potentially while offering other benefits not available from gas turbines. Energy storage could also simultaneously support renewable generation and provide ancillary services. In the electricity market structure relevant to this work, these services include frequency regulation, provided by operating generators that can make small, rapid adjustments

in their power output to stabilize grid frequency; responsive reserve (RRS) from facilities operating below their maximum capacity; and non-spinning reserve (NSRS) from offline generators that can startup quickly [9]. The optimal services provided from a given energy storage system depend on the technology used. For example, compressed air energy storage (CAES) relies on turbines and expanders to operate, and is thus well-suited to longer duration and response time services, such as NSRS [10]. Because electrochemical storage technologies can respond quickly, they are suitable for frequency regulation [11–14]. Electrochemical storage devices can respond faster than any existing ancillary service providers, and the Federal Electricity Regulatory Commission (FERC) has ordered system operators to develop new market mechanisms to enable these storage devices to monetize their superior capabilities [11].

Vehicles designed to rely on the electric grid to charge their batteries might eventually create operational complications for utilities and system operators. These plug-in electric vehicles (PEVs) vary widely in range and capabilities, and include both plug-in hybrid electric vehicles (PHEVs) and battery electric vehicles (BEVs). PHEVs have both an electric drivetrain and an internal combustion engine (ICE), allowing them to drive a limited distance relying exclusively on the onboard battery, typically between 10 and 40 miles, before requiring the ICE to continue driving. On the other hand, BEVs have a larger battery than PHEVs but do not have an ICE to extend their range once the battery is exhausted, thus limiting their range to, in most cases, less than 100 miles. While more familiar hybrid electric vehicles (HEVs) also have a large onboard battery pack, it is typically much smaller than those found in PEVs. The motive power for an HEV is primarily provided by the ICE, and the battery is primarily intended to facilitate more efficient operation of the ICE. PHEVs and BEVs of varying electric range and price are commercially available nationwide and offered by several automakers.

Various market conditions, including zero-emission vehicle mandates, tax incentives, and decreasing battery and power electronics costs [15,16] all foreshadow future growth in PEV sales. Widespread availability of PEVs has led to questions about the electric grid impacts associated with light duty vehicle (LDV) fleet electrification. PEV charging is currently a large unscheduled residential load, similar to an air conditioner, and thereby yields additional energy sales for utilities. Scheduled charging could defer the need for distribution network infrastructure upgrades [17,18], and there is ample generating capacity to support scheduled overnight charging [16,19–21], but most PEV owners do not currently receive any incentives to schedule their charging for the middle of the night. Without charge scheduling, PEV charging load could increase peak electricity demand [16, 19, 22, 23]. Because charge scheduling is not currently widespread, it is as yet unclear whether system operators, as PEVs increase in popularity, will need to accommodate unscheduled charging in their long-term planning and operations strategies.

Conversely, PEVs could offer benefits otherwise requiring dedicated grid-scale energy storage, such as increasing the capacity factors of renewable generators [1], providing ancillary services [24], and reducing emissions from thermal generators [21,25]. Such a scheme is typically referred to as V2G. This use of PEVs could be a boon for storage on the grid, as the capital costs of existing grid-scale storage technologies are prohibitive for most applications [26,27]. Vehicle owners would receive compensation for allowing their vehicles' batteries to be used to provide grid services by participating in a V2G program [17,28,29]. This compensation would need to be sufficient to offset any battery lifetime effects caused by participation in a V2G program and might even offer additional income for vehicle owners. Since a single vehicle's battery is too small to participate in an electricity market, it is widely proposed that vehicles should be pooled by a third-party entity, often referred to as a "vehicle ag-

gregator" [30]. The aggregator would use the combined resources of all the vehicles in its program, potentially across an entire region, to offer energy and ancillary services in an electricity market. Payments received for providing these resources would then be shared with program members. Providing frequency regulation is often cited as the best revenue opportunity for vehicle owners because they can be compensated for being available to provide regulation and the power and energy requirements upon deployment might be small if the total vehicle resource is sufficiently large [24, 30].

1.2 Scope and organization

This work is intended to augment the existing literature to facilitate a better understanding of the characteristics of individuals' vehicle use choices, the ancillary service impacts of the interaction between PEVs and the electric grid, and the potential revenue opportunities for vehicle aggregators and owners in a V2G system. To that end, this dissertation is structured around the following research objectives:

1. Identify temporal patterns in vehicle availability using Global Positioning System (GPS)-based travel data and compare those patterns with load and wind generation in ERCOT,

2. Using Monte Carlo methods, examine the magnitude and variability of vehicle use and battery charging as a function of total fleet size and time of day to assess the impact of plug-in electric vehicles on electric load and frequency regulation procurements,

3. Determine potential revenue for vehicle owners and aggregators using a two-stage stochastic linear program to model a vehicle-to-grid system operating in ERCOT, and

4. Evaluate the value to the consumer of a plug-in hybrid electric vehicle or a battery electric vehicle when participating in a vehicle-to-grid system using decision analysis methods.

The analysis performed to respond to each of these objectives relied on a wide array of data. The details of these data are provided in the relevant sections, along with descriptions of how the data were prepared for subsequent analysis. The data fall into two primary classes: vehicle data and electric grid data, and are drawn from four sources:

- **Puget Sound Regional Council** [31] — GPS-based travel data collected between November 2004 and May 2006 from 293 households with 436 unique vehicles
- **US Department of Transportation** [32] — travel survey data collected from individuals in all 50 US states for travel completed on an assigned day between March 2008 and May 2009, totaling 1.17 million trips
- **Pecan Street Research Consortium** — vehicle charging data from 33 households in Austin, Texas from April 2012 to April 2013
- **Electric Reliability Council of Texas (ERCOT)** [33] — a range of parameters describing day-ahead and real-time market prices, day-ahead ancillary service offers, and real-time operational conditions from 2009 and 2011

The remainder of this dissertation is divided into six primary sections. Chapter 2 provides background information and a review of the literature on subjects relevant to this work, including PEV charging and control strategies, vehicle use patterns, grid-scale energy storage simulation, and V2G system modeling. The four

subsequent chapters review the work completed to cover the primary research objectives. Chapter 3 details the identification of temporal patterns in vehicle availability using GPS-based travel data and comparison of those patterns with load and wind generation in ERCOT. Chapter 4 explores the variability in available vehicle battery capacity as a function of total fleet size and time of day to assess the impacts of electric vehicles and vehicle charging on total electric load and frequency regulation. Chapter 5 describes a stochastic optimization approach to assess the potential revenue for vehicle owners and aggregators in a V2G system in ERCOT. Chapter 6 details the impact of V2G program participation on individual PEV owners. Each of these chapters covers the relevant methodology, results obtained, and conclusions. Finally, Chapter 7 reviews conclusions drawn from all four research objectives and identifies opportunities for extension of this work.

Chapter 2

Background

This chapter provides information relevant throughout this work, and is divided into six sections:

1. Plug-in electric vehicles
2. Historical vehicle use data
3. Plug-in electric vehicle charging models
4. Vehicle-to-grid system models
5. Stochastic modeling approaches
6. Decision analysis concepts

Sections 2.1 and 2.2 provide background information on the types of vehicles being modeled, the characteristics of vehicle charging equipment, and the travel survey data used in this work to estimate the times when vehicles are parked and when they are charging. Section 2.2 also details how this work improves upon common approaches to employing vehicle use data in the literature. Sections 2.3 and 2.4 provide a review of the literature on PEV charging and V2G system modeling and simulation. These sections also include discussion of how this work augments the existing literature. Finally, Sections 2.5 and 2.6 provide background information on stochastic programming and decision analysis methods relevant to the models presented in Chapters 5 and 6, respectively.

2.1 Plug-in electric vehicles

Plug-in electric vehicles are now available for purchase across the United States. Most major manufacturers offer at least one PHEV or BEV, and often sell the same vehicle in international markets. These vehicles range widely in price, body style, and range. Because the onboard ICE extends the range of PHEVs, most have an electric driving range of only 10 to 40 miles, provided by a battery pack with between 3.2 and 16 kWh capacity [34]. Since BEVs are limited to the range offered by the battery alone, their batteries are much larger than in PHEVs, between 16 and 85 kWh, but offer driving ranges from 62 to over 250 miles [35]. Since January 1, 2010, the US Internal Revenue Service has offered a federal income tax credit for the purchase of a PEV, ranging from $2,500 to $7,500, depending on the capacity of the vehicle's battery [36]. These incentives will be phased out on a per-manufacturer basis beginning when 200,000 qualifying vehicles (PHEVs and BEVs) have been sold by a given manufacturer [36]. Some states, counties, municipalities, and utilities offer additional incentives or rebates for the purchase of a PEV, or offer other indirect or non-financial benefits, such as solo high-occupancy vehicle (HOV) lane access, free parking, or reduced electricity tariffs for vehicle charging.

Though PEVs typically have a "convenience charger" that allows recharging anywhere a standard 120 V outlet is available, most manufacturers recommend that PEV buyers also purchase a more powerful, 240 V charging station (EVSE) for installation at their home. While the convenience charger might be sufficient for the smaller batteries in PHEVs, vehicles with larger batteries will likely require these dedicated home EVSEs. Typical EVSEs designed for home applications are rated at 3.3 or 6.6 kW, but some are designed to provide in excess of 10 kW. The EVSEs marketed for use with different PEVs typically have power ratings appropriate to the

size of the PEVs battery. Most EVSEs have a charging cable with a conductive charging plug designed to comply with the Society of Automotive Engineers (SAE) J1772 standard [37], though some manufacturers have chosen to develop proprietary connectors and provide adapters for the J1772 plug, since that standard is used for public charging equipment. Higher power, direct-current public charging infrastructure that provides short recharge times for PEVs is forthcoming.

2.2 Vehicle use data

The characteristics of vehicle use have been the focus of significant study in civil engineering for decades. As LDVs have become more pervasive, the importance of understanding individuals' decisions regarding travel mode, route, timing, and extent has grown dramatically. Studies to quantify these individual travel characteristics are dependent upon historical travel data. Cities, Metropolitan Planning Organizations (MPOs) and the US Department of Transportation (DOT) have all collected these data, which serve various purposes, including measuring the effect of new infrastructure or travel management regimes, assessing transport-related areas of greatest need for remediation and augmentation, and supporting academic research. Large-scale travel data collection typically relied on self-reporting of travel in diaries. This approach is simple and (relatively) low-cost, but the travel logged by individuals is not always accurate. There exist well-established biases in the data reported by surveyed individuals that require correction when the data are used for modeling or prediction [38, 39]. With the advent of civilian access to GPS satellites, the precision of vehicle travel data was improved considerably. Because of the cost of these devices and their relatively recent availability, GPS-derived travel data often include only a small sub-sample from a travel study, and are limited in prevalence and availability to the academic community. Regardless of the method used, vehicle travel data col-

lection efforts have emphasized capturing the largest and most representative sample of vehicle travel given available funds. As a result, few travel surveys include any longitudinal time-series observations of vehicles or households.

In the absence of publicly available empirical charging data, travel survey data have been used by researchers to model electricity demand for vehicle charging. These travel data are based on vehicles other than PEVs, and thus do not include any information about charging times or patterns. Transforming travel data into vehicle charging is not central to the interests of many researchers, and thus they have chosen to use whatever data are readily available or can be quickly transformed to suit their needs. Accurate modeling of the interactions between electric vehicles and the electric grid, whether for vehicle charging or to simulate a V2G system, requires an understanding of the temporal variations in vehicle use. For example, studies of the emissions impacts of PEV charging have highlighted the importance of accounting for the time-varying features of vehicle use [21]. The literature is replete with research on the variability of wind generation and electricity demand over various time intervals and in various regions [4,5,40,41], but studies involving PEV charging often lack such detail. Early work examining the potential resource size and revenue opportunities from V2G typically ignored temporal variations in vehicle use, and instead selected an availability fraction, which was held constant throughout the modeling period [42–44]. Dallinger et al. [24] compared this constant vehicle availability assumption with the results of a Monte Carlo simulation of vehicle use and demonstrated that such assumptions are entirely inconsistent with vehicle use data.

Recent studies of revenue opportunities for vehicle owners and aggregators have begun to account for time-of-day variations in vehicle availability, but authors have not generally undertaken a close examination of underlying vehicle use trends.

As a result, many models use a single daily profile to represent all days of the year [28, 45–47]. Further, most studies do not address the uncertainty present in vehicle use shares by time of day, but rather treat them as deterministic, exogenous model parameters [25, 45–47]. Pearre et al. [48] did perform a close study of longitudinal GPS derived travel data, but focused primarily on characterizing daily driving distances.

Some researchers have used stochastic methods, but they often rely on small vehicle use data sets. While using these data is an improvement over the simplistic assumptions in earlier publications, they are often only one day in duration and rely on a small sample vehicle pool. Markel et al. [17] used stochastic selection from a set of individual vehicle use profiles from California, but assumed that all days of the year have similar vehicle use patterns and focused primarily on distribution level effects. Similarly, Sandels et al. [30] used stochastic simulation of vehicle use but worked from a small set of vehicle use profiles synthesized from a German travel study. Other publications have used stochastic methods as well, but have been constrained by a focus on V2G for specific applications, such as smoothing wind generation or penalizing aggregators who cannot fulfill their ancillary service deployments [1,24,49].

To address the shortcomings in many studies of PEV charging and V2G systems with respect to the characterization of demand for electricity to support PEV charging, this dissertation includes a close examination of the travel data employed for modeling. Both GPS- and diary-based travel data are reviewed closely to identify patterns present in vehicle use and determine whether any anomalous conditions exist in the data that might require correction. PEV charging data prepared from the analyzed travel data and used in various models are verified against empirical charging data. These data review and preparation efforts are conducted prior to the

development of models or simulations to ensure that inputs created for those models are consistent with both the underlying travel data and real-world PEV charging decisions made by vehicle owners.

2.3 Plug-in electric vehicle modeling

Even if vehicle use patterns are well-characterized, remaining uncertainty can create operational challenges on the grid. Extensive research has sought to identify strategies to avoid the possible detrimental consequences on grid operations of PEV proliferation, largely focusing on the development of charging control strategies. These strategies rely on control approaches that are either centralized, where a central authority explicitly manages the charging activity of vehicles in the system, or decentralized, where individual EVSEs respond independently to data or command signals from an existing centralized entity (e.g. a utility or independent system operator). Fettinger et al. [50] propose that centralized coordination of vehicle charging reduces communication bandwidth requirements as compared to a decentralized system where EVSEs autonomously coordinate charging. Aabrandt et al. [51] develop optimal charging control strategies using a centralized approach based on prior work from Sundstrom and Binding [52]. Kristoffersen et al. [53] and Acha et al. [54], among others, also use a centralized control conceptual framework for their vehicle charge management simulations. These centralized control methods appear most commonly in research that addresses multiple objectives or is developed with the intention of expanding the approach to support secondary objectives. In contrast, research that proposes methods to reduce peak system demand often use decentralized control approaches. For example, Su and Chow [55] develop an algorithm that could be deployed within EVSEs and provide distributed automated control of charging. Mets et al. [22] also develop approaches for decentralized control but focus specifically on

home charging, where charge control is provided by a Home Energy Management System (HEMS).

Vehicle charging models developed in the literature have sought to estimate opportunities to improve system performance and to address a wide range of operational challenges. Reducing charging costs for individual vehicle owners has been an area of particular interest [46, 51, 53, 56], since an electric vehicle is the largest single load in a home and, unlike air conditioning or heating loads, is persistent year-round. It is anticipated that the introduction of these large electric loads will put significant strain on the distribution system [50, 54], and researchers have found that even limited charge scheduling might be sufficient to delay or avoid distribution equipment replacement [17, 18, 57]. Assessments of the available generating capacity to support overnight charging [16, 19–21] have led many researchers to examine opportunities for peak load reduction and load shifting [22, 55, 56]. Though the efforts applied to these manifold objectives are themselves diverse, the results might be applicable across objectives. For example, reductions in peak demand might, at the distribution level, reduce operation and maintenance costs, and shifting charging to overnight hours can, assuming vehicle owners are exposed to time-of-use or near real-time electricity prices, reduce charging costs. Most authors, however, have not explicitly quantified these cross-objective benefits, possibly because not all of these objectives can be measured in monetary terms, or the benefits might not accrue to the same entities. This work focuses on services provided by PEVs that can be monetized, such that some portion of those earnings can be paid back to vehicle owners.

A wide array of modeling approaches have been applied to the study of vehicle charging. Unit commitment methods borrowed from the electric grid modeling literature and often used to model large-scale energy storage have seen limited

adoption [46, 58]. Other optimization-based methods developed specifically to assess control approaches for vehicle charging demand, subject to various objectives, enjoy widespread application. Most optimization approaches used by researchers are deterministic, that is, the models require fixed, known values for all of the parameters at every interval in the model. Linear programs are occasionally used [53, 56], but the characteristics of the objective function or constraint equations employed in many studies require the use of quadratic programming (QP) [22, 53, 57]. There appears to be little consistency among studies using QP methods as to which constraints or elements in the objective are the source of the quadratic term(s). Given that the results of these optimal charging models could be sensitive to marginally predictable input parameters, such as market prices and driver behavior, some models attempt to mitigate the uncertainty of those parameters or build robustness into their charging schedules [51, 53]. Fettinger et al. [50] propose a two-part semi-heuristic "stochastic" formulation to reduce distribution system costs, estimated based on related parameters. More generalized stochastic programming approaches have not been adopted in the literature.

To mitigate the potential variability in vehicle charging or availability to provide V2G services, instead of using stochastic programming methods that quantify uncertain system elements, some researchers have proposed that vehicle owners simply notify aggregators of their schedules in advance. This approach ensures that vehicle use can be considered as a deterministic, exogenous component of their model [25, 45–47]. For example, Han et al. [59] assumed that vehicle use behavior cannot be readily quantified and that V2G program participants would need to notify the aggregator before using their vehicles. Stein et al. [60] developed a model for PEV charging that requires user input so that charging loads can be scheduled optimally. Both Sioshansi and Denholm [25] and Saber and Venayagamoorthy [58]

assume that vehicles can be optimally scheduled for V2G services. The approach in this work builds upon these previous efforts by avoiding the use of deterministic methods to describe PEV use or electricity demand for vehicle charging and avoiding the assumption that vehicle charging loads can be made deterministic [25, 45, 47].

With adequate charging control in place, generating capacity otherwise idled overnight can serve charging loads, even for very large PEV fleets [16, 19–21], and distribution infrastructure upgrades might be deferred or avoided [17, 18, 57]. Most PEV owners are not currently exposed to real-time electricity prices, and it is possible that even with time-of-use or real-time pricing for retail customers, drivers will not change their charging behavior. Pre-scheduling of charging would require additional effort on the part of PEV owners, and they might reject mandatory participation in a charging control program operated by their utility or a third-party entity. Because of these uncertainties in the willingness of PEV owners to modify their charging choices, this work seeks to explore the effects of unscheduled charging. Since these effects are well-established with respect to overall electricity use and peak-hour demand [16, 19, 22], this study focuses on ancillary services, particularly frequency regulation procurements.

2.4 Vehicle-to-grid system models

Some PEV charging models use the language "vehicle-to-grid" to describe methods that halt or slow the rate of charging subject to a particular control strategy. For the purposes of this work, V2G refers exclusively to operational cases where vehicles can charge at varying rates and also discharge, returning power back to the grid when it is needed. This electricity sent onto the grid could be used for any purpose that is served by electricity from traditional sources, though there might be

some services that benefit more from V2G, as an energy storage-based provider, than others [13, 14]. Determination of when vehicles should be charged or discharged is typically modeled as handled by a centralized entity, referred to as a "vehicle aggregator." While vehicle charging control could be decentralized, where HEMS or EVSEs could respond to price signals or coordinate with other such devices, centralized management of a V2G system is required because sending power onto the grid requires participating in an electricity market, but in US electricity markets, current rules specify minimum capacity thresholds that preclude PEVs from participating as individuals. The aggregator serves as the representative in an electricity market for a large group of PEV owners, and is thus responsible for market participation decisions, determining which vehicles will provide the requested services at a given time, and providing compensation (if applicable) to the PEV owners in its program. In this work, it is assumed that participation in a V2G program would be optional, and not instituted by a utility as a condition of charging a PEV in their service area, since mandatory programs might frustrate first-time PEV owners and alienate potential buyers.

In a V2G (or vehicle charging) model, vehicle use patterns and vehicles themselves can be modeled using a few different strategies. A few authors have developed profiles for archetypal vehicles and use those generalized profiles to create a fleet of the desired size and composition [17]. Generalized vehicle profiles allow for changes in the number of PEVs modeled using a single parameter, but sacrifice detail present in the original vehicle travel data. Treating vehicles individually offers greater accuracy, but can make adding vehicles difficult, particularly if drawing a small fleet from a large sample of vehicles. Within the model formulation itself, vehicles can be either pooled or discretized. Modeling a set of discrete vehicles enables analysis of the impacts of specific V2G operational strategies on a per-vehicle basis, which can be useful for

assessing the experience of program participants. Discrete vehicles can also be readily grouped by household to determine the impact of PEVs or V2G on a per-household basis [61]. As the size of the fleet increases, the number of non-zero variables in a model with discrete vehicles will grow rapidly, which might result in an unacceptably large problem. Pooling the vehicles and applying constraints to the collective resource accurately represents the operational perspective of the vehicle aggregator. It also reduces the number of variables in the model and can help control computational expense in otherwise complex models. By pooling the resource of all the vehicles available, it is possible that the operational strategies for the aggregator indicated by the model might not satisfy the vehicle use needs of individual participants, and the extent of that effect cannot be readily determined.

There are several entities that could capture the monetary benefits generated from a V2G program, including individual vehicle owners, electric utilities, system operators, and vehicle aggregators. If the aggregator operates as a third-party in the system, much like an independent power producer, they will require compensation to offset their operational costs. There might be additional costs outside of the aggregator's operations, such as upgrades to communications infrastructure to reduce latency or replacement of EVSEs with models that support V2G communications, that would need to be covered using revenue generated from V2G operations, but these costs are not well-explored in the literature. It is likely that PEV owners will require some inducement to enter a V2G program, which might take the form of a financial incentive, and many researchers have envisioned that a V2G program would generate sufficient revenue that participants would be compensated regularly [28, 47, 62, 63]. Many V2G models, regardless of their methods, have focused on quantifying these payments to program participants. As a further simplification, most studies that estimate these V2G revenue opportunities to date treat electricity market prices as

exogenous [7, 28, 45, 47, 62–64]. Some authors have noted that this assumption might have a significant effect on expected payments to V2G participants, though none has attempted to address it comprehensively. Sioshansi et al. [7] stated that using a price-taker assumption will tend to overestimate earnings from market participation for large-scale grid storage systems. White and Zhang [47] noted that in scenarios with high V2G participation, revenues for arbitrage and regulation services could change, but they only quantify the revenue impacts on V2G arbitrage. They also note that the number of vehicles that can provide frequency regulation without affecting prices is far smaller than most of the scenarios in their study.

Several authors have specifically quantified the potential earnings that a PEV owner could expect from participation in a V2G system. The estimated profits from publications that examine PEV participation in the frequency regulation market, like the model in this work, are summarized in Table 2.1. All of the publications in the table assume all revenues earned from market participation are directed to PEV owners and make some attempt to account for battery degradation costs, which have already been deducted from the annual profits indicated. These publications also assume that the aggregator's offers in the day-ahead market have no effect on prices; only White and Zhang [47] recognize the potential weakness of this assumption in their discussion. The profits calculated in these publications can be approximated using typical day-ahead market conditions, since the revenues for providing frequency regulation services are primarily derived from the day-ahead market. Day-ahead markets have hourly intervals, and in each interval, prices for frequency regulation p in ERCOT are usually on the order of \$100/MW, but for the markets and years examined in these publications, prices were one order of magnitude smaller, typically between \$20/MW and \$50/MW [47, 64]. For the results shown, PEVs have a 10 kW or larger EVSE, and their average daily availability x is between 0.5 and 0.96.

Table 2.1: Several publications calculate the per-vehicle profits available from a V2G system providing frequency regulation services. All of the publications tabulated assume the aggregator does not influence prices in the day-ahead market. Each study also includes an estimate of battery degradation costs, which have already been deducted from the profits shown. The result from White and Zhang is from the most conservative case considered for frequency regulation: the shortest battery life and lowest earnings of the three years analyzed.

Authors	Market	EVSE size (kW)	Annual profit ($)
Kempton and Tomic [64]	CAISO	15	2,554
White and Zhang [47]	NYISO	10	2,410
De Los Rios, et al. [63]	ISO-NE	19.2	1,400
Quinn, et al. [62]	CAISO	10	407

Multiplying these values together yields estimated per-vehicle revenue of $\$87.6\,p\,x$. Frequency regulation is divided into two services, thus doubling the total effective revenue to $\$175.2\,p\,x$. At an average price p of \$35/MW and availability x of 0.9, estimated annual V2G revenue for a single PEV is \$5518.80. Since the results from the studies in Table 2.1 include battery degradation costs in their profit estimates, this estimated revenue compares reasonably well with the literature.

The V2G aggregator model developed in this dissertation compares the assumption of exogenous ancillary service prices with an endogenous price approach to determine the effect of the exogenous price assumption made by many authors. To control model size and computational expense, the model uses a set of vehicle use scenarios developed using a clustering technique, conceptually similar to the approach used by Kristoffersen et al. [53]. The model is framed from the perspective of the vehicle aggregator, and treats the storage resource available at any given time as a pooled resource. Further detail on the modeling approach is provided in Chapter 5.

2.5 Modeling systems with stochastic elements

Many systems can be modeled using optimization methods, such as linear programs (LPs) or mixed-integer programs (MIPs). These systems are often subject to uncertainty, which is not readily captured in traditional LP or MIP model formulations. There are a range of methods that can be applied to model the stochastic elements of the system of interest. Sensitivity analysis, "worst-case" analysis, and scenario analysis are popular methods to characterize a system with respect to unknown future conditions. With sensitivity analysis, a deterministic mathematical program is developed to assess whether the solution is relatively invariant to changes in the values assigned to the model parameters. A model solution that changes little subject to a range of values for the parameters of interest is considered optimal under uncertainty [65,66]. It can be shown, however, that such a solution can be optimal for a limited number of cases, perhaps only those analyzed, while performing poorly or becoming infeasible in other conditions [65]. Worst-case analysis and scenario analysis are variations on the sensitivity analysis approach, examining combinations of parameter values that the modeler believes are possible, observing the behavior of the objective function, and using a heuristic approach to adjust the model parameters or tune the solution until the performance is satisfactory [65, 66]. More sophisticated methods can be used to find better solutions to stochastic problems. These methods include dynamic programming, chance-constrained programming, and stochastic programming, among others.

Dynamic programming (DP) methods are commonly used to solve problems where stochastic elements arise, including reservoir operation, inventory planning, and (financial) portfolio optimization. Apart from these application areas, this modeling approach is particularly effective for those problems that can be readily decomposed

into a subproblem that can be solved recursively and have a finite number of possible states [67]. DP can also be applied to systems where there are multiple stages in the system or process, but only when decisions at each stage are independent of previous stages [67]. DPs suffer from the "curse of dimensionality" and thus grow exponentially with the number of possible states [68]. In a V2G model, the stochastic variables are naturally treated as non-negative continuous, and the system cannot readily be defined by a small set of discrete states, thus dynamic programming is not an ideal modeling approach.

Chance-constrained programming offers an appealing alternative to dynamic programing because it does not suffer from the same dimensionality issues. Chance-constrained programs are developed by modifying the underlying LP formulation by relaxing constraints that contain the stochastic parameter of interest. These constraints are relaxed according to a measure that defines what fraction or percent of the time they must be satisfied (the "chance"). The constraint equations can then be reformulated to account directly for variability using the selected chance measure with the cumulative density function of the stochastic parameter. This modeling approach has the advantage of being able to be solved as an LP, and is particularly appealing in situations where the operational objective is specified in such performance terms (e.g. 93% of flights should depart on time). For the purposes of a V2G model, the preselection of a chance measure might be arbitrary or weakly justified. Further, chance constrained optimization models address the lowest-cost $n\%$ conditions and avoid the remaining $(1-n)\%$ conditions, even though the additional cost to address those cases might be low [65]. Avoiding the $(1-n)\%$ cases that are the most difficult could also yield a solution that is quite poor for the overall problem [65]. Chance-constrained programming is thus not suitable for capturing the full range of the conditions of interest.

SP[1] transforms traditional deterministic models, which have a single stage in which all information about the system is known, into a model with two or more stages, where information becomes known after the completion of each stage (or prior to the start of the next stage). These stages essentially represent states of knowledge. At any point in a deterministic model with stochastic parameters where there is a change in the state of knowledge about those parameters, there should be a transition to a new stage. Models are typically either two-stage or multi-stage (many more than three stages). Two-stage models are characterized by a decision that requires some commitment before a single event, quite often a purchase or investment decision subject to uncertain demand. Once demand is realized, the remainder of the problem can be solved. Multi-stage models typically describe ongoing operations of a system, though multi-stage models can sometimes be reformulated as two-stage models. Solution methods for two-stage models are well-established in the literature, and the ability of this approach to more fully characterize the stochastic elements of a system make it a suitable candidate for modeling a V2G system. In the case of a vehicle aggregator, the day-ahead and real-time components of an electricity market are well-suited to a two-stage program, where vehicle battery availability and state-of-charge and ancillary service deployments are unknown in the day-ahead market. Further discussion regarding the two-stage SP structure and the development of the stochastic V2G model is detailed in Chapter 5.

[1] Stochastic programming should not be confused with stochastic optimization, which is a class of iterative solution methods for mathematical problems, such as simulated annealing and genetic algorithms.

2.6 Conceptual approaches to decision analysis

The methodological approach applied to decision analysis in this work is one drawn broadly from a microeconomic perspective of individual decision-making [69]. When applying these conceptual frameworks to discrete decisions, individuals seek to maximize the expected value of their decision, based on their assessed valuation of the possible outcomes (or "prospects" [70]) and the probabilities of each of those prospects. Such individuals, often termed "rational actors" in the microeconomic literature, make internally consistent decisions with regard to a range of prospects and accept the transitivity of prospects between decisions. Howard [70] refers to these fundamental characteristics of decision-makers as the "five rules of actional thought," which are prerequisites for this approach to decision analysis. The assessment of probabilities drawn from information held by the decision-maker or the values and preferences that underlie the decision-maker's risk attitude are both important components of the decision [71]. An advantage to this framework for decision analysis is that while these components — information, values, and preferences — are important, the *basis* for values and preferences or the *sources* of information are only relevant insofar as parameters material to the decision are drawn from them.

Outside of this discipline, there are a range of approaches on a continuum between individual and social "frames," referring to, at their extremes, exclusionary decision-making centered around the preferences and interests of single individuals, and decisions made by individuals, but where the agency of the individual is curtailed by social constructs [69]. These approaches include direct, empirical study of decision-makers' preferences; technology diffusion models; and sociological assessments of external influences that can affect behaviors, preferences, and attitudes [69]. While these approaches need not be applied to the specific cases studied here, espe-

cially since these decisions are prospective, they are relevant to understanding the factors that influence consumer preferences with respect to vehicle use patterns and vehicle type choice.

A method of particular interest from decision making under uncertainty is options analysis, sometimes referred to as "real options." In the context of a decision, an option creates an opportunity for a decision in the future, following the availability of new information relevant to the original decision [70]. While there exist financial instruments called options that have this characteristic, there are more accessible examples of options encountered in daily life. Car insurance (apart from the minimum liability coverage required by statute) is an option that many people purchase. Car insurance is purchased in anticipation of a possible collision in the future. The policy has some premium (the option price) that is paid regularly, and if the driver has an accident, they can draw on the policy to recover their damages, less a deductible (the exercise price). Individuals who purchase car insurance have assessed the probability of having an accident and the potential costs associated with such an event and determined that they justify the cost of the policy they select. Extended warranties and other types of insurance have similar option structures. In the particular case of PEVs, BEVs have a finite driving range available from their battery. Once the battery is exhausted, it must be recharged before the vehicle can be used again. PHEVs have an ICE drivetrain that eliminates this range limitation, but also typically cost more than a comparable BEV and have a shorter electric driving range. The price differential between a BEV and a PHEV is the option price, where the option is provided by the ICE, and the exercise price is the cost of operating the ICE once the battery is depleted.

2.7 Summary

As discussed in this chapter, this work augments the existing literature in several key ways. In Chapter 3, vehicle use patterns are explored to improve the understanding of how those patterns vary in time. This understanding is critical to the remainder of this work and can impact the validity of simulation efforts in the literature. Given the focus in the literature on developing charging control strategies, Monte Carlo simulations in Chapter 4 are used to better understand the impact of unscheduled vehicle charging on frequency regulation procurements and overall electricity demand. A novel SP formulation of a V2G aggregator providing frequency regulation in an electricity market is developed in Chapter 5. This approach considers the potential effect of aggregator market power on day-ahead regulation prices, uncertainty in vehicle battery availability and real-time regulation deployments on the aggregator's day-ahead offers, and revenue diversions to cover payments to program participants, which are all typically excluded from modeling efforts in the literature. Finally, based on results from the SP V2G model, Chapter 6 applies decision analysis methods to assess the potential for PEV owner participation in a V2G program.

Chapter 3

Temporal vehicle use analysis and evaluation of the interaction between plug-in electric vehicle charging and electric load and wind generation

Accurate V2G models require an understanding of the temporal variations in vehicle availability. Current literature on V2G systems could benefit from an improved understanding of the temporal variations in vehicle availability. To that end, this work explores recent vehicle use data to determine the presence and timing of vehicle use patterns, the time of day when vehicle use peaks, and whether location has a significant impact on vehicle use patterns. Further, because frequency regulation procurements are highest during periods of significant change in electricity demand or wind generation [4], the assessment of vehicle use patterns is followed by a comparison of net load and vehicle use to determine whether vehicles might be available when they are most needed to provide grid services.

3.1 Vehicle use assessment

3.1.1 Driving data sources

In this study, driving data collected by the Puget Sound Regional Council (PSRC) were used to estimate vehicle use patterns [31]. These data were collected by PSRC using Global Positioning System (GPS) vehicle tracking devices on 436 vehicles from November 2004 through April 2006. During that period, PSRC conducted a study to explore the effect of various tolling strategies on route choice decisions

among study participants. These data with tolling influence effects were excluded from this study, leaving eight unique months of available data — January through June, November, and December. The PSRC data were chosen for this analysis over other available GPS traffic study data or national travel survey data because this analysis includes an examination of seasonal variations in vehicle use, and long-term data from a constant set of vehicles were desired for that portion of the study.

The PSRC data were made available by the National Renewable Energy Laboratory (NREL) through their Transportation Secure Data Center (TSDC) online repository of GPS study data. To protect the privacy of study participants, NREL converted the raw GPS data obtained from PSRC into "tours" reported with minute-by-minute resolution. Tours are individual vehicle trips grouped by a common purpose. For example, an individual might drive from their home to a grocery store, then to a pharmacy, a gas station, and a hardware store before returning home. These five trips could be grouped together into a single tour because they are in series and are all devoted to household errands. Along with the tour data, NREL released a subset of the demographic data collected by PSRC.

3.1.2 Transformation of GPS study data

For the purposes of this work, the processed PSRC data were converted into the parameter "vehicle use," denoting the fraction of the vehicle fleet being driven at any given time. Approximately 130,000 unique tours were recorded by PSRC, excluding the tolling influence portion of their study. In these data, PSRC recorded tours as long as 159 days, thus the data were first modified to exclude tours of excessive length, which was defined as those tours longer than 36 hours. Reporting anomalies in February, March, and April 2005 resulted in minimum vehicle use values well above the maximum reported values throughout the rest of the study, which led

to the exclusion of tours from those months. Because vehicles entered the study pool gradually over the course of the first month of the study, November 2004, the changing size of the pool made accurately assessing vehicle use difficult, thus these data were excluded as well. With the remaining 127,500 tours, a count of vehicles in use was generated. The number of participating vehicles was not constant throughout the study period, so the number of vehicles in each week of the study was calculated from the week of the first and last time each household ID recorded a tour. The count data were then shifted to a minute-by-minute time basis and converted to vehicle use by dividing by the number of vehicles in the study at that time. To determine the time intervals on which substantive variations in vehicle use were present and thus develop the results presented in the next section, vehicle use was averaged across several groups: for each day of the week and each month, each day of the week in all months, each month for all days of the week, all weekdays and weekends for each month, and all weekdays and weekends in the study. These results are shown in Figures 3.1 through 3.3.

3.1.3 Patterns identified from vehicle use data

Figure 3.1 shows for two months how vehicle use varies among weekdays, weekends and all days in the month. In each case, weekdays have two distinct local maxima, one around noon and another around 6 pm. In addition, weekdays show a rapid increase in vehicle use during the early morning hours, reflective of the morning commute to work, followed by a gradual increase towards peak use around midday. Weekend vehicle use is markedly different, characterized by use increasing later in the morning, a single peak around midday and a more gradual decrease in vehicle use throughout the afternoon and evening hours. The standard deviation of both the weekday and weekend data, shown as a grey area around the average, indicates

that there is minimal variation in the magnitudes and timing of vehicle use within a month's weekdays and weekends. Though the timing and magnitude of the midday peak is similar for weekdays and weekends, averaging all the days of any month together yields a smoother use curve, associated with the lack of a rapid initial increase in use during the early morning hours and the absence of a second peak during the evening commute on weekend days. The mismatch between weekday and weekend vehicle use is also apparent from the standard deviation of the monthly averaged data, which does not follow the average line throughout the day. These findings indicate that any studies of PEV charging or V2G systems should avoid yearly or monthly averaged vehicle use profiles, as weekday and weekend vehicle use are markedly different.

Without further examination, the apparent conclusions from Figure 3.1 could be misleading, particularly with respect to the separation of days of the week into weekdays and weekends. Figure 3.2 shows the average of each weekday over the course of the study period. As can be seen in these figures, the trends are similar regardless of the day of the week, supporting the earlier separation of data into weekday and weekend groups. It appears that there is a slight reduction in vehicle use throughout the first and last day of the work week, which could reflect alternative work schedules.

Having separated the data into weekday and weekend groups, vehicle use was then compared across all the months in the study to determine whether use was affected by seasonal changes in weather or daylight hours. It was anticipated that meteorological conditions, both weather and the duration of daylight hours, could affect vehicle use in three primary ways: by encouraging or forcing individuals to stay home all day, to travel earlier or later than they otherwise would, and to slow down while driving, increasing congestion and travel times. These effects would appear in

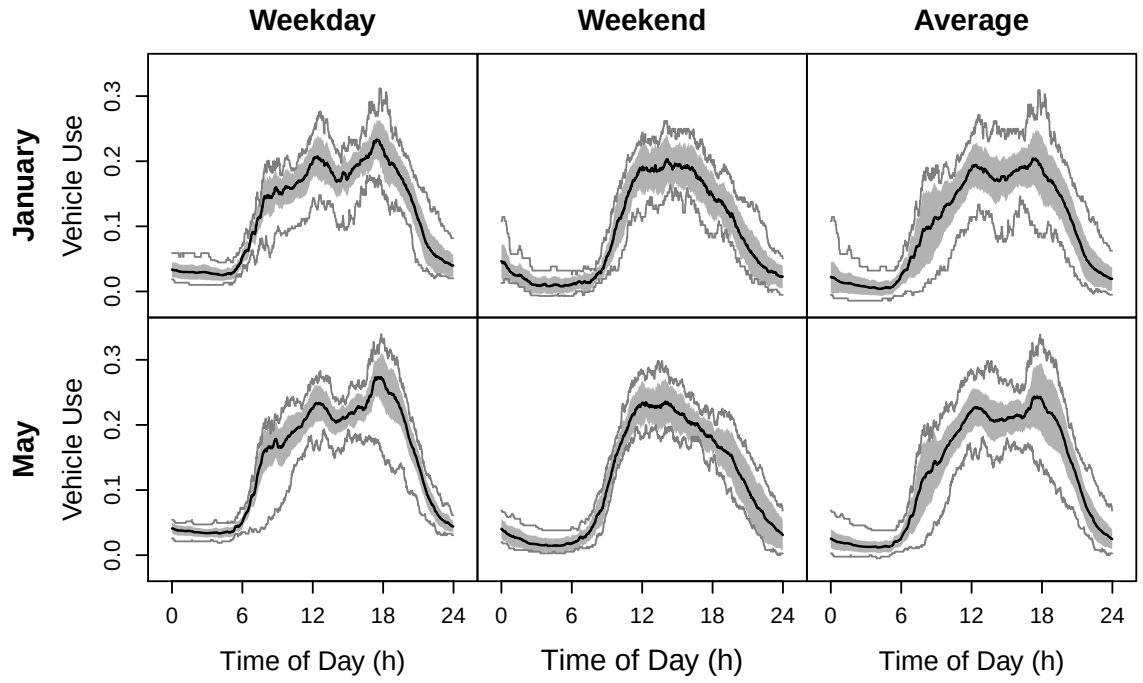

Figure 3.1: Weekday and weekend driving patterns differ significantly, regardless of the time of year, and averaging weekday and weekend data yields estimated vehicle use that is not reflective of either weekday or weekend patterns. The shaded grey area in each subplot shows the range of one standard deviation above and below the average. Grey lines indicate the maximum and minimum vehicle use values.

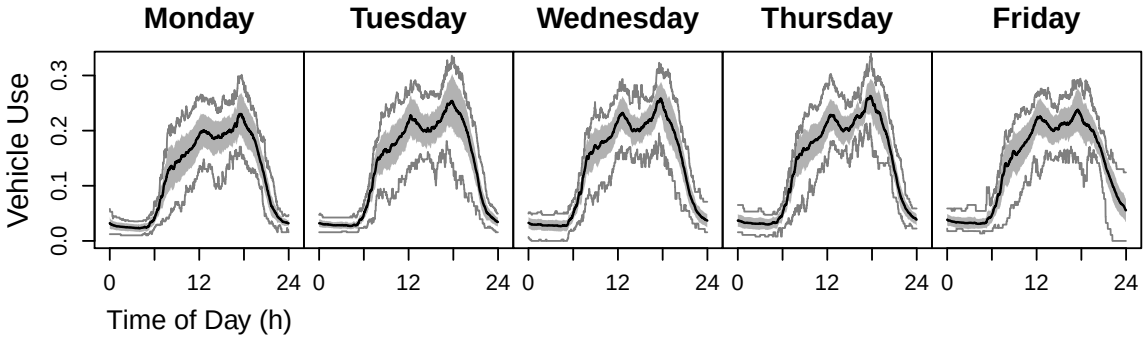

Figure 3.2: Comparing weekdays across all months in the study reveals that differences in vehicle use between weekdays are small.

the data as a reduction in vehicle use throughout the day, a shift of or change in the duration of the daytime vehicle use period, and a widening and smoothing of the daytime vehicle use curve, respectively. A selection of months in the study, shown in Figure 3.3, were examined for the potential appearance of these effects. These results do not show significant variation in the timing of peaks across months, indicating that weekday usage times are primarily affected by work or schedule requirements and not daylight or weather conditions. This indication is consistent with the minimal standard deviation that appears in Figure 3.2. Widening, smoothing or shifting of daytime vehicle use periods does not appear to occur between months or within the data in each month.As noted in the US National Research Council's Transportation Research Board (TRB) Highway Capacity Manual, vehicle use patterns might be affected more dramatically in regions that experience severe winter weather conditions that make driving difficult [72].

There also exists the possibility of weather conditions impacting both renewable generation and vehicle use simultaneously. Such a common-mode change would require that a significant quantity of renewable generation be located near a concentration of PEVs such that they experience same weather conditions. (Large weather phenomena and natural disasters, such as hurricanes or earthquakes, are a notable exception.) In Texas, most wind generation is located far from major load centers and likely concentrations of PEVs. This geographic separation, as well as the lack of strong dependence of vehicle use on the time of year, as shown in Figure 3.3, allowed the combination of vehicle use data with wind generation and load data from ERCOT.

Generally, the results in Figures 3.1 through 3.3 are consistent with the TRB Highway Capacity Manual, in that urban vehicle use is relatively constant between

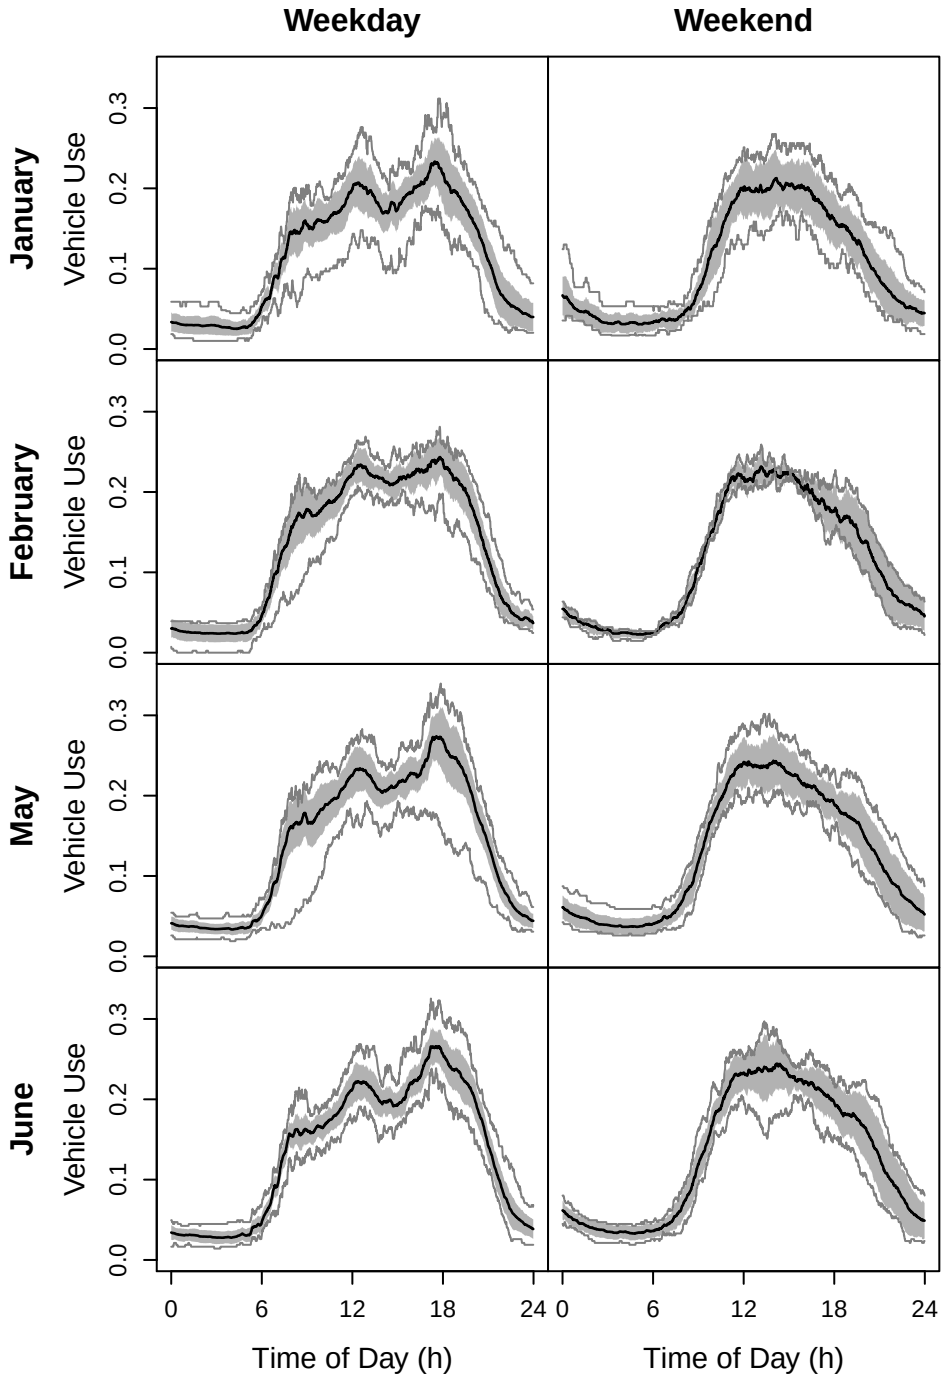

Figure 3.3: Trends in vehicle use on weekdays (left panel) or weekends (right panel) appear similar across months, suggesting that changes in weather or daylight hours are not significant factors.

weekdays but varies between weekdays and weekends, intra-day vehicle use is roughly bimodal, with peak use on weekdays in the afternoons around 5 or 6 pm, and variations in vehicle use between the same hours on different days of the month are small but not insignificant [72].

3.2 Electricity market and battery availability analysis

3.2.1 Transformation of vehicle use into battery availability

In anticipation of comparison with data from ERCOT, vehicle use values derived from PSRC data in the previous section were transformed to represent aggregate vehicle battery availability, or the total energy stored in the PEV fleet that is connected to the grid. This transformation requires modifying vehicle use, denoted by x_t, to reflect battery charge depletion, Q, as a consequence of vehicle use during the day. Because PEV batteries have been aggregated for the entire fleet, an average distance driven can reasonably be substituted for a more complex distribution of tour lengths. The US average daily distance driven of 29 miles was used as a starting point for estimating charge depletion [73]. For a vehicle with a 24 kWh battery, such as the Nissan LEAF [74] if 40% charge depletion is assumed, 9.6 kWh would be depleted over the course of a tour. Assuming an approximate energy use of 0.34 kWh/mile [75, 76], 28 miles would be traversed in a tour, comparable to the US average daily driving distance.

For the example given, a 3.3 kW 240 V EVSE would require a minimum of 2.9 hours (174 minutes) to recharge 9.6 kWh, and to allow for some deviation from the maximum charge rate, an average total recharge time, τ, of 3.2 hours (192 minutes) was assumed. Implicit in this approach is the assumption that vehicles will be recharged only at the end of each tour. While it is possible that some drivers will have access to a charging station at work or while conducting business away from

home, attempting to account for the potential availability of public or workplace EVSEs is outside the scope of this work.

Vehicle use is transformed into aggregate vehicle battery availability with Equations 3.1 through 3.4. This analysis begins by equating the aggregate "PEV storage use" fraction to the previously calculated "vehicle use" fraction, x_t, for all time t. As shown in Equation 3.1, the difference between PEV storage use in each period t and the previous period t–1 yields the variable δ_t. Each period in the model is one minute in duration.

$$\delta_t = x_t - x_{t-1} \tag{3.1}$$

The total charge depleted (as a consequence of driving activity) from the batteries of vehicles completing tours and reconnecting to the grid in each period t is denoted by q_t^d, and is calculated as the charge depletion fraction, Q, of the change in PEV storage use, δ_t, as shown in Equation 3.2. For example, if the PEV storage use fraction changes from 0.4 to 0.35 in a single period, δ_t will equal -0.05 and the total charge depleted, q_t^d, will be 0.02, assuming charge depletion, Q, is 40%. In periods where PEV storage use increases (vehicles are starting tours), q_t^d is zero.

$$q_t^d = \begin{cases} |\delta_t| Q & \text{if } \delta_t < 0 \\ 0 & \text{else} \end{cases} \tag{3.2}$$

The quantity of charge restored in each period (minute), q_t^r, once vehicles' tours have ended, is described by Equation 3.3. This quantity is simply the total battery depletion for period t that requires recharging, q_t^d, divided by the number of periods (minutes) τ required for recharging.

$$q_t^r = \frac{q_t^d}{\tau} \tag{3.3}$$

Modifying PEV storage use, x_t, with the parameters from Equations 3.2 and 3.3 to reflect the effects of charge depletion yields "adjusted PEV storage use,"

denoted by y_t. In each period where PEV storage use decreases (δ_t is negative; vehicle tours are completed), y_t is calculated by adding the charge depleted q_t^d from each of the last τ periods and then subtracting the fraction of the charge restored q_t^r in each period, multiplied by the number of periods that have elapsed since the charge depletion event occurred. By including these two terms, adjusted PEV storage use thus reflects the state-of-charge of vehicles that have recently completed driving tours, at the time of their return (second term), as well as the time required to recharge their partially depleted batteries (third term). After τ periods, when the batteries of vehicles that completed tours in period t are fully charged, the latter two terms in Equation 3.4 drop out, as indicated by the summation range on those terms.

$$y_t = x_t + \sum_{i=t-\tau}^{t} q_i^d - \sum_{i=t-\tau}^{t} q_i^r (t-i) \tag{3.4}$$

The effect of charge depletion is especially evident in the evening hours when most tours end. Because of the charging time required, the cumulative effect of these charge-depleted PEVs completing their tours at the end of the day increases apparent vehicle use throughout the evening and early morning hours. It should be noted, as shown in Figure 3.4, that partially depleted batteries can offer frequency regulation down, when a reduction in generation or an increase in load is needed to correct grid frequency, in addition to regulation up and spinning reserves. For later use with ERCOT data, the complement of charge depletion-adjusted PEV storage use was calculated following Equation 3.5. This parameter, z_t, is referred to here as "battery availability."

$$z_t = 1 - y_t \tag{3.5}$$

In addition to averaging the data to investigate patterns in vehicle use, various statistics on the study data were assessed to compare with national driving statistics, including the number of miles driven by each vehicle during the study period, the

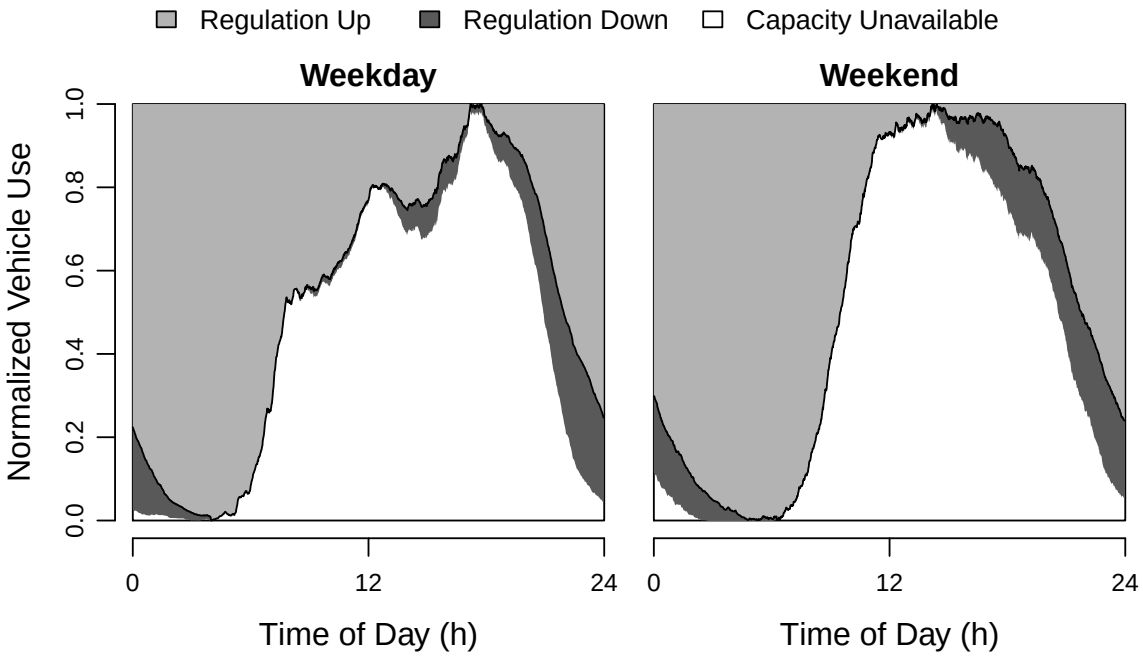

Figure 3.4: Normalized regulation up capacity is highest when most vehicles are stationary and fully charged in the early morning hours, before most tours begin, while regulation down capacity is highest in the evenings, when vehicles' tours are completed and their batteries have been partially depleted.

number of tours taken, and the number of days the vehicle was used. These data were largely consistent with the National Household Travel Survey (NHTS), conducted by the US Department of Transportation, except that annual vehicle mileage was lower in the PSRC study [77]. This difference in vehicle mileage could be a consequence of several factors, such as the geographic distribution of PSRC study participants or demographic differences between the NHTS and PSRC samples. Despite the difference in average annual vehicle miles traveled, comparison of the results against the TRB Highway Capacity Manual indicates that temporal trends observed in the vehicle use data are consistent with driving data collected in other municipalities [72]. This consistency indicates that while the PSRC data might underestimate volumes in other cities, the timing and trends in use are comparable, thus no further manipulation of the data was performed, as the primary focus of this work is on the temporal characteristics of vehicle use.

It should be noted that initially, PEV usage patterns will likely differ from this data set. For example, two-car households with one PEV might rearrange their vehicle use such that shorter trips are all taken using the PEV and less-frequent, longer trips are completed using their other car. The details of how people will change their driving choices might be able to be estimated through close examination of driving patterns filtered with demographic parameters, along with forthcoming data on PEV use from early adopters, but this effort is subject to significant uncertainty and is outside the scope of this work. Further, it is anticipated that as PEVs become more widespread, PEV owners are no longer exclusively early adopters, and concerns like range anxiety are resolved through consumer education and improved public charging infrastructure, PEV use patterns will approach current vehicle use patterns in the general population.

PEV use patterns might differ from current use patterns as a result of the financial incentives created by a V2G program. This effect is highly dependent on the magnitude of the revenue opportunities for vehicle owners; however, as the number of vehicles providing V2G increases and thus the importance of any behavior change increases, the revenue potential from providing V2G will probably diminish. Further study regarding the elasticity of departure time choice and the effect of increasing V2G participation on ancillary service capacity prices could illuminate the value of accounting for changes in driver behavior, but such analysis is outside the scope of this work.

3.2.2 Comparison of ERCOT data and battery availability

Wind generation and electric load data from 2010 were obtained from ERCOT to compare with the vehicle use data described previously. These data were reported by ERCOT in 15 minute intervals for each month, with generators grouped by fuel type, making wind generation and total load discernible. Data corresponding to the months missing from the PSRC data, July through October, were removed, and the remaining data were averaged in the same groups as vehicle use in preparation for subsequent comparisons.

Wind generation and electric load in ERCOT were compared with battery availability to determine whether battery availability and peaks in load were aligned. Net load, l_t, which is load minus wind generation, was used for this comparison. Both battery availability and net load were normalized to fall between zero and one for the average weekday and average weekend of each month. The resulting parameter, a_t, that describes the comparison is denoted by "availability-load alignment" and was calculated by multiplying battery availability and net load.

$$a_t = l_t z_t \tag{3.6}$$

Peaks in the availability-load alignment curve reflect periods when vehicle batteries are most available to provide grid services and net load is relatively high. This approach assumes that neither the user nor the utility schedules charging of PEVs, but rather that vehicles will charge when owners plug them into an EVSE. As noted previously, it is assumed that charging will primarily take place at home, which is the terminus of most of the tours in the PSRC data. Depending on the extent to which net load and battery availability are aligned, PEVs might be available during peak demand hours when ancillary services are often crucial, or vehicle charging might not require widespread scheduling to ensure primarily nighttime charging. Figure 3.5 shows availability-load alignment for January. Representative of cooler months in Texas, wind generation comprises a greater portion of generation, and electric load is bimodal, with morning and evening peaks associated with residential activity and limited daylight hours. Alignment between battery availability and net load is best when availability-load alignment and net load are similar. Apart from overnight hours when net load is lowest, battery availability and net load in Figure 3.5 show the greatest coincidence in the early morning hours, between 6 and 8 am, and in the late evening hours, after 9 pm. These periods of alignment arise from net load peaking when people are preparing to depart for work or errands in the morning, just before using their vehicles, and again when people complete their evening activities, having just used their vehicles for their last tour of the day.

As noted by many authors in the literature, in ERCOT, diurnal variations are present in both wind and load, and these variations change seasonally [4]. To examine seasonal variations, Figure 3.6 shows battery availability, net load and availability-load alignment for June. Representative of hotter months in Texas, air conditioning loads yield high electricity demand during the afternoon and evening hours. Vehicle use shows a similar trend, with use peaking during the afternoon and evening hours.

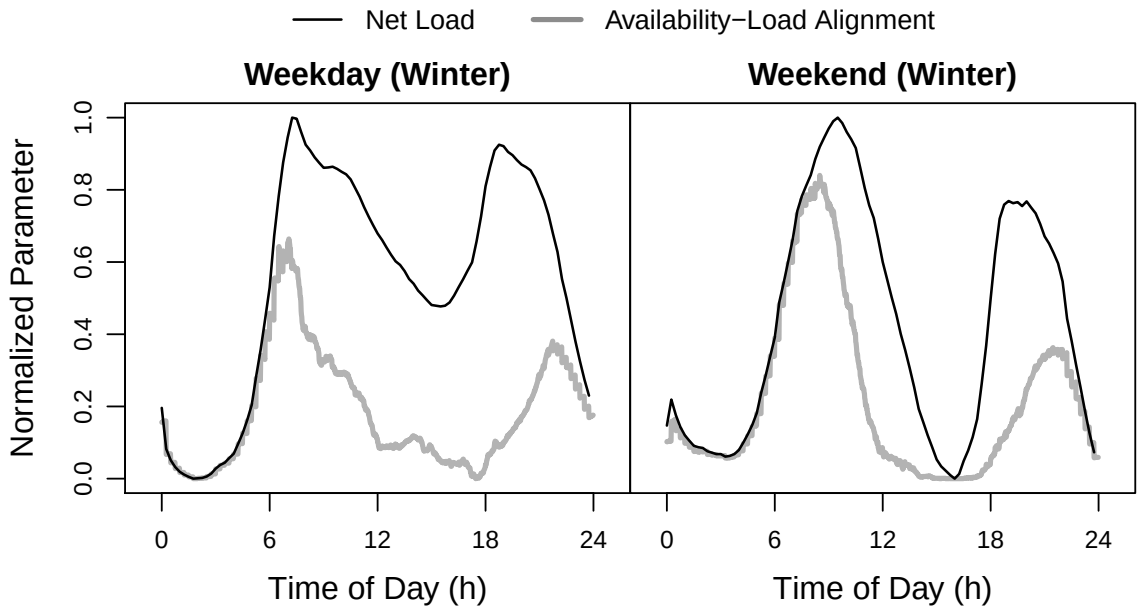

Figure 3.5: In the winter, hours of highest vehicle utilization in the morning and afternoon are bounded by periods of high electricity demand, before people leave home in the morning and after they arrive home in the evening, leading to strong alignment between availability and load during those hours.

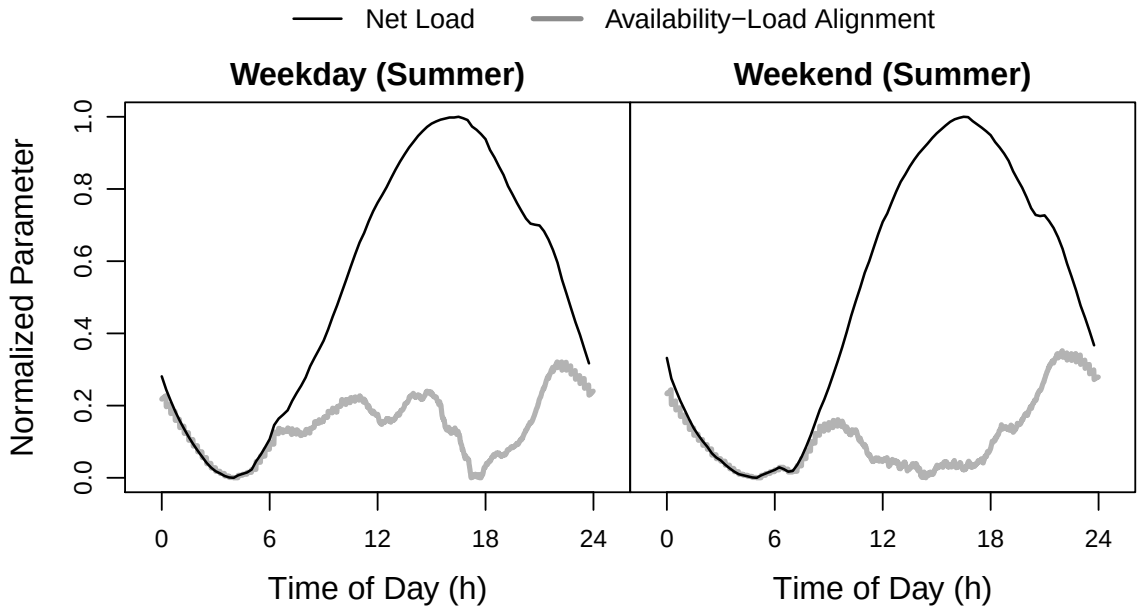

Figure 3.6: In the summer, battery availability is roughly the inverse of electricity load, leading to minimal alignment between the two parameters except on a limited basis midday and again in the late evening.

Because trends in vehicle use and net load are closely aligned, rising throughout the day, beginning around 6 am, and falling again in the evening, starting at around 6 pm, the availability of vehicles to provide load shifting, peak shaving, or valuable ancillary services, is far more limited. Further, at times when vehicle use is lowest, net load is also near its minimum, and the potential value of V2G services is thus diminished.

3.3 Conclusions

The first part of this analysis sought to assess the presence of patterns in vehicle use across various time intervals — within days, between days of the week, and between months of the year. The results yielded two primary findings. First, diurnal driving patterns vary significantly between weekdays and weekends, and second, ve-

hicles are often used during hours when grid services are in high demand, especially during the summer months, contrary to prior V2G studies. Examining GPS vehicle use data from the Puget Sound region revealed that weekdays and weekends show significantly different vehicle use profiles. Weekdays have three distinct periods — a rapid increase in vehicle use during the early morning hours, a midday peak and an afternoon peak. Weekends, on the other hand, have a single peak in vehicle use around midday, with lower use in the morning and late evening hours. Examination of the data for each day of the week revealed that variations within weekdays and weekends are comparatively minor. Identifying this difference between weekdays and weekends is crucially important, as it indicates that V2G studies that use average driving profiles should be careful to not conflate weekday and weekend driving data, as doing so could result in under-prediction of vehicle usage, particularly in the early morning and late evening hours. The data also show that variation between months is limited, but the presence of seasonal variations could be a function of climate. In particular, in regions where winters are especially severe, limited daylight and poor weather conditions could restrict mobility, yielding lower overall vehicle use, and possibly a slight narrowing of the hours of peak vehicle use. The households monitored in the PSRC study drove somewhat less than the average urban US household on an annualized basis, thus it is expected that these results and those for other regions differ only in the magnitudes of vehicle use.

This study also sought to compare the relationship between battery availability and net load. In ERCOT, battery availability appears to align best with net load during cooler months, when net load is bimodal and electricity use occurs primarily in the hours just before vehicles are used, between 6 and 8 am, and after tours are completed in the evening, beginning around 9 pm. Overnight, between those times, vehicles remain available to provide ancillary services while they recharge. In the

summer, significant air conditioning loads in ERCOT yield a mismatch between net load and battery availability, suggesting V2G services might be limited during those months. Given the regional dependence of wind (or other stochastic renewable) generation and electric load, and the potential for some variation in vehicle use between regions, it is important that researchers interested in performing V2G studies use regional data and, if possible, perform a long-term analysis to be able to account for seasonal variations in wind generation, electricity load and vehicle use.

Chapter 4

Assessment of the impact of plug-in electric vehicle charging and variability in charging load on electric load and frequency regulation procurements

The importance of maintaining the supply-demand balance in the electric system motivates efforts to mitigate and forecast sources of variability. Since PEVs could be a significant new source of variability, a better understanding of the stochastic characteristics of vehicle use is desirable. Monte Carlo simulation methods can be applied to assess the effect of the PEV fleet size on charging load variability. Dallinger et al. [24] performed such an analysis, but their study relied on German vehicle use data and was limited by characteristics of the German electricity system that constrained the use of vehicles for ancillary services. This effort expands on that study in several key ways: by accounting for demographic shifts with PEV fleet growth, the effect of fleet size on total variability, and the impact of variability on ancillary service needs.

4.1 Estimate the variability in vehicle use with Monte Carlo methods

Many efforts, at both the regional and national levels, have been made to collect historical vehicle use data in the United States for the purposes of transportation infrastructure planning. In 2009, the US Department of Transportation (DOT) re-

leased the current NHTS [32]. The NHTS data offer extensive details on 1.17 million trips reported via telephone surveys between March 2008 and May 2009 [32]. A telephone auto-dialer was used to generate random telephone numbers, thus effectively randomizing the sampled population [32]. The trip data were collected for the assigned day's travel activity, including trips taken using private automobiles, carpools, public transportation, and other personal travel modes.

Using individual vehicle use profiles created from the NHTS data, Monte Carlo simulations were performed to assess possible variability in vehicle use and how that variability changes with the time of day and the size of the PEV fleet. To examine the effect of growth in the fleet of PEVs in a region, eight different fleet sizes were tested: 100, 500, 1,000, 5,000, 10,000, 50,000, 100,000, and 500,000 vehicles. These simulations were performed by sampling directly from the pool of individual profiles, assuming each is equally probable, with a uniform random variable on the interval [1, q], where q denotes the number of profiles available. Based on observations of vehicle use patterns detailed in Chapter 3, all of the results presented in this chapter, except for those shown in Figure 4.15, are limited to only weekday vehicle use data; weekend data were not subject to detailed study.

It is anticipated that, consistent with other unfamiliar technologies, as prices come down and consumers' familiarity with PEVs improves, typical buyers will no longer be early adopters. With this shift away from early adopters, the demographic profile of PEV owners will move from more educated, higher income households towards the average household [78]. To observe whether these changes in the demographic characteristics of the vehicle owner pool can significantly influence vehicle use variability as a function of vehicle fleet size, additional Monte Carlo simulations were performed. The procedure paralleled the method previously described, except that

prior to sampling individual vehicle use profiles, the pool of candidate profiles was adjusted according to the anticipated demographics of the relevant fleet size — smaller fleets included only profiles with demographics characteristic of early adopters, while larger fleets included profiles closer to the average vehicle owner. These results and those from the previous Monte Carlo analysis were compared to determine whether demographic changes have a significant effect on vehicle use patterns.

4.1.1 NHTS data preparation procedure

Prior to performing the simulations, the NHTS data were converted into vehicle use profiles, following the procedure outlined in Figure 4.1. The individual vehicle and trip data were collected by DOT through phone surveys [32]. Survey respondents were provided diaries to record their travel, but were not required to have used them unless more than one week elapsed between their assigned travel day and when they reported their travel over the phone [32]. This data collection approach can result in misreporting of travel times, reflected in the NHTS data as an overabundance of travel events that occur on five-minute and quarter-hour intervals, likely because those times come readily to mind [38, 39]. Since the times when vehicles are in use are of interest throughout this work, preprocessing of the NHTS data was required to minimize the impact of these reporting errors.

NHTS survey data are available online in several comma separated variable (CSV) files. Each of the files includes various parameters, such as the individual trip or vehicle, along with extensive demographic information. The CSV files all have unique entries (trips, vehicles, etc.) recorded in each row, with the details of each record given in columns. To reduce the size of the data to be manipulated, the data recorded on individual trips and vehicles were reduced to include only those columns with demographic details relevant to the planned simulations. Because the

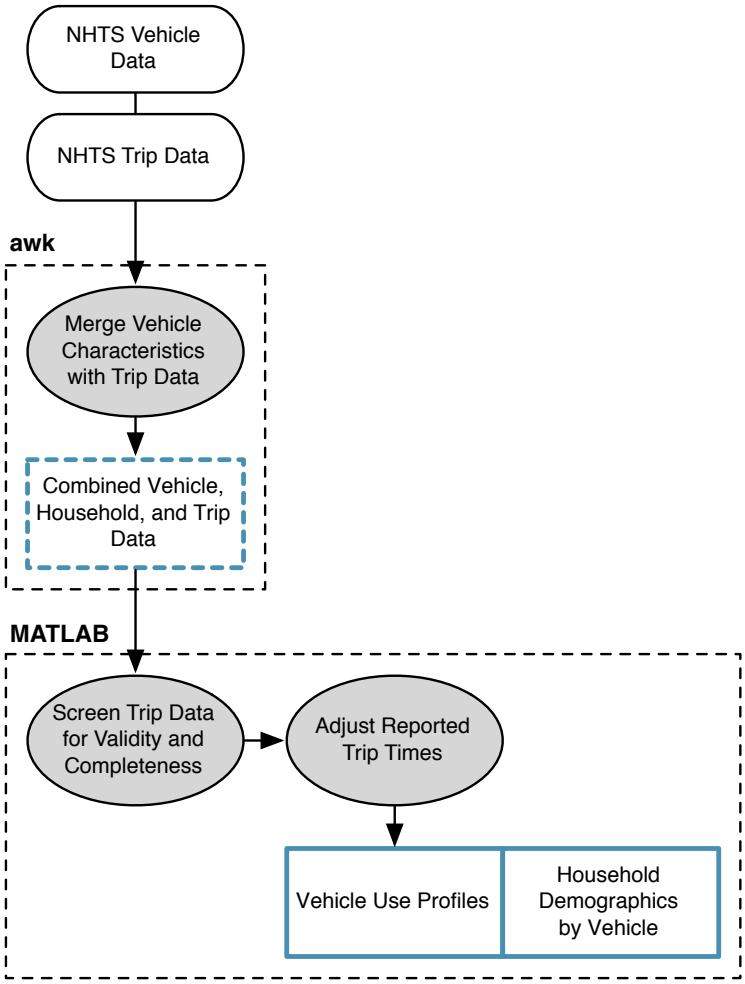

Figure 4.1: Trip and vehicle data reported in the NHTS were converted into vehicle use profiles in two stages. The data were initially reduced and combined using *awk*, yielding a single array containing only those data of interest, and eliminating trips that were taken using a mode other than a private car. Subsequent screening steps removed trips with error values or incomplete data, and any vehicles that did not have an annual mileage estimate. The remaining trips were grouped by vehicle to create vehicle use profiles, which required correction of any overlapping trips reported, as well as shifting of all trips to reduce the effect of rounding in reported travel times, as shown in Figures 4.2 and 4.3.

data included 1.17 million recorded trips (each row in the data constitutes one trip), the interpreted scripting language *awk* was used to select the desired columns. Given that the vehicle data file was also extensive, including more than 300,000 vehicles, *awk* was used to remove unnecessary columns from that file as well.

The resulting files detailing each trip taken and each vehicle in all participating households were then combined. To protect the privacy of survey respondents while maintaining the usability of the data collected, each household is identified by a unique 8-number code, and each vehicle is identified by a number in each household (e.g. 1, 2, or 3). These household and vehicle numbers were used to combine vehicle data with travel data. When the two data sets were combined, any trips that did not have a vehicle number coded could not have any vehicle data matched to them, thus these trips were discarded. There were a total of 220,719 trips discarded. Because the survey includes all trips taken by the respondents, including those completed by modes other than personal vehicles, not all of the trips are relevant to an analysis of individual vehicle use patterns. Non-personal vehicle travel accounted for more than 75% of the trips removed. The remaining trips removed can be attributed to efforts undertaken to prevent problematic data from entering the simulation.

As part of the process of combining the trip and vehicle data recorded in the NHTS data into a single matrix detailing each trip, any trips that could create problems for the simulations were removed. All trips that had error codes (values less than zero) recorded for start times, end times, or trip mileage, traversed zero miles, or had equal start and end times, were removed. All of these trips were either physically impossible, in the case of the trips that started after they ended, lacked sufficient data to be fully characterized, or were not valid trips, in the case of those that covered zero miles. The "BESTMILE" annual miles traveled estimate was used

for these studies. Any vehicles that did not have an annual mileage estimate recorded were also eliminated.

With the vehicle-specific parameters defining annual miles driven and the vehicle type[1] combined with the trip data, individual trips were restructured into vehicle use profiles. Each vehicle's use profile defines when, according to the survey respondent, the vehicle is being driven, and the distance traversed at the end of each trip. Because information about individual trips was not collected immediately following each trip, problems with the trip times reported had to be corrected. Figure 4.2 shows how the coincidence of reported trip start and end times yielded large peaks in vehicle use at certain times, where vehicle use is shown as a fraction of the total number of vehicles in the survey. The inset in Figure 4.2 reveals that the largest modes occur on the hour and half-hour, while smaller modes appear on the quarter-hours (0:15 and 0:45). These problems were corrected by jittering the start and end times.

This process of adjusting trip start and end times sometimes moved trips outside the range of 00:00 and 23:59. These trips were then modified depending on whether their start times, end times, or both, were outside the allowable range. Trips with only start times or end times out of range were shifted forward to start at 00:00 or backward to end at 23:59, while retaining the same total duration. Trips that were beyond the end of the day were shifted to the start of the day, such that a trip starting at 00:07 and ending at 00:42 the following day was shifted back a full day. Conversely, trips that occurred before the start of the day were shifted to the end of the day, meaning a trip that started at 22:05 and ended at 22:16 the previous day would be shifted forward a full day. Shifting trips that were out of range yielded adjusted trips that sometimes overlapped existing trips. Generally, these trips were corrected by

[1] Vehicle types coded in the NHTS data were defined as: car, van, SUV, pickup, other truck, RV, motorcycle, golf cart, and other.

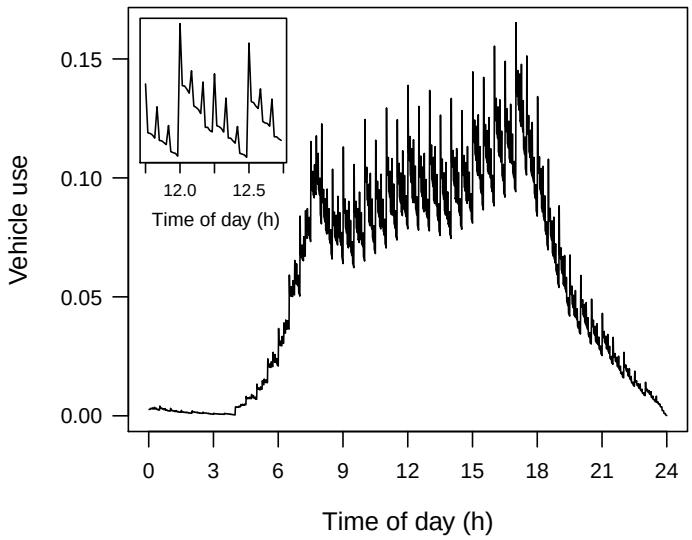

Figure 4.2: The times when trips are reported as taking place in the original NHTS survey data create an unusual number of coincident trip starts and ends that, when combined yield a vehicle use curve with large spikes, are inconsistent with actual vehicle use patterns. The inset region highlights the effect of centering on certain start and end times, with apparent spikes in vehicle use at each five minute interval. This effect is strongest on the hour and half-hour, with smaller peaks on the quarter-hours.

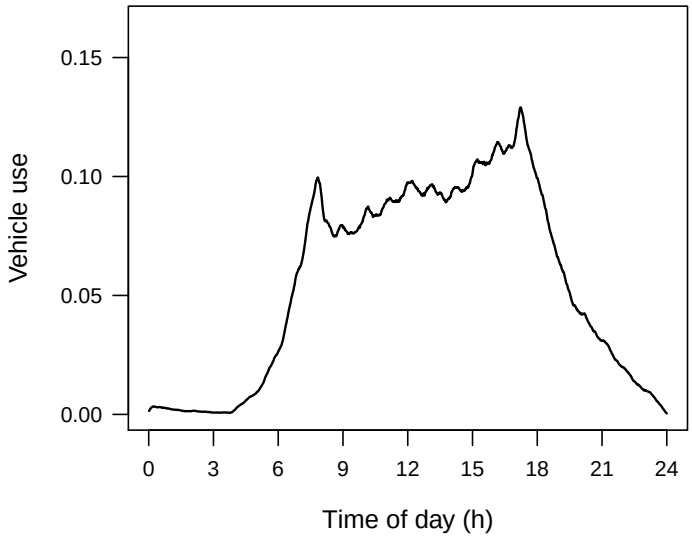

Figure 4.3: After completing the extensive trip correction process, the vehicle use curve is much smoother, indicating that the corrections undertaken were successful in reducing the effect of centering in reported trip start and end times.

shifting the overlapping trip forward or backward, depending on whether the start or end time, respectively, was overlapping. If a trip was longer than an existing trip, it overwrote the existing trip, while any trips shorter than existing trips were discarded. Each of the cases of possible overlap were identified and addressed according to the criteria detailed in Table 4.1.

These corrective measures minimize the modes that appear in the original NHTS trip survey data, as shown in Figure 4.3. Based on the results in Witlox [39], it is assumed that survey respondents are not subject to centering biases when reporting trip distances; for example, they are no more likely to report four miles versus five miles when the actual distance is somewhere in between. Because the duration of each trip was not changed significantly, the miles reported were not adjusted. The reported miles traversed for each trip, along with the start and end times, were recorded in a matrix.

Table 4.1: All of the conditions where new trips might overlap existing start and end times were identified. For each trip that satisfies one of these conditions, the corrective action identified was applied to either the new trip or the existing trip, whichever is appropriate.

Condition	Corrective action
Both start and end times overlap exactly	Discard trip
Start time overlaps with the end time of the previous trip	Adjust trip start time
End time overlaps with start time of the next trip	Adjust trip end time
New trip begins before and ends after existing trip	Overwrite existing trip
Start time overlaps with another start time and	
current trip ends after existing trip	Overwrite existing trip
current trip ends before existing trip	Discard trip
End time overlaps with another end time and	
current trip starts after existing trip	Discard trip
current trip starts before existing trip	Overwrite existing trip
Start time is within another trip and	
current trip ends after existing trip	Adjust trip start time
current trip ends before existing trip	Discard trip
End time is within another trip and	
current trip starts before existing trip	Adjust trip end time
current trip starts after existing trip	Discard trip

4.1.2 Demographic trip screening

In the first few years of widespread PEV availability in the United States, most buyers will likely fit the profile of early adopters [69, 78, 79]. Like the first buyers of other new-to-market technologies, these individuals are willing to accept significant uncertainty with respect to whether PEVs will perform to their expectations, just to be able to be one of the first to have one. These early adopters tend to have higher than average household incomes and educational attainment [78]. Early adopters of PEVs also likely live in urban areas, own other vehicles, and are employed [78, 80, 81]. As public charging infrastructure expands and the number of PEV owners increases, public awareness and familiarity with PEV capabilities, benefits, and drawbacks will also increase [81], moving the demographic characteristics of PEV buyers toward the average vehicle buyer.

Using the extensive demographic information reported alongside the vehicle and trip data included in the NHTS sample, the pool of candidate vehicle use profiles selected during some Monte Carlo simulations was tailored to reflect the characteristics of PEV buyers at each fleet size. Results from the Monte Carlo simulations with either all vehicle use profiles or the smaller set of use profiles selected based on key demographic parameters will reveal whether changing demographics as PEVs become more popular will have a substantial impact on vehicle use patterns. Nine of the demographic parameters available in the NHTS data were selected to be used to reduce the candidate set of vehicle use profiles:

1. total household income,
2. number of people in the household (household size),
3. number of vehicles in the household,
4. educational attainment of household members,

5. population of the city the household is in,
6. whether the household owns or rents their home,
7. type of dwelling the household resides in,
8. total annual mileage of each vehicle, and
9. the vehicle type.

These parameters were selected based on analyses of the characteristics of PEV buyers in the literature [78–81]. Each of these parameters is coded in the NHTS data as either an exact quantity, in the case of household vehicle count, household size, and annual mileage, or as numeric codes for categories defined by DOT.

Since each trip is recorded separately, any vehicles that were driven by multiple members of a household on the surveyed day might have different educational attainment codes assigned to each trip. To create the vehicle use profiles, all the trips for a single vehicle were combined. Instead of attempting to retain all the educational attainment codes for the drivers of each vehicle, all vehicles in each household were assigned new values for educational attainment based on the highest level attained by any member of the household. It is implicitly assumed that one of the primary decision-makers in the household is the person with the highest level of education.

Tables 4.2 and 4.3 show, for each of the demographic parameters selected, the constraints applied as a function of the simulated fleet size. The use profiles included are constrained to those parameters greater than or less than the values identified depending on which direction is appropriate for the given parameter. While the specific constraint values selected for each fleet size are somewhat arbitrary, the values and trends are based on studies of vehicle owners' preferences and demographic characteristics. For example, it is anticipated that initial buyers of PEVs will be more educated than average [82, 84]. Because of the cost and range limitations of some

Table 4.2: Several assumptions about the characteristics of households acquiring PEVs at various points along the technology adoption curve, indicated by the fleet size, are implicit in these tabulated selection criteria. In particular, it is expected that higher income households, particularly those who own (denoted by O) detached single-family homes (denoted by S), are likely to acquire PEVs early, as they have the ability to install a charging station at home. Those that rent (denoted by R) or live in duplexes (denoted by D) or multi-family dwellings (denoted by M) might not have ready access to a charging station or electrical outlet near home or work, though that will change as workplace and public charging infrastructure grows. Households where at least one person holds a graduate or professional degree (denoted by G) or a bachelor's degree (denoted by B) are expected to be earlier adopters of PEVs than those with a high school equivalent education (denoted by H).

Fleet size	Dwelling type			Home ownership		Household income	Household size	Educational attainment			City population
	S	D	M	O	R			G	B	H	
100	✓			✓		100,000	2	✓			1,000,000
500	✓			✓		100,000	2	✓			1,000,000
1,000	✓	✓		✓		100,000	2	✓			1,000,000
5,000	✓	✓		✓		100,000	2	✓	✓		500,000
10,000	✓	✓	✓	✓	✓	80,000	1	✓	✓		500,000
50,000	✓	✓	✓	✓	✓	80,000	1	✓	✓		200,000
100,000	✓	✓	✓	✓	✓	70,000	1	✓	✓	✓	200,000
500,000	✓	✓	✓	✓	✓	60,000	1	✓	✓	✓	0

PEVs, it is anticipated that households with multiple members, higher incomes, and multiple vehicles will be more likely to acquire PEVs initially [80, 81]. The range limitations and operating cost benefits associated with PEVs will likely yield initial adoption among city dwellers who want to use them for commuting purposes [80, 82]. Individuals who own single-family homes will be most able to install an EVSE at home to charge their PEV overnight [81]. Vehicles that are passenger cars and have lower annual mileage are likely more representative of the vehicles that could be supplanted by PEVs, especially since nearly all PEVs widely available are cars, and not vans, SUVs, or trucks.

Table 4.3: Vehicles that are passenger cars (denoted by C) driven a limited number of miles each year in a multi-vehicle household are anticipated to be the best candidates for replacement with a PEV. As PEVs grow in popularity, prices fall, and their capabilities increase, it is anticipated that they will be more likely to replace higher annual mileage vehicles and be found in single-vehicle households. Also, with time, manufacturers are likely to introduce PEVs that are vans or SUVs (denoted by V), and eventually, pickup trucks (denoted by T) as well. These assumptions guided the selection criteria indicated.

Fleet size	Vehicle count	Annual mileage	Vehicle class		
			C	V	T
100	2	10,000	✓		
500	2	10,000	✓		
1,000	2	15,000	✓		
5,000	2	15,000	✓		
10,000	1	20,000	✓	✓	
50,000	1	25,000	✓	✓	
100,000	1	30,000	✓	✓	✓
500,000	1	40,000	✓	✓	✓

4.1.3 Monte Carlo simulation procedure

A predictive vehicle model can quantify anticipated vehicle use and the uncertainty in use patterns. A single estimate of vehicle use can be readily calculated from the NHTS data, as shown in Figure 4.3. While such an approach demands little computational expense, a single estimate could over- or under-predict vehicle use at a given time of day. With only a single estimate, little can be known about the quality of the estimate. To better understand the range of possible outcomes, bounding estimates could be calculated or a series of probable scenarios could be developed, though these estimates of electricity demand could be limited by incorrect assessments of possible outcomes. Instead, in this work, a Monte Carlo approach is applied to predict vehicle use. A modified version of this approach, described in Section 4.2.1, is used to predict vehicle charging loads.

Generally, Monte Carlo simulations generate random estimates of a parameter or set of parameters, and the resulting model objective value is calculated and recorded. After repeating this process many times, there exist a corresponding number of estimates of the variable of interest. These can be examined to determine the characteristics of their distribution. The resulting values might be tightly clustered or widely spread, show multiple modes, or have non-zero higher-order standardized moments. The characteristics of the distribution can indicate the likely realizations of the variable of interest. In simulating vehicle use, the vehicles comprising the NHTS data can be used as the pool of potential outcomes from which a subset can be drawn.

Each Monte Carlo simulation begins by first applying any relevant constraints to reduce the candidate set based on the demographic parameters discussed in Section 4.1.2, thus narrowing the pool of potential vehicle use profiles. Each of the eight simulated fleet sizes n were simulated both with and without these demographic constraints. Vehicle use profiles for the vehicles in the simulated fleet are pulled from the pool of q possible profiles by generating n random numbers between 1 and q. These random numbers are the column indices of the profiles selected for that simulation. The sum of the selected profiles is recorded and a new set of random column indices is generated.

4.1.4 Results

The 500 simulation runs performed for each fleet size were summarized using descriptive statistics to facilitate exploration of the variability present in vehicle use patterns as a function of PEV fleet size and determine whether imposing demographic constraints that are a function of fleet size affects the nature of the results. These results also reveal whether vehicle use variability is different at various times of day and whether demographic narrowing of the candidate vehicle fleet should be continued

in subsequent studies. These profiles included only those whose data were collected on weekdays; weekend data were excluded.

Figure 4.4 summarizes the results from the simulations performed with all binary vehicle use profiles as candidate profiles during the sampling process. There appears no obvious relationship between characteristics related to the time of day, such as periods with greater rates of change in vehicle use (e.g. morning or evening hours), and increased variability, given by the standard deviation. In these results, fleet size is the dominant parameter affecting vehicle use variability. With smaller fleets, the standard deviation varies slightly throughout the day. Larger fleets, beyond 50,000 vehicles, show negligible variability in vehicle use, regardless of the time of day. Further, the minimal variability in vehicle use with larger fleet sizes suggests that V2G aggregators, system operators, and electric utilities might not face significant operational problems created by such variability. The results in Sections 4.2.3 and 4.3.2 examine these issues more closely.

Changing the pool of candidate vehicle profiles using demographic constraints affects vehicle use patterns throughout the day. Figure 4.5 shows that vehicle use with demographic screening is increased during the morning hours, regardless of fleet size, but is decreased slightly throughout midday. For larger, less volatile vehicle fleets, use is also increased during the evening hours. If these demographic constraints are consistent with future PEV buyers, those drivers tend to use their vehicles on weekdays in a work commute-type pattern: more in the morning and evening, and less midday. Further, as shown in Figure 4.6, this trend persists with increasing fleet size for these simulations; even when the candidate vehicle use profile pool is larger and is demographically closer to the general population, the prevalence of morning and evening commute-dominated vehicle use only increases. This effect could be

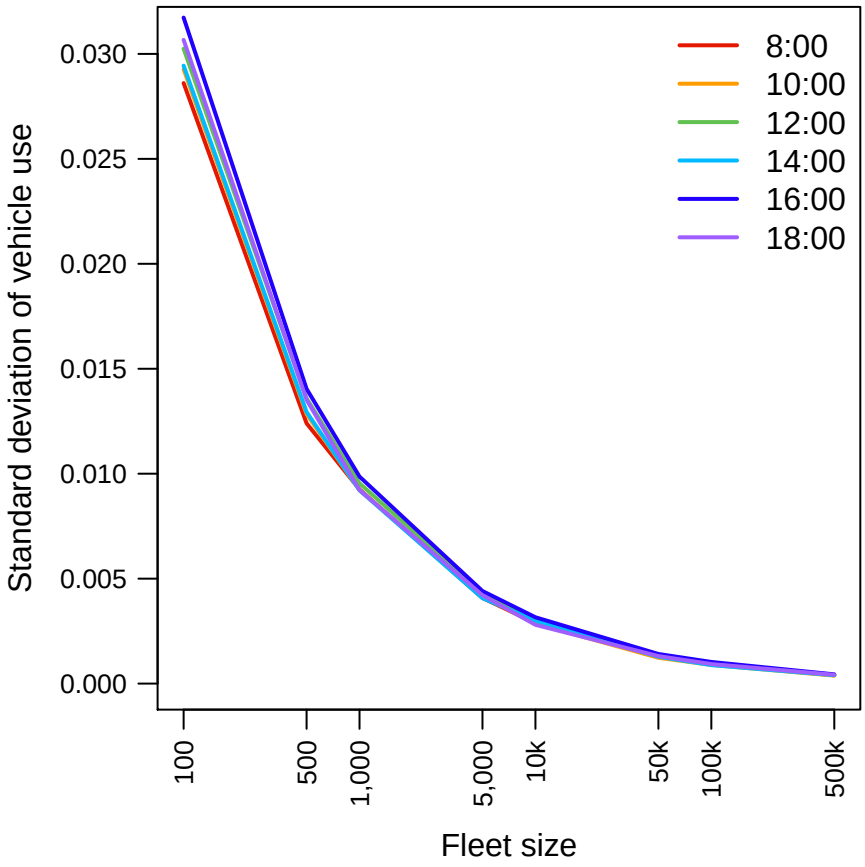

Figure 4.4: The six times of day plotted were selected such that they would be uniformly spaced across the majority of the hours of the day when vehicle use is high, based on Figure 4.3. The standard deviation values shown for the selected times summarize the range of outcomes from the 500 simulation runs performed for each fleet size. It appears that the primary driver of variability in vehicle use is fleet size, with larger fleets showing significantly reduced standard deviation and spread in reported standard deviation as a function of time of day.

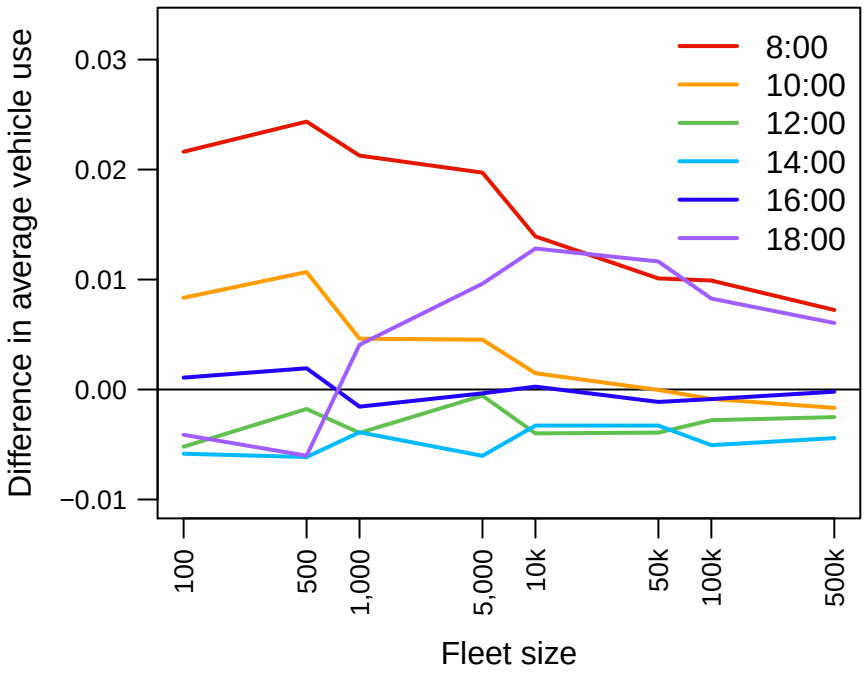

Figure 4.5: The results shown reflect the difference in average vehicle use reported from the simulations performed for both cases: including all vehicles and restricting the pool of candidate vehicles based on demographics. The differences indicated are an order of magnitude smaller than reported vehicle use values, but the average vehicle use is not consistently higher in one simulation case. Around midday, when vehicle use is highest, results from the two cases are most similar.

caused by demographic constraints that, even when they are most relaxed, with the largest fleet size simulated, still remove enough vehicles from the pool of candidate profiles to have a noticeable effect. Figure 4.6 also shows that volatility in vehicle use between periods decreases with increasing fleet size, following the same trend as volatility in vehicle use within a given period, shown in Figure 4.4.

These results suggest that variability in vehicle use, and thus, the potential for variability in vehicle charging, might not be significant with larger PEV fleets. Further, changes in the demographic characteristics of PEV owners are likely to affect PEV use patterns as their numbers grow, thus subsequent analyses include both

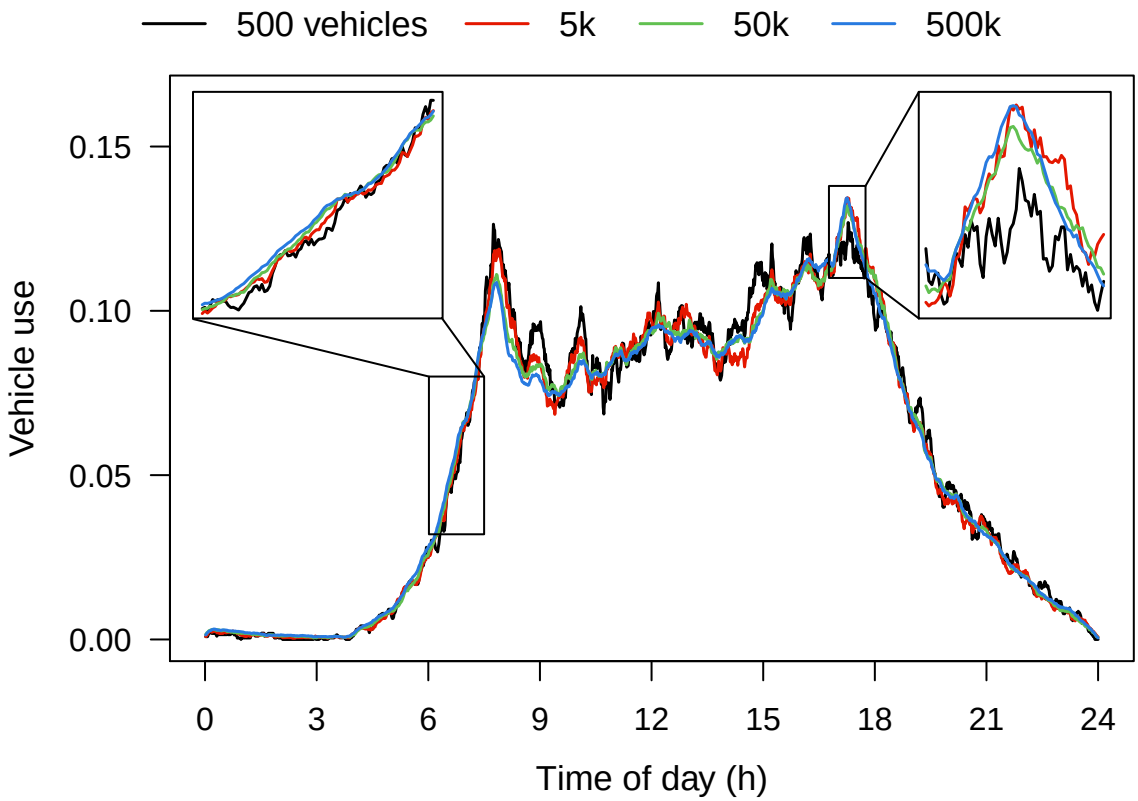

Figure 4.6: Examining average vehicle use from the simulations with demographic constraints for four fleet sizes, it appears that as the fleet size increases, vehicle use increases during morning commute hours (left inset) and evening peak driving hours (right inset). Also, not unsurprisingly, the inter-temporal volatility of vehicle use decreases with increasing fleet size.

general and demographic constrained simulation results. These upcoming analyses incorporate more detailed vehicle driving data to observe the temporal characteristics of vehicle charging loads, power capacity, and fleet state-of-charge.

4.2 Simulate unscheduled PEV charging loads

The simulation results from Section 4.1.4 highlight the effect of fleet size on the stochasticity of vehicle use measures, but for a system operator facing unknown and unfamiliar demand from PEVs, uncertainty in the total power demand from PEVs connected to the grid might be of greater interest. For this analysis, it is assumed that charging behavior is not scheduled or controlled by the vehicle aggregator or the system operator. The results detailed in Section 4.2.3 improve upon previous work that has attempted to calculate the power available from PEVs on the grid because the methodology applied here avoids the use of deterministic methods to create measures of variability [25, 45, 47].

4.2.1 Simulation procedure for estimating variability

Following the procedure detailed in Section 4.1.1, the use profile of each vehicle in the NHTS survey included all trips, with the start time, end time, and intermediate periods coded using unique identifiers. The miles traversed for each trip were recorded at the end of the trip. These profiles were modified to examine the variability in the energy stored in grid-connected PEV batteries and the power demanded for battery charging. The three parameters of interest for this analysis, all functions of the time of day, were: the state-of-charge of each vehicle while it is connected to an EVSE, the power draw for charging demanded by each vehicle, and the potential available power from each vehicle whenever it is parked, regardless of whether it is connected to an EVSE. The procedure used to develop these parameters is outlined in Figure 4.7.

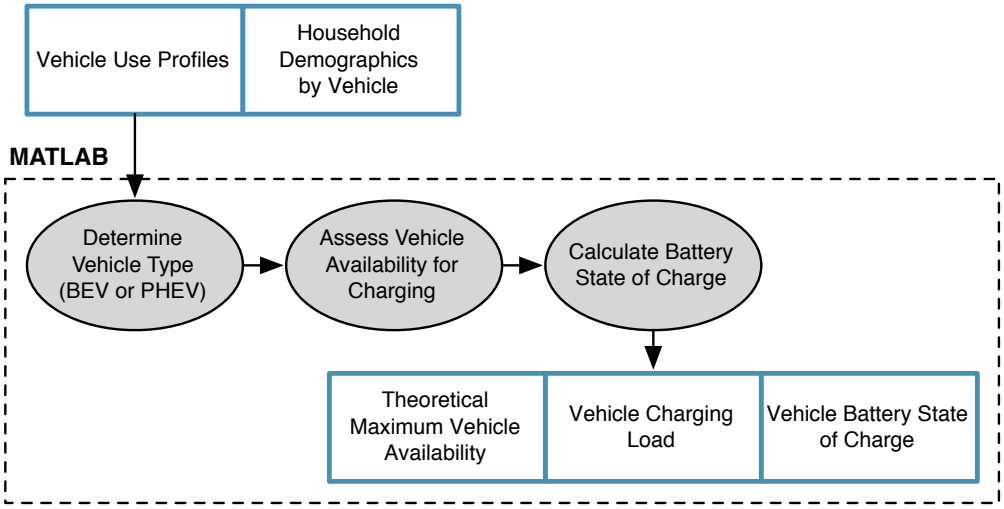

Figure 4.7: The vehicle use profiles and demographic characteristics prepared from the NHTS data, as described in Section 4.1.1, were used to develop three parameters: vehicle charging power demand, vehicle battery state-of-charge, and maximum theoretical vehicle power capacity if all idle vehicles were connected to the electric grid. Based on the number of miles traveled on the surveyed day, each vehicle was assigned to be a BEV, PHEV, or neither. All non-PEVs were discarded, and the remaining vehicles' travel data were converted into battery state-of-charge, availability, and power demand based on assumed battery and EVSE sizes, battery depletion rate per mile driven, and estimated charging probabilities.

Each of these parameters was calculated by separately restructuring each vehicle's travel data. The time between each trip and the number of minutes required to recharge after each trip were calculated for each vehicle. The time required to recharge the vehicle's battery was calculated based on a nominal EVSE power rating and PEV energy depletion rate per mile driven of 5.5 kW and 0.34 kWh/mile, respectively. Because it is unlikely that PEV owners will want or be able to charge their vehicles at the end of every trip, whether each trip would end with being plugged into an EVSE was determined using the probabilities in Figure 4.8. Implicit in these probabilities are several assumptions: vehicles used during commuting times are more likely to be plugged into an EVSE at work than those used for mid-day errands, the availability of public charging stations will in the near term remain somewhat limited, free public charging will not be commonplace, and midday driving is dominated by short trips that do not lend themselves to recharging following each trip. Uniformly-distributed random numbers between zero and one for each trip end were combined with the EVSE use probability distribution to create a vector of charging status values equal to one when a vehicle would charge at the end of a trip, and zero otherwise. The charging status vector was then modified with the assumptions that a vehicle stopped less than 15 minutes would never be plugged into an EVSE, and the last trip of the day would always be followed by charging. It is possible that in the future, mandatory residential demand response programs, or critical peak pricing or real-time pricing programs for residential customers might encourage or require PEV owners to make specific charging choices. Given that these programs are not yet widespread, it is assumed here that charging behavior is not scheduled or controlled by the vehicle owner or a third-party, such as the owner's electric utility or the system operator in the region.

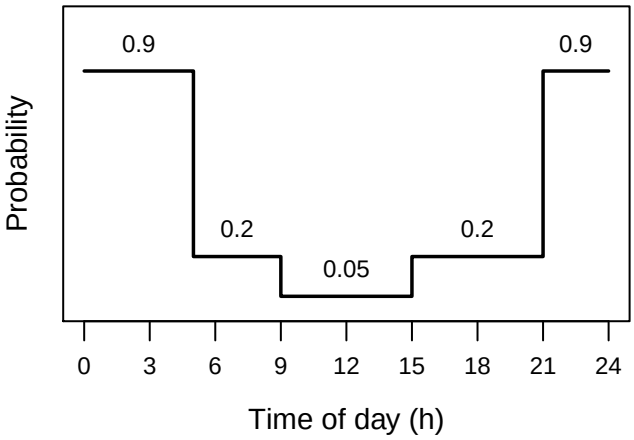

Figure 4.8: In converting vehicle driving profiles derived from NHTS data into charging profiles, it is assumed that the probability that a vehicle will plug into an EVSE at the end of a trip is a function of the time of day.

To be able to generate the three parameters of interest, each vehicle's battery capacity had to be determined. The vehicle type, and hence, the battery capacity for each use profile, was determined based on the total miles covered in each reported profile. Vehicles driven less than 80 miles on the sample day were assumed to be a mix of 40% BEVs and 60% PHEVs, vehicles driven between 80 and 200 miles were assumed to all be PHEVs, and vehicles driven more than 200 miles were assumed to be neither PHEVs nor BEVs. Average battery capacities of 8.5 kWh and 27.2 kWh were selected for PHEVs and BEVs, respectively, yielding driving ranges of 25 and 80 miles, based on the assumed nominal energy use rate of 0.34 kWh/mile [35, 76]. For simplicity, it is assumed that these and the other preset characteristics of the vehicle fleet — PEV energy use per mile and average EVSE power rating — are not a function of the fleet size. These assumptions are implicit in the procedure used to modify the vehicle use profiles prior to performing the Monte Carlo simulations, and are made primarily to manage the complexity of the simulations.

For all vehicles with electric drivetrains (already identified as either BEVs or PHEVs), charging events were transformed into a battery state-of-charge profile. The duration of recharging following all trips was recalculated using the charging status vector. If the vehicle was not plugged in, the state-of-charge would remain the same while the vehicle was stopped. If the vehicle was plugged in, the battery state-of-charge during charging was calculated using the trip end state-of-charge to determine the minutes of recharging required. In cases where the vehicle could not fully recharge before the start of another trip, the final state-of-charge was calculated based on available time to charge. Following the final trip of the day, recharging was permitted to extend beyond the end of the day, if needed, to ensure complete overnight replenishment of the battery. In these cases, the remainder of the state-of-charge profile was recorded at the beginning of the day.

4.2.2 Verification of simulation approach

The results from these Monte Carlo simulations were validated against empirical charging data collected as part of the Pecan Street Smart Grid Demonstration Project ("Pecan Street"). Pecan Street currently includes over 375 homes instrumented to measure their energy, water, and natural gas usage [85]. The Pecan Street PEV charging data are comprised of one-minute interval average power measurements of the circuit dedicated to the in-home EVSE. As of April 2013, there were 33 PEVs in the Pecan Street study area with charging data acquired from their circuit monitoring equipment, of which 25 are PHEVs. EVSE circuit monitoring began on April 23, 2012, but data acquired during the first 34 days, when fewer than 10 homes' EVSE circuits were being measured, were excluded from this study. The last recorded data included were reported on April 4, 2013.

Using the demographic parameters in the NHTS data noted in Section 4.1.2, a subset of the vehicle charging data, with characteristics corresponding to those households with monitored EVSE sub-circuits in Pecan Street, were selected for use in the Monte Carlo simulations. Demographic characteristics of the homes in Pecan Street were collected via online surveys conducted in 2011 and again in 2012. Of the 33 homes with available EVSE data, 27 responded to one or both of the surveys. The responses between the two surveys were typically quite consistent. For the handful of households that responded to both surveys and indicated changes to the parameters of interest, the answers recorded in the later survey superseded earlier responses. Of the nine parameters identified in Section 4.1.2, four were observed from the survey responses: 1) household income, 2) educational attainment, 3) number of vehicles owned, and 4) number of people in the household. In the surveys, incomes were

given as a range; in the interest of simplicity, the midpoint of the range was used to calculate the average income of the 27 households.

The responding households had an approximate average of 2.1 vehicles and 2.3 individuals living in the home. At least one member of most of the households held one or more postgraduate or professional degrees, and their average gross income was approximately $150,000 per annum. The remaining five parameters used to select the candidate charging profiles were assigned values based on the characteristics of the Pecan Street study area: most own their homes, most of their homes are single-family units, they live in an urban area[2] with a population greater than 200,000, they typically drive less than 20,000 miles per annum, and the car being supplemented or replaced with a PEV is a car, van, or SUV.

For the purposes of this verification, the charging probabilities shown in Figure 4.8 were modified to differentiate the probabilities for PHEVs and BEVs. These separated charging probabilities are shown in Figure 4.9. The Pecan Street charging data indicates that PHEV owners often charge at least once in the middle of the day, possibly to ensure that the vehicle never runs short of available energy, since the range of most PHEVs is close to or less than the average daily miles traveled for drivers in the US [73]. Further, relying primarily on the battery for motive power is cheaper, and likely one of the primary reasons for purchasing a PHEV over an HEV. Regardless of the extent of these midday charging events, almost all PHEV drivers will need to charge every night if they want to minimize their use of the ICE, and hence, their operating costs. Conversely, the range limitation of BEVs mean that owners are likely confident they can comfortably complete their daily driving within that range. The typical driver covers only 29 miles in a day [73], thus some BEV

[2] It was assumed that urban versus rural was a more important differentiating factor for PEV ownership than the population of the urban area.

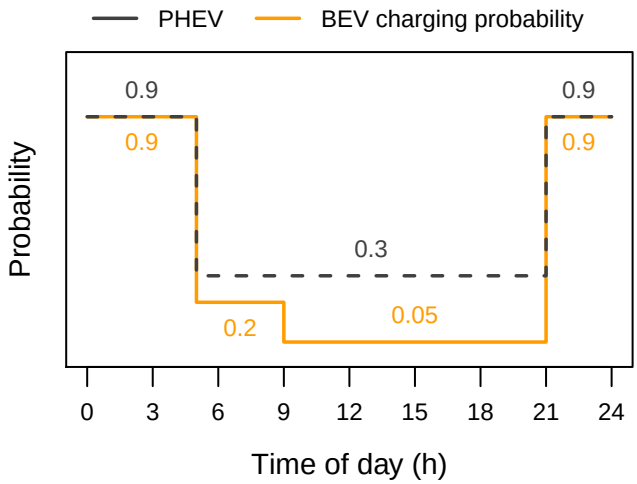

Figure 4.9: To simulate power and energy demand from PEV charging, data defining the travel of vehicles in the NHTS sample were modified using the probabilities shown, which define, as a function of the time of day, the likelihood that a vehicle would plug into an EVSE at the end of a trip. Separate probability measures were defined for PHEVs and BEVs. As indicated, PHEVs are more likely to charge throughout the day, while BEVs are less likely to plug in than PHEVs, especially during the afternoon and evening hours. It is assumed that all vehicles charge following their final trip of the day.

drivers might even charge every other day and still be comfortably within the range of the vehicle at the end of the second day. PHEV owners in Pecan Street charge an average of 1.6 times per day, while BEV owners charge half as often, an average of less than once per day. Because the modeling approach used here spans a single day, the simulations cannot directly account for the possibility that BEV drivers might charge their vehicles overnight every other day, but they can easily account for the correlated effect that those drivers will be less likely to pursue daytime charging. Other assumptions applicable to the probabilities in Figure 4.8 were carried over to the development of the revised charging probabilities in Figure 4.9.

The Monte Carlo simulations used for this validation effort include the application of the BEV and PHEV charging probabilities and the demographic char-

acteristics from the Pecan Street study area. Figure 4.10 reveals that there is good agreement between normalized average charging load from the Monte Carlo simulations and empirical Pecan Street data. This agreement suggests that the assumptions made to convert the NHTS data from vehicle use data to vehicle charging load profiles are consistent with real-world charging behavior. In particular, Figure 4.10 indicates that assumptions regarding the probability that vehicles would be connected to an EVSE at the end of a trip, summarized in Figure 4.9, were reasonable. It should be noted that those quantitative probability measures were predicated on speculative, qualitative assumptions about EVSE use. As public EVSE availability increases, it is possible that their use will change the frequency and duration of home vehicle charging events. Conversely, PEV owners might realize that the driving range of the vehicles they have purchased are well-suited to their needs and thus they find little need for midday charging, in which case, the charging load will continue to follow the temporal profile shown in Figure 4.10.

Figure 4.10 does show some differences between the simulated results and empirical data. These differences are shown explicitly on the bottom of Figure 4.10, where "deviation" denotes empirical data subtracted from simulated charging load. The measured deviation indicates that the model predicts charging behavior to within 7% of the empirical data except at a few specific times. In particular, average charging loads increase slightly later in the afternoon in the Pecan Street data; beginning around 15:00, there is a delay of about one hour that declines to 30 minutes by the time the peak is reached at 19:00. Also, from the peak, charging loads in the Pecan Street study area decrease more quickly initially than in the Monte Carlo simulation results, and also show near-constant charging demand for roughly an hour around 21:00. This period of constant charging demand is probably a period when, because of their much larger batteries, most of the BEVs in the system are still charging,

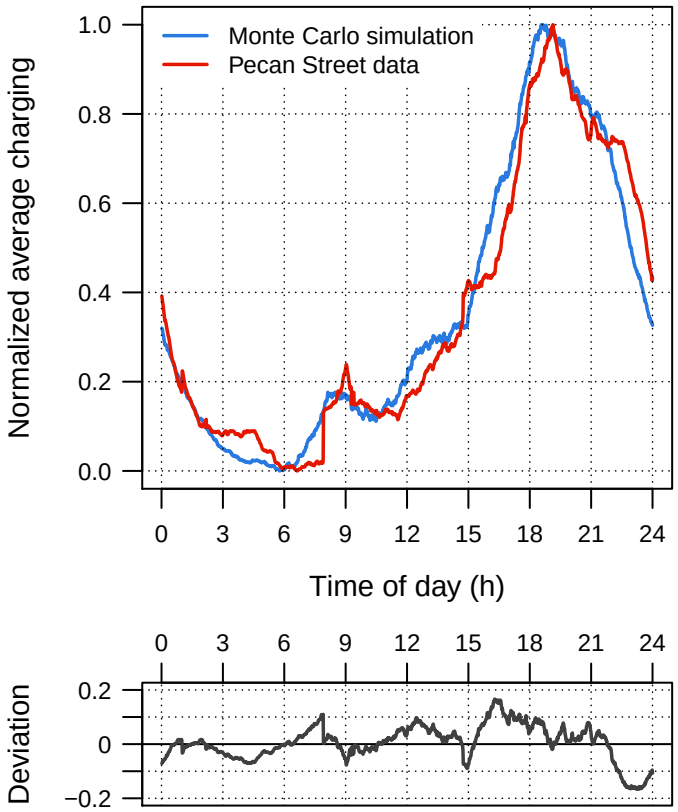

Figure 4.10: Average PEV charging load was determined using Monte Carlo simulation and compared with charging data from Pecan Street. The trends in demand are similar, including a noticeable increase in demand around 08:00, but the peak is slightly later in the Pecan Street sample, and the shape of the decline in demand in the late evening hours is somewhat different as well.

while many of the PHEVs are completing their charging for the evening because they need only charge for a few hours to be fully replenished. This effect is visible, but less pronounced, in the simulation data as a change in the slope of the average charging curve around 21:00. The effect of the battery size differential is likely magnified in the Pecan Street data by the prevalence of EVSEs rated at 3.3 kW, lower than the simulated average 5.5 kW EVSE, and the observation that many BEV owners in Pecan Street have discovered that they do not need to charge their vehicles every night, and thus charge for approximately twice the time every other night.

4.2.3 Results

The simulations detailed in Section 4.2.1 enable calculation of vehicle charging loads, battery state-of-charge, and potential power capacity for a range of fleet sizes. These data were examined to determine whether significant variability arises, especially in the power demanded to serve vehicle charging needs, as the PEV fleet grows. Whether charge scheduling or significant capacity expansion might be needed to serve charging loads from larger vehicle fleets was also of interest.

Figure 4.11 compares normalized average charging demand with vehicle use and electric load to show the comparative timing of events with those parameters. The average charging demand is derived from the 10,000 vehicle simulation without demographic constraints. Charging demand peaks just after vehicle use. It is likely that the peak does not occur later because as vehicles complete their last trips of the day, other vehicles are nearly finished charging or taking their first trip of the day. The timing of charging demand does not align well with electric load which, as shown, is the average of 15 minute ERCOT total generation for all of the weekdays in June 2010. Consistent with the results in Chapter 3, when vehicle charging is not scheduled, it is highest during the same hours when summer electric load is already high. As

a result, charging loads would exacerbate already high peak demand throughout the summer months, suggesting that charge scheduling might be valuable once the PEV fleet in a given region or interconnection is large.

In Figure 4.12 and subsequent results in this section, "variability" or "variation" is calculated as the difference between the maximum and minimum values reported from the simulations for each fleet size. Variability in vehicle charging loads, shown in Figure 4.12, appears to be correlated to the magnitude of charging demand (shown in Figure 4.11). While charging load variability increases with increasing fleet size, Figure 4.13 shows that the increases in variability are outpaced by growth in total charging load. As a result, the cumulative effect of variability is diminished significantly at larger PEV fleet sizes. Further, the magnitude of peak average charging demand, even with the largest fleet size simulated, is approximately 630 MW. Peak power demand in the summer months in ERCOT is slightly less than 70 GW [86], two orders of magnitude larger than peak average charging demand. This difference suggests that charge scheduling might be significantly less important than Figure 4.11 implies. On the other hand, it is possible that as battery costs decline, the capacity of vehicle batteries will increase and EVSE power ratings will keep pace to prevent increases in charging times. Such changes in nominal EVSE power demand would have a multiplicative effect on these results: a doubling of EVSE power demand would double variability as well. This change could dramatically affect the charging demand variability indicated in Figure 4.12.

4.3 Assess changes in ancillary service provision and requirements resulting from vehicle use variability

This work is predicated on the hypothesis that PEVs might be able to provide usable ancillary services in electricity markets, but the physical infrastructure

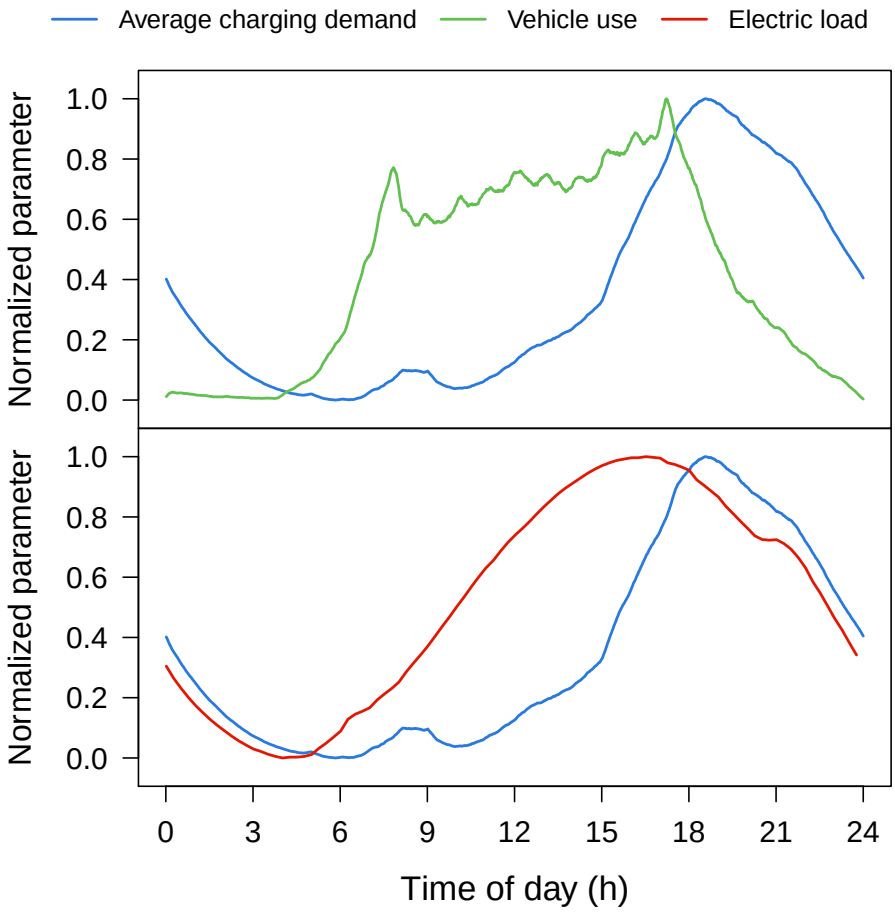

Figure 4.11: Comparing normalized average charging demand with vehicle use reveals that charging demand peaks just after vehicle use and then, as expected, declines steadily throughout the evening hours as vehicles recharge overnight. A comparison of charging demand with ERCOT summer electric load shows that without charge scheduling, peak charging load will occur during peak electricity demand hours, typically reaching its maximum less than two hours after peak electric load.

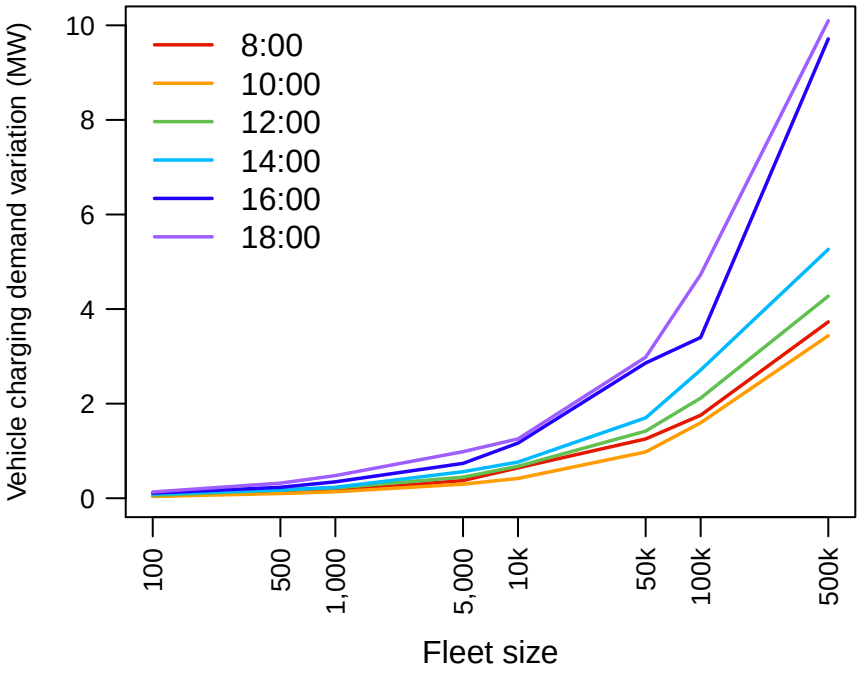

Figure 4.12: Charging demand variation represents the difference between the maximum and minimum charging loads simulated at each of the six times of day shown. Comparing the variability with the magnitude of charging demand shown in Figure 4.11 indicates that variability is greatest during periods when vehicle charging load is also high.

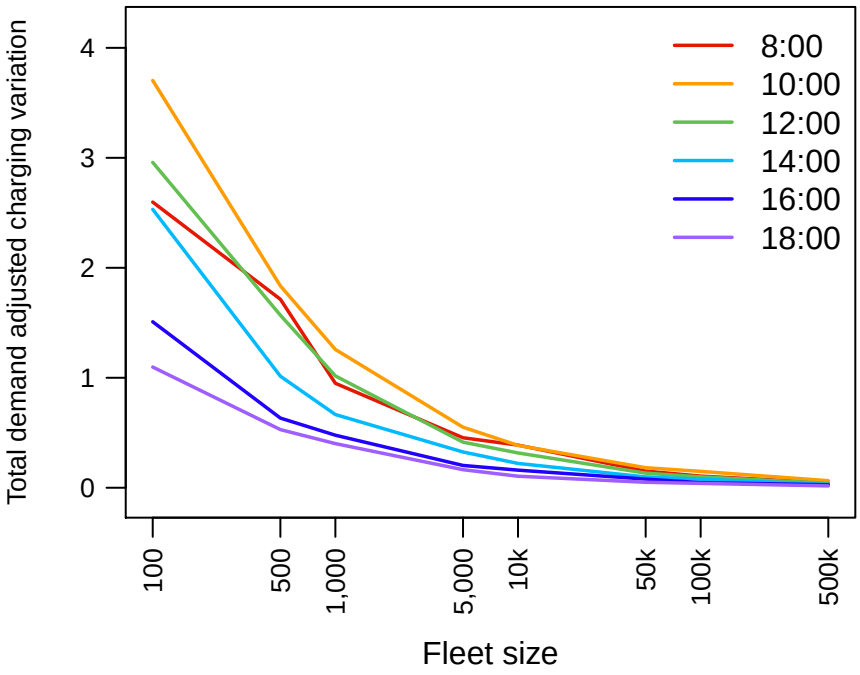

Figure 4.13: Dividing charging demand variability by average charging demand shows that as fleet size increases, while variability increases, it is outpaced by the growth in average charging demand, yielding lower relative variability with larger PEV fleets.

and communications backhaul required to support V2G will not be in place before utilities and system operators are confident of consumer acceptance of PEVs. Prior to the rollout of the necessary infrastructure, stochastic, unscheduled charging loads could create a need to procure additional ancillary services. To determine whether changes might need to be made to frequency regulation procurements to accommodate variability in PEV charging loads, the Monte Carlo simulation approach detailed in Section 4.1.3 was modified following the procurement estimation approach suggested by GE Energy in their report on the impact of wind generation on ancillary services in ERCOT [87].

4.3.1 Modifications to the Monte Carlo simulation approach

ERCOT currently calculates regulation procurements for every upcoming month by taking, for each hour in the day, the 98.8th percentile value of the deployments in the previous month and the same month the previous year [87]. The greater of these two values is used to set regulation procurements for each hour. To apply this methodology to as yet unobserved changes in load or generation in ERCOT, the change in charging loads between periods is used as a proxy to estimate the variability that would affect regulation needs. For each hour, the charging demand values from the 500 simulations were sorted in ascending order and the 98.8th percentile entry was selected. This entry in the sorted list was selected using the nearest rank method, shown in Equation 4.1, where the position r of the 98.8th percentile entry is derived by rounding the value calculated from the desired percentile c and the total number of entries n in the ranked list.

$$r = \left\| \frac{c\,n}{100} + 0.5 \right\| \tag{4.1}$$

The determination of additional regulation procurements needed in each hour to support variability in charging loads without accounting for other unpredictable system elements in ERCOT, existing load (apart from charging load) and wind generation, implies that vehicle use patterns are independent of these elements. Based on the results detailed in Chapter 3, the relationship between net load, measured as load minus wind generation, and vehicle use changes seasonally. Since there are some times during the year when net load and charging loads are correlated, simulated charging loads were combined with a full year of ERCOT net load.

The charging demand data developed in the previous section included only weekday data. Before combining charging load and net load from ERCOT, the simulation procedure was repeated to generate weekend vehicle charging profiles. ERCOT net load data on one-minute intervals were used to match the vehicle charging data. Since the results in Chapter 3 indicate that there are no significant seasonal patterns in vehicle use, the simulated charging loads were applied to all times throughout the year. Weekday and weekend charging profiles were randomly assigned to each day, as appropriate, to create a full year of charging loads. This load profile was combined with net load to develop net load with vehicle charging. This new net load parameter was then examined using the approach detailed previously, by sorting the net load for each hour in each month and selecting the 98.8th percentile value. This procedure yielded estimated regulation procurement needs for each hour in each month for each fleet size tested. With these results, which account for seasonal changes in intra-day correlations between load, wind generation, and vehicle charging, the effect of vehicle charging on regulation procurements with respect to increasing PEV fleet size can be studied. These results can also be compared with the results examining vehicle charging loads independent of wind generation and electric load to determine whether the correlations between net load and charging loads are significant.

4.3.2 Results

Initially, vehicle charging patterns were simulated for a range of fleet sizes without regard to the interaction between charging load and total net load. The regulation procurement quantities for each hour are shown in Figure 4.14. These results are applicable for the entire year, as it was observed in Chapter 3 that vehicle use trends, and hence, vehicle charging loads, are not dependent on the time of year. Also, to be consistent with the results detailed in Chapter 3, the results shown in Figure 4.14 are only applicable to weekdays.

Figure 4.14 indicates markedly different trends in the changes in procurements for regulation up and down as a function of the time of day. Additional regulation up is needed primarily during the daytime, increasing sharply during the late afternoon and early evening, between 16:00 and 20:00. It was assumed in the formation of the vehicle charging profiles that after the final trip of the day, all vehicles will begin charging, and it is during this period that most vehicles will complete their final trips of the day. As a result, there will be significant changes, as well as significant minute-to-minute variability, in charging loads during these hours. The sharply increased need for regulation up during the late afternoon and early evening is likely a result of this variability. The apparent timing of vehicle charging is consistent with the timing indicated by "average charging demand" in blue in Figure 4.11. In contrast, additional regulation down is needed primarily in the overnight hours, between the 19:00 and 03:00. This need for regulation down overnight is also likely motivated by overnight charging loads, where the time when a vehicle completes charging depends on the time when it completes its last trip of the day and the state-of-charge of the battery at that time. Variability present in these two parameters will tend to create variability in the timing of charge completion and hence, variability in the timing

of necessary decreases in the output of generators serving those loads. The times when additional regulation up and regulation down need to be procured are similar for all the fleet sizes simulated. Increases in fleet size only increase the magnitude of additional regulation procurements.

Vehicle charging load profiles were combined with net load to determine whether the interaction between them might be complementary and to explore whether seasonal changes in net load would change these interaction effects. Figure 4.15 shows, for each month of the year, the change in regulation up procurements needed to accommodate the fleet size indicated. The quantity of additional regulation needed is highest during the daytime hours, especially in the late afternoon and early evening, consistent with the charging load-only simulation results in Figure 4.14. Interestingly, with larger vehicle fleets, the quantity of regulation up needed decreases slightly overnight. This effect is likely attributable to the large and persistent charging load during those hours acting to swamp smaller minute-to-minute changes in load. This effect becomes more pronounced with larger fleet sizes. Comparing the results across months does not reveal any appreciable seasonal effects. Changes to regulation down are also consistent with the results in Figure 4.14, with reductions in the need for regulation during the midday hours, and no evidence of any seasonal trends.

To explore the influence that these changes in the quantity of frequency regulation needed will have on day-ahead market clearing prices, the regulation quantities shown in Figure 4.14 can be compared to actual regulation procurements. In 2011, the actual average regulation up clearing quantity was 529 MW, with a standard deviation of 141 MW, and the actual average and standard deviation of regulation down clearing quantities were 459 MW and 98 MW, respectively. The quantities shown in Figure 4.14 are less than 1% of average regulation procurements. While the average

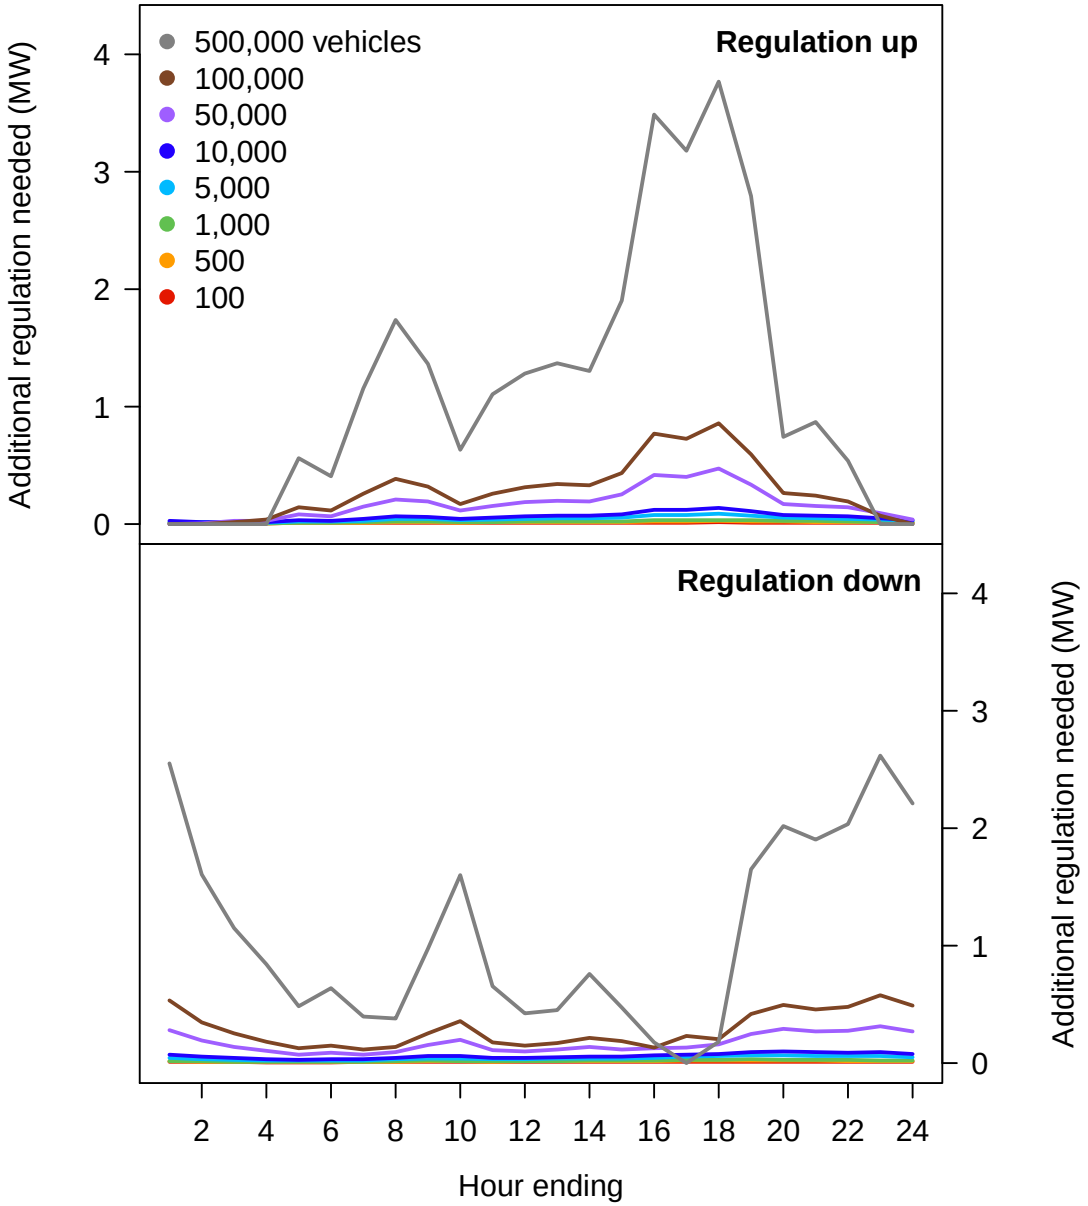

Figure 4.14: Simulations of weekday vehicle charging loads indicates that additional regulation would be needed to support the variability of those loads. On weekdays, additional regulation up is most needed during the daytime hours, roughly mirroring times of high vehicle use, while conversely, regulation down is most needed overnight. Changes in regulation requirements scale with fleet size, but as noted in Section 4.2.3, the relationship is non-linear.

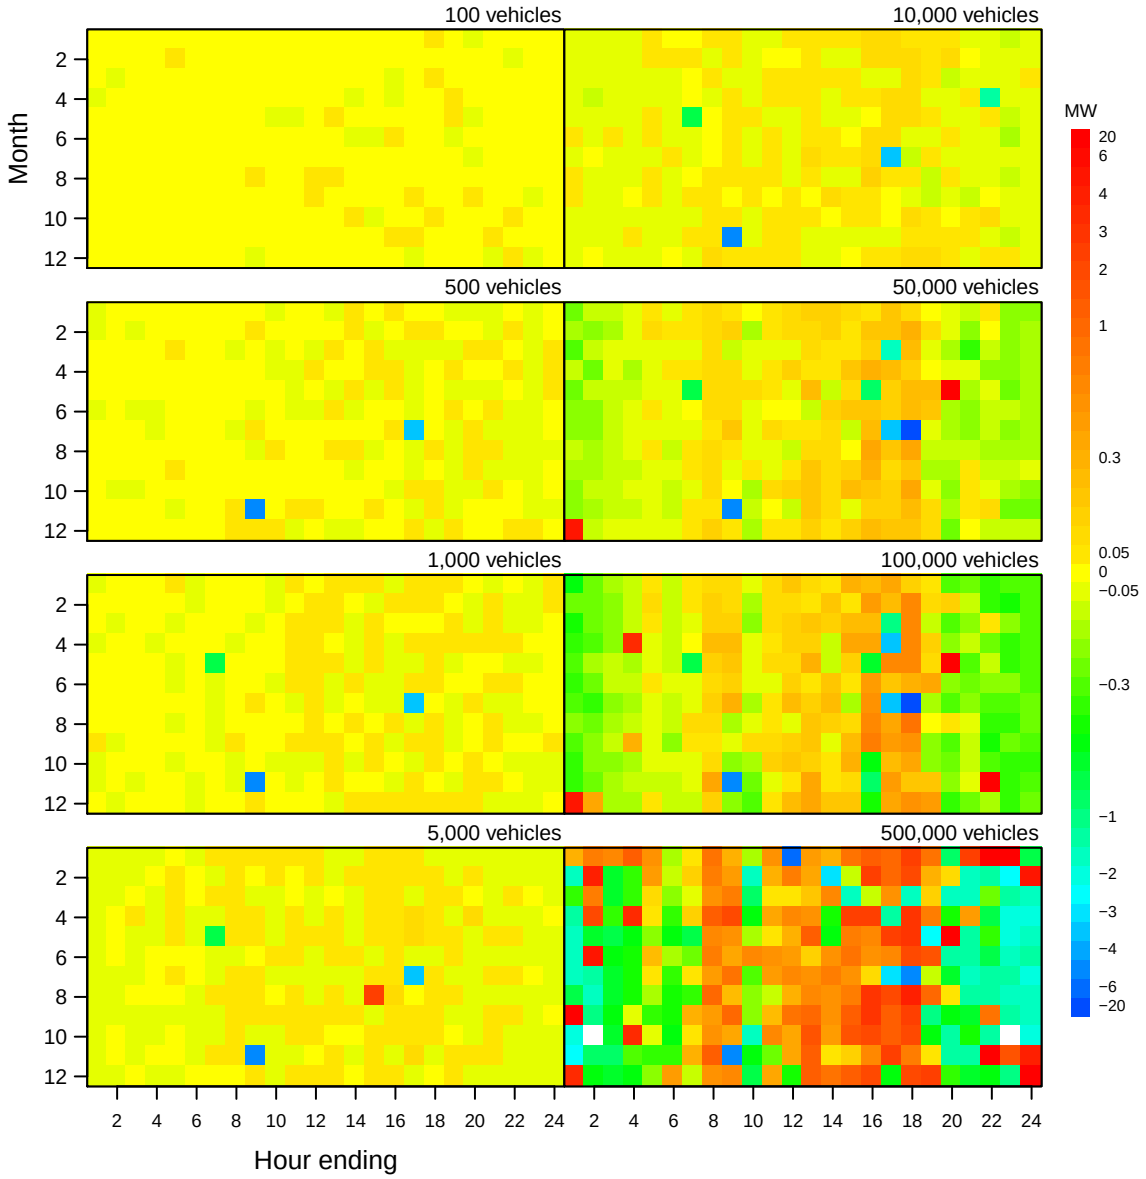

Figure 4.15: Combined net load and vehicle charging results are shown for each fleet size, where baseline requirements are denoted by yellow cells, while increasing requirements move toward red, and decreasing requirements move toward blue. Results from these simulations indicate that the trends shown in Figure 4.14 remain: additional regulation up is most needed during the daytime hours. Net load stabilization from vehicle charging loads suggests that regulation requirements might be reduced during the overnight hours between 20:00 and 3:00, indicated by the green and cells that become especially apparent beginning with the simulated 50,000 vehicle fleet. White cells indicate values that exceed the range indicated in the legend.

regulation procurement quantities are two orders of magnitude larger than the additional regulation needed to support the variability presented by vehicle charging loads, changes in day-ahead prices might not scale linearly with procurement quantities. By combining the changes in procurement quantities indicated by the combined net load and vehicle charging load simulations with day-ahead offer quantity and price data from 2011, the change in prices arising from changes in procurement quantities can be readily assessed. The largest fleet size tested represents the worst-case changes in procurement quantities. With that PEV fleet size, the average increases in day-ahead prices are 0.38% and 3.3×10^{-4}% for regulation up and regulation, down respectively. On an hourly average basis, the highest observed increase in prices was 6.8% and 0.39% for regulation up and regulation down. These results indicate that while prices will increase during some hours, the net impact on prices will be small.

It should be noted that for the fleet sizes simulated, the effect of charging loads on prices in the energy market is less certain. Texas represents approximately 6% of the light-duty vehicle fleet in the United States [88]. If the national PEV fleet reaches one million vehicles by the end of the current decade, Texas' share of the market will be approximately 50,000 to 60,000 vehicles, indicating a maximum additional procurement needed of approximately 0.4 MW for both regulation up and down. In comparison, the maximum possible charging load for 50,000 vehicles, based on the results in Section 4.2.3, is 64 MW. In ERCOT, where system peak summer load (worst-case) is approximately 67 GW [86], this worst-case charging load will, on average, only have a small effect on prices, even though unscheduled charging peaks just after electric load, as shown in Figure 4.11. For the largest fleet simulated, while only 3 or 4 MW of additional regulation might be needed, the maximum charging load is 630 MW. Around peak times, this increase in load could yield an increase in prices, though the operational consequences would likely be limited.

4.4 Conclusions

From travel survey data collected by the US Department of Transportation, this work employs Monte Carlo simulation methods to examine the effects of PEV charging on the electric grid as a function of the PEV fleet size in a given region. Initial simulations focused on the influence of the demographic characteristics of drivers on the magnitude of vehicle use at various times of day. These simulations were used to determine whether subsequent analyses should account for potential changes in the demographics of PEV owners as the vehicle fleet grows. Results from these simulations indicated that the demographic parameters used to identify potential PEV buyers as a function of fleet size had an influence on vehicle use, most notably by increasing vehicle use during morning and evening commuting times.

The next set of simulations used the same travel survey data, converted to vehicle charging data, to estimate electricity demand associated with PEV charging and the variability in those charging loads. Results from these simulations show that the variability in charging loads presented by PEVs is small for even the largest fleet sizes tested, thus the relative variability in charging load becomes quite small relative to the total charging load. Further, even with PEVs comprising approximately 5% of the light duty vehicle fleet in Texas, average charging demand is several orders of magnitude smaller than typical peak demand during summer months in ERCOT, representing a worst-case average increase of approximately 1% of peak demand. Unfortunately, the timing of peak charging demand is nearly coincident with existing peak electricity demand. While there might be sufficient generating capacity to address a 1% increase in on-peak demand, charge scheduling could still take advantage of lower electricity prices and increase the utilization of generating capacity otherwise idled overnight.

Comparing results from the simulations with empirical charging data collected as part of the Pecan Street Smart Grid Demonstration Project indicate that the effort undertaken to process and prepare the data from the NHTS was successful in developing results that are consistent with actual vehicle charging patterns. From this result, though the NHTS data on personal vehicle use is based on fossil-fueled vehicles, for utilities and grid operators that do not have access to empirical PEV charging data, applying the methodology detailed to NHTS data could reasonably be used to anticipate temporal trends in future PEV charging loads. The simulation results could thus be useful for utilities and grid operators as part of their long-term strategic planning in lieu of empirical PEV charging data.

Again using the Monte Carlo simulations of PEV charging, the effect of variability in charging loads on ERCOT frequency regulation procurement quantities was assessed. These simulations indicate that additional regulation up is required to support uncertainty in the times when vehicles will begin charging, especially in the afternoon, when many drivers complete their last trips of the day. In contrast, additional regulation down is required to accommodate variability in the times when vehicles finish charging in the evening and early morning hours, since in this model, charging is dominated by overnight home charging. Combining vehicle charging load with net load revealed trends similar to the results without net load. Further, any seasonal changes in wind generation and electric load were dominated by diurnal patterns in those parameters and in vehicle charging loads. Simulation results indicate that minimal procurement changes, on the order of 1% of current procured quantities, will be required to accommodate unscheduled PEV charging, even for a region with more than 500,000 PEVs. Moreover, these increases in the quantity of regulation procured will increase day-ahead regulation prices by no more than 10%, depending on the service (regulation up or down) and time of day. While implementation of charge

pre-scheduling or optimal charging control could minimize the effect of PEV charging loads on the grid, these results indicate that such sophisticated control methods will likely offer only minimal benefits in the frequency regulation market.

Chapter 5

Modeling and evaluation of a vehicle-to-grid system for the provision of frequency regulation ancillary services

In this chapter, the development of a model to simulate a V2G program coordinator or "vehicle aggregator" is pursued. The intention of this model is to assess the market potential of providing frequency regulation services using V2G storage under circumstances where the aggregator is subjected to real-world operational conditions. These operational conditions include:

- uncertainty with respect to PEV use and state-of-charge,
- day-ahead prices that vary as a function of the offer quantity decisions made by the aggregator,
- revenue diversions to cover payments to V2G program participants, and
- uncertainty in real-time frequency regulation deployment quantities.

Modeling a V2G program or a vehicle aggregator is not itself a novel concept, but no known work exists that captures these complicating operational conditions. Such a model will yield a more realistic assessment of the revenue potential for a V2G system based on participation in frequency regulation markets.

The objective of this model is to observe the revenue maximizing decisions of the vehicle aggregator. These decisions can be readily posed in terms of variables and constants, which permits their representation in a mathematical programming for-

mulation. The vehicle aggregator maximizes their revenue by carefully choosing the quantity of frequency regulation to offer, but these decisions occur in the day-ahead market, when vehicle use and regulation deployments are unknown. Any solution approach should capture the uncertainty in those factors, since they affect the quality of the aggregator's decisions under real-time conditions. Capturing the separation between decisions made in the day-ahead market subject to uncertain future conditions and the consequences of those decisions in the real-time market is fundamentally analogous to a two-stage stochastic programming approach [65, 89]. In a two-stage SP, decision variable(s) are set in the first stage of the model and the second stage is then executed. In a typical problem, there exist recourse variables in the second-stage that can be used to mitigate the impact of the original decision(s) and increase the value of the objective function. This problem structure is commonly used in a variety of decision-making applications that require robust strategies [90, 91]. The structure of a two-stage SP with recourse is shown in Figure 5.1a.

In most electricity markets, after ancillary service offers are submitted for the day-ahead market, there are limited opportunities to later adjust those offers. Each market participant should have the ability to provide, in full, any ancillary services offered. Because typically there are limited recourse opportunities, this model is formulated as a two-stage SP with "simple recourse" [65, 89]. Simple recourse is the imposition of a penalty function for the failure to satisfy second-stage constraints due to first-stage decisions, without an opportunity to take other corrective action. This model thus requires a programming approach similar to the classical newsvendor problem [65, 89]. This modeling structure is illustrated in Figure 5.1b.

This vehicle aggregator model has time intervals dividing the modeled period. Each of these time intervals might appear as though it should be considered a sep-

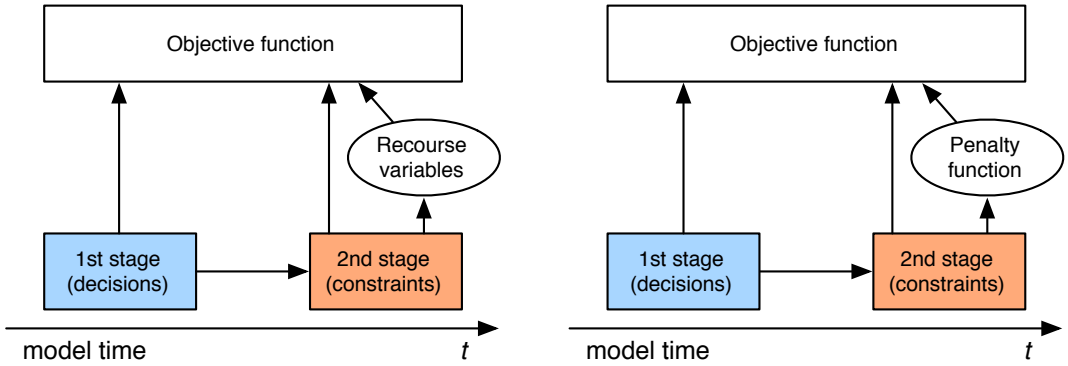

(a) recourse (b) simple recourse

Figure 5.1: Two-stage stochastic programs are characterized by a first stage, where some decisions are made while facing uncertainty about some parameters in the model, and a second stage, when the uncertain parameters become known and the consequences of the first stage decisions can be assessed. Often, there exist opportunities to recover from these consequences experienced in the second stage. These opportunities take the form of recourse variables. In the vehicle aggregator model, traditional recourse is not available to the aggregator, and the costs of the first stage decisions are assessed with a penalty function.

arate stage, thus creating a multi-stage SP [89]. A multi-stage model enables the observation of changes in recourse decisions during the temporal evolution of real-time conditions (as each stage is completed). Multi-stage models can also include time-dependent stochastic parameters, where the values of those parameters in a given interval can influence their subsequent values. Because recourse is not modeled for the vehicle aggregator and the stochastic parameters are not time-dependent, a multi-stage approach would add unnecessary complexity to this model.

5.1 Two-stage stochastic model formulation

Identifying the central features that define a vehicle aggregator and a V2G system led to the development of the model parameters set out in Section 5.1.1. The SP formulation of the system was created, establishing the form of the objective

function and model constraints that fully define the operation of a V2G system from the perspective of a vehicle aggregator.

5.1.1 Nomenclature

Index	Description
i	piecewise constant revenue curve segments
q	probabilistic scenarios
t	time intervals for modeled period

Constant	Units	Description
		Scalars
c^p	\$/MW	penalty price for units short of deployment
E	MWh	energy storage capacity per vehicle
L	hr	length of a single time step
N		number of vehicles in the system
η	MWh/mi	vehicle drivetrain efficiency
ρ	MW	average nominal vehicle charge rate (EVSE capacity)
		Vectors
c_t^{d+}	\$/MW	day-ahead regulation up procurement price (MCPC)
c_t^{d-}	\$/MW	day-ahead regulation down procurement price (MCPC)
c_t^r	\$/MWh	real-time energy price (MCPE)
p_q		joint probability distribution on scenarios q
		Matrices
$b_{t,q}$	mi	average distance driven by vehicles completing trips at time t; probabilistic
$F_{i,t}^+$, $F_{i,t}^-$	\$/MW	slopes for each segment in the piecewise curve
$f_{t,q}$		regulation deployment fraction; probabilistic
$g_{t,q}$		binary indicator set to 1 when $f_{t,q}$ is positive and 0 otherwise
$T_{i,t}^+$, $T_{i,t}^-$	MW	boundaries of each piecewise segment
$t'_{t,q}$		average duration (as a number of intervals) of trips completed at time t; probabilistic
$x_{t,q}$		vehicle departures, as a fraction of the total fleet; probabilistic

Variable	Units	Description
$a^+_{t,q}$, $a^-_{t,q}$	MW	slack variables for regulation provided short of deployed quantity
C^{p+}_t, C^{p-}_t	\$	total revenue from procurements
$\delta^+_{i,t}$, $\delta^-_{i,t}$	MW	width or partial width of segment i traversed
$d^+_{t,q}$, $d^-_{t,q}$	MW	regulation quantities deployed
$h^+_{t,q}$, $h^-_{t,q}$	MW	regulation quantities deployed in excess of requested quantity
$n_{t,q}$		number of vehicles connected in each period
r^+_t, r^-_t	MW	regulation capacities offered (decision variable)
$s_{t,q}$	MWh	total energy stored in the batteries of connected vehicles in each period
$v_{t,q}$	MW	total charging (non-regulation activity) of all connected vehicles
z	\$	value of the objective function

5.1.2 Governing equations

This model was designed with a focus on identifying available revenue opportunities for the vehicle aggregator in a frequency regulation ancillary service market, thus it has a profit-maximizing objective function. The objective function, given by Equation 5.1, is divided into the components for each stage. The first summation in the objective reflects the capacity payments for regulation up and regulation down, which occur in the first stage of the model. The second summation is the expected value of the second stage of the model, which represents real-time market conditions experienced by the aggregator. The superscript + and − symbols indicate whether a variable corresponds to regulation up or down, respectively.

$$\begin{aligned} \max \quad & z = \sum_t (c^{d+}_t r^+_t + c^{d-}_t r^-_t) L + \mathcal{Q}_q(a, d, h, v), \\ \text{where} \quad & \mathcal{Q}_q(a, d, h, v) = \sum_{t,q} c^r_t p_q (d^+_{t,q} - a^+_{t,q} + d^-_{t,q} + a^-_{t,q}) L \\ & \qquad - c^p p_q (a^+_{t,q} + a^-_{t,q}) - c^r_t p_q v_{t,q} L + c^p p_q (h^+_{t,q} + h^-_{t,q}) \end{aligned} \tag{5.1}$$

The second-stage model parameters and variables are all functions of the scenario variable q. The set of joint scenarios q arise from the combinations of the

stochastic input parameters, which are initially defined over smaller sets of scenarios. The stochastic parameters $b_{t,q}$, $t'_{t,q}$, and $x_{t,q}$ are initially defined over the same set of q_2 scenarios. The stochastic parameter $f_{t,q}$ is defined by a different set of q_1 scenarios. Each of the q_1 and q_2 scenarios has a corresponding probability. The joint set of scenarios q is all the possible combinations of the q_1 and q_2 input scenarios, thus there are a total of $q_1 q_2$ joint scenarios q. All of the second stage variables are thus defined over the scenarios q, corresponding to the possible combinations of input parameters.

The four terms in the summation for the second stage of the model represent, in order of appearance: 1) expenses or revenue from energy flows to and from the storage resource associated with providing regulation, 2) penalties for unmet regulation demand, 3) payments for energy used to charge vehicles, and 4) penalties for over-deployment of regulation. It is anticipated that vehicle owners will continue to pay for any energy used to charge their battery, but it is unclear exactly how they will be compensated for energy drawn from their battery to serve regulation up deployments. While developing this model, testing payments to the aggregator by participants for energy used to charge vehicle batteries revealed a condition where the aggregator acted to artificially inflate demand for battery charging to maximize these payments, even at average retail prices for electricity. To ensure that charging takes place during periods when prices are low and to avoid this distortionary aggregator behavior, these payments are not included explicitly in the model.

In the interest of simplicity, it is assumed that regulation offered by the vehicle aggregator is procured in full. This assumption is apparent in the substitution of offer quantities, the decision variable in the model, for procured regulation quantities. In so doing, it is assumed that the offer will be of a size and price that ensures that the aggregator will never be the marginal service provider. If it is further assumed

that the vehicle aggregator's offer is always the lowest priced offer, the quantity of regulation that must be served by the aggregator will be either the total capacity offered (and, as noted previously, procured) or the total quantity deployed, whichever is smaller.

This model is designed to treat the batteries of the vehicles parked and connected to the grid during any period t as a pooled resource. As such, Equation 5.2 ensures that the total energy available in those connected batteries, $s_{t,q}$, is adjusted to account for any conditions that could affect the size of that resource between periods. The parameters that affect the size of the connected storage resource are regulation deployments $d_{t,q}$, vehicle departures $x_{t,q}$, vehicle arrivals $x_{t-t',q}$, and battery charging $v_{t.q}$. In this model, vehicle departures occur whenever the vehicle is unplugged from the grid, regardless of whether it is driven immediately. Vehicle arrivals are modeled indirectly using trip duration and distance. Vehicles arriving in interval t are, in aggregate, vehicles that started a trip at $t - t'$, where t' is a function of the interval t and scenario q. Vehicles completing a trip in interval t can thus be denoted by combining the departing vehicle fraction $x_{t,q}$ and trip duration t'. These vehicles drove an average of $b_{t,q}$ miles during their trip. The energy remaining in their batteries upon completion of their trips is calculated based on their total energy at departure, $s_{t-t',q}$, reduced by the energy used to complete their trips, calculated based on average trip distance $b_{t,q}$ and the scalar parameters for drivetrain efficiency, η, and the number of vehicles in the simulation, N.

$$s_{t,q} = s_{t-1,q} - L(d^+_{t,q} - a^+_{t,q} + d^-_{t,q} + a^-_{t,q}) - x_{t,q}s_{t,q} + x_{t-t',q}(s_{t-t',q} - \eta b_{t,q}N) + Lv_{t,q} \quad \forall t, q \quad (5.2)$$

In Equation 5.2, the lagging storage variable $s_{t-1,q}$ is written in GAMS[1] with a circular operator, which forces the quantity available at the end of the day to be the same as at the beginning of the day, thus ensuring that the storage resource is not left completely depleted at the end of the simulation. This constraint also ensures the storage resource is never forced to meet an artificially created set point or threshold, such as ending the day at 50% state-of-charge or starting the day with the entire fleet fully charged. Using a circular variable and running the model for a single day at a time limits the ability of the aggregator to act in anticipation of market conditions days in advance, but since vehicle travel characteristics are relatively stable between weekdays, with the exception of Friday night, the state-of-charge of the storage resource should be similar during the overnight hours of each weekday. The arriving vehicle fleet fraction and initial state-of-charge variables, $x_{t-t',q}$ and $s_{t-t',q}$, respectively, are also circular with respect to t', so that trips ending at the beginning of the day (just after midnight) can start prior to midnight.

Equations 5.3a and 5.3b constrain the capacity available to serve regulation deployments in the real-time market, using $s_{t,q}$ to determine the stored energy and depleted capacity of the pooled storage resource. Equation 5.3a indicates that the energy capacity available to provide regulation up, either a reduction in load or an increase in generation, cannot exceed the total energy stored in the aggregated battery in any given period. Similarly, Equation 5.3b limits the energy capacity available to provide regulation down, either an increase in load or decrease in generation, to no more than the capacity depleted in the connected storage resource.

$$s_{t,q} \geq L(d^{+}_{t,q} - a^{+}_{t,q}) \quad \forall t, q \tag{5.3a}$$

[1] General Algebraic Modeling System (GAMS) is a high-level programming language specifically designed to solve large-scale optimization problems.

$$n_{t,q}E - s_{t,q} \geq L(-d^-_{t,q} - a^-_{t,q} + v_{t,q}) \quad \forall t, q \tag{5.3b}$$

The quantity of regulation deployed in the real-time market, denoted by $d^+_{t,q}$ and $d^-_{t,q}$ for regulation up and down, respectively, is defined in Equations 5.4 as the fraction $f_{t,q}$ of the quantity offered in each direction. Positive $f_{t,q}$ denotes regulation up and negative $f_{t,q}$ denotes regulation down. The inequality constraints in Equations 5.4 allow the regulation deployment quantity variables $d^+_{t,q}$ and $d^-_{t,q}$ to be any value below the deployment fraction $f_{t,q}$ multiplied by the quantity offered. If equality relations were used, these equations would force regulation deployments when they are counter-productive. The result of these inequality relations is that in periods when regulation up is requested, Equation 5.4b only ensures that the quantity deployed for regulation down is below whatever positive value is present on the right-hand side of the equation. A similar effect occurs with Equation 5.4a when regulation down is requested. To control these conditions, Equation 5.5 requires that the regulation up deployment variable must be non-negative, while the regulation down deployment variable must be non-positive, thus ensuring that regulation deployments are in the correct domain. The underline and overline notation in Equation 5.5 indicates the lower bound and upper bound of a variable, respectively.

$$d^+_{t,q} \geq f_{t,q} r^+_t \quad \forall t, q \tag{5.4a}$$

$$d^-_{t,q} \leq f_{t,q} r^-_t \quad \forall t, q \tag{5.4b}$$

$$\underline{d^+_{t,q}}, \overline{d^-_{t,q}} = 0 \tag{5.5}$$

As noted previously, the scenario-dependent parameter $x_{t,q}$ indicates the *fraction* of the total fleet departing at a given time t, while the duration of trips being completed in interval t, measured in model intervals, are denoted by $t'_{t,q}$. From these

parameters, the number of vehicles that can be counted as available in the pooled resource, $n_{t,q}$ can be directly calculated for all periods t in each of the joint scenarios q. The first term in Equation 5.6, $n_{t-1,q}$, is suppressed when $t = 1$, thus the initial value of $n_{t,q}$ can be optimized by the model for each scenario q. With that initial value, Equation 5.6 then calculates the deterministic changes to the number of vehicles in the system for each interval in each scenario.

$$n_{t,q} = n_{t-1,q} + N(x_{t-t',q} - x_{t,q}) \quad \forall t, q \tag{5.6}$$

Using the number of vehicles idle and plugged into the electric grid, denoted by $n_{t,q}$ and calculated in Equation 5.6, the real-time constraints on power flow to or from those vehicles can be readily calculated. Equations 5.7 use the assumed nominal EVSE power rating ρ to calculate the total power capacity of the pooled storage resource, where the quantity demanded, $d^+_{t,q}$ or $d^-_{t,q}$, is adjusted using the slack variables $a^+_{t,q}$ and $a^-_{t,q}$ if any demand could not be served. It is implicitly assumed in these equations that V2G-capable EVSEs are able to provide any amount of power, up to their rated maximum, in either direction (into or out of a vehicle's battery). Further, this constraint addresses only the power limitations associated with the EVSEs to which vehicles are connected, not the energy limitations associated with the "state-of-charge" of connected vehicles' batteries, which are addressed by Equations 5.3.

$$d^+_{t,q} - a^+_{t,q} \leq n_{t,q}\rho \quad \forall t, q \tag{5.7a}$$

$$-d^-_{t,q} - a^-_{t,q} + v_{t,q} \leq n_{t,q}\rho \quad \forall t, q \tag{5.7b}$$

In the existing model structure, regulation up deployments are essentially energy sales. As a result, it might be profitable to offer additional capacity and then use some excess available capacity to ensure that the requested deployment is met

and other constraints are simultaneously satisfied. To minimize these manipulative real-time market participation strategies, marked by regulation over-provision and provision of regulation opposite the requested direction, Equations 5.8 calculate these quantities, denoted by $h^+_{t,q}$ and $h^-_{t,q}$, which are then penalized in the objective function.

To illustrate the function of Equations 5.8, recall that $d^+_{t,q}$ is defined by Equation 5.4a as the quantity of regulation up deployed in only those periods when regulation up is requested. The parameter $g_{t,q}$ is equal to 1 when regulation up is requested, so if $d^+_{t,q}$ is exactly equal to the quantity requested, the two terms on the right-hand side of Equation 5.8a will be equal, and $h^+_{t,q}$ will be zero. Further, $h^+_{t,q}$ will equal any deployment beyond the requested quantity $f_{t,q}r^+_t$. In the same period, Equation 5.8b will assign a positive value (since Equation 5.5 forces $d^-_{t,q}$ to be non-positive) to $h^-_{t,q}$ for the provision of any regulation down. Equations 5.8a and 5.8b function similarly in periods when regulation down is requested, except that Equation 5.8b controls excessive provision of regulation down, while Equation 5.8a minimizes provision of uncommanded regulation up.

$$h^+_{t,q} = d^+_{t,q} - g_{t,q} f_{t,q} r^+_t \quad \forall t, q \tag{5.8a}$$

$$h^-_{t,q} = -d^-_{t,q} - (g_{t,q} - 1) f_{t,q} r^-_t \quad \forall t, q \tag{5.8b}$$

The PEV fleet size N is the independent variable used to scale the constraints in the model. Equation 5.9 defines the maximum available regulation capacity to be offered in the day-ahead market to no more than the number of vehicles N multiplied by the assumed V2G power capacity ρ of a single EVSE. Neither of these constraints define the real-time limits of these variables, which is important if the model is to yield viable solutions, but they are useful in ensuring that the physical bounds of the simulated system are not exceeded. Real-time energy and power constraints are handled by Equations 5.3 and 5.7, respectively.

$$\overline{r_t^+}, \overline{r_t^-} = N\rho \quad \forall t \tag{5.9}$$

$$a_{t,q}^+, a_{t,q}^-, n_{t,q}, r_t^+, r_t^-, s_{t,q}, v_{t,q} \geq 0 \quad \forall t, q \tag{5.10}$$

Finally, Equation 5.10 indicates the variables that are defined as non-negative and have not been specified as such in a previous equation. In the GAMS code, the special case of variable declaration, "positive variable," is used to impose non-negativity on these variables, obviating the need for separate equations.

5.1.2.1 Controlling day-ahead market variables

This model simulates elements of both a day-ahead market and a real-time market, but because it has been formulated as a two-stage SP, only one execution of the model is required to optimize all the variables. The operation of the two market stages simulated is defined using specific variables and equations to control each stage. The time scales on which these two markets operate, however, is not reflected in the formulation, which has only one apparent time index t. This time index reflects the number of intervals in the real-time market, which is chosen to balance computational expense and fidelity to the real-time deployment of regulation in ERCOT. Regardless of the number of intervals of t, the day-ahead market is operated on hourly intervals. Since the length of a time interval in hours is given by L[2], any day-ahead parameters should exist L^{-1} times for each hour. For parameters that apply to the day-ahead market, c_t^d, $T_{i,t}$, and $F_{i,t}$, the input values are repeated for all the intervals in each hour of the day. Day-ahead market variables r_t^+ and r_t^- are forced to repeat L^{-1} times within the model itself using Equation 5.11. Though Equation 5.11 is shown only once, to constrain each of the day-ahead market variables, it appears twice in

[2] This procedure is only applicable to models with time intervals less than or equal to one hour.

the code. The notation mod in Equation 5.11 denotes the modulo operator, which is the remainder of the division of the left term by the right term.

$$r_t = r_{t+1} \quad \forall t,\, t-1(\mathrm{mod} L^{-1}) < t(\mathrm{mod} L^{-1}) \tag{5.11}$$

An equation in the form of Equation 5.11 is only required for variable r_t. Equation 5.14 ensures that if r_t repeats, $\delta_{i,t}$ will also repeat over the same intervals, and Equation 5.13 does the same for the total revenue variable C_t^p. In both the exogenous and endogenous objective function forms (detailed in Section 5.1.2.2), the repetition of the day-ahead procurement payments L^{-1} times in each hour indicates that revenue will be over-predicted by the same factor, thus the time step length variable L is multiplied by the day-ahead revenue terms in both objective functions.

Instead of using Equation 5.11 to correct for the difference between the number of intervals in the model and the hourly intervals in the day-ahead market, variables and parameters that define elements of the day-ahead market could have been defined on a separate time interval set t^*. Such an approach would have eliminated the need for Equation 5.11, but would have required the introduction of a new variable to transform regulation procured on t^* intervals to the real-time market t intervals so that the quantity of regulation procured could be used as a constraint for variables defining the real-time market. Compared to the approach ultimately adopted, including another interval set t^* did not appear to offer appreciable benefits by reducing the number of variables and equations or otherwise simplifying the model formulation.

5.1.2.2 Incorporating endogenous prices

Most studies of energy storage on the electric grid, regardless of the capacity of the device studied, assume exogenous capacity prices for the ancillary service markets in which the devices participate [7, 28, 45, 47, 62, 64]. This exogenous capacity price

approach is equivalent to assuming that the storage system is a price taker. Depending on the power and energy capabilities of the device, and the particular market in which it is participating, this assumption might not be entirely valid. Some authors have noted that this assumption is one of the primary limitations of their results [7, 47]. Additional equations to support endogenous prices have been developed for this model to determine whether including these prices has a significant impact on expected revenue and market participation strategies for the vehicle aggregator. This treatment of prices as a function of the vehicle aggregator's offer quantities simulates a case where the aggregator has market power.

Including capacity prices for ancillary services directly in the model would introduce non-linear terms into the objective function because the capacity offered would be multiplied by the capacity price, where the price is a function of the offer quantity. To ensure a linear model, ancillary service price data are prepared following the procedure in Section 5.3.1. The resulting piecewise linear revenue curves are included in the model with a formulation derived from the approach used by Carrión and Arroyo to include non-linear production cost curves as piecewise linear functions in their unit commitment model [92]. The segments of the piecewise revenue curve are indexed by set i and there exist unique piecewise revenue curves for each period t. Figure 5.2 shows a sample piecewise revenue curve and a corresponding revenue curve implied by the exogenous price model approach.

This approach to calculating the revenue obtained from offering frequency regulation, where prices are not treated as exogenous components of the model, requires that the objective function be modified. The updated objective function is given in Equation 5.12.

$$\max \quad z = \sum_{t} (C_t^{p+} + C_t^{p-})L + \mathcal{Q}_q(a, d, h, v) \tag{5.12}$$

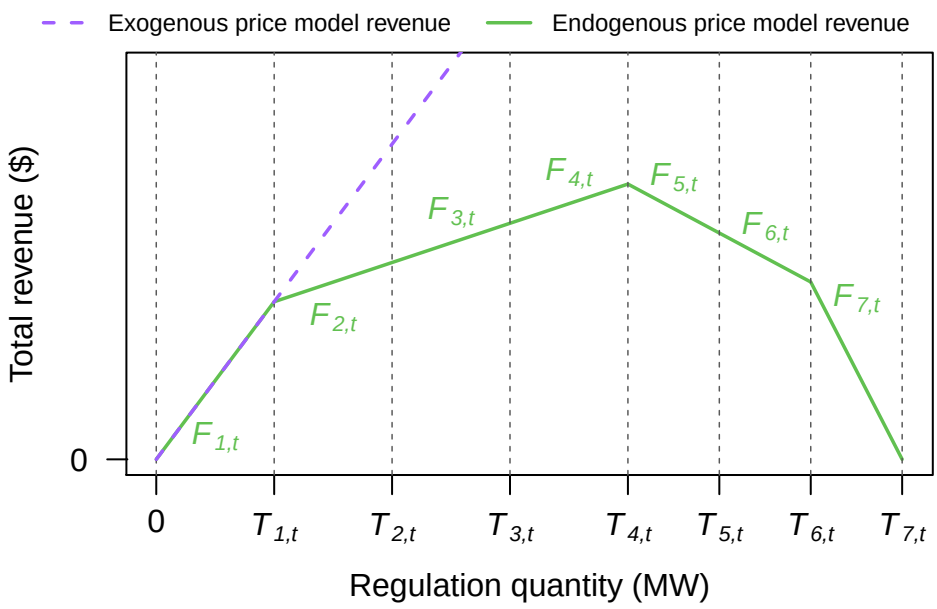

Figure 5.2: Endogenous prices are included in the model using piecewise linear revenue curves. These curves are parametrized by their slopes, $F_{i,t}$ and the boundaries, $T_{i,t}$, dividing the curve into I piecewise segments. In the exogenous price model, changes in the offer quantity have no effect on prices, so the total revenue generated in each hour of the day-ahead market is a linear function of the quantity offered.

Equations 5.13 through 5.16 symbolically represent the parameters that define piecewise revenue curves and indicate how the total revenue is updated based on the quantity of regulation up or down offered (and procured) in each period. Each of these equations is duplicated in GAMS such that regulation up and regulation down capacity are handled using separate equations.

Equation 5.13 determines the total revenue for each period t in the model by multiplying the quantity included from each segment i by the slope $F_{i,t}$ of that segment, thus determining the contribution from each section to the total revenue.

$$C_t^p = \sum_i F_{i,t}\delta_{i,t} \quad \forall t \tag{5.13}$$

Equation 5.14 determines the total contributions of the consecutive total and partial segments such that the sum of those segments matches the procured quantity.

$$r_t = \sum_i \delta_{i,t} \quad \forall t \tag{5.14}$$

Each segment i should be included sequentially, such that if the procured quantity r_t is within the third segment, the contributions from segments 1 and 2 should be included in full, but none of the fourth segment should be captured. With a revenue maximizing objective function, the segments will only be used sequentially if the slopes $F_{i,t}$ of the piecewise revenue curve are strictly monotonically decreasing. The form of the revenue curve is discussed in detail in Section 5.3.1.2. The term $\delta_{i,t}$, calculated in Equation 5.15, is equal to the whole or some part of each segment, depending on its contribution to the total procured quantity, and hence, the total revenue.

$$\begin{aligned} \delta_{i,t} &\leq T_{i,t} - T_{i'-1,t} \quad \forall i, t, \\ \text{where} \quad & i' \subset i, i' \in 2, \ldots, I \end{aligned} \tag{5.15}$$

$$\delta_{i,t} \geq 0 \quad \forall i, t \tag{5.16}$$

Finally, Equation 5.16 requires that $\delta_{i,t}$ be greater than or equal to zero for all sub-segments and all periods.

5.1.2.3 Incorporating compensation payments

If there are sufficient revenue opportunities for an aggregator to at least cover the costs of their own facility, staff, and operations, there are a multitude of approaches they could employ to compensate V2G program participants for the use of their vehicles' batteries to provide ancillary services. Capacity payments from the frequency regulation day-ahead market are often identified as the primary revenue source for a V2G program, thus any compensation strategy would likely, though not necessarily, include a portion of the capacity payments. A few compensation approaches were considered for this work:

- payment per unit of regulation provided on a per-vehicle basis
- payment per unit of regulation provided plus a capacity payment based on a ratio of the total grid-connected time for the vehicle to the total time for all the vehicles in the system over the operational period (e.g. one month)
- payment per unit of regulation provided plus a capacity payment based on a fraction of the total fleet $1/N$
- constant payment to each participant in each operational period (e.g. one month)

In the interest of limiting the size of the model, individual vehicles are not modeled discretely, thus any compensation payments cannot be calculated independently

for each of the N participating vehicles. Instead, the total cost of compensatory payments is captured in the objective function. Fixed compensation for all participants, unless the number of vehicles N becomes a variable, should reduce the objective function value, but it is unclear whether it will affect the decisions made by the model. Ultimately, the compensation approach modeled assumes that the aggregator retains some margin m of the total revenue indicated by the objective function.

$$\begin{aligned} \max \quad & z = \sum_{t} m(C_t^{p+} + C_t^{p-})L + \mathcal{Q}_q^*(a,d,h,v), \\ \text{where} \quad & \mathcal{Q}_q^*(a,d,h,v) = \sum_{t,q} c_t^r p_q \left[m(d_{t,q}^+ - a_{t,q}^+) + d_{t,q}^- + a_{t,q}^- \right] L \\ & \qquad - c^p p_q (a_{t,q}^+ + a_{t,q}^-) - c_t^r p_q v_{t,q} L + c^p p_q (h_{t,q}^+ + h_{t,q}^-) \end{aligned} \tag{5.17}$$

The objective function used in this model, given by Equation 5.17, is intended to maximize aggregator revenue. Compensatory payments are thus subtracted from the objective. If compensation is provided to each V2G participant based on a share of the total earnings passed on by the aggregator, the revenue devoted to these payments can be represented in their inverse, as the margin m, or fraction, of the revenue retained by the aggregator. Revenue is derived from payments for energy dispatched when regulation up is provided and from capacity in the day-ahead market, thus the margin is applied to both terms, as shown in Equation 5.17.

5.2 Solving the stochastic program

The model presented in Section 5.1 takes the form of a stochastic program. The stochastic parameters in the model are defined for a set of discrete scenarios q, which converts those parameters to a deterministic form. With this conversion of the stochastic parameters, the model being solved is referred to as the "deterministic equivalent program," and can be solved using existing optimization software packages that can solve LPs. The model was written using GAMS and solved using the IBM

Table 5.1: The scalar parameters in the model are set to the values indicated. The fleet size, N, is varied between 1,000 and 5,000,000 vehicles to explore the effect of fleet size on aggregator revenue and other model variables. Later, case studies will examine the effect of other values for the vehicle battery size, E, and the EVSE power rating, ρ.

Scalar parameter	Units	Value
c^p	\$/MW	100
E	MWh	0.024
L	hr	1/240
η	MWh/mi	0.00034
ρ	MW	0.0055

ILOG CPLEX optimizer. Final values for model variables were exported by GAMS to CSV files.

The GAMS code was run on a high performance computing system with GAMS 24.1.2 and CPLEX 12.5.1. The code is structured such that it solves both the exogenous and endogenous price models for a given set of inputs. Both models had a runtime limit of 25,000 seconds (almost 7 hours). For a single fleet size on a single day, runtimes typically varied between one and two hours per model, depending on whether other users simultaneously had active processes. Because of the relatively consistent runtimes, few cases reached the runtime limit. Any cases that reached the runtime limit typically encountered some sort of error during the solution process and never found a feasible solution. These cases were identified during the results summary process and were run again.

5.3 Stochastic model input data preparation

Four categories of data comprise the inputs to the stochastic linear program described in Section 5.1. These inputs are, as shown in Figure 5.3, 1) vehicle trip

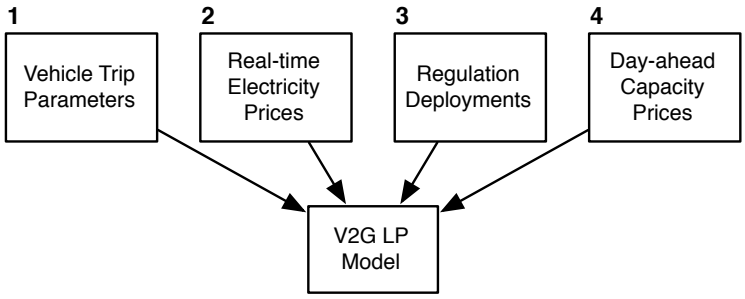

Figure 5.3: Four primary data sets, indicated by boxes 1–4, are required by the SP V2G system model. Each of these boxes belies manifold preparatory steps required to transform original data into inputs of the form needed by the model. These preparatory steps are detailed for boxes 1–4 by Figures 5.6, 5.10, 5.8, and 5.4, respectively, and also in the discussion in the sections surrounding these figures.

parameters, 2) real-time electricity prices, 3) real-time regulation deployments, and 4) day-ahead capacity prices. As detailed in Chapter 2, this model was developed with the intention of capturing uncertainty in vehicle availability and regulation deployments, as these parameters can be forecast but not directly controlled by the vehicle aggregator and could have a significant effect on the aggregator's revenue and exposure to penalties imposed by the system operator. Additionally, the effect of including prices directly in the model, where the aggregator's quantity offered in the day-ahead market influences clearing prices was of interest. To support this capability and facilitate comparison between deterministic and uncertain day-ahead capacity prices, both constant day-ahead prices and piecewise linear revenue curves were prepared for the model. Parameters in each of the categories of model inputs were prepared for the model using processes specific to the format and characteristics of the underlying source data. These procedures are detailed in the remainder of this section.

The stochastic model is designed to run single-day simulations. To capture seasonal variations in electricity market conditions, especially day-ahead and real-time prices, multiple dates were simulated. Twenty-six dates were selected to capture

conditions approximately every two weeks throughout the study year. These dates apply to only price data, as the remaining input parameters are not considered unique to any particular day or subject to seasonal variations. All the price data used with the model were from 2011, thus dates from that calendar year were used. Starting with Monday on the first full week in 2011 (January 3), every 15th day was selected until a Friday was reached. Following a Friday, the process would repeat beginning with the Monday occurring two weeks after the Friday, exactly 10 days later. For example, the first Friday selected was March 4, falling in the 9th week of 2011, and the next date selected was the Monday occurring 10 days later, in the 11th week, March 14. This pattern was repeated until the end of the year, except for the dates in November. To avoid the impact of changes in schedules throughout the week of the Thanksgiving holiday, Tuesday, November 15 was substituted for Thanksgiving day, Thursday, November 24. To retain the same number of Tuesday, Wednesday and Thursday dates sampled, Wednesday, November 9 was replaced with Thursday, November 10 and Tuesday, October 25 was replaced with Wednesday, October 26. A complete accounting of the 26 dates used in the model is given in Table 5.2.

5.3.1 Frequency regulation capacity prices

Day-ahead ancillary service market data were acquired from ERCOT for calendar year 2011. These data included clearing quantities, as well as all tendered offer prices and quantities, for each hour and for all the ancillary services in ERCOTs ancillary service market structure.[3] Deterministic day-ahead frequency regulation prices included in the model are referred to as "exogenous" prices, since the prices are determined outside of the model and are not affected by market participation decisions

[3] Additional day-ahead ancillary service market data not required for this work were included in the files provided by ERCOT.

Table 5.2: The V2G aggregator model is set up to run single-day simulations. To limit the total number of simulations while also capturing seasonal changes in day-ahead and real-time market conditions, biweekly dates were drawn from calendar year 2011. In case there might be some small, unobserved differences between conditions on different days of the week, dates were drawn from all five weekdays. To be consistent with vehicle use patterns, weekends were omitted from this study.

Monday	Tuesday	Wednesday	Thursday	Friday
January 3	January 18	February 2	February 17	March 4
March 14	March 29	April 13	April 28	May 13
May 23	June 7	June 22	July 7	July 22
August 1	August 16	August 31	September 15	September 30
October 10	November 15	October 26	November 10	December 9
December 19				

made by the vehicle aggregator. In contrast, "endogenous" prices are the price data used to create piecewise constant revenue curves, where those revenue curves facilitate changes in the apparent market clearing price subject to the aggregator's decisions.

5.3.1.1 Exogenous capacity prices

Since the market clearing prices were not provided by ERCOT, they were determined by identifying the price associated with the clearing quantity in each hour. Because the model has 15 second periods, the clearing prices for each hour were repeated 240 times. Repeating the offer quantity and clearing price 240 times would yield a corresponding overestimation of each simulated day's revenue. This overestimation is corrected by multiplying the day-ahead revenue term in the objective function, Equation 5.1, by the scalar parameter L. The same correction factor is applied to the day-ahead market revenue term in the objective function formulation with piecewise linear revenue curves, Equation 5.12.

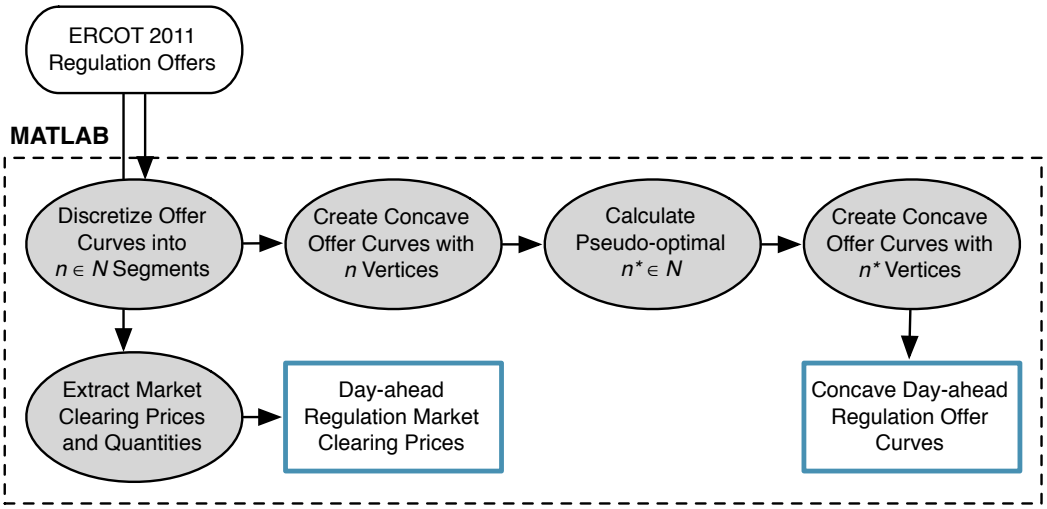

Figure 5.4: Day-ahead frequency regulation offer quantities and prices from ERCOT for 2011 were converted into two different model inputs, one for the exogenous price model and one for the endogenous price model. Clearing prices and quantities were extracted to serve the exogenous price model. The offer curves for the endogenous price model were prepared in a multi-step process using all the offers below the clearing price.

5.3.1.2 Endogenous capacity prices

According to ERCOT nodal market protocols, day-ahead frequency regulation offers must be submitted with a single price for the entire quantity offered [93], which yields a piecewise constant offer curve. The offer curve is constructed from the price and quantity data received by ERCOT. Including piecewise constant offers for frequency regulation directly in the model formulation would require binary or integer variables. Further, including the capacity prices in the model would introduce non-linear terms into the objective function, with the capacity price multiplied by the capacity offered. To avoid these non-linear terms, prior to importing the regulation offer data into the model, the capacity prices are multiplied by the regulation quantities offered, leading ultimately to the curve shown in the bottom panel of Figure 5.5. This piecewise linear revenue curve is included in the model with a formulation based

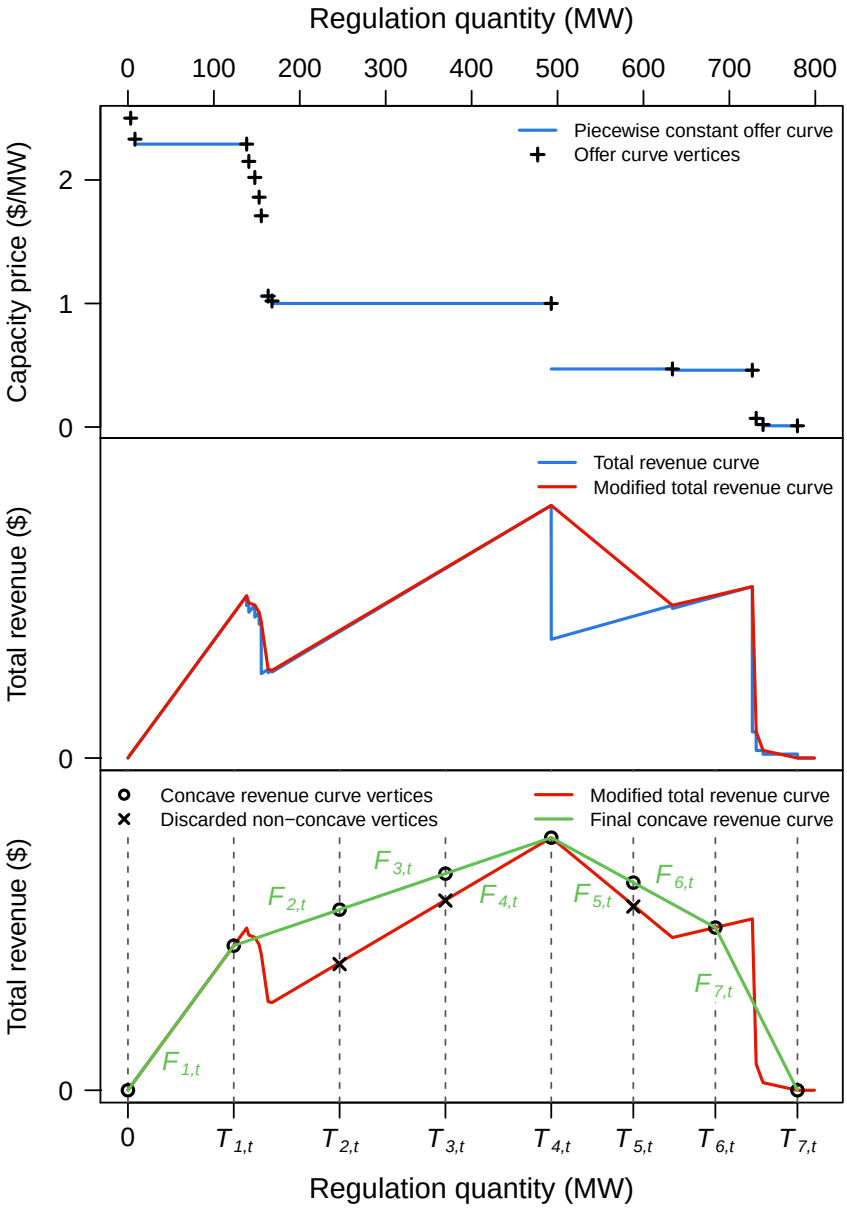

Figure 5.5: The quantity and price-defined offers submitted to ERCOT for frequency regulation are reversed, such that the marginal offer price is set to a quantity of 0, yielding a piecewise constant capacity price curve (top panel). To avoid non-linear terms in the objective function, this capacity price curve is multiplied by the offer quantity to develop a total revenue curve (middle panel). This curve is converted to a piecewise linear revenue curve, parameterized by its slopes $F_{i,t}$ and the vertices $T_{i,t}$ that define the edges of each discrete region in which those slopes exist (lower panel).

on the approach used by Carrión and Arroyo [92] to include non-linear production cost curves as piecewise linear functions in their unit commitment model.

These offer curves are (indirect) deterministic inputs to the model, which implies that the vehicle aggregator has enough experience participating in the electricity market in its region to have a sense of how prices will change subject to their behavior and other market activity. While it is certainly reasonable to assume that there would be some uncertainty associated with market prices, due to the reaction of other market participants to the aggregator's offers and other market conditions that affect prices, it is anticipated that this uncertainty and its effects are small relative to the change in prices arising from the aggregator's offers themselves.

The piecewise linear revenue curves in the model, illustrated by the green line in the lower panel of Figure 5.5, must not only be piecewise linear, but also have *strictly* monotonically decreasing slopes. Stated differently, the piecewise segments must together form a concave function. This requirement stems from the structure of the equations in Section 5.1.2.2, where day-ahead market revenue is calculated from the slopes and vertices of the $i \subseteq I$ curve segments. According to Equation 5.14, the model will determine the portions of each of the segments $\delta_{i,t}$ that will be included, such that their sum is equal to the offer quantity set by the aggregator. For a given t, this equation should "fill" each segment $\delta_{i,t}$ sequentially, beginning at the origin, until the equality condition is satisfied. In a linear model, this "filling" procedure cannot be enforced directly with a constraint equation, but without such a constraint, if the slopes are not strictly monotonically decreasing, the model will maximize revenue by filling the segments with the largest slopes first, even if those segments are non-consecutive.

To comply with the requirement that the revenue curves have strictly monotonically decreasing slopes, for any offer quantities larger than the clearing quantity, the revenue will drop below 0. This last segment in the revenue curve has a slope slightly less than $F_{I-1,t}$ and an upper quantity bound (not shown in Figure 5.5) of $T_{I,t} = 20{,}000$, which is the maximum simultaneous power output of 3.6 million vehicles. Though the largest fleet size simulated has 5 million vehicles, because market clearing quantities are far lower than 20,000 MW, it was assumed that the effective offer quantity limit was sufficiently high to not affect aggregator decisions or estimated revenue.

Preparing piecewise linear and strictly monotonically decreasing revenue curves from the original data requires a multi-stage process. Each step in this process is performed identically for both regulation up and regulation down data. First, for a given day, quantities below the clearing quantity and their corresponding prices are extracted for each hour. The prices and quantities are then reversed such that the prices are in decreasing order, beginning with the clearing price. The clearing price is thus associated with a quantity of 0. The form of the offer curve generated by this procedure is illustrated in the top panel of Figure 5.5. This restructured curve represents the anticipated clearing prices for the vehicle aggregator, subject to the quantity they offer. The price-quantity vertices in the restructured curve are then converted to total revenue vertices by multiplying prices and quantities together. The piecewise linear curve created from these points is shown in blue in the middle panel of Figure 5.5.

The final concave revenue curves can be interpolated from the modified total revenue curves using the desired number of vertices for the new curves. The modified total revenue curve vertex with the highest revenue identifies the inflection point

of the final concave revenue curve (in Figure 5.5, the inflection is at the point at $T_{4,t}$ in the lower panel), separating the curve into two regions, between 0 and the inflection point, and from the inflection point to the maximum allowed quantity. The x coordinates for the final concave revenue curve vertices are uniformly spaced in the regions before and after the inflection point. The number of points in each region is based on the fraction of the curve that appears before the inflection point. If the inflection point is at vertex $i = a$, then this fraction is equal to $T_{a,t}/T_{I,t}$. There are a total of $I + 1$ vertices, thus the number of vertices that appear before the inflection point is given by $\lfloor (I + 1)(T_{a,t}/T_{I,t}) \rfloor$. Any remaining vertices are used in the region after the inflection point. The y coordinates of the final concave revenue curve vertices are then determined by interpolation from the modified total revenue curves.

With the final concave revenue curve vertices calculated, the points are checked to ensure that they form a strictly monotonic set. All the points appearing after zero and before the inflection point are checked, from left to right, by determining whether they lie above the line connecting the previous point, located at (x_{k-1}, y_{k-1}),[4] and the inflection point. The slope of this evaluative line is denoted by σ. If any points (x_k, y_k) appear below that line, their y coordinates (y_k) are adjusted using Equation 5.18. When needed, Equation 5.18 recalculates the y coordinate by first calculating the y coordinate of the point that would lie on the line between the previous point and the inflection point (calculated by the term $\sigma(x_k - x_{k-1}) + y_{k-1}$), then increasing the value by a small fixed amount (the term $1 + 1 \times 10^{-8}$), which ensures that the point will be on a monotonically concave curve, and by an amount that decreases as k increases (the term $\frac{1}{1 \times 10^4 \ln(k-0.5)} + 1$), which makes the point part of a strictly monotonic set. A similar procedure is applied to check the final concave revenue curve

[4] Notation used in these equations does not correspond with or supersede the notation for the SP formulation introduced in Section 5.1.

vertices that appear to the right of the inflection point, except that these points ($x_{k'}$, $y_{k'}$) are traversed from right to left, and the adjustment equation for points that lie below the line between the previous point (to the right) and the inflection point is given by Equation 5.19.

$$y_k = \left(\frac{1}{1\times 10^4 \ln(k-0.5)} + 1\right)(\sigma(x_k - x_{k-1}) + y_{k-1})(1 + 1\times 10^{-8}) \tag{5.18}$$

$$y_{k'} = (1\times 10^{-7}\exp(k') + 1)(\sigma(x_{k'+1} - x_{k'}) + y_{k'+1})(1 + 1\times 10^{-8}) \tag{5.19}$$

Final concave revenue curves, like that shown in green in Figure 5.5, are calculated for integer $n \in N = [7, 12]$ vertices. To determine the number of vertices to use in the input to the GAMS model, the sum of squared errors (SSE) s_n between the original curve, shown in Figure 5.5 in red, and the final curve, shown in green, is calculated for each of the vertex counts n. There then exists an SSE value for each hour and vertex count. The sum of the SSE values for each vertex count is calculated and then adjusted based on the number of vertices, following the form of Equation 5.20.

$$s_{n*} = \min s_n n^{0.2} \quad n \in N \tag{5.20}$$

The "optimal" vertex counts $n*$ that minimize the RHS of Equation 5.20 for regulation up and regulation down are averaged, and the floor of the average is used as the number of vertices for the given date. The new vertex values corresponding to the "optimal" vertex count identified are used to calculate slopes $F_{i,t}$. The x coordinates of the vertices, corresponding to the model parameter $T_{i,t}$, and the slopes are exported for use in the GAMS model.

5.3.2 Vehicle travel data

To simulate the uncertainty in PEV charging activity faced by a vehicle aggregator, parameters defining vehicle travel in the model are scenario-dependent.

Because detailed travel data from PEVs were not readily available, and there exists little compelling evidence that PEVs will be driven in patterns significantly different from the ICE vehicles they replace, GPS-based ICE vehicle travel data were used as a proxy for PEVs. The PSRC data discussed at length in Chapter 3, were used again here. From these data, vehicle travel in the model was parametrized by departures $x_{t,q}$, tour duration $t'_{t,q}$, and tour distance $b_{t,q}$. The procedure used to develop these parameters is shown in Figure 5.6. As shown, a two step process was used, first preparing the data using the same procedures applied to the PSRC data used in Chapter 3, and then clustering those intermediate parameters in R to develop "profiles" or "scenarios" of the final, desired parameters.

The PSRC data used in Chapter 3 were employed here, but different outputs were desired. The four intermediate parameters needed as inputs to the clustering procedure were trip departure counts, trip distances, trip durations, and the number of cars active in the study on each day. These parameters were formulated such that each day was a separate record. The clusters leading to the desired parameters were then created by combining similar days from these intermediate parameters. It is believed that seasonal changes in vehicle use patterns are not significant compared to diurnal changes, thus the intermediate parameters prepared for clustering include data from throughout the year. To be consistent with the other data prepared for the model, only weekday data from the PSRC study are included in the MATLAB outputs. Vehicle participation, or the number of vehicles in the study on any given day, was determined using the procedure in Chapter 3. Trip departure counts were drawn from the starting timestamps associated with all the trips that start on a given day and are simply counts of the number of trips started in each minute of that day. Trip durations were calculated as the number of minutes elapsed for each trip, and were recorded at the trip end time. All the trips ending in a given minute had their

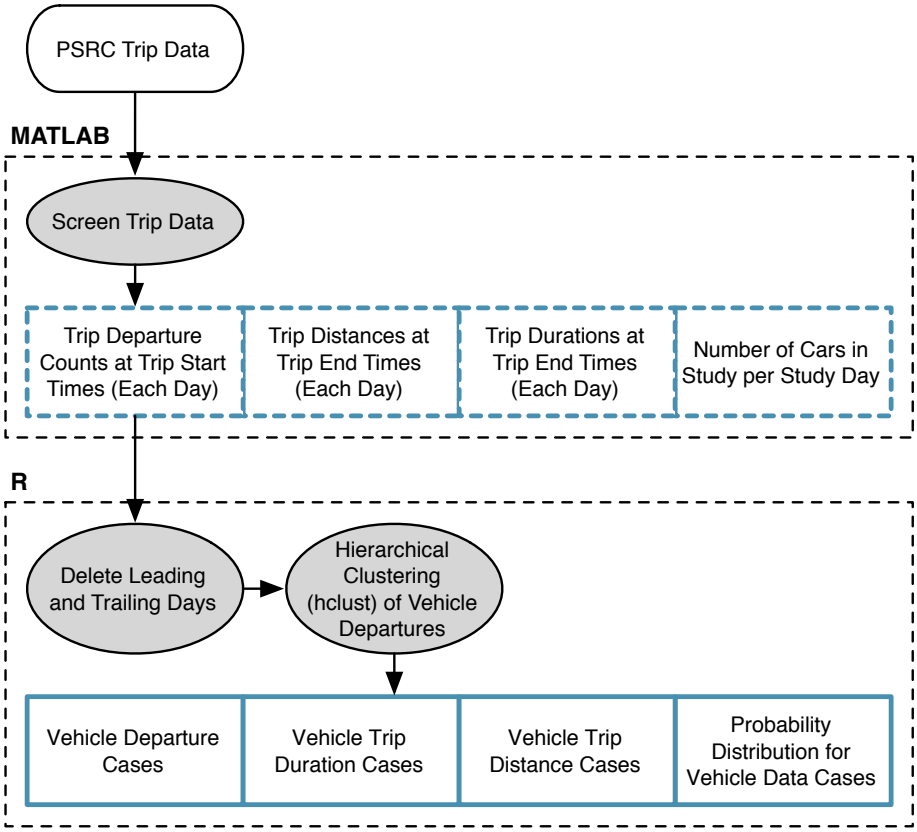

Figure 5.6: Three stochastic parameters define vehicle use patterns observed by the vehicle aggregator: vehicle departure counts for each minute in a 24-hour period, trip durations, measured in minutes, and trip distances, measured in miles. Each of these parameters is described by a set of scenarios that define the possible outcomes of the parameters, and is subject to a probability distribution over these scenarios. Each scenario is defined for an entire 24-hour period for each of the three parameters, such that they form a matched set. The scenarios are developed from the GPS-based travel data collected by the PSRC. These data are prepared in MATLAB following a rubric similar to that used with the same data in Chapter 3, whereupon four new intermediate parameters are produced. These intermediate parameters correspond to each of the three final stochastic parameters, with a separate vector of 1-minute data for each day in the PSRC study (except those removed during the screening process), plus a measure of the number of vehicles participating in the study on each day. In R, hierarchical clustering is applied to the intermediate departure fraction parameter to determine which days belong in each cluster. The remaining vehicle travel parameters are grouped using the same cluster membership. Finally, the probability distribution for the stochastic parameters are determined by dividing the number of days in each cluster by the total number of travel days in the original data.

durations averaged, such that each day would have a single duration value reported in each minute. Trip distances were calculated similarly, taking the calculated distances (in miles) and the trip end times for each day and averaging all the distances for a given minute such that each minute of each day had a single corresponding trip distance value. Using trip departure counts based on trip start times, while trip distances and durations are based on trip end times, was pursued when it became evident that using values all based on trip start times or trip end times would not lead to a viable set of governing equations and would yield poorly fitting clusters.

Probabilistic parameters that characterized single-day vehicle use patterns were derived from the four intermediate parameters prepared in MATLAB using hierarchical clustering. This approach to cluster analysis draws from the initial data and identifies pairs, working from greatest to least commonality, until all of the data have been combined. This data clustering method can be shown graphically in a "dendrogram." Figure 5.7 shows an example of a simple dendrogram developed from clustering 16 days of data. Fundamental to this approach was the assumption that the days of travel data drawn from the PSRC study can be treated as a characteristic sample of vehicle travel days, and can thus be used to develop a set of typical vehicle use scenarios and assess the probabilities of those scenarios. In this case, the input data were vehicle departure counts divided by the number of vehicles participating on each day. The R function *hclust* was used to cluster these data. The first time the vehicle departure counts were clustered, the first 54 days and seven of the last 10 days displayed use trends sufficiently inconsistent with the remaining data and previously observed use patterns that they were screened from the set. The remaining 88 days were clustered again. To determine the cluster averages, or "centers," the desired number of clusters had to be determined. Cluster counts between 4 and 20 were tested to identify a balance between sufficient diversity between the clusters and

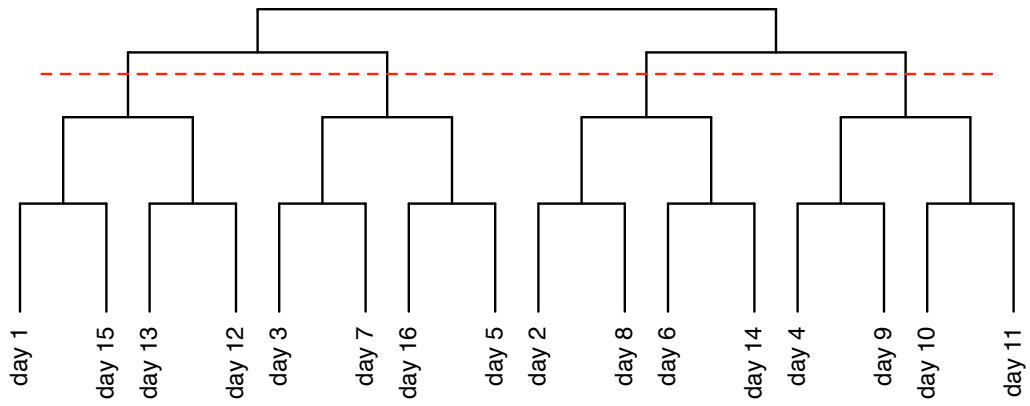

Figure 5.7: The process of identifying similar data using hierarchical clustering can be illustrated using a dendrogram. The dendrogram shown is highly simplified for illustrative purposes. With hierarchical clustering, data reported for each day are combined based on similarity. At the bottom of the dendrogram, data within clusters are more similar than they are toward the top. For example, days 1 and 15 are more similar to each other than they are to days 4 and 9. This is shown in the figure as days 1 and 15 being paired at the bottom of the dendrogram, while days 1 and 15 are paired with days 4 and 9 only at the highest level. Clusters can be identified by cutting the dendrogram horizontally at any point. The dotted red line shows the cut that yields four clusters.

model input size, since each additional cluster would add another column to each of three inputs to the SP model. Ten clusters was chosen as an acceptable initial compromise, subject to later reduction if solution times were found to be excessive. With the cluster count chosen, and thus membership in each cluster determined, cluster centers were calculated, excluding zero values reported in the original data, and trip duration and distance were averaged in clusters using the same cluster membership as the vehicle departure counts data. The probability of each cluster was calculated as the number of members (number of days) in each cluster, divided by the total number of days clustered. This clustering procedure thus developed the parameters needed to characterize driving patterns for a set of archetypal days, along with the probability of each of these days.

Because the SP model is structured on 15 second intervals and the PSRC data were provided with one-minute fidelity, the cluster centers had to be converted to the shorter intervals, quadrupling the number of entries in each cluster. Departure fractions are repeated for each 15 second interval in a minute and then divided by four to ensure that in a given minute, an equivalent total number of departures occur. Trip durations were repeated and then multiplied by four to account for the quadrupling in the number of intervals (e.g. a 15 minute trip spans 60 15-second intervals). Trip distances are replicated for the intervals within each minute but not adjusted, since the distance traversed in a trip is not affected by the change in the number of intervals.

5.3.3 Frequency regulation deployments

In addition to uncertainty about the times when vehicles would be idle or in use, which is handled by parameters prepared following the procedure outlined in Section 5.3.2, the aggregator also faces uncertainty with respect to frequency regulation deployments in the real-time market. To capture this uncertainty, regulation deployments are modeled as a stochastic parameter based on historical measurements of area control error (ACE). The regulation deployment scenarios were created by first preparing the original ACE data in MATLAB, screening weekend data and shifting the original data onto common time intervals set at the SP model's interval length. These data were then converted to deployment fractions before clustering in R by dividing the magnitudes by the maximum or minimum values on each day. Because the ACE data were combined into clusters, thus largely erasing their their year- or date-specific characteristics, it was not critical that the original data be from the same year as the day-ahead and real-time market frequency regulation price data. Since ERCOT ACE signal data from 2009 were readily available, they were used to create

the input to the SP model. This procedure for preparing the regulation deployment data is shown graphically in Figure 5.8.

ACE is a parameter used by ERCOT to calculate the appropriate system response to a frequency event. It is measured as MW of generation to be reduced or added by generators providing frequency regulation. Deviations in frequency, measured as Δf, from the system nominal 60 Hz are converted to ACE based on Equation 5.21 [94].

$$ACE = \frac{(667\,\Gamma)\Delta f}{10} \tag{5.21}$$

At maximum system load (maximal demand for electricity) in ERCOT, each 0.1 Hz deviation, either up or down, from the nominal frequency corresponds to an ACE of 667 MW. This ACE value indicates that 667 MW of generation is desired as a response to the frequency event. The original ERCOT ACE data from 2009 were nominally reported on four second intervals, but the timestamps were not spaced exactly four seconds apart. Further, on each day, the times were not exactly the same and some four second intervals were missed entirely. For example, the first value after 08:00:00.0000 reported on 27 March was at 4.31700 seconds, but on 11 April, it was at 0.32600 seconds, followed by a value reported at 8.18501 seconds. As with vehicle use data, the scenarios that define the stochastic parameter would be developed using hierarchical clustering in R, which would require that each day's ACE data were on uniform time intervals, with the same total number of values reported for each day. The original ACE data were thus simultaneously prepared for clustering and shifted onto the GAMS model's 15 second time basis using a customized nearest (temporal) neighbor placement approach.

The nearest neighbor adjustment approach began with several data preparation steps. To be consistent with other inputs to the SP model, all ACE data reported

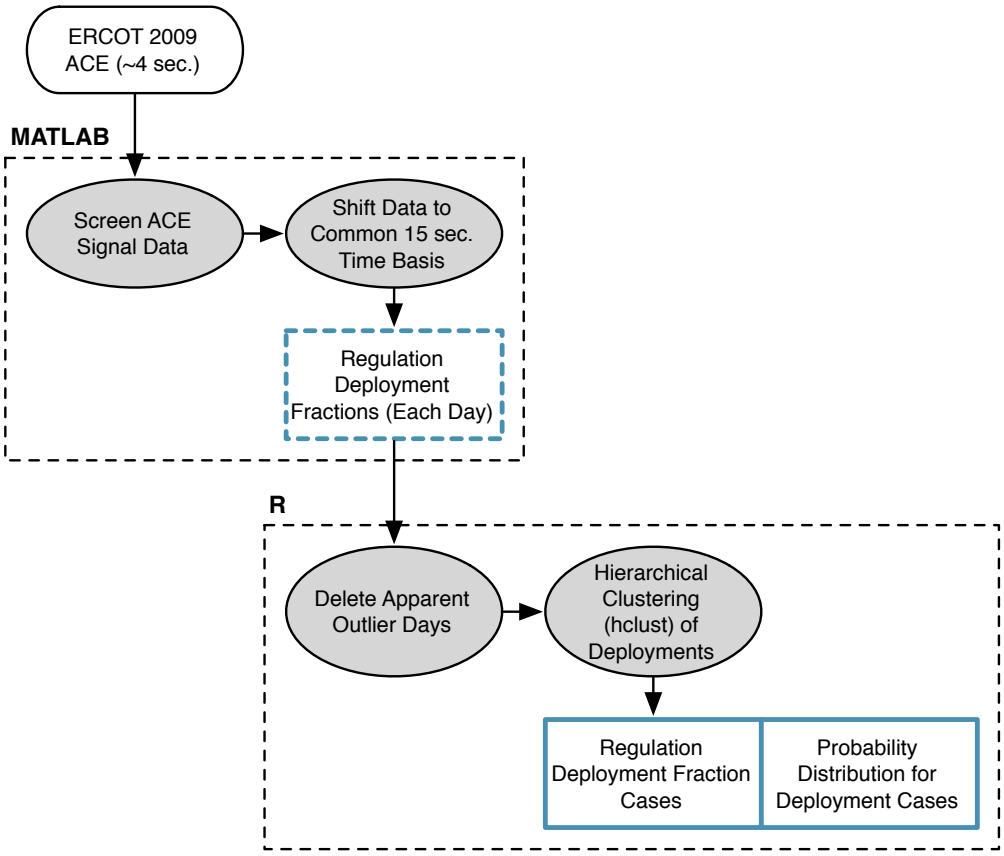

Figure 5.8: Regulation deployments are treated as stochastic parameters in the model formulation, similar to the vehicle use parameters. Accordingly, a similar procedure is used to develop the regulation deployment data for the SP model, where the data are first prepared in MATLAB and then some intermediate data are sent to R for hierarchical clustering into scenarios. Initial manipulation of the data in MATLAB involves removing weekend data and correcting for inconsistencies in the interval on which the data were originally reported. At the same time, the data are shifted onto the 15 second SP model interval. The data are then converted from deployment magnitudes to fractions, based on the minimum and maximum values reported on the given day, and passed to R for clustering. Two outlier days were removed from the input data before applying hierarchical clustering to divide the days into six scenarios. The final regulation deployment fraction scenarios are calculated as the cluster centers, or the average of the days included in each cluster, and the probability of each of these scenarios is determined based on the number of days included in the scenario.

on weekend days were removed. Combined date and time stamps were converted into separate date numbers (1–365) and timestamps (in seconds from midnight). Part of the way through the year, the reported times changed from seconds to fractions of a second, and these times were retained as originally reported, without rounding. With the dates prepared, each day's regulation values were modified using the custom nearest neighbor function. In this procedure, illustrated in Figure 5.9, "nearest neighbor" refers to the retention of data reported at the time nearest to each of the new, common 15 second time intervals. To determine the ACE value nearest each common time, the number of seconds between the time of each reported ACE value and every common time was calculated and stored in a "difference matrix." Because of the limited number of original and common times, the resulting difference matrix was sufficiently small that methods to minimize its size or avoid calculation of all of the differences to limit computational cost were not necessary.[5] In the difference matrix, original times ahead of a common time would appear as negative values, while original times behind common times would be positive. As a tie-breaker, it was decided that shifting values forward would be favored over shifting backward, thus all negative values in the difference matrix were multiplied by a factor of 1.5 and the absolute value of the matrix was taken. The minimum of the adjusted differences was then used to identify which of the original entries would be assigned to each of the common times.

Final assignment of ACE values to the common time basis first required verifying that no original data were assigned to two common times and correcting any such duplicate assignments. Figure 5.9 shows a duplicate assignment, where the original

[5] Computational time was reduced primarily by using a high-performance computing cluster. If such a system had not been available, using a lower-level language like Python would have yielded a significant decrease in computation time.

common times	0	15	30	45	60
original times 0	0	15	30	45	60
4	-4	11	26	41	56
7	-7	8	23	38	53
20	-20	-5	10	25	40
37	-37	-22	-7	8	23
41	-41	-26	-11	4	19
45	-45	-30	-15	0	15
49	-49	-34	-19	-4	11
53	-53	-38	-23	-8	7
58	-58	-43	-28	-13	2

	0	15	30	45	60
0	0	15	30	45	60
4	6	11	26	41	56
7	10.5	8	23	38	53
20	30	7.5	10	25	40
37	55.5	33	10.5	8	23
41	61.5	38	16.5	4	19
45	67.5	45	22.5	0	15
49	73.5	51	28.5	6	11
53	79.5	57	34.5	12	7
58	87	64.5	42	19.5	2

Figure 5.9: In the nearest neighbor procedure performed in MATLAB, for each weekday, a difference matrix is used to determine which original data are assigned to each of the common times. Times are measured in seconds from midnight. The original times here are simply for illustrative purposes. The difference matrix on the left shows the actual difference values calculated as the difference between the common times and the original times. On the right, the adjusted values are shown, where all the negative values have been multiplied by 1.5, and the absolute value of the matrix has been taken. These adjusted values are used to assign ACE values associated with the original times to each of the common times. Assignment is based on the minimum value in each column, indicated here with green boxes. The corresponding row indicates the appropriate time, and thus, the appropriate ACE value for each of the common times. In this case, one row is assigned to two columns. Such duplicate assignments are corrected during the assignment process.

time at 20 seconds is the selected as the best fit for the common times at both 15 seconds and 30 seconds. For each row in the difference matrix that is a best fit to more than one column, the column with the best fit to that row is identified, and all other columns matched to that row are identified for assignment to a new row (where rows correspond to original times and their associated ACE data). In the example in Figure 5.9, the original time of 20 seconds is a better fit to the common time at 15 seconds, so a new original time match must be found for the common time at 30 seconds. A smaller, temporary difference matrix is created for just the columns that need to be reassigned. In Figure 5.9, only the 30 second column would appear in the temporary difference matrix. From this temporary difference matrix, all rows that are within the discard limit (equal to the interval length) for each column (each common time) and that are not currently assigned to another column are identified. The difference between the common time and any rows that meet these requirements is calculated and the best fit row is assigned to that column. For the common time at 30 seconds in Figure 5.9, the only original time not currently assigned and within the discard limit of 15 is at 37 seconds. The ACE data originally reported at 37 seconds would thus be associated with the common time interval of 30 seconds. In cases where all the original data that satisfy the discard limit are already matched to other common times, a non-numeric (NaN) value is temporarily assigned to indicate that there is no viable match to that common time from the original data.

With the assignments completed, a few final preparatory steps were performed prior to clustering. Any non-mumeric or otherwise empty entries where no viable match was available needed to be filled. It was assumed that in these cases, the previous deployment reported persisted until a new command for regulation was sent. Finally, the reported regulation deployment magnitudes (ACE data) were converted to "fractional" deployments. All negative deployments (regulation down) were di-

vided by the minimum deployment for that day, while positive deployments were divided by the maximum deployment. Conversion to fractional deployments ensured that clustering would group days with similar diurnal deployment profiles or trends, rather than similarity in deployment magnitudes *and* trends. Deployments must also be in fractional form to be compatible with the model formulation presented in Section 5.1.2. The fractional deployments for the 261 weekdays were reduced to 258 days after removing three days that had three or more leading zeros, which indicated that no ACE data had been reported for at least the first full minute of that day.[6]

The number of days in the MATLAB output had to be reduced to manage the size, and thus the runtime, of the final SP model. The 258 individual weekdays of ACE data from 2009 prepared in MATLAB were combined into a set of regulation deployment scenarios using hierarchical clustering. As with the vehicle driving data, the clustering function *hclust* in R was used to identify which days would be placed in each cluster. Examining the clustering results for cluster counts ranging from 4 to 12 showed that days were well-distributed for either six or ten clusters. In the interest of limiting the size of the SP model, six clusters were used in the final formulation. To improve the quality of the clusters, two dates (June 29 and August 8) were screened from the data. Both of these weekdays were dissimilar from each other and the remaining data due to unusual deployments reported on those days. Again, cluster centers were calculated by averaging the ACE values reported for the days included in each cluster, thus creating a set of six frequency deployment scenarios. Unlike the vehicle travel data, which are quite similar on most weekdays, even with several separate clusters of regulation deployments, the averaging of ACE data used to create the cluster centers will tend to minimize extreme deployment events. It is

[6] Two of these three days did not have reported data for the first several hours of the day.

assumed that there exists a uniform probability of each of the days clustered, thus the probability of each cluster occurring was calculated as the number of days in each cluster divided by the total number of days (256).

5.3.4 Real-time electricity prices

Real-time electricity prices are handled as deterministic inputs to the model. Like day-ahead prices for frequency regulation, market clearing prices for electricity (MCPEs) are drawn from ERCOT in 2011. Price data were prepared for each of the dates shown in Table 5.2. A single price for each market settlement period was calculated by averaging prices from the four hubs, North, West, South, and Houston. These average prices were then converted from the 15 minute market settlement interval to the 15 second model interval by repeating each MCPE 60 times. Since the SP model runs one day at a time, these price data were paired with the corresponding day-ahead market prices, prepared following the procedure in Section 5.3.1, thus creating a separate set of prices for each day.

5.4 Stochastic model results

The GAMS model was run for eight fleet sizes N, between 1,000 and 5,000,000 and for each of the 26 dates identified in Table 5.2. The results from all of the dates were averaged together and converted to annualized quantities by multiplying the average revenue estimates by 365. The revenues shown in these results do not include the cost of vehicle charging included in the objective function, as it is assumed that these costs will be paid by the vehicle owner, independent of their participation in a V2G program.

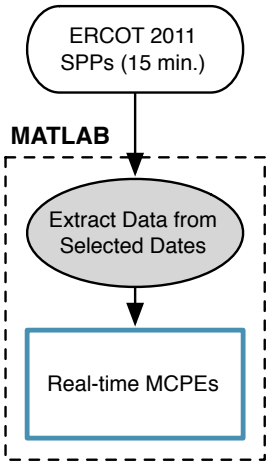

Figure 5.10: Real-time price data were drawn from ERCOT settlement point prices (SPPs) for the year 2011. Prices reported on the 26 dates in Table 5.2 were extracted and shifted to 15 second model intervals from the 15 minute market settlement interval by repeating each reported price 60 times. These prices were then passed to the GAMS model along with the corresponding day-ahead frequency regulation prices.

5.4.1 Effect of endogenous capacity prices

One of the primary objectives of this V2G model was to examine the effect of incorporating frequency regulation day-ahead market prices into the model such that changes in the offer quantities chosen by the vehicle aggregator would affect those prices. As noted in Chapter 2, most studies of grid-scale energy storage, regardless of the storage device, assume that the entry of that device in the market has no effect on capacity prices. While this assumption might reasonably be valid for small storage devices, a network of such devices, or a single larger facility, could affect prices [45, 47, 53]. This effect will be especially prominent in cases where the storage system participates in the market assuming negligible operating costs [45].

The effect of endogenous frequency regulation capacity prices ("endogenous day-ahead prices") on aggregator revenue is first analyzed excluding the influence of changes in the aggregator's retained revenue margin. Accordingly, the figures in

this section are drawn from the modeling case where $m = 1$, in which no payments are made to V2G program participants. Figure 5.11 compares annualized revenue between the exogenous and endogenous price models.

Because the V2G model is linear, annualized revenues in the exogenous price model are linearly dependent on fleet size. As the fleet size increases, so too does the aggregator's annual revenue. As expected, when the offer quantities affect day-ahead prices, annual revenue declines with increasing fleet size. In the endogenous price model, smaller fleet sizes show nearly linear increases in revenue with increasing fleet size because the aggregated storage resource is not large enough to maximize revenues during the day. The small difference in revenues between the exogenous and endogenous price models for these small fleet sizes can be attributed to the method detailed in Section 5.3.1.2 to convert the offer quantities and prices into concave piecewise linear revenue curves. As the PEV fleet grows, for the endogenous price model, the storage resource becomes large enough to capture the maximum available revenue in an increasing number of model intervals, and the gap between revenues in the exogenous and endogenous price models widens. For the largest fleet size tested, predicted annual revenue in the endogenous price model is 85% lower than in the exogenous price model.

The annualized revenues from Figure 5.11 can also be examined on a per vehicle basis, as shown in Figure 5.12. The linear relationship between fleet size and revenue means that on a per vehicle basis, revenues in the exogenous price model are constant. On the other hand, consistent with the total revenue trend in Figure 5.11, the endogenous price model yields sharply decreasing per vehicle revenue beyond the first three fleet sizes simulated. On a percentage basis, the difference in per vehicle revenues between the exogenous and endogenous price models are the same as the

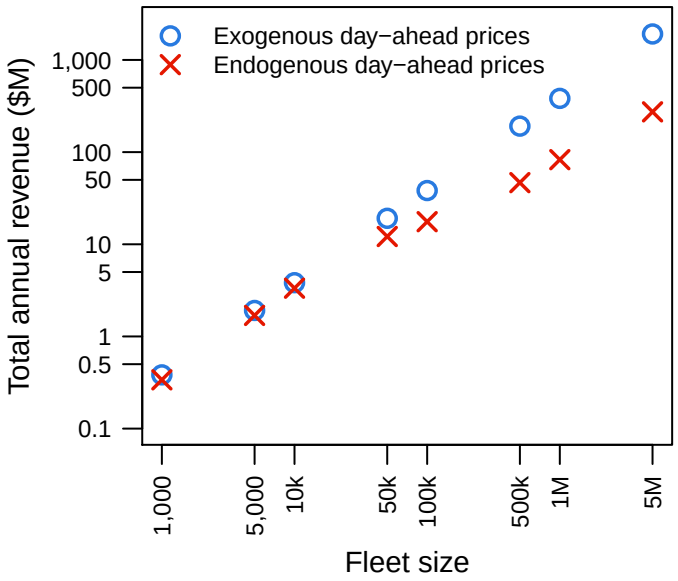

Figure 5.11: Consistent with the SP model formulation, with exogenous day-ahead prices for frequency regulation, annual aggregator revenue is a linear function of the number of vehicles in the program. On the other hand, for the endogenous price case, beyond a certain number of vehicles, the marginal benefit of additional vehicles declines steadily. While total revenue increases with any increase in fleet size (for the fleet sizes modeled here), the gap between the exogenous and endogenous models' annual revenue estimates widens steadily.

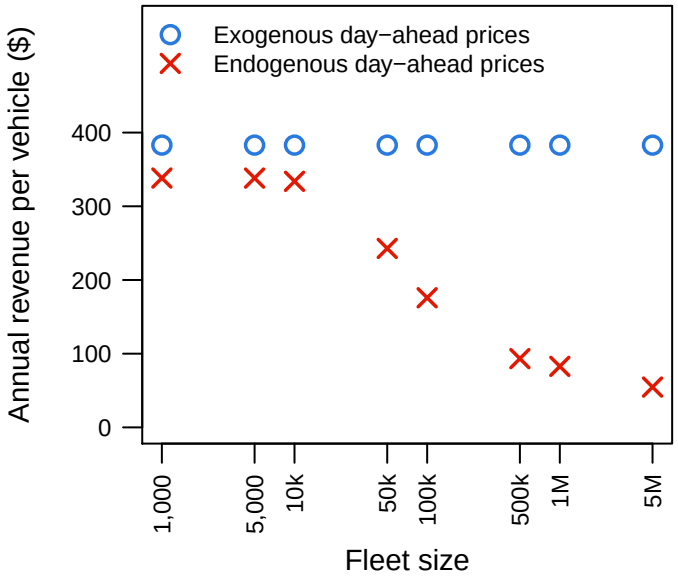

Figure 5.12: As seen in Figure 5.11, in the exogenous price model, annual revenue is a linear function of fleet size, thus the annual revenue *per vehicle* is constant. With endogenous day-ahead prices for frequency regulation, with a fleet size of only 10,000 vehicles, the effect of fleet size on revenue per vehicle is already observable. For fleets with more than 10,000 vehicles, the fleet size has a profound effect on per vehicle revenues, though the impact does diminish somewhat for PEV fleets with more than 1 million vehicles.

total revenues, thus for the largest fleet size tested, the per vehicle revenue declines by 85% when the effect of offer quantities on market prices is captured by the model. The effect of increasing fleet size on per vehicle revenues is important because revenue generated by the aggregator from its participation in the frequency regulation market must be sufficient to cover both its own expenses and any payments to V2G program participants. The per vehicle revenues are thus reasonable estimates of an upper bound on payments to PEV owners in a V2G program.

The impact of fleet size on V2G system revenue in the endogenous price model belies the changes in day-ahead offer quantity decisions made by the aggregator when subjected to offer quantity-dependent prices. As the number of vehicles increases, the

size of the aggregate storage resource increases. In the exogenous price model, these storage capacity increases lead to concomitant increases in revenues, as the aggregator increases the size of their offers to correspond to additional available capacity. When day-ahead regulation prices vary as a function of the offer quantity, as in the endogenous price model, in any given hour in the day-ahead market, once there is sufficient capacity for the aggregator to offer the revenue maximizing quantity, further increases in the fleet size will not change the offer quantity. As a result, for any fleet sizes above the optimal capacity for that hour, the quantity offered in the exogenous price model will be larger than in the endogenous price model. This effect is shown for both frequency regulation up and down in Figure 5.13. The smallest three fleet sizes show nearly identical offer quantities between the two price models, indicated in Figure 5.13 as a difference of almost zero. As the fleet size increases, the difference between the offer quantities grows, such that at very large fleet sizes, the regulation down offers in the endogenous model are quite small when compared to the exogenous model quantities. In Figure 5.13, the difference in offer quantities between the two models is larger for regulation down than it is for regulation up, which might be a result of the aggregator being required to pay for energy used to meet regulation down deployments in the real-time market. It could also indicate that offers in the day-ahead market for regulation down tend to be lower, which itself could be a result of any generator deployed for regulation down being required to absorb the loss of revenue from the energy market when they reduce their output.

5.4.2 Effect of revenue margins

With the effect of changes in fleet size established for both price models, the revenue margin parameter m, introduced in Section 5.1.2.3, was used to assess how a vehicle aggregator's decisions might change when some of the revenue generated

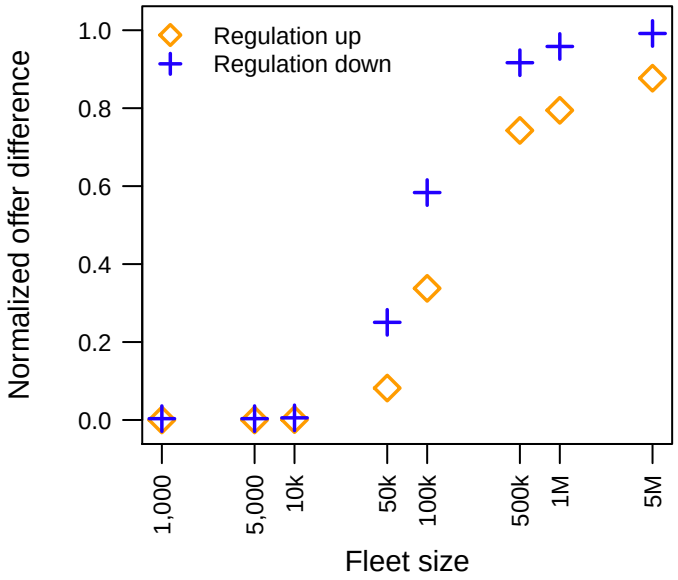

Figure 5.13: For small PEV fleet sizes, the difference between the regulation service quantity offered in the exogenous and endogenous price models is small because the full storage capacity available at at a given time t from the vehicle fleet can be offered without affecting day-ahead prices in the endogenous price model. As the fleet size increases, the entire available capacity cannot be offered without reducing the total revenue earned from the offer, thus the offered quantity is reduced relative to the exogenous price model (the difference between the two models increases). This effect, whereby the optimal quantity offered in the endogenous price model is reduced relative to the maximum available capacity at a given t, increases as fleet size increases, such that for the largest fleet sizes modeled, hardly any of the available capacity is offered. This effect is more pronounced for regulation down services, which might have occurred because the energy to provide regulation down must be paid for in the real-time market (a real aggregator would likely have program participants cover some or all of these costs).

in the frequency regulation market was not retained. These revenue margin tests simulate the effect that high internal operating costs or other external costs not directly captured in the model might have on the aggregator's market participation decisions. In particular, these external costs could represent the use of some revenues to pay program participants. Results from these simulations can also reveal whether the revenue impacts of fleet size might be insignificant relative to the bulk loss of revenue as the revenue margin parameter is reduced. Revenue margins between 0.4 and 1 in increments of 0.1 were examined. Lower margins indicate a greater portion of revenue is "lost" to outside expenses, as indicated in Equation 5.17. Revenue margins less than 0.4 were excluded under the assumption that the aggregator would retain at least that fraction of their revenues for operating expenses. Because the results from the revenue margin extrema of 0.4 and 1 bound the results from the intermediate cases, Figures 5.14 and 5.15 show the revenue range, where the lower edge of each bar indicates revenues for the $m = 0.4$ case.

Figure 5.14 is analogous to Figure 5.11, except that Figure 5.14 reflects the range of results from the revenue margin cases. As first noted in Figure 5.11, for the exogenous model, revenue is linearly increasing with fleet size for all of the revenue margin cases, thus on a percentage basis, the differential between the largest and smallest revenue margin cases is the same. Reducing the revenue margin parameter by 60% has the effect of reducing overall revenues by 69%. Revenues for the endogenous model also show a similar trend between revenue margin cases, but the percentage of revenue lost varies non-monotonically with fleet size. At smaller fleet sizes, revenue is reduced by approximately 71%, comparable to the exogenous model results, increasing slightly for the two intermediate fleet sizes, before declining to approximately 65% for the three highest fleet sizes. The impact of reductions in the revenue margin for

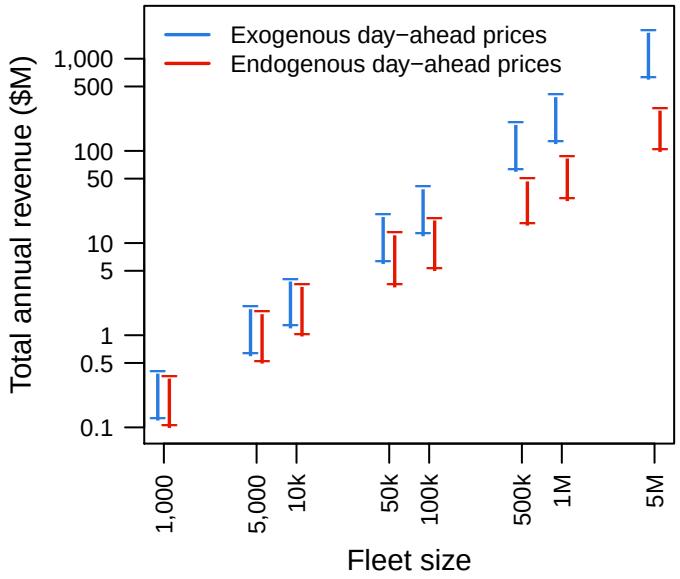

Figure 5.14: If the aggregator retains only a fraction m of its gross revenues, as the fraction decreases, so too does the annual revenue indicated in the model. The bars shown reflect the range of revenues for the indicated price model and fleet size. The top of the bars indicates the annual revenue for the given fleet size if all of the revenue is retained ($m = 1$) and the bottom of the bars indicates the revenue for the lowest margin case simulated ($m = 0.4$). For the smallest fleet sizes, both the magnitude and the change in revenue between the highest and lowest margin case are comparable between the exogenous and endogenous price models. As the fleet size increases, the revenue range becomes smaller for the endogenous price model than the exogenous price model, and as observed in Figure 5.11, the difference in revenue generated between the two models increases.

the endogenous price model are thus greatest at smaller fleet sizes, countering slightly the effect of fleet size on revenues in the endogenous price model.

The effect of revenue margins on per vehicle revenue, shown in Figure 5.15, further highlights the difference in revenues between the two price models. The magnitudes of revenues in both price models are comparable across the range of revenue margins m at small fleet sizes. As noted previously, the small difference between the price cases can be attributed to the conversion of day-ahead offer quantities and prices into concave piecewise constant revenue curves for the endogenous model. Consistent

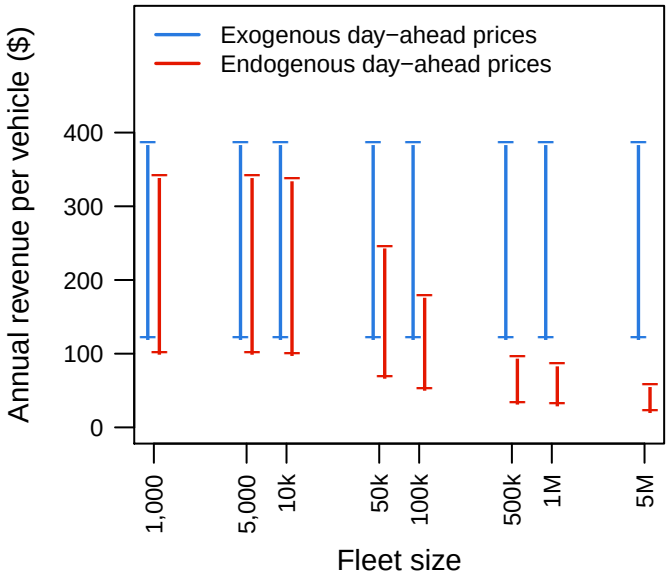

Figure 5.15: As in Figure 5.14, each of the bars indicates results over the revenue margin range $m = 0.4, \ldots, 1$. In the exogenous price model, reductions in the revenue margin have a dramatic effect on per vehicle revenue. The effect is similarly pronounced for smaller fleets in the endogenous price model. This is consistent with Figure 5.14, where the two price models show similar total revenue magnitudes and ranges for small fleet sizes. Also consistent with Figure 5.14, for the endogenous price model, larger fleet sizes show a much smaller change in total revenue as the margin decreases.

with the results in Figure 5.14, the revenues in the endogenous price model decrease with increasing revenue margin, but as a percentage of maximum revenue for a given fleet size, this effect is reduced somewhat in larger fleet sizes. Comparing the per vehicle revenue range between the two price models, at the largest fleet sizes, the highest anticipated revenue from the endogenous model, with $m = 1$, is lower than revenue in the exogenous price model at the lowest revenue margin case, $m = 0.4$.

The range of price models and revenue margin cases were tested to observe their effect on offer quantities, in addition to revenues. Figure 5.16 shows the average quantity offered in a single day-ahead market interval for both regulation up and down. As implied by Figures 5.11 and 5.14, offer quantities in the exogenous price

model increase linearly with increasing fleet size. In the endogenous price model, beyond the smallest fleet sizes, increases in the offer quantities are only observed for regulation up. In each case, decreases in the revenue margin lead to a decrease in the average offer quantity. Because of the method of plotting, the effect of changes in the revenue margin appears small for the exogenous price models, but the regulation up offer quantities are reduced by 23%, while the regulation down offers are reduced by 6%. Reductions in that range are observed for regulation down offer quantities in the endogenous price model as well. In contrast, regulation up offers in the endogenous price model are significantly reduced as the revenue margin decreases, up to a maximum of 74% for the largest fleet size.

Payments to V2G program participants can be readily estimated for the revenue margin cases simulated. In these compensation estimates, it is assumed that the diverted revenue is used exclusively to pay PEV owners in the V2G program. As the revenue margin decreases, the magnitude of diverted funds increases, thus increasing compensation. Figure 5.17 shows the effect of fleet size and revenue margin on participant compensation for the endogenous price model. For smaller fleet sizes, program participants could anticipate earning up to $210 each year. As more vehicles enter the program, potential compensation decreases, such that for programs with more than 100,000 participants, it is unlikely participants could expect annual compensation in excess of $100. For these large fleet sizes, it seems unlikely that payments from the aggregator will be sufficient to induce new participants to enter the program. Moreover, program participants can anticipate that their earnings will deteriorate until reaching a settling point where existing participants are satisfied with their compensation, but that compensation is insufficient to attract new members faster than the rate of replacement for current members.

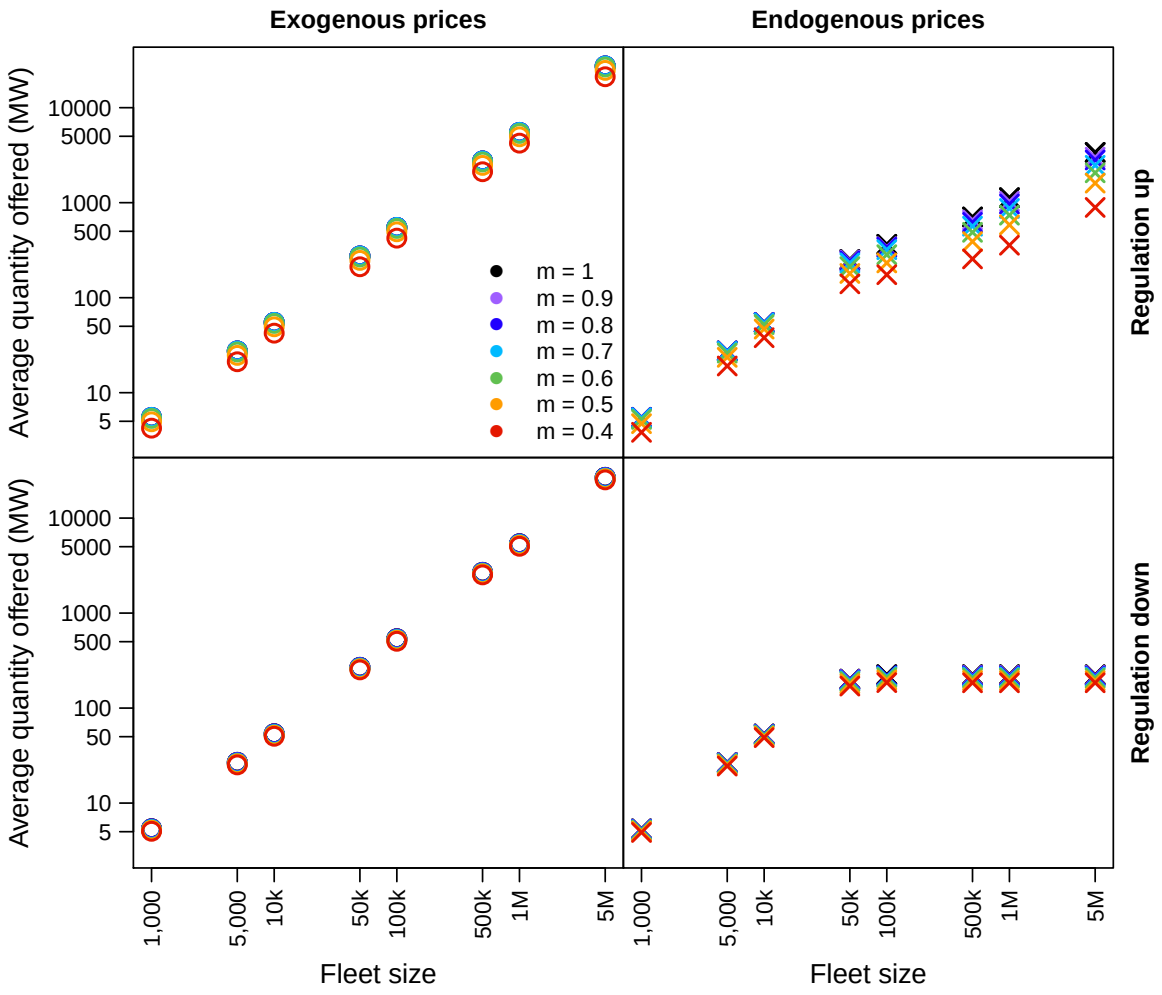

Figure 5.16: Changes in the offer quantities for regulation up as a function of fleet size are similar to the trends in overall revenue shown in Figure 5.14. Reductions in the revenue margin appear to have a significant effect for only the endogenous price model, where larger fleet sizes have more than half an order of magnitude difference between the quantity offered in the highest and lowest revenue margins. In the case of regulation down, for both price models, the effect of changes in the revenue margins appear to be smaller for regulation down than regulation up. As observed in Figure 5.13, regulation down has a lower apparent value than regulation up in the model, likely because the cost of electricity used to serve regulation down deployments in the real-time market is paid by the aggregator, thus limiting the growth in regulation down offers beyond a fleet size of 50,000 vehicles.

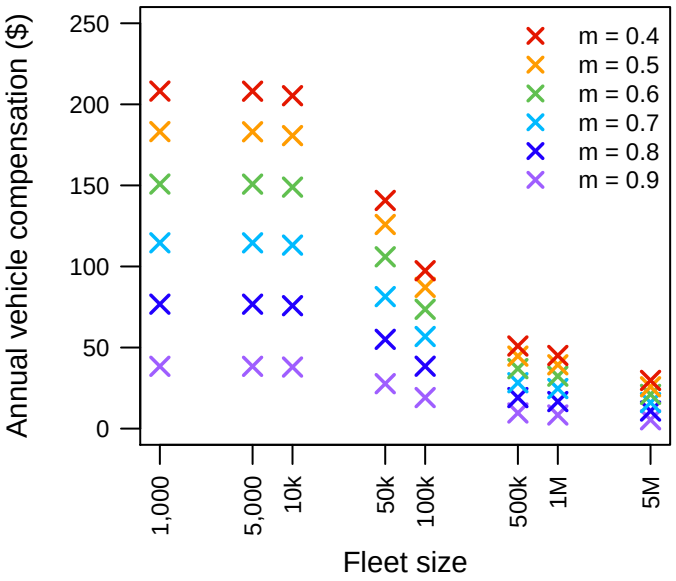

Figure 5.17: In the revenue margin cases simulated, it is assumed that earnings from market participation are transferred to V2G program participants. Thus as the revenue margin decreases, the earnings available to be paid to vehicle owners increases. For small fleet sizes, these decreases in the revenue margin have a significant impact on the annual compensation paid to vehicle owners. On the other hand, for larger fleet sizes, the quantity paid to vehicle owners, even in the most favorable scenarios modeled here, does not exceed $60 per year.

5.4.3 Vehicle parameter case studies

Results from the SP V2G model have been presented for a range of fleet sizes. These results show that when accounting for changes in day-ahead prices that arise from increasing PEV fleet size, aggregator revenue per vehicle declines considerably. As a consequence, potential compensation payments made to PEV owners for their participation in the V2G program necessarily decline. Because the model is structured such that the storage resource of all vehicles plugged into an EVSE at any given time is treated as a pooled resource, compensation payments appear uniform across the PEV fleet, but in a real-world V2G program, it is likely that participants would be compensated according to their ability to support the aggregator's needs. How often a vehicle is connected to an EVSE, extent of battery depletion from driving, vehicle battery size, and EVSE power rating will likely all affect compensation. Vehicle battery availability (to the aggregator) and vehicle use for driving are governed primarily by the daily needs and choices of individual drivers, and while they might be subject to incentives offered by the aggregator, it would be difficult to model the effect of those incentives directly, thus it is assumed that these parameters will be stochastic but broadly consistent with historical trends. Fleet-average battery size and EVSE power rating are dependent upon the equipment selected by PEV owners, but changes to these parameters, via equipment replacement or updates to product designs, will persist until new vehicles or EVSEs are purchased. Because of this persistence, and the relative ease with which changes in these parameters can be incentivized or legislated (as compared to vehicle owners' driving decisions), the sensitivity of the results to these two parameters is explored further. The simulations to support these case studies were performed only with the highest revenue margin, $m = 0.4$, as this yielded payments of sufficient size to vehicle owners in the base case

that changes in aggregator revenue and owner compensation should be clearly shown as the parameter of interest is changed.

5.4.3.1 Average PEV battery size

As battery prices fall and energy densities increase, battery capacities will likely increase to address concerns about the range limitations of current PEVs. Vehicles currently available have batteries ranging from 3.2 kWh to 85 kWh [95]. Four battery sizes were tested, 12, 18, 30, and 36 kWh, in addition to the 24 kWh size used in earlier results. These will reveal the general effect of changes in fleet-average battery size and, in particular, whether increases in the average battery size will increase revenues and compensation payments.

Results shown in Figure 5.18 for the five battery sizes tested shows that the battery size has little effect on aggregator revenue. In the exogenous price case, revenues increase only 6.5% between the smallest and largest battery sizes. The effect is more variable in the endogenous price case, where revenues rise by approximately 10% in most cases, but by 16-17% for the 50,000 and 100,000 vehicle fleet size cases. Trends favoring these intermediate fleet sizes don't appear in the other conditions examined. It is possible that there is some optimal battery size during certain times of day that can only be achieved with the appropriate number of batteries with the needed capacity and SOC, and these two cases have that optimal configuration.

Because aggregator revenue does not grow significantly with increasing battery size, payments to program participants, shown in Figure 5.19 also change little as battery size increases. As with Figure 5.17, only the payments in the endogenous price model are shown, since that model accurately reflects the effect that changes in the fleet size has on day-ahead market prices, and hence, aggregator revenue and participant payments. As the fleet size increases, the minimal benefit conferred by

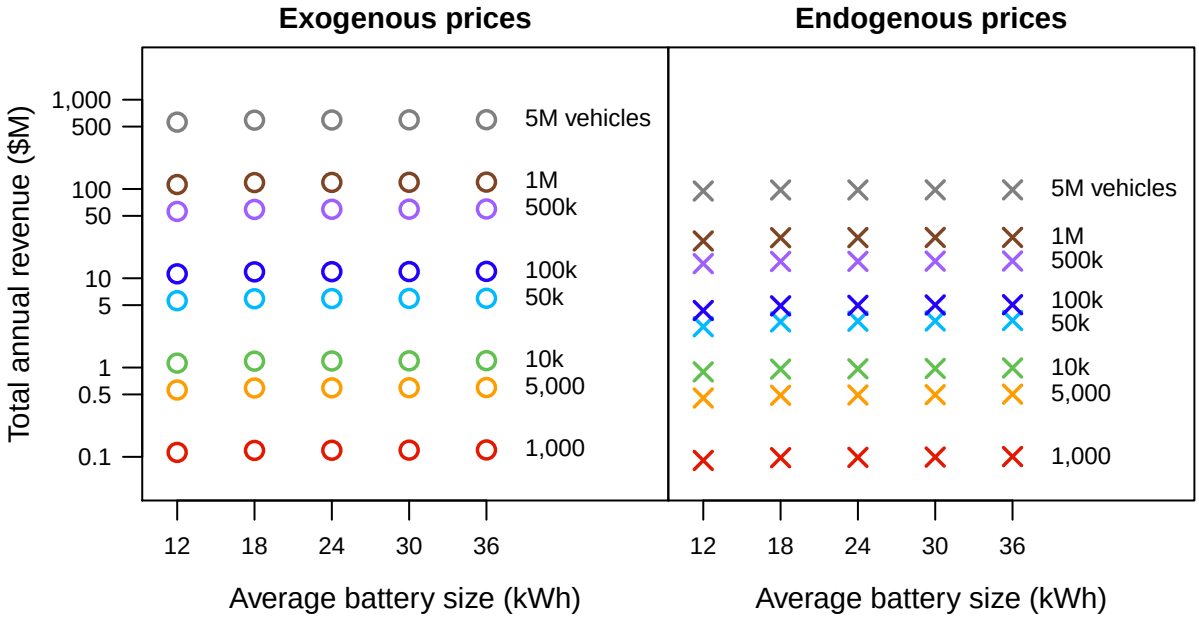

Figure 5.18: Varying battery sizes has a minimal effect on total aggregator revenue in either the exogenous or endogenous price models. The largest difference between the battery sizes examined appears with the endogenous price model in the intermediate fleet sizes. It is possible that this improvement for the 50,000 and 100,000 vehicle fleet sizes arises from some operational flexibility conferred by the larger batteries that can't be realized at the smaller fleet sizes. Other than those two fleet sizes, revenue does not increase by more than 11% with either price model.

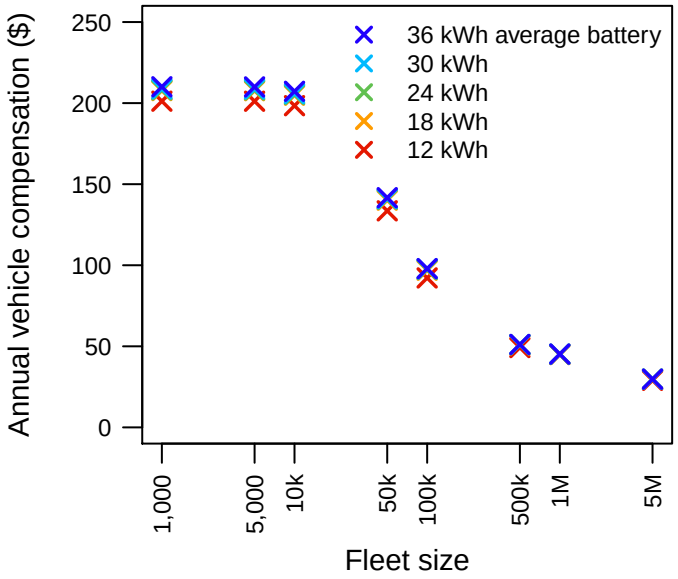

Figure 5.19: Figure 5.17 showed that for the endogenous price model, payments to program participants will decline dramatically with increasing fleet size. This trend persists across the range of battery sizes simulated. Increased battery size has a minimal effect on annual payments to vehicle owners. At large fleet sizes, the difference in payments among battery sizes nearly disappears.

larger batteries in PEVs is lost due to excess capacity available for all the battery sizes simulated.

5.4.3.2 EVSE power rating

Though increases in the average size of PEV batteries might yield minimal benefits for V2G aggregators or vehicle owners, larger batteries will likely motivate increases in EVSE ratings to prevent increases in recharge times. Even if PEV battery sizes do not change significantly, EVSEs might still move toward higher power ratings to reduce recharge times. Four new EVSE power ratings were tested, 2.2, 3.3, 4.4, and 10 kW. The nominal EVSE power rating for the preceding simulations was 5.5 kW. The lower power ratings were selected because 5.5 kW was perhaps optimistic in the near term, given that PHEV owners might be able to rely on only their onboard

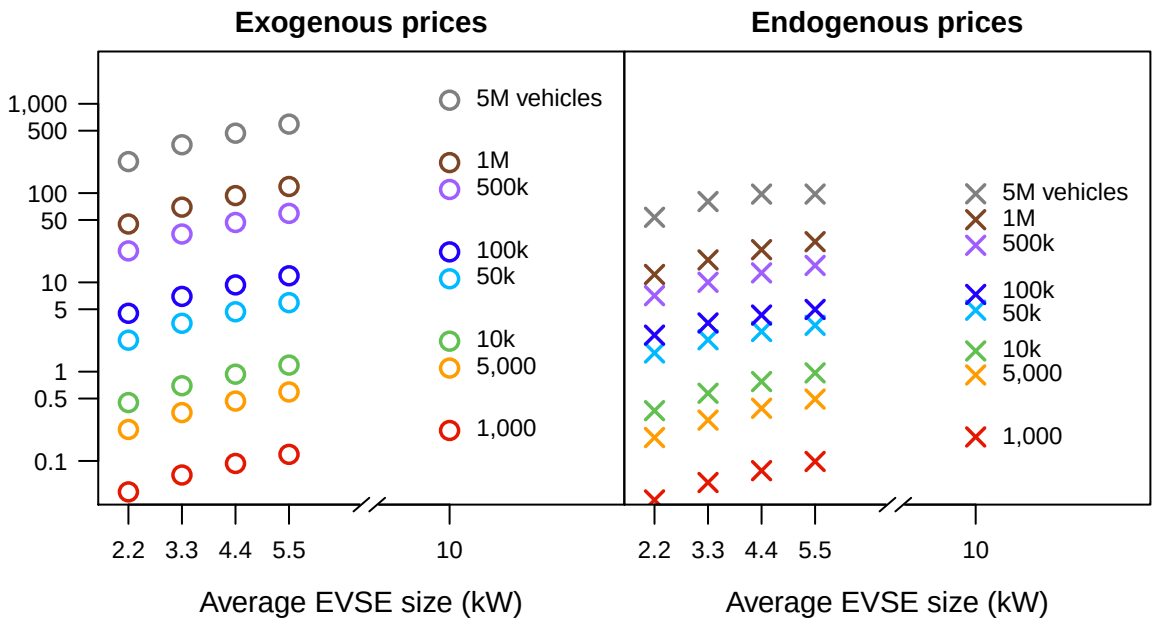

Figure 5.20: Increasing the average EVSE power rating from 2.2 kW to 10 kW increases aggregator revenue by almost half an order of magnitude in the exogenous price model and for the smaller fleet sizes with endogenous prices. The revenue trend observed in Figure 5.14 for the endogenous price model remains, but with increasing fleet size, the benefits diminish more for larger EVSE sizes, such that the revenue difference between EVSE sizes shrinks considerably at the largest fleet sizes simulated.

120 V charger, while the higher power rating was tested to observe the effect of the introduction of high power home EVSEs for vehicles with large batteries. All of these cases assume an average of 24 kWh for PEV batteries.

Results in Figure 5.20 show that EVSE size has a significant effect on aggregator revenues. Larger EVSEs enable the aggregator to take advantage of more revenue opportunities, particularly with a smaller vehicle fleet, suggesting that power, and not SOC (energy), is the primary limiting factor for revenue. These results are consistent with the findings of both White and Zhang [47] and Quinn et al. [62] that increasing the EVSE power capacity can increase total revenues significantly. In the endogenous price model, eventually the PEV fleet becomes sufficiently large that neither

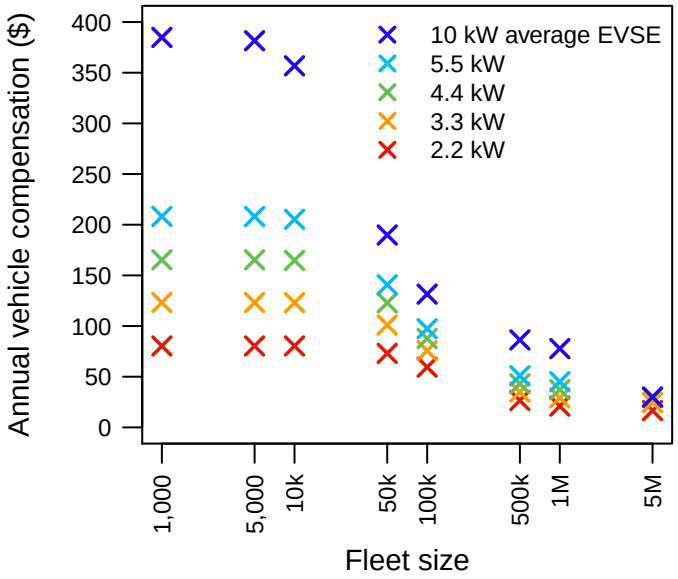

Figure 5.21: EVSEs with higher power ratings can significantly increase total revenue, which in turn increases the revenue available for compensatory payments to program participants. The increase in payments between the 5.5 kW and 10 kW cases appears large because there is a difference of 4.4 kW capability between these two cases, whereas the other cases are all 1.1 kW apart. This difference diminishes considerably beyond the three smallest fleet sizes tested, such that in the 5 million vehicle scenario, consistent with Figure 5.20, larger EVSEs hardly increase participant payments.

available energy nor power are limitations to market participation, at which point additional revenue opportunities become quite limited. This phenomenon is evident in the greater rise in revenue, over the fleet sizes examined, for the 2.2 kW EVSE than the 10 kW EVSE. The differential between the two cases in the endogenous price model is 408% with the smallest PEV fleet examined, but only 83% with a simulated fleet of 5 million vehicles. The 2.2 kW case closes the revenue gap on the larger EVSEs simulated when the fleet is large enough that power is not a limitation, in spite of the EVSEs lower power rating.

As observed with revenue from the endogenous price model shown in Figure 5.20, payments to program participants rise significantly with increasing EVSE

power rating, but the benefits are primarily realized for small fleet sizes. With larger EVSEs, available revenue opportunities are captured with smaller fleet sizes. Beyond the saturation point, payments decline precipitously, eventually reaching near parity with the smaller EVSE cases. While larger EVSEs are beneficial to the aggregator and vehicle owner for small fleet sizes, it is likely that these more powerful EVSEs won't become commonplace until much later, at which point the pool of candidate PEVs for a V2G program in a given region would likely exceed 10,000 vehicles. Further, the introduction of large EVSEs pose even greater challenges at the distribution level than is anticipated from near term growth in home EVSE infrastructure. As it is, an EVSE is typically the largest standalone power draw in a home. Increasing the power flow between a home and the substation in either direction by an average of 10 kW would probably require extensive distribution system upgrades in some of the fleet size scenarios envisioned, where several homes on a single feeder line might have such EVSEs in their homes.

5.5 Conclusions

In this chapter, a two-stage SP model of a vehicle aggregator was created. The structure of the model allowed the aggregator to offer frequency regulation services in the day-ahead market and observe the real-time operation of their system subject to regulation deployments and the use of vehicles in their program. The aggregator is subject to penalties if it cannot satisfy its day-ahead commitments to provide frequency regulation in the real-time market, but faces uncertainty about vehicle availability and SOC in real-time conditions. For the penalty imposed in the model, the aggregator chooses to offer frequency regulation services, even with the uncertainty in its available capacity to meet real-time deployments. The revenue opportunities presented by the frequency regulation market generate revenues sufficient

to promote participation, and estimated per-vehicle earnings approached $400. The model was structured to facilitate the comparison of two different frequency regulation pricing regimes. The "exogenous price model" assumed that prices were not a function of the offers made by the vehicle aggregator in the day-ahead market, and the "endogenous price model" adjusted revenue as a function of the offer quantity selected by the aggregator in each day-ahead market interval. Eight different PEV fleet sizes were tested, ranging from 10,000 to 5 million vehicles. With a small PEV fleet, the results from both price models were comparable, and the results compare well with other results in the literature. As fleet size increased, per-vehicle revenues in the endogenous price model declined dramatically, indicating that the market for frequency regulation is not large enough to absorb available capacity for those fleet sizes without experiencing a significant drop in capacity prices. Further, these results show that ignoring the effect of aggregator offers in the day-ahead market when modeling a V2G system could lead to grossly overestimated revenues. For a V2G aggregator with 100,000 vehicles in their system, current conditions in the frequency regulation market in ERCOT suggest that the aggregator likely wouldn't be able to compensate its members more than $100 per year.

Initial results generated from the SP model assumed constant values for several vehicle parameters. Following the analysis of the exogenous and endogenous day-ahead pricing model structures, case studies were executed to determine the influence of two of these parameters — EVSE power rating and PEV battery size — on revenues. A total of five battery sizes and EVSE ratings were tested. Increases in battery size increased revenues, but even with a spread of 24 kWh between the smallest and largest battery size modeled, the change in revenues was between 3 and 17%. Increasing EVSE size also increases revenues, but the range of revenues is a stronger function of the fleet size. For the largest fleet sizes examined in this work,

the absolute differential in compensation is still small. Generally, these results show that while V2G revenues are sensitive to PEV battery size and EVSE power rating, their influence on revenues are insufficient to support a vehicle aggregator making any programmatic or policy change to incentivize participants purchasing vehicles with larger batteries or installing uprated EVSEs.

Chapter 6

Decision analyses of plug-in electric vehicle owners in a voluntary vehicle-to-grid participation paradigm

In the United States, with the exception of a few urban areas where vehicle ownership is impractical due to space constraints, pervasive refueling infrastructure and short refueling times mean that ICE vehicles afford owners unparalleled mobility. In comparison, BEVs offer a much shorter cruising range, and the lengthy required recharge times mean that long or unexpected trips might not be possible, depending on the owner's prior driving and recharging choices. BEVs have lower operating costs than their ICE counterparts, but the limited range and comparable purchase prices are seen as a barrier to adoption by many vehicle owners.

The range limitations posed by BEVs are overcome by PHEVs, but the on-board range-extending ICE adds cost and complexity, and PHEVs have shorter electric driving ranges than BEVs. The inclusion of an ICE in PHEVs can be viewed as an option that allows a vehicle owner to always be able to take a trip, instead of being range constrained and potentially forced to find alternate means to complete a desired trip. The option price, p_o, is the difference in purchase price between a comparable BEV and PHEV. Accordingly, the option would be exercised if a situation presented itself where, without the ICE, the vehicle owner would not be able to complete their desired trip, and thus the exercise price, p_e, is the cost of the gasoline required to complete that trip. These differences in purchase price, cost of ownership, and vehicle

use flexibility between BEVs and PHEVs could influence a vehicle owner's experience participating in a V2G program.

There are several decisions that a PEV owner or prospective owner might consider with respect to a V2G program. For a PEV owner faced with the opportunity to enter a V2G program, the potential payments for their participation could be enticing. There is some possibility that because of the aggregator's market participation decisions, a participant in a V2G program might be unexpectedly left with less charge than they anticipated. A vehicle plugged in with a partially discharged battery and subject to V2G operations could be fully discharged before being recharged again by the aggregator. It would be difficult for an aggregator to guarantee more than a preset minimum state-of-charge to participants if they had no foreknowledge of their participants' travel intentions. Vehicle owners and the aggregator could avoid these conditions with vehicle use pre-scheduling, but it might be difficult for an aggregator to provide incentives sufficient to offset the inconvenience of advance travel planning. Given the increased probability of being unable to complete future travel due to the aggregator's use of a vehicle's battery, the decision to participate in a V2G program should balance any potential compensation against this risk. A battery with insufficient range for upcoming travel is a particular problem for a BEV driver. With a PHEV, an empty battery simply invokes the use of the ICE engine to complete any travel. The probabilities of a BEV owner being stranded or a PHEV owner being required to use the range-extending ICE are different, and thus the differences between the vehicle types will affect the experience of participating in a V2G program. Vehicle type is thus another important decision with respect to the interaction between a V2G system and PEV owners.

6.1 Preparation for decision analyses

Decision analysis methods can be used to assess the decisions faced by individuals considering whether to enter a V2G program or the appropriate vehicle type for participation. The results from the GAMS model of a V2G aggregator in Chapter 5 can be used to describe the operating conditions created by the aggregator and experienced by vehicle owners. Using these results to assess the probabilities relevant to the decisions, evaluating a range of values for the remaining parameters can reveal the conditions under which joining a V2G program is reasonable.

For the purposes of the forthcoming analysis, it is assumed that all prospective V2G program participants are risk neutral "deltapeople" [70]. Let there exist an individual who will accept either a sure quantity q or a deal with $i \subset 1, \ldots, n$ outcomes q_i of probability p_i, as shown on the left side of Figure 6.1. Let the certain quantity then be increased by some quantity δ. The individual decision-maker is considered a "deltaperson" if they would only exchange the new certain equivalent $q + \delta$ for an updated deal where each of the outcomes q_i are also increased by the same δ (for the range of payoffs $q_i + \delta$), illustrated on the right side of Figure 6.1. This property requires that the slope of the curve quantifying the individual's risk attitude be invariant with respect to the magnitudes of the outcomes. Only linear and exponential risk curves satisfy this requirement. Since the V2G program participants are also assumed to be risk neutral, their underlying risk curve is linear. Risk neutrality necessarily satisfies the delta property. Assuming that program participants are risk neutral also allows the direct commingling of assessed probabilities and estimated dollar values ascribed to particular outcomes without first translating those monetary measures through each individual's risk curve. These assumptions greatly simplify the decision assessment calculations in Sections 6.2 and 6.3.

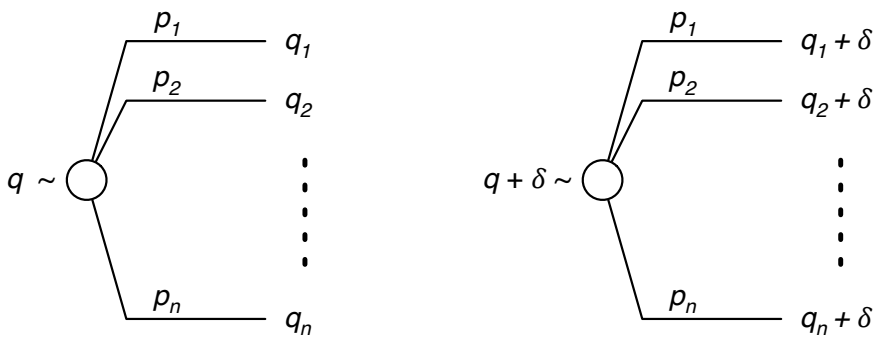

Figure 6.1: An individual faced with a choice between certain acquisition of some wealth (monetary or otherwise) and a deal between several possible outcomes and associated probabilities can define the deal and/or the certain quantity such that they would be willing to accept either. The equivalence relationship between these two possible conditions is shown on the left. Given that they have achieved such substitutability between the certain quantity and the deal, if the certain quantity is increased in value by some amount δ, the individual is a deltaperson (satisfies the conditions of the delta property) if the deal they consider equivalent to this new certain quantity is one in which each outcome has a value that is also increased by δ. This condition is shown on the right. While being risk neutral is not a condition of being a deltaperson, risk neutrality satisfies the requirements of the delta property. Adapted from Howard [96].

Estimating the probabilities p, p', q, and q' is pivotal to completing the decision analyses. These probabilities were calculated based on the state-of-charge from the V2G stochastic optimization model described in Chapter 5 and the vehicle travel clusters developed as inputs to that model. There were 1,456 revenue margin cases examined with that model. To limit the extent of the modeling effort required to generate the needed probabilities, only the smallest revenue margin ($m = 0.4$) results were analyzed, as these results represent the largest potential payments to vehicle owners from the margin cases. The probabilities were assessed using Python, and were initially developed as separate probabilities for each of the 208 applicable model cases (26 dates and 8 fleet sizes). For the GAMS model cases, there were 11 vehicle travel clusters. Each of the 80 days of travel data used in the clustering analysis is a member of one of those clusters. Because there were six regulation deployment clusters, each of the travel data clusters is associated with six scenarios of results from the GAMS model. Each of the 80 days was analyzed individually with the appropriate set of scenarios for that day's cluster membership. The impact of V2G system operation or only vehicle charging for the given GAMS results case was assessed for each of the 436 vehicles used during those days. Vehicle charging was determined based on charging decisions made by the vehicle aggregator. Sometimes more charge was available than was needed for a single vehicle, thus the available charge in an interval, v_t, was restricted to no more than the (depleted) state-of-charge at that time, s_t, using Equation 6.1. The $[(s_t+v_t) > 0]$ term acts as a logical operator to limit the application of Equation 6.1 to only those vehicles that could be over-charged in interval t. The charging provided to each vehicle in interval t is then added to that and all subsequent intervals for the day, as defined by t^*.

$$s_{t^*} = s_{t^*} + (s_t + v_t)[(s_t + v_t) > 0] - s_t \quad \text{where } t^* \in t, \ldots, T \tag{6.1}$$

For vehicle use and charging with and without V2G system operation, the range of state-of-charge for each vehicle was assessed. The total number of vehicles with insufficient charge to complete their day's travel with and without V2G were calculated for each of the days in the vehicle data clusters. Both vehicle counts were divided by the number of vehicles in use on that day, yielding the probability of a vehicle having insufficient charge on that day. These probabilities were multiplied by the probability weights from the frequency regulation deployment clusters that correspond to the six scenarios and summed to get a single probability estimate for each day. The average of the probabilities across all the days were recorded for the two operating conditions — with and without V2G — as the elemental probabilities for the given GAMS model results. This procedure yielded a total of 208 probabilities for each of the two operating conditions.

Probabilities calculated using this methodology were compared for several battery sizes, since forthcoming decision cases consider PHEVs and BEVs, which typically have markedly different battery capacities. For a single operating condition, either with or without V2G participation, probabilities varied little among the 208 fleet size and date combinations. For a single fleet size, there was no apparent trend in probabilities from month-to-month, or evidence of a consistent seasonal trend. Similarly, there was no consistent change in the estimated probabilities with respect to fleet size. The primary differentiation in the probabilities was between operating conditions and with batteries of differing sizes. Table 6.1 shows, for four battery sizes and the two operating conditions, an average of the 208 estimated probabilities calculated from the Python code. The largest battery shown, 27.2 kWh, corresponds to 80 miles of range, assuming an energy conversion rate of 0.34 kWh/mile. Most drivers travel less than 80 miles a day, especially when accounting for additional range provided by charging available during the day, thus the probability of having insuf-

Table 6.1: The probability of being without sufficient charge to complete a day's planned travel was assessed for the vehicles in the PSRC study across the 80 days of vehicle travel data included in the clusters used as an input to the V2G model, as described in Chapter 5. Probabilities were estimated for two operating cases, with V2G operations and without. Each of the dates and fleet sizes analyzed with the V2G model had a probability value calculated for both operating conditions, for a total of 208 probabilities per condition. The probabilities tabulated are averages of the calculated values. With a large battery, the probability of not being able to complete future travel due to insufficient charge is low. At an energy conversion rate of 0.34 kWh/mile, a 27.2 kWh battery provides a range of 80 miles. Decreases in battery size yield non-linear increases in the probability, and the non-linearity increases when V2G operations are introduced.

	Probabilities	
Battery Size (kWh)	Without V2G	With V2G
8.5	0.479	0.589
12	0.293	0.384
24	0.045	0.069
27.2	0.027	0.046

ficient charge is less than 3%. As battery size decreases, the probability of having insufficient charge increases non-linearly, such that a vehicle with a 25 mile range has a nearly 50% probability of depleting its battery before completing the day's travel. This non-linearity likely arises because of the distribution of daily travel distances, where the total miles traveled in a day, even when accounting for available charging, is between the range of a 12 and 24 kWh battery, or approximately 35 and 70 miles of driving range, respectively. This effect is compounded by the introduction of V2G operations. The activity of the V2G aggregator increases probabilities regardless of battery size. On a percentage basis, the increases are smaller with smaller batteries, but the non-linear trend in the probabilities is more pronounced with V2G operations.

In a multi-vehicle household, if one of the vehicles is replaced with a PHEV or BEV, the members of that household might be able to switch vehicles to accommodate

the anticipated needs of each driver. This switching could take place on a per-day basis, where vehicles are selected based on driving needs at the beginning of each day, or on a per-tour basis, if drivers in a household have travel patterns that facilitate switching during the day. These probabilities do not account for the possibility of exchanging travel between vehicles in a household.

6.2 Participation in a V2G program

The prospect of compensation for ancillary service provision through a V2G program would likely encourage many PEV owners to consider participating in such a system. If a vehicle owner does not participate in a V2G program, with knowledge of their vehicle's last reported range (and approximate charge rate, if the vehicle is plugged in), they should be able to estimate the vehicle's state-of-charge and feasible range before they begin another trip. For owners of range-limited BEVs, it is likely that with experience they would learn to avoid trips they could not comfortably complete, thus minimizing the probability, p, of not being able to complete a trip. On the other hand, for vehicle owners who have chosen to enter a V2G program, the vehicle aggregator might use their PEV to serve ancillary service deployments while the vehicle is parked, and thus they could not reasonably expect any particular state-of-charge prior to their next trip. As a consequence, there is a higher probability, p', that a given vehicle owner will be unable to complete their desired trip.

The tree representing the decision of whether to participate in a V2G program, specifically for a BEV owner, is shown in Figure 6.2. In each case where sufficient charge remains in the vehicle battery to complete a desired trip (denoted in Figure 6.2 by "Suff. Charge"), the trip will be completed, where the value of the trip is given by variable v. Compensation paid by the V2G provider for the use of a vehicle's

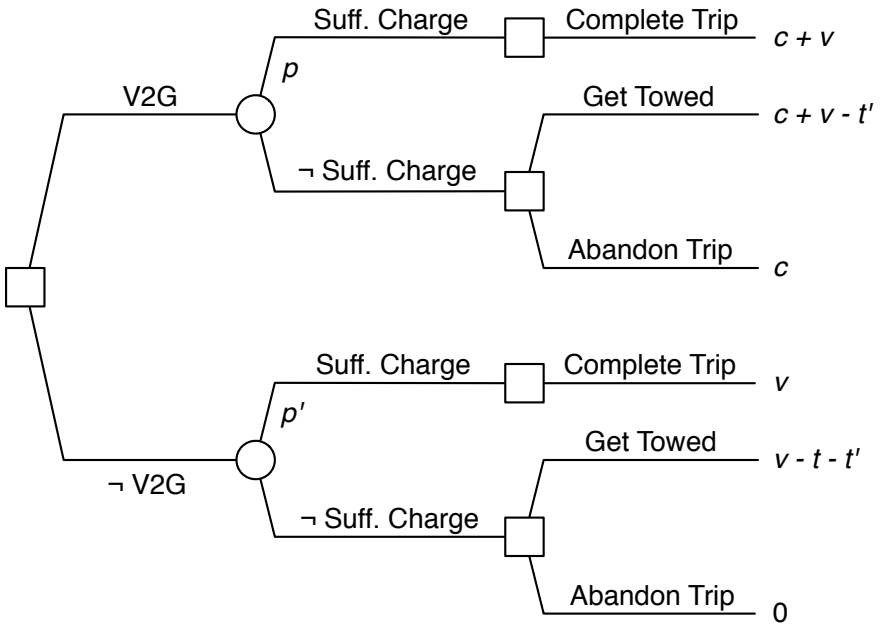

Figure 6.2: This decision tree shows the primary alternatives of whether or not to participate in a V2G program and the follow-on consequences with respect to one's ability to reliably complete future trips on-demand. As shown, it is assumed that the driver has a BEV. The endpoint probabilities and costs can be readily applied to facilitate a decision between the two primary alternatives.

battery is denoted by c. If insufficient charge is present in the vehicle's battery to complete the desired trip, assuming the vehicle is a BEV, an individual can either abandon the trip, which could simply mean getting a ride with a friend or choosing an alternate mode to reach their destination, or they can have their vehicle towed, incurring towing cost t and inconvenience cost t'. It is assumed that a V2G provider will cover any towing expenses for participants. Having established the probabilities p and p' of being unable to complete trips as a function of whether or not an individual is a member of a V2G program, the anticipated values associated with each of the outcomes in the decision tree can be combined with the endpoint probabilities. The decision is thus framed in a quantitative form.

The decision tree in Figure 6.2 can be simplified, eliminating some of the decisions shown, by making a few assumptions. The value of the resulting equations can then be calculated, where the "best" choice is whichever has the highest expected value. Figure 6.2 shows a decision where, upon learning that their BEV does not have enough charge to complete their planned travel, the driver might choose to abandon the trip. For the purposes of this analysis, it is assumed that decision to travel has already been made, and that having insufficient charge is not a deterrent to completing the trip by some other means. While it is certainly possible that an extraneous trip for casual purposes might be abandoned, it is assumed in this analysis that travel of any significant distance is of sufficient importance that it will not be readily abandoned, thus this second stage decision has been removed from the tree. From the perspective of sociological decision theory, the decision to travel discretized in Figure 6.2 might not be an explicit decision as much as a consequence following from ancillary decisions regarding social interactions or assumed responsibilities [69].

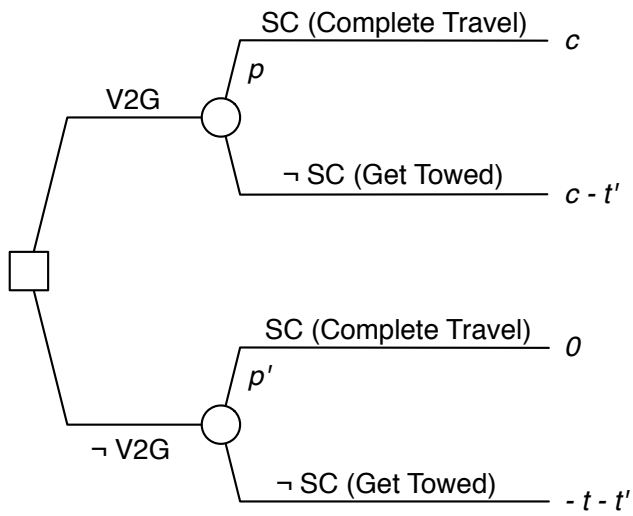

Figure 6.3: By assuming that travel decisions are committed before knowing the state-of-charge of the PEV, the "abandon trip" option in Figure 6.2 is removed, simplifying the decision tree. All of the remaining outcomes involve completion of any intended travel, thus including the value of the completed trip, v, it can be removed from the decision as well. The label "SC" denotes "sufficient charge."

As a result of removing the possible outcome of not completing the intended travel, the remaining outcomes all included the value of the travel, v. This parameter accounts only for the value inherent in being able to complete the intended travel, not any reductions in value associated with delays in reaching the planned destination, which are captured by the inconvenience penalty parameter t'. Because the trip value term appears in every outcome, it uniformly changes the magnitude of all the potential outcomes and can thus be ignored. The resulting reduced decision tree is shown in Figure 6.3.

Each of the legs of the decision tree in Figure 6.3 can be converted to an equation. The value of the resulting Equations 6.2 and 6.3 can be compared to determine the best choice. In those equations, the probabilities are known, following the procedure detailed in Section 6.1. The quantities used in these equations apply to

a single day because the decision of whether or not to participate in a V2G program and the potential accompanying compensation payments c are not associated with single trips. While it is unlikely that a PEV owner could decide to change their membership in a V2G program from day to day, the per-year quantities for the known parameters would be annualized from single day simulation results, thus the preferred outcome would be unchanged whether the time scale were daily or annual. There are two unknown parameters in Equations 6.2 and 6.3: the inconvenience penalty t' and the cost of towing or alternative transportation t. Several values are tested for both parameters to observe how they affect the decision to participate in a V2G program for a BEV driver. The experience for a PHEV owner is somewhat different, and is discussed in Section 6.3.

$$pc + (1 - p)(c - t') \tag{6.2}$$

$$-(1 - p')(t' + t) \tag{6.3}$$

Results

Participation in a V2G program provides a BEV owner an opportunity to earn money in exchange for making their vehicle's battery available to the aggregator, but exposes a BEV owner to an increased probability of being left without sufficient charge to complete their planned travel. The participation decision is represented quantitatively by Equations 6.2 and 6.3. The probabilities p and p' were assessed using the methodology described in Section 6.1, and compensation is based on the quantities developed in Chapter 5. The remaining unknowns in the decision equations are the travel disruption penalty, t' and the travel completion cost, t. Creating an equality relation between these two equations and rearranging the terms yields formulas that can be plotted to explore the relationship between unknown parameters. The two unknown parameters in this decision are compared in Figure 6.4.

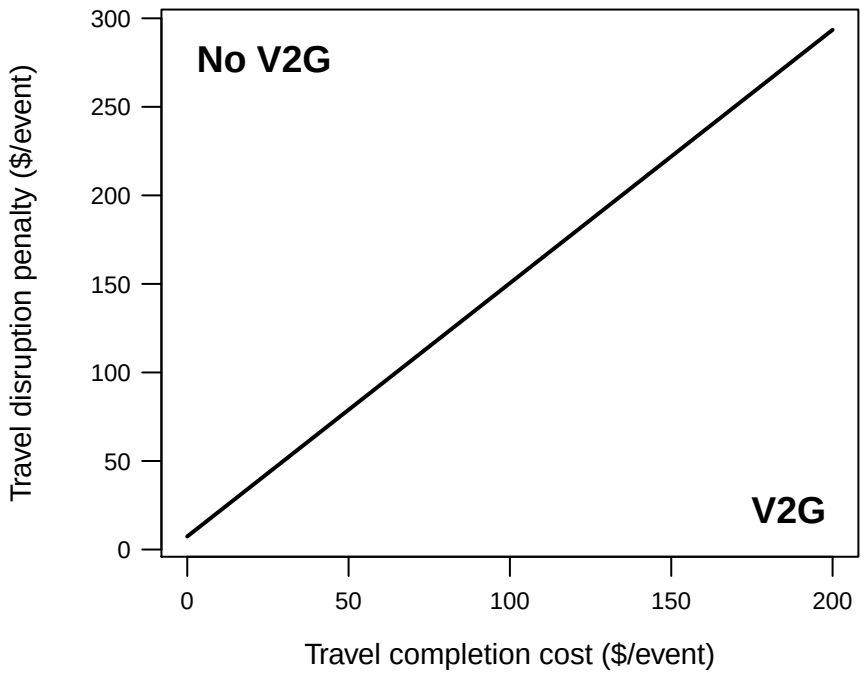

Figure 6.4: Comparing the effects of the two unknown parameters in the decision to participate in a V2G program, assuming a PEV fleet of 500,000 vehicles, yields the relationship shown. When travel completion costs (e.g. towing, changing travel modes) are high or the driver places a value on the disruption to their travel plans less than the cost to complete them via an alternate mode, participating in a V2G program is preferred. Otherwise, payments made to the vehicle owner by the aggregator are insufficient to balance the additional certainty afforded by not participating in a V2G program.

The payments made by a vehicle aggregator to drivers participating in their V2G program would need to be sufficient to offset the increased probability of having insufficient charge to meet their travel needs. This compensation also implicitly includes the avoided costs otherwise faced by the vehicle owner to get their vehicle towed or otherwise transport themselves to their destination in situations where their battery is depleted prior to the conclusion of their day's travel. Figure 6.4 shows that the payments and the avoided costs, represented by the added expenses, t, for travel completion faced by the non-participant, are insufficient to always offset the travel disruption penalty. For V2G to be favorable, penalties associated with travel delay must be limited relative to the travel completion cost. Given that the revenue generated by the aggregator is limited, compensation thresholds for participation in a V2G program might provide a guideline to program viability.

For a constant travel completion cost, t, the compensation required for an individual BEV owner to join a V2G program can be determined from their assessed travel disruption penalty. The relationship between the penalty and annual vehicle compensation required is shown in Figure 6.5 for three different travel completion costs: \$50, \$100, and \$150. Due to the increased probability that a vehicle owner will be left with insufficient charge as a result of V2G participation, as the penalty value assessed by a driver for such an event increases, so too does the compensation necessary for it to be favorable for them to participate. Interestingly, for lower penalty values, depending on the particular travel completion cost scenario, individuals should be willing to pay to participate in a V2G program. This condition arises from the value of the travel completion costs paid by the aggregator that would otherwise have to be borne by the vehicle owner when their battery is depleted, despite the increase in the probability of such an event. As a result, for a given penalty value, as the travel completion cost increases, the value of the aggregator increases as well. Given the

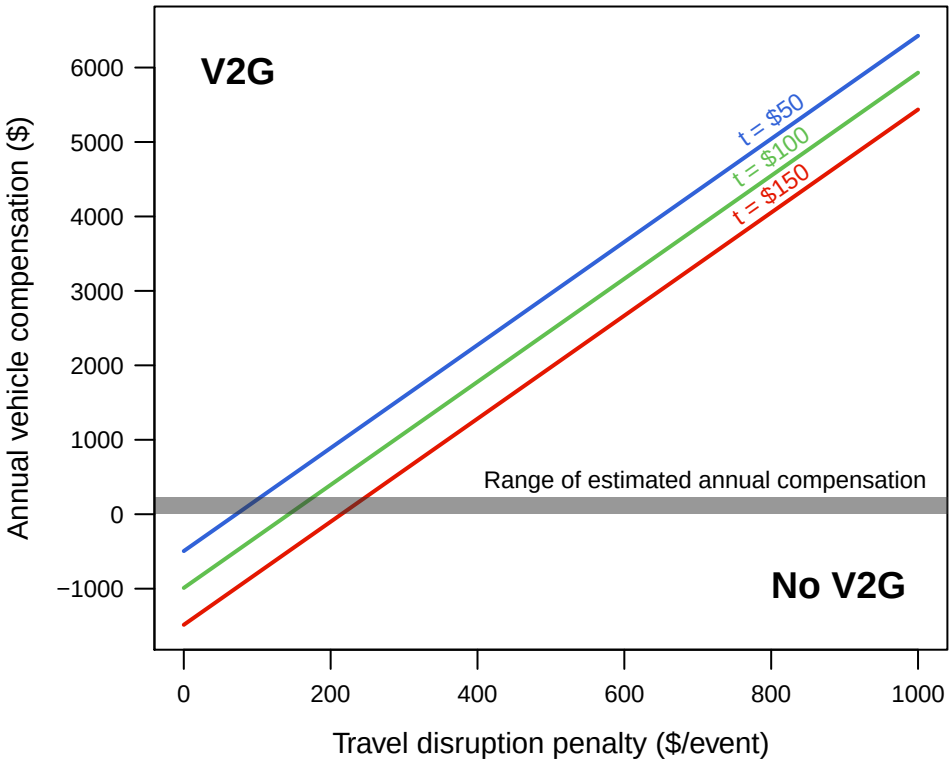

Figure 6.5: If an individual can assess their travel disruption penalty, it can be determined what compensation would be required from the aggregator to favor their entry into a V2G program, subject to the estimated value of travel completion costs (covered by the aggregator if participating in their program). Predictably, compensation payments increase as the individual's assessed disruption penalty increases. At low penalty values, the value of the travel completion costs paid by the aggregator are sufficient that vehicle owners should pay to enter a V2G program. Compensation indicated by the results in Chapter 5 is less than $1 per day, even for the largest revenue margin and smallest fleet size studied. As a result, individuals who assess a travel disruption penalty higher than $250 per event are probably not candidates for a V2G program unless there is some reason that their travel completion costs are anticipated to be higher than the cases shown.

magnitude of the annual payments indicated by the results in Chapter 5, individuals who assess their travel disruption penalty higher than approximately \$250 should not join a V2G program unless they believe the cost of travel completion in those cases when they do not have sufficient charge to be much higher than \$150 per event. Moreover, the compensation thresholds indicated in Figure 6.5 do not include any battery degradation costs.

6.3 Option value of a PHEV versus a BEV if participating in a V2G program

As noted in the previous section, V2G participants cannot readily predict their vehicle's state-of-charge in the future, and thus, the vehicle's state-of-charge might be insufficient to complete upcoming travel planned by the vehicle owner but unknown to the aggregator. If the vehicle is a PHEV, this risk is mitigated. The option value of a PHEV as compared with a BEV for individuals participating in a V2G program can thus be calculated, where the difference between the purchase price of the PHEV and BEV denotes the option price, p_o, and the additional cost of completing a trip using gasoline is the exercise price, p_e. Again, v denotes the dollar value ascribed to the trip, t' denotes the inconvenience associated with being towed, and it is assumed that all towing costs are paid by the vehicle aggregator. The results from the V2G aggregator model in Chapter 5 indicate that vehicles with different battery sizes would be compensated slightly differently. PHEVs typically have shorter electric driving ranges and smaller batteries than BEVs, thus they will receive unequal compensation. The payments for PHEV and BEV owners are denoted by d and d', respectively.

Figure 6.6 depicts the complete decision tree, which can be reduced to simplify the analysis. As with the previous decision, it is assumed that neither the PHEV nor

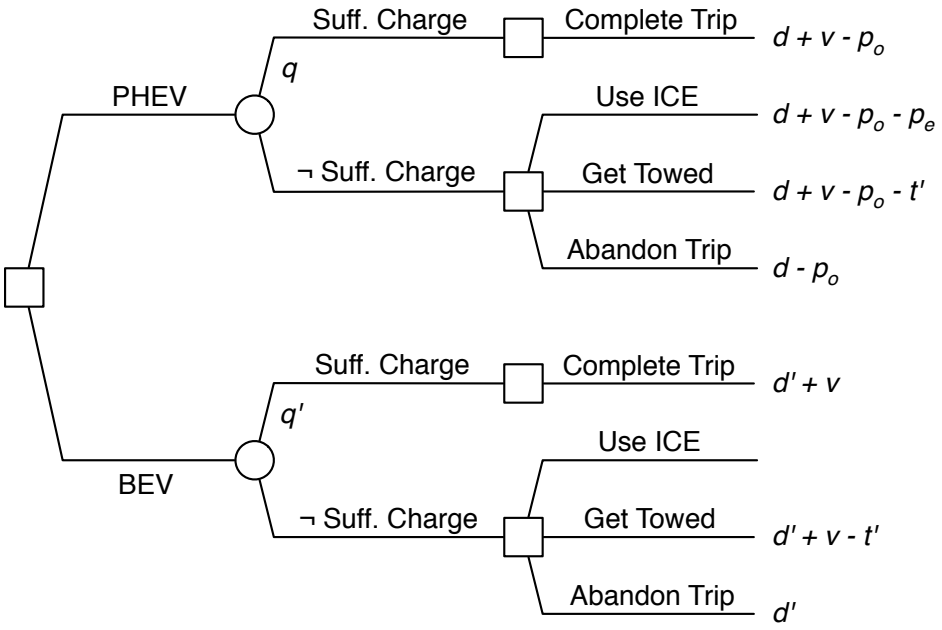

Figure 6.6: The decision for a prospective V2G participant hinges upon the type of vehicle they own, where the ICE in a PHEV provides an additional option reflected in the alternatives available for each vehicle type and their respective endpoint value measures. Again, the label "SC" denotes "sufficient charge."

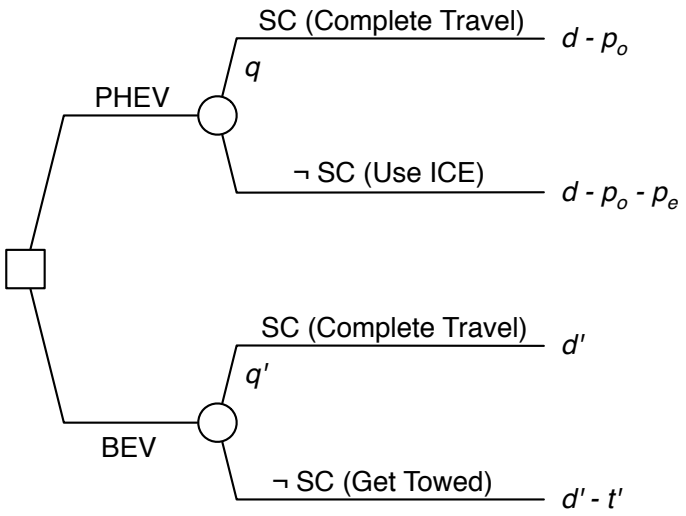

Figure 6.7: If it is assumed that drivers will not readily abandon their planned travel and that a PHEV owner will always use their ICE rather than be towed to their destination, then the decision described in Figure 6.6 can be reduced to the form shown. For each vehicle type, the uncertainty faced is between having sufficient charge to complete their upcoming travel or having to rely on the incumbent alternative, either the on-board ICE or towing (or another mode of transport).

the BEV driver would abandon their planned travel, because the point at which the driver learns they do not have sufficient charge in their battery is after the decision to travel has already been made, at which point the driver would probably not readily abandon the effort. The two remaining alternatives shown for the PHEV when the battery has been depleted are to use the ICE to extend the vehicle's range or to get towed. Given that the primary advantage of a PHEV over a BEV is the lifting of the range limitation on the vehicle via the ICE drivetrain, it is assumed that a PHEV owner would not intentionally be towed. Removing the "abandon trip" alternative for both vehicles, and the "get towed" alternative for the PHEV leaves only alternatives that include the completion of planned travel. As a result, all of the terms include the value of travel v, which can thus be removed from all of the terms. The simplified decision tree is shown in Figure 6.7.

From the reduced decision tree shown in Figure 6.7, equations can be readily developed to describe the decision between purchasing a PHEV and a BEV when participating in a V2G program. The procedure for developing these equations mirrors that discussed in Section 6.2. Equations 6.4 and 6.5 give the sum of the values of each of the alternatives multiplied by their endpoint probabilities. The primary unknowns in these equations are the option price, p_o, exercise price, p_e, and the value of the disrupted travel, t'.

$$q(d - p_o) + (1 - q)(d - p_o - p_e) \tag{6.4}$$

$$q'd' + (1 - q')(d' - t') \tag{6.5}$$

Results

Since not all of the terms in Equations 6.4 and 6.5 are known, the analysis of the decision between the two vehicle types when participating in a V2G program relies on examining the relationships between the unknown parameters to determine their comparative importance. From these relationships, conclusions can be drawn about the decision itself. For the purposes of the comparisons here, the parameter t', which represents the travel disruption penalty, is held as the independent variable, and its effect on the decision for a range of option prices and exercise prices. The results shown in Figure 6.8 are based on Equation 6.6, which defines the relationship between the independent variable and the option and exercise prices.

$$p_o = (q - 1)p_e + d - d' - (q' - 1)t' \tag{6.6}$$

Equation 6.6 was developed by rearranging the terms from an equality relation between Equations 6.4 and 6.5. Analogies are made between the option and exercise prices and parameters salient to the vehicle owner's decision, such as fuel prices and

discount rate, to clarify the decision. These results provide a guide for the decision faced by an individual driver, but their vehicle use choices could significantly influence the results shown, and the decision would be improved by tailoring the probabilities q and q' to the individual driver.

Figure 6.8 shows the relationship between the three unknown parameters: the option price, exercise price, and "travel disruption penalty", or t'. The travel disruption penalty is equivalent to the value ascribed by the driver to tolerating the inconvenience created by the vehicle aggregator's operations. Using an assumed ICE fuel economy and nominal fuel price, the exercise price can be converted into the number of miles that must be driven using the ICE due to aggregator operations. The option price observed by the vehicle buyer is the total price differential between a BEV and a comparable PHEV. As shown in Figure 6.8, the option value represents the benefits associated with the presence of the ICE in a PHEV. These benefits accumulate during the entire period the PHEV is owned. To determine the present value of these benefits when the vehicle is purchased, the benefits are discounted using a real rate of 5% per year over a 6-year ownership period, which is roughly correlated with the product lifecycle of a typical vehicle model. The range of option values shown in Figure 6.8 indicate that this price differential must be small, or the assessed travel disruption penalty must be high, to support the purchase of PHEV. Since the ICE engine in a PHEV affords avoidance of any travel disruptions due to a depleted battery, only drivers who place an extremely high value on the ability to maintain their travel plans should purchase a PHEV.

The dividing line between independent and dependent variable values favoring PHEVs or BEVs is shown for four exercise prices, ranging from \$1 to \$7. These prices are converted to distance measures more easily interpreted in the context of vehicle

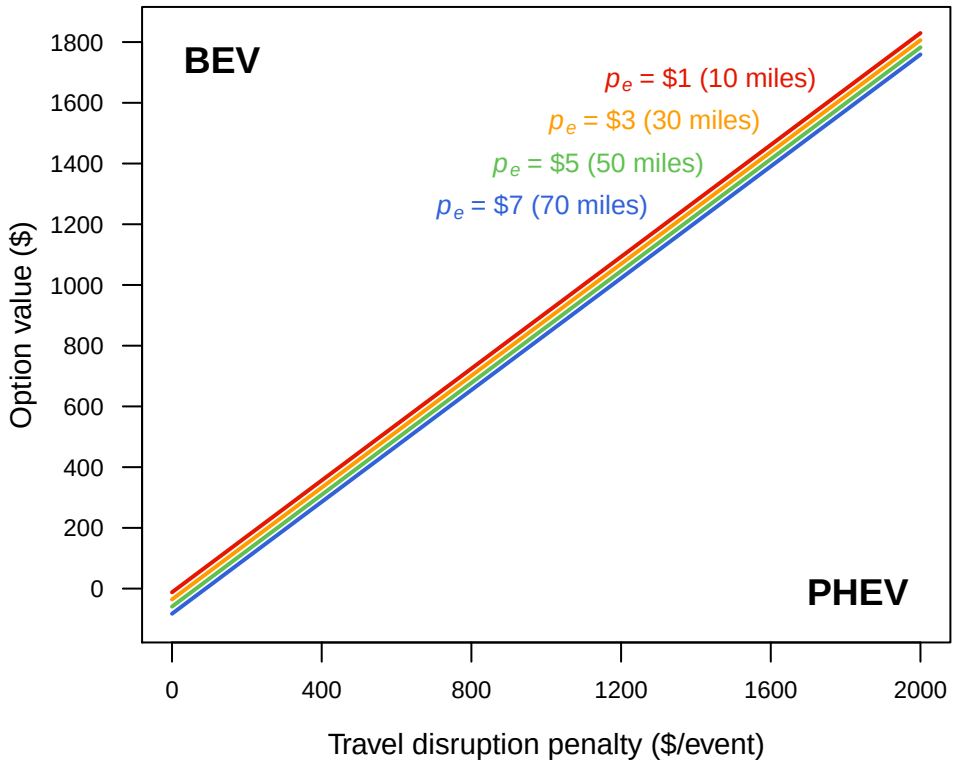

Figure 6.8: The relationship between the option price, p_o, and the inconvenience value associated with the difficulty of completing upcoming planned travel using an alternate vehicle or mode, t', is shown for four different exercise prices, ranging from \$1 to \$7. Both vehicle types considered in this analysis, PHEVs and BEVs, are subject to the charging and discharging decisions made by the vehicle aggregator. High option prices and low assessed values for travel disruption favor BEVs, as indicated with the "BEV" label, while converse conditions favor PHEVs. At a fixed option price, as the disruption penalty increases, the preferred vehicle type changes from a BEV to a PHEV. Increasing the exercise price has a corresponding effect on the travel disruption penalty threshold where a PHEV is preferred. Since the benefit of the ICE engine in a PHEV accrues to the driver gradually over the period of ownership, the option values shown represent the present value of these benefits at the time of vehicle purchase, assuming a 6-year ownership period and a real discount rate of 5% per year. The estimated mileages shown for each of the exercise prices represent the miles driven using the ICE drivetrain, and assume a combined fuel economy of 37 mi/gal and a fuel price of \$3.75/gal. Other ownership costs are not included, with the assumption that those costs are roughly constant between a comparable PHEV and BEV.

use. The conversion was calculated by multiplying the exercise price by the vehicle's average fuel economy and dividing by fuel price, with assumed values of 37 mi/gal and \$3.75/gal, respectively. While there are other costs associated with vehicle ownership, most of these costs are assumed to be equivalent between comparable PHEVs and BEVs. There might be differences in maintenance costs due to the ICE drivetrain in the PHEV that should be captured in the exercise price, but there exists insufficient data on PHEV and BEV operation to quantify these costs without speculation. Because the exercise price applies to only PHEVs, the distances indicated represent the distance driven beyond the PHEVs range. In this comparison, PHEVs and BEVs were assumed to have electric driving ranges of 25 and 80 miles, respectively.

The relationship between the option price and the travel disruption penalty is strongly dependent upon the differential between the compensation to PHEV and BEV owners in Equation 6.6. Figure 5.19 shows that compensation is not strongly dependent on battery size. Because compensation differs little between battery sizes at all fleet sizes, the trends shown in Figure 6.8 are applicable across the range of fleet sizes studied in Chapter 5. If the compensation paid to BEV owners, d', was significantly greater than for PHEV owners, d, much higher assessed travel disruption penalty values would be needed for the decision to favor the PHEV. Compensation payments vary more with respect to EVSE power rating, reflected in Figure 5.21, and the differential is greater than \$50 per year for PEV fleets smaller than 50,000 vehicles. While PHEV and BEV owners can use similar EVSEs, the larger batteries in BEVs necessitate EVSEs with higher power ratings than are needed by PHEVs. Increases in BEV range will be accompanied by increases in battery size, and to limit increases in required charging times, EVSEs with higher power ratings will likely become more commonplace. Even with EVSEs of at least 10kW, however, the compensation differential will increase the travel disruption penalty by less than \$20.

6.4 Conclusions

The model developed in Chapter 5 was used to determine the optimal operation of a V2G system subject to uncertain market and vehicle conditions. To determine the effect of these operations on the participants of a V2G program, the battery state-of-charge results from the model were used to estimate the probability that a driver would be left with insufficient charge to complete their future travel. The estimated probabilities were used to determine whether a BEV owner should join a V2G program and what PEV type is preferred for V2G participation. Given the finite driving range of a BEV, V2G operations were shown to increase the probability of having insufficient charge to complete future travel. Determining participation for a BEV owner is function of their assessment of the value of being delayed due to having insufficient charge, but given the compensation available for participating in a V2G program, it is unlikely that individuals would choose to participate subject to the additional risk. The participation decision changes when the vehicle owner has a PHEV, which offers them the option to use the ICE drivetrain when their battery is depleted. At a travel disruption penalty value that would indicate that a vehicle owner should not participate in a V2G program with a BEV, a PHEV might be a better choice. Because the benefits of the ICE drivetrain accrue during the entire ownership period of the PHEV, while the incremental cost of the drivetrain is incurred on the first day of ownership, this incremental cost must be on the order of $1,000. This result suggests that without some sort of foreknowledge of vehicle use or a sophisticated predictive algorithm that can anticipate upcoming travel, a vehicle aggregator might struggle to retain BEV owners in their program. If the aggregator could provide some vehicle purchase incentives to increase the acceptable option value, they might have better operational success and participant satisfaction with PHEV owners.

Chapter 7

Conclusions

The research detailed in this dissertation can be broadly divided into two phases. The first phase of the research, featured in Chapters 3 and 4, focused primarily on understanding the characteristics of vehicle use patterns and observing how those patterns, when transformed into vehicle charging load, affected existing conditions on the electric grid, with a particular focus on frequency regulation. The second phase of this work, covered in Chapters 5 and 6 focused on modeling a V2G system from the perspective of a vehicle aggregator, and observing the effect of that system's operations on the aggregator and vehicle owners participating, or considering participating, in that system. This work relied on vehicle travel data from several sources, as well as electric grid data from ERCOT, and employed a range of analytical methods to respond to the research objectives framed in Chapter 1.

The first research phase began with an effort to understand vehicle use patterns across various time intervals — within days, between days of the week, and between months of the year. These patterns could reveal characteristics of vehicle travel relevant to any research that relies on travel data to simulate the effect of electric vehicles on the grid, especially whether travel patterns indicate that peak electricity demand would be increased as a result of vehicle charging loads and whether there exist seasonal variations in driving patterns that parallel trends present in electricity demand. Data collected by PSRC were used for this work because they were longitudinal, and could thus show the presence of any seasonal characteristics. Analysis of the vehicle

travel data showed that there is a consistent diurnal pattern among weekdays, and a markedly different pattern on weekends. Weekday vehicle use is characterized by two distinct peaks, one around noon and another in the late afternoon, while weekend vehicle use only peaks once around midday. Variations within weekdays and weekends are comparatively small. This difference between weekday and weekend vehicle use is critically important. Averaging weekday and weekend travel together to create a single vehicle use profile is common in the literature, and will underrepresent vehicle use in the early morning and late evening hours. The data also show limited variation between months, which affects the relationship with electric load and wind generation in ERCOT. During cooler months, net load, which is electric load less available wind generation, has two peaks. These peaks occur just before the increase in vehicle use in the morning, and just after vehicle use declines again in the evening. The summer months in ERCOT are characterized by an afternoon peak in demand associated with residential air conditioning load. Vehicle use decreases just after the peak, indicating that vehicle charging loads might exacerbate peak demand, and that electric vehicles might not be available to provide ancillary services when they are most needed. These results underscore the importance of modeling any V2G system for a full year of electricity market conditions in any region where those conditions are subject to significant seasonal changes.

From this understanding of vehicle use data and the interaction between vehicle use and electric load, a larger set of travel survey data were used to simulate unscheduled PEV charging loads on the grid as a function of the number of PEVs in a region. Previous comparisons of vehicle use data with net load in ERCOT indicated that PEV charging loads might be coincident with summer peak electricity demand. The results of these simulations indicated that if PEVs comprised 5% of the light duty vehicle fleet in Texas, average peak demand during the summer months

would increase by approximately 1%. These results suggest that while the effort to develop optimization methods to manage charging might be useful for reducing the cost of PEV charging or controlling distribution system operation and maintenance costs, unscheduled charging will have a limited impact on peak demand in ERCOT. In addition, variability in PEV charging loads is small and not significantly affected by the number of PEVs present, which makes the effective variability quite small relative to total charging load for the larger fleet sizes considered. This variability in charging loads, however limited, was analyzed to determine its effect on frequency regulation procurements. The results of these simulations indicate that additional regulation procurements are required as a function of dominant trends in vehicle use — additional regulation up is needed to compensate for uncertainty in exactly when drivers will begin charging their vehicles in the afternoon and evening, while more regulation down was needed to accommodate uncertainty about the times when vehicles would finish charging, throughout the overnight hours, as their batteries were replenished. The increase in regulation required to support unscheduled charging is on the order of 1%, consistent with initial results showing small variability in charging loads, regardless of fleet size. Finally, the estimated PEV charging data from these simulations were validated using empirical charging data. Simulated and empirical charging patterns were found to be comparable, thus verifying the procedure used to convert the original travel survey data into charging data, despite the survey data being derived from ICE vehicles, not PEVs.

The second phase of this work began with the development of a two-stage stochastic linear programming formulation of a vehicle aggregator participating in the ERCOT market for frequency regulation. This model was developed to address several issues not fully explored in the literature: whether participation in frequency regulation markets would be economically favorable for the aggregator when faced

with uncertainty in the real-time availability and state of charge of PEVs in their system, the effect of changes in day-ahead regulation prices due to the aggregator's market power, the revenue generated by the aggregator and available to compensate participating vehicle owners, and the influence of vehicle-specific parameters on the aggregator's revenue. To be consistent with the results from the first phase of this research, a range of days selected from the year were modeled to capture any seasonal characteristics that might arise from electricity market conditions, and again, only weekdays were included in these dates. Due to limitations in the model formulation, it cannot be shown conclusively that there is sufficient stability in vehicle availability and state of charge to enable an aggregator to participate in the day-ahead frequency regulation market, but the results suggest that it is possible for an aggregator to bid in that market, even with that uncertainty present. Two different models for day-ahead market prices were developed to determine the effect of the aggregator's market power. The results from these models show that as the number of vehicles in a given region increases, per-vehicle revenues, and hence, compensation provided to vehicle owners, are reduced by as much as 85%. For the largest fleet sizes tested, annual compensation did not exceed $100 per year. While these results are highly dependent on prices in the frequency regulation market, and thus are not strictly applicable to every market in the United States or from year-to-year in ERCOT, the difference between the models does indicate that a vehicle aggregator will have sufficient market power to influence frequency regulation prices, and as such, these effects should be incorporated into V2G modeling efforts in the literature. Finally, the effect of PEV battery size and EVSE power rating on aggregator revenues were tested. Over the range of battery sizes tested, increasing size yields a corresponding increase in compensation, but the overall effect is quite small, and nearly vanishes at large fleet sizes, suggesting that the added cost associated with a vehicle's larger battery would not provide a

corresponding compensation benefit. EVSEs with higher power capabilities, which might be correlated with PEVs with larger batteries, yield a much larger increase in compensation, likely due to the increase in effective resource availability, which suggests that the resource provided by vehicle batteries for frequency regulation is primarily power-limited, not capacity- or duration-limited.

The results of the stochastic linear program were then used to assess two decisions: 1) whether an individual BEV owner should choose to participate in a V2G program, and 2) whether the additional cost of a PHEV could be justified when participating in a V2G program to avoid the risk of being delayed by a BEV with insufficient charge. Using the state of charge of the pooled vehicle resource and the number of vehicles available to provide V2G services at each time interval from the stochastic LP results, the probability of an individual vehicle owner being unable to complete anticipated future travel due to insufficient charge was assessed with and without V2G operations and for several battery sizes. The results from these simulations are dependent on the penalty associated with the travel delay that arises from a depleted battery. This value can be assessed by an individual as a prospective V2G member or PEV buyer, but in general, the compensation offered to participate in a V2G program is likely insufficient as an incentive to BEV owners, due to the increased probability of having an unexpectedly depleted battery. For PHEV owners, the on-board ICE drivetrain ensures that they will not be stranded if the vehicle aggregator uses their battery for regulation services, thus they are better suited to participating in a V2G program. Since the benefit of the ICE drivetrain accrues over the vehicle's entire ownership period, the incremental cost of the ICE at the time of purchase must be on the order of $1,000.

Though there might be sufficient compensation to retain PHEV owners in a V2G program, and improvements to the aggregator's operational capabilities could reduce the increase in the probability of a BEV owner having insufficient charge if they participate, it is unclear whether these incentives would be sufficient to offset potential costs associated with battery degradation or induce PEV owners to join a V2G program. Moreover, this compensation is likely inadequate as a means to encourage PEV adoption. Perhaps more importantly, it is unclear whether after-purchase financial incentives can play a major role in vehicle choice decisions, and if those incentives have greater influence with new or established vehicle types. It is possible that even if the results from these analyses showed that it was favorable for all PEVs to participate and that vehicle owners would receive $1,500 per year in compensation, such incentives might not have a material effect on car buying decisions.

Future work

This work leaves open several avenues for further development of the underlying data and modeling approaches employed. Existing work on electric vehicle charging and V2G systems rely primarily on travel survey data due to the lack of publicly available vehicle charging data. These survey data are a viable proxy for charging data, but historical charging data would be superior, and would obviate the complex conversion of travel data into charging data. It is not entirely clear how these data could be obtained, but they would be a valuable resource to the research community. From these data, sophisticated models to predict charging loads could be developed. These data could also underlie regression models that provide insight into minimum charging requirements as a function of time of day and vehicle owner

characteristics. The results from such a model could be useful to a vehicle aggregator and potentially as an input to a V2G system model.

There are several opportunities for expansion of the stochastic LP formulation of a V2G system developed in Chapter 5. Creating a computationally efficient method for discretizing the pooled storage resource would enable more accurate assessment of operating costs for the aggregator and could reveal whether specific vehicle use characteristics might affect the probability that a vehicle will have sufficient charge for upcoming travel. Discrete vehicle modeling could also be used to determine whether the aggregator could increase its revenues and reduce insufficient charge events by defining an n minute notification threshold for drivers about to use their vehicles, though introducing the notification threshold might require reformulating the model as a multi-stage program. Managing the cost of V2G operations with respect to battery degradation could also be readily included in a stochastic LP formulation with discrete vehicles.

One of the major shortcomings of the grid-scale energy storage literature is that many of the models developed focus on the use of storage for a specific single objective, such as relieving transmission congestion, smoothing or firming wind generation, or performing energy arbitrage. As shown by Townsend [10], energy storage typically has a range of capabilities, but capitalizing on only a single function can miss significant revenue opportunities. This dissertation focused on frequency regulation because it has been shown to be the highest value ancillary service for electrochemical energy storage to provide and it likely has the least impact on battery lifetime, but the stochastic LP formulation could be expanded to include other ancillary service markets. In addition to existing market structures, the effect of modifications to the compensation for certain ancillary services proposed by ERCOT [94] and FERC [11]

could be examined. It is possible that by including these additional revenue generation opportunities, aggregator revenue and participant payments could increase considerably. There are other objectives, such as maximizing charging from renewable sources or managing distribution system conditions, that have been the focus of other vehicle charging and V2G studies, but operational benefits in these areas might not be readily monetized for the vehicle aggregator and PEV owners.

Bibliography

[1] Bri-Mathias S Hodge, Shisheng Huang, Aviral Shukla, Joseph F Pekny, and Gintaras V Reklaitis. The Effects of Vehicle-to-Grid Systems on Wind Power Integration in California. In S. Pierucci and G. Buzzi Ferraris, editors, *20th European Symposium on Computer Aided Process Engineering*. Elsevier B.V., 2010.

[2] Department of Energy. Bottling Electricity: Storage as a Strategic Tool for Managing Variability and Capacity Concerns in the Modern Grid. Technical report, DOE Electricity Advisory Committee, 2008.

[3] Energy Information Administration. Shares of electricity generation from renewable energy sources up in many states. http://www.eia.gov/todayinenergy/detail.cfm?id=5750, 2012.

[4] R.A. Walling, L.A. Freeman, and W.P. Lasher. Regulation requirements with high wind generation penetration in the ERCOT market. In *2009 IEEE/PES Power Systems Conference and Exposition*, pages 1–7. IEEE, March 2009.

[5] Erik Ela, Michael Milligan, Brian Parsons, Debra Lew, and David Corbus. The evolution of wind power integration studies: Past, present, and future. In *2009 IEEE Power & Energy Society General Meeting*, pages 1–8. IEEE, July 2009.

[6] Patrick Sullivan, Walter Short, and Nate Blair. Modeling the Benefits of Storage Technologies to Wind Power. *Wind Engineering*, 32(6):603–615, December 2008.

[7] Ramteen Sioshansi, Paul Denholm, Thomas Jenkin, and Jurgen Weiss. Estimating the value of electricity storage in PJM: Arbitrage and some welfare effects. *Energy Economics*, 31(2):269–277, 2009.

[8] Paul Denholm, Erik Ela, Brendan Kirby, and Michael Milligan. The Role of Energy Storage with Renewable Electricity Generation. Technical report, National Renewable Energy Laboratory, 2010.

[9] Chioke Harris, Jeremy P. Meyers, and Michael E. Webber. A unit commitment study of the application of energy storage toward the integration of renewable generation. *Journal of Renewable and Sustainable Energy*, 4(1), 2012.

[10] Aaron K. Townsend. *A Grid-Level Assessment of Compressed Air Energy Storage in ERCOT*. PhD thesis, The University of Texas at Austin, 2013.

[11] Federal Energy Regulatory Commission. Frequency Regulation Compensation in the Organized Wholesale Power Markets. http://www.ferc.gov/whats-new/comm-meet/2011/102011/E-28.pdf, 2011. 137 FERC P. 61064, Order No. 755.

[12] Jim Eyer and Garth P. Corey. Energy Storage for the Electricity Grid: Benefits and Market Potential Assessment Guide. Technical Report February, Sandia National Laboratories, 2010.

[13] Y V Makarov, S Lu, J Ma, and T B Nguyen. Assessing the Value of Regulation Resources Based on Their Time Response Characteristics. Technical Report 500, Pacific Northwest National Laboratory, 2008.

[14] A Oudalov, D Chartouni, C. Ohler, and G. Linhofer. Value Analysis of Battery Energy Storage Applications in Power Systems. In *2006 IEEE PES Power Systems Conference and Exposition*, pages 2206–2211. IEEE, 2006.

[15] Department of Energy. One Million Electric Vehicles by 2015: February 2011 Status Report. http://www1.eere.energy.gov/vehiclesandfuels/pdfs/1_million_electric_vehicles_rpt.pdf, 2011.

[16] D M Lemoine, D M Kammen, and A E Farrell. An innovation and policy agenda for commercially competitive plug-in hybrid electric vehicles. *Environmental Research Letters*, 3(1):1–10, January 2008.

[17] Tony Markel, Michael Kuss, and Michael Simpson. Value of Plug-in Vehicle Grid Support Operation. In *2010 IEEE Conference on Innovative Technologies for an Efficient and Reliable Electricity Supply*, pages 325–332, Waltham, September 2010. Institute of Electrical and Electronics Engineers.

[18] Luis Pieltain Fernandez, Tomás Gomez San Roman, Rafael Cossent, Carlos Mateo Domingo, and Pablo Frias. Assessment of the Impact of Plug-in Electric Vehicles on Distribution Networks. *IEEE Transactions on Power Systems*, 26(1):206–213, February 2011.

[19] P Denholm and W Short. An Evaluation of Utility System Impacts and Benefits of Optimally Dispatched Plug-In Hybrid Electric Vehicles. Technical report, National Renewable Energy Laboratory, Golden, 2006.

[20] K Parks, P Denholm, and T Markel. Costs and Emissions Associated with Plug-In Hybrid Electric Vehicle Charging in the Xcel Energy Colorado Service Territory. Technical report, National Renewable Energy Laboratory, 2007.

[21] Tammy M Thompson, Carey W King, David T Allen, and Michael E Webber. Air quality impacts of plug-in hybrid electric vehicles in Texas: evaluating three battery charging scenarios. *Environmental Research Letters*, 6(2):1–11, April 2011.

[22] Kevin Mets, Tom Verschueren, Wouter Haerick, Chris Develder, and Filip De Turck. Optimizing smart energy control strategies for plug-in hybrid electric vehicle charging. *2010 IEEE/IFIP Network Operations and Management Symposium Workshops*, pages 293–299, 2010.

[23] Chioke B Harris and Michael E Webber. Quantifying the Effect of Plug-in Electric Vehicles on Future Grid Operations and Ancillary Service Procurement Requirements. In *Proceedings of the ASME 2013 International Mechanical Engineering Congress & Exposition*, 2013.

[24] David Dallinger, Daniel Krampe, and Martin Wietschel. Vehicle-to-Grid Regulation Reserves Based on a Dynamic Simulation of Mobility Behavior. *IEEE Transactions on Smart Grid*, 2(2):302–313, June 2011.

[25] Ramteen Sioshansi and Paul Denholm. Emissions impacts and benefits of plug-in hybrid electric vehicles and vehicle-to-grid services. *Environmental Science & Technology*, 43(4):1199–1204, February 2009.

[26] Aidan Tuohy and Mark OMalley. Pumped storage in systems with very high wind penetration. *Energy Policy*, 39(4):1965–1974, April 2011.

[27] D Rastler. Electricity Energy Storage Technology Options. Technical report, Electric Power Research Institute, Palo Alto, 2010.

[28] Scott B. Peterson, J.F. Whitacre, and Jay Apt. The economics of using plug-in hybrid electric vehicle battery packs for grid storage. *Journal of Power Sources*, 195(8):2377–2384, 2010.

[29] Ramteen Sioshansi and Paul Denholm. The Value of Plug-In Hybrid Electric Vehicles as Grid Resources. *The Energy Journal*, 31(3):1–23, 2010.

[30] Claes Sandels, Ulrik Franke, Niklas Ingvar, Lars Nordstrom, and Roberth Hamren. Vehicle to Grid Monte Carlo simulations for optimal Aggregator strategies. In *2010 International Conference on Power System Technology*, pages 1–8. IEEE, October 2010.

[31] National Renewable Energy Laboratory. NREL Transportation Secure Data Center. http://www.nrel.gov/vehiclesandfuels/transportationdata/index.cfm, 2012.

[32] U.S. Department of Transportation. 2009 National Household Travel Survey User's Guide, Version 1. http://nhts.ornl.gov/2009/pub/UsersGuideV1.pdf, 2011.

[33] Electric Reliability Council of Texas. About ERCOT. http://www.ercot.com/about/, 2013.

[34] Environmental Protection Agency. New Plug-in Hybrids. http://www.fueleconomy.gov/feg/phevsbs.shtml, 2013.

[35] Environmental Protection Agency. Electric Vehicles — Compare Side-by-Side. http://fueleconomy.gov/feg/evsbs.shtml, 2013.

[36] Internal Revenue Service. Plug-in Electric Drive Vehicle Credit (IRC-30D). http://www.irs.gov/Businesses/Plug-In-Electric-Vehicle-Credit-(IRC-30-and-IRC-30D), 2010.

[37] Society of Automotive Engineers. SAE Electric Vehicle Conductive Charge Coupler. http://www.arb.ca.gov/msprog/zevprog/stakeholders/infrastructure/finalsaej1772.doc, 2009.

[38] Piet Rietveld. Rounding of Arrival and Departure Times in Travel Surveys: An Interpretation in Terms of Scheduled Activities. *Journal of Transportation and Statistics*, 5(1):71–82, 2002.

[39] Frank Witlox. Evaluating the reliability of reported distance data in urban travel behaviour analysis. *Journal of Transport Geography*, 15(3):172–183, May 2007.

[40] J Decarolis and D Keith. The Costs of Wind's Variability: Is There a Threshold? *The Electricity Journal*, 18(1):69–77, January 2005.

[41] J C Smith, Stephen Beuning, Henry Durrwachter, Erik Ela, David Hawkins, Brendan Kirby, Warren Lasher, Jonathan Lowell, Kevin Porter, Ken Schuyler, and Paul Sotkiewicz. Impact of variable renewable energy on US electricity markets. In *IEEE PES General Meeting*, pages 1–12. IEEE, July 2010.

[42] W Kempton and J Tomic. Vehicle-to-grid power implementation: From stabilizing the grid to supporting large-scale renewable energy. *Journal of Power Sources*, 144(1):280–294, June 2005.

[43] Willett Kempton and Amardeep Dhanju. Using electric vehicles as storage for large offshore wind power. In *Copenhagen Offshore Wind 2005*, pages 1–15, 2005.

[44] Willett Kempton, Victor Udo, Ken Huber, Kevin Komara, Steve Letendre, Scott Baker, Doug Brunner, and Nat Pearre. A Test of Vehicle-to-Grid (V2G) for Energy Storage and Frequency Regulation in the PJM System. Technical report, 2008.

[45] Casey Quinn, Daniel Zimmerle, and Thomas H. Bradley. The effect of communication architecture on the availability, reliability, and economics of plug-in hybrid electric vehicle-to-grid ancillary services. *Journal of Power Sources*, 195(5):1500–1509, March 2010.

[46] Ramteen Sioshansi. Modeling the Impacts of Electricity Tariffs on Plug-In Hybrid Electric Vehicle Charging, Costs, and Emissions. *Operations Research*, 60(2):1–11, February 2012.

[47] Corey D. White and K. Max Zhang. Using vehicle-to-grid technology for frequency regulation and peak-load reduction. *Journal of Power Sources*, 196(8):3972–3980, April 2011.

[48] Nathaniel S. Pearre, Willett Kempton, Randall L. Guensler, and Vetri V. Elango. Electric vehicles: How much range is required for a days driving? *Transportation Research Part C: Emerging Technologies*, 19(6):1171–1184, February 2011.

[49] Sekyung Han, Soohee Han, and Kaoru Sezaki. Estimation of Achievable Power Capacity From Plug-in Electric Vehicles for V2G Frequency Regulation: Case Studies for Market Participation. *IEEE Transactions on Smart Grid*, 2(4):632–641, December 2011.

[50] Nathan Fettinger, Chee-Wooi Ten, and Chunxiao Chigan. Minimizing residential distribution system operating costs by intelligently scheduling plug-in hybrid electric vehicle charging. *2012 IEEE Transportation Electrification Conference and Expo (ITEC)*, pages 1–8, June 2012.

[51] Andreas Aabrandt, Peter Bach Andersen, Anders Bro Pedersen, Shi You, Bjarne Poulsen, Niamh O'Connell, and Jacob Ostergaard. Prediction and optimization

methods for electric vehicle charging schedules in the EDISON project. In *2012 IEEE PES Innovative Smart Grid Technologies (ISGT)*, pages 1–7. IEEE, January 2012.

[52] Olle Sundstrom and Carl Binding. Charging service elements for an electric vehicle charging service provider. In *2011 IEEE Power and Energy Society General Meeting*, pages 1–6. IEEE, July 2011.

[53] Trine Krogh Kristoffersen, Karsten Capion, and Peter Meibom. Optimal charging of electric drive vehicles in a market environment. *Applied Energy*, 88(5):1940–1948, May 2011.

[54] Salvador Acha, Tim C Green, and Nilay Shah. Effects of optimised plug-in hybrid vehicle charging strategies on electric distribution network losses. In *IEEE PES T&D 2010*, pages 1–6. IEEE, 2010.

[55] Wencong Su and Mo-yuen Chow. Performance Evaluation of an EDA-Based Large-Scale Plug-In Hybrid Electric Vehicle Charging Algorithm. *IEEE Transactions on Smart Grid*, 3(1):308–315, March 2012.

[56] Alexander Schuller, Jens Ilg, and Clemens van Dinther. Benchmarking electric vehicle charging control strategies. In *2012 IEEE PES Innovative Smart Grid Technologies (ISGT)*, pages 1–8. IEEE, January 2012.

[57] K. Clement-Nyns, Edwin Haesen, and Johan Driesen. The Impact of Charging Plug-In Hybrid Electric Vehicles on a Residential Distribution Grid. *IEEE Transactions on Power Systems*, 25(1):371–380, February 2010.

[58] Ahmed Yousuf Saber and Ganesh Kumar Venayagamoorthy. Intelligent unit

commitment with vehicle-to-grid A cost-emission optimization. *Journal of Power Sources*, 195(3):898–911, February 2010.

[59] Sekyung Han, Soohee Han, and Kaoru Sezaki. Development of an Optimal Vehicle-to-Grid Aggregator for Frequency Regulation. *IEEE Transactions on Smart Grid*, 1(1):65–72, June 2010.

[60] Sebastian Stein, Enrico Gerding, Valentin Robu, and Nicholas R Jennings. A Model-Based Online Mechanism with Pre-Commitment and its Application to Electric Vehicle Charging. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*, pages 4–8, Valencia, 2012.

[61] Mobashwir Khan and Kara M Kockelman. Predicting the Market Potential of Plug-in Electric Vehicles using Multiday GPS Data. In *91st Annual Meeting of the Transportation Research Board*, Washington, D.C., 2012.

[62] Casey Quinn, Daniel Zimmerle, and Thomas H. Bradley. An Evaluation of State-of-Charge Limitations and Actuation Signal Energy Content on Plug-in Hybrid Electric Vehicle, Vehicle-to-Grid Reliability, and Economics. *IEEE Transactions on Smart Grid*, 3(1):483–491, March 2012.

[63] A. De Los Rios, J Goentzel, K E Nordstrom, and C W Siegert. Economic analysis of vehicle-to-grid (V2G)-enabled fleets participating in the regulation service market. In *2012 IEEE PES Innovative Smart Grid Technologies (ISGT)*, pages 1–8. IEEE, January 2012.

[64] W Kempton and J Tomic. Vehicle-to-grid power fundamentals: Calculating capacity and net revenue. *Journal of Power Sources*, 144(1):268–279, June 2005.

[65] Alan J King and Stein W Wallace. *Modeling with Stochastic Programming*. Springer Series in Operations Research and Financial Engineering. Springer, 2012.

[66] Gerd C Infanger. *Planning under uncertainty: solving large-scale stochastic linear programs*. Boyd & Fraser, 1993.

[67] Harry M. Kaiser and Kent D. Messer. *Mathematical Programming for Agricultural, Environmental, and Resource Economics*. John Wiley and Sons, Inc., 2011.

[68] K. D. W. Nandalal and Janos J. Bogardi. *Dynamic Programming Based Operation of Reservoirs: Applicability and Limits*. Cambridge University Press, 2007.

[69] Charlie Wilson and Hadi Dowlatabadi. Models of Decision Making and Residential Energy Use. *Annual Review of Environment and Resources*, 32(1):169–203, November 2007.

[70] Ronald A. Howard. Options. In Richard J. Zeckhauser, Ralph L. Keeney, and James K. Sebenis, editors, *Wise Choices: Decisions, Games, and Negotiations*, chapter 5. Harvard Business Review Press, 1996.

[71] J. E. Matheson and R. A. Howard. An Introduction to Decision Analysis. In *Readings in Decision Analysis*. Decision Analysis Group, 1976.

[72] William P Coleman. Chapter 3: Modal Characteristics. In *TRB Highway Capacity Manual*. October 2010.

[73] Bureau of Transportation Statistics. National Household Travel Survey: Daily Travel Quick Facts. http://www.bts.gov/programs/national_household_travel_survey/daily_travel.html, 2012.

[74] Nissan Motor Company. Nissan LEAF Specifications. http://www.nissan-zeroemission.com/EN/LEAF/specs.html, 2012.

[75] Environmental Protection Agency. 2011-12 Electric Vehicles — Compare Side-by-Side. http://fueleconomy.gov/feg/evsbs.shtml, 2012.

[76] Don Anair and Amine Mahmassani. State of Charge: Electric Vehicles' Global Warming Emissions and Fuel-Cost Savings Across the United States. Technical report, Union of Concerned Scientists, February 2012.

[77] Department of Transportation. 2009 National Household Travel Survey Data — Version 2.1. http://nhts.ornl.gov/download.shtml#2009, February 2011.

[78] Pacific Gas & Electric. Electric Vehicle Penetration Study Using Linear Discriminant Analysis. http://www.energy.ca.gov/2012_energypolicy/documents/2012-02-23_workshop/comments/Pacific_Gas_and_Electric-Electric_Vehicles_Penetration_Study_2012-03-01_TN-63900.pdf, 2012.

[79] Ecotality. Who are the participants in The EV Project? http://www.theevproject.com/cms-assets/documents/128842-80098.devproj.pdf, 2013.

[80] Chandra R. Bhat, Sudeshna Sen, and Naveen Eluru. The impact of demographics, built environment attributes, vehicle characteristics, and gasoline prices on household vehicle holdings and use. *Transportation Research Part B: Methodological*, 43(1):1–18, January 2009.

[81] Sashank Musti and Kara M. Kockelman. Evolution of the household vehicle fleet: Anticipating fleet composition, PHEV adoption and GHG emissions in Austin, Texas. *Transportation Research Part A: Policy and Practice*, 45(8):707–720, October 2011.

[82] Sangho Choo and Patricia L Mokhtarian. What type of vehicle do people drive? The role of attitude and lifestyle in influencing vehicle type choice. *Transportation Research Part A: Policy and Practice*, 38(3):201–222, March 2004.

[83] Binny M Paul, Kara M Kockelman, and Sashank Musti. The Light-Duty-Vehicle Fleet's Evolution: Anticipating PHEV Adoption and Greenhouse Gas Emissions Across the U.S. Fleet. In *Transportation Research Board 90th Annual Meeting*, pages 1–25, Washington, D.C., 2011. Transportation Research Board.

[84] Abolfazl Mohammadian and Eric Miller. Empirical Investigation of Household Vehicle Type Choice Decisions. *Transportation Research Record*, 1854(1):99–106, January 2003.

[85] Joshua D. Rhodes, Charles R. Upshaw, Chioke B. Harris, Colin M. Meehan, David A. Walling, Paul A. Navrátil, Ariane L. Beck, Kazunori Nagasawa, Robert L. Fares, Wesley J. Cole, Harsha Kumar, Roger D. Duncan, Chris L. Holcomb, Thomas F. Edgar, Alexis Kwasinski, and Michael E. Webber. Experimental and Data Collection Methods for a Large-Scale Smart Grid Deployment: Methods and First Results. *Energy*, in review, 2013.

[86] Electric Reliability Council of Texas. 2013 ERCOT Planning: Long-Term Hourly Peak Demand and Energy Forecast. http://www.ercot.com/gridinfo/load/2013_Long-Term_Hourly_Peak_Demand_and_Energy_Forecast.pdf, 2012.

[87] Reigh A. Walling, Venkat Banunarayanan, Amanvir Chahal, Lavelle Freeman, Jerry Martinez, Nicholas Miller, Devin Van Zandt, and Mark Walling. Analysis of Wind Generation Impact on ERCOT Ancillary Services Requirements. Technical report, GE Energy, Schenectady, 2008.

[88] U.S. Federal Highway Administration. *Statistical Abstract of the United States*, page 689. U.S. Census Bureau, 2012.

[89] John R Birge and François Louveaux. *Modeling with Stochastic Programming*. Springer Series in Operations Research and Financial Engineering. Springer, 2012.

[90] Allen R Ferguson and George B Dantzig. The Allocation of Aircraft to Routes–An Example of Linear Programming Under Uncertain Demand. *Management Science*, 3(1):45–73, October 1956.

[91] Joyce W Yen and John R Birge. A Stochastic Programming Approach to the Airline Crew Scheduling Problem. *Transportation Science*, 40(1):3–14, February 2006.

[92] Miguel Carrión and Josée Arroyo. A Computationally Efficient Mixed-Integer Linear Formulation for the Thermal Unit Commitment Problem. *IEEE Transactions on Power Systems*, 21(3):1371–1378, August 2006.

[93] Electric Reliability Council of Texas. ERCOT Nodal Protocols, Section 4: Day-Ahead Operations. http://www.ercot.com/content/mktrules/nprotocols/current/04-080113-Nodal.doc, 2013.

[94] Electric Reliability Council of Texas. Storage Participation in ER-

COT. http://www.ercot.com/content/gridinfo/etts/keydocs/ERCOT_Storage_Issues_Whitepaper.doc, 2010.

[95] US Department of Energy. Fuel Economy. http://www.fueleconomy.gov/, 2013.

[96] Ronald A. Howard. Decision Analysis, 2005. Unpublished manuscript.