

DISSERTATION

MODELING AND SIMULATION TO INVESTIGATE THE ELECTRIFICATION POTENTIAL
OF MEDIUM- AND HEAVY-DUTY VEHICLE FLEETS

Submitted by

David A. Trinko

Department of Systems Engineering

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Spring 2023

Doctoral Committee:

Advisor: Thomas H. Bradley

Jason C. Quinn

Steven Simske

James Hurrell

Copyright by David A. Trinko 2023

All Rights Reserved

ABSTRACT

MODELING AND SIMULATION TO INVESTIGATE THE ELECTRIFICATION POTENTIAL OF MEDIUM- AND HEAVY-DUTY VEHICLE FLEETS

This project involves developing and integrating new modeling tools to simulate the dynamics of electric medium- and heavy-duty fleet vehicle adoption. A technical and economic modeling tool, combining a data-driven hardware cost model with a cost-optimal charging strategy microsimulation, enables tailored analysis of the costs and benefits of electrifying individual fleets. Next, a novel text synthesis process, applied to a curated corpus of literature, quantifies trade-offs between technical, economic, and other factors in the fleet vehicle procurement decision. The outcomes of these tasks combine with knowledge from recent literature on fleet decision processes to specify the vehicle procurement model used by fleets in an agent-based model of the medium- and heavy-duty electric vehicle market. This model embodies an especially disaggregated approach to adoption modeling, internalizing factors and dynamics that conventional adoption models externalize. In particular, explicitly modeling the formation and diffusion of opinions among agents enables experiments that conventional models cannot support. Demonstrations show, for example, that increasing the extent of interactions between populations with different proclivities to electric vehicles has an asymmetrical outcome. High-proclivity electric vehicle adoption is generally unaffected as interactions increase, but low-proclivity adoption is accelerated. By representing individual fleets' requirements and costs at a high level of detail, incorporating an adoption decision model informed by a wide body of empirical research, and broadening the array of variables and dynamics available for experimentation, this integrated model offers a new way to understand the urgent challenge of eliminating emissions from the most emissions-intensive transportation sectors.

ACKNOWLEDGEMENTS

This material is based upon work supported by the U.S. Department of Energy's Office of Energy Efficiency and Renewable Energy (EERE) under the Vehicle Technologies Office, award number DE-EE0009664. Additional funding and support were provided by the Electric Power Research Institute and Lightning eMotors. Chapter 3 was a collaborative effort with Fletcher Ouren, another student in Systems Engineering. Informal support has also been graciously provided by the EPRI Electric Transportation team as a whole, fellow students from the Bradley and Quinn research groups, and friends and family.

DEDICATION

This dissertation is dedicated to my parents, Sue and Tony, and my sister, Olivia. They're really great people.

TABLE OF CONTENTS

ABSTRACT	ii
ACKNOWLEDGEMENTS	iii
DEDICATION	iv
LIST OF TABLES	viii
LIST OF FIGURES	ix
Chapter 1 Introduction	1
1.1 Motivation	1
1.2 Existing literature and models	2
1.3 Funding and support	4
1.4 Overview	4
1.5 Research questions	6
1.6 Research contributions	8
1.7 Organization of chapters	9
Chapter 2 Electric vehicle and fleet economics modeling	11
2.1 Abstract	11
2.2 Introduction	11
2.3 System model of Charging-as-a-Service	16
2.4 Functions of charging-as-a-service	18
2.4.1 Economic functions	19
2.4.2 Technical functions	21
2.4.3 Managerial functions	22
2.4.4 Societal functions	22
2.5 Business cases	24
2.5.1 Detached residence case	26
2.5.2 Multi-unit dwelling case	28
2.5.3 Business parking case overview	31
2.5.4 Commercial fleet case overview	35
2.6 Conclusion	40
Chapter 3 Decision and preference modeling	42
3.1 Abstract	42
3.2 Introduction	42
3.2.1 Contributions	43
3.3 Brief overview of text mining	44
3.3.1 Performance metrics	46
3.4 Methods	47
3.4.1 Corpus description	49
3.4.2 Data wrangling/preprocessing	49
3.4.3 LDA implementation	50

3.5	Verification	51
3.5.1	Split-corpus testing	51
3.5.2	Model repeatability	51
3.6	Results	52
3.6.1	Validation results	52
3.6.2	Final saliency scores	53
3.7	Discussion	54
3.7.1	Suitability/compatibility	54
3.7.2	Capital cost	55
3.7.3	Sustainability and long-term cost	55
3.8	Future work	56
3.8.1	Synonym normalization	57
3.8.2	Model comparison	57
3.8.3	Corpus amendments	57
3.8.4	Industry corpus	58
3.9	Conclusion	58
Chapter 4	Agent-based model development and application	60
4.1	Abstract	60
4.2	Introduction	60
4.3	Adoption process model	62
4.4	Learning and opinion formation	65
4.5	Verification and validation	71
4.5.1	Submodule verification	71
4.5.2	Comparison with other projections	73
4.5.3	Peer review	74
4.6	Demonstration and sensitivity assessment	77
4.6.1	Sensitivity analysis	78
4.7	Simulating more heterogeneous populations	85
4.7.1	Weight class heterogeneity	87
4.7.2	Preference heterogeneity	90
4.7.3	Range option heterogeneity	93
4.8	Conclusion	96
Chapter 5	Future work	98
Chapter 6	Conclusions	101
Appendix A	Electric vehicle and fleet economics modeling	123
A.1	Details of business cases	123
A.1.1	Detached residence	123
A.1.2	Multi-unit dwelling	124
A.1.3	Business parking	127
A.1.4	Commercial fleet	128
Appendix B	Decision and preference modeling	132

Appendix C	Agent-based model development and application	140
C.1	Vehicle attributes and calculation methods	140
C.1.1	Range Suitability	140
C.1.2	Total cost of ownership	140
C.1.3	Purchase price	141
C.1.4	Maintenance cost	141
C.1.5	Fuel economy	141
C.1.6	Resale value	142
C.1.7	Familiarity	142
C.1.8	Driver acceptance	142
C.1.9	Fueling price stability	143
C.1.10	Fuel security	143
C.1.11	Refueling convenience	143
C.1.12	Fuel cost	144
C.1.13	Reliability	144
C.1.14	Brand image and environmental sustainability	144
C.2	Adoption process	145
C.2.1	Awareness	145
C.2.2	Consideration	145
C.2.3	Decision	147
C.2.4	Implementation	148
C.2.5	Confirmation	149
C.3	Observation abstractions	149
C.4	Observational inference process details	151

LIST OF TABLES

2.1	Summary of economic, technical, managerial, and social functions of charging-as-a-service.	19
2.2	Selected assumptions used in the single-family-home business case. A complete version of this table can be found in A.1.1.	26
2.3	Selected assumptions used in the MUD business case. A complete version of this table can be found in A.1.2.	30
2.4	Selected assumptions used in the business parking case. A complete version of this table can be found in A.1.3.	32
2.5	Selected assumptions used in the commercial fleet business case. A complete version of this table is provided in A.1.4.	36
4.1	List of vehicle and technology attributes that fleet agents incorporate in their adoption decisions.	66
4.2	List of major peer-review instances and outcomes. Acronyms: Colorado State University (CSU, author’s PhD committee), U.S. Department of Energy Vehicle Technologies Office (DOE VTO), Electric Power Research Institute (EPRI), Lightning eMotors (LeM), National Renewable Energy Laboratory (NREL), Renewable and Sustainable Energy Reviews journal (RSER), Sacramento Municipal Utility District (SMUD). . . .	76
4.3	Google Maps queries used to find sites for simulated fleets, of five different vocations, in the Los Angeles metropolitan area.	78
4.4	Summary of scenarios considered in the weight class experiment.	88
4.5	Summary of scenarios considered in the decision preference experiment.	90
4.6	Summary of scenarios considered in the range option experiment.	93
A.1	Assumptions used in the single-family-home business case.	124
A.2	TOU-EV-9 rate schedule [159] used in cases 2 and 4.	125
A.3	Assumptions used in the multi-unit dwelling business case.	126
A.4	General Service TOU Option R rate schedule [160] used in case 3.	127
A.5	Assumptions used in the business parking business case.	128
A.6	Assumptions used in the commercial fleet business case.	131
B.1	Index of keywords (left column) and the categories to which they were assigned. . . .	139
C.1	Methods by which fleet agents make observations, and the relative importance weight assigned to those observations	150

LIST OF FIGURES

1.1	Conceptualization of information flow between the modeling tools.	6
2.1	Graphical abstract.	12
2.2	System diagram illustrating the roles Charging-as-a-Service can play at the interface between the transportation and electricity systems. ChaaS can serve as a unified conduit for electric, economic, and information flows, and can help manage risks for transportation system and electricity system entities by assuming EVSE ownership responsibilities. Abbreviations: Distributed Energy Resources (DERs); Electric Vehicle Service Equipment (EVSE); Transmission & Distribution (T&D).	17
2.3	Breakdown of costs paid by the customer in the detached residence example, expressed both as a net present cost and as a levelized cost of charging (LCOC).	28
2.4	Breakdown of costs paid by residents in the MUD example, expressed both as a net present cost for a group of 15 vehicles and as a levelized cost of charging (LCOC).	31
2.5	Charging sales price, in \$/kWh, and utilization, in hours per day, required for a charging system host to achieve a net profit over the system’s lifetime. In the “No ChaaS” scenarios, peak demand is maximized (all chargers are assumed to be used simultaneously at least once per month). In the “ChaaS” scenarios, peak demand is minimized via load pooling, shifting via DERs, or other means, achieving a theoretical lower bound to electricity cost. The savings realized by ChaaS substantially reduce the price that must be assessed to achieve a profit, particularly at low utilization levels.	33
2.6	Breakdown of value streams paid and accrued by station hosts in the business parking case, expressed as a levelized revenue per kWh sold, when the charging system is utilized 2.7 hours per day and charging is sold at \$0.31/kWh.	34
2.7	Power flows throughout a representative day for the commercial fleet business case, for populations of 3, 10, and 25 fleets. The baseline charging load is shown in the top row. The ChaaS power flows, which include the integration of a PV array and a battery storage system, and the accrual and sale of LCFS credits, are shown in the lower rows.	37
2.8	Breakdown of value streams paid and accrued by groups of fleets in the commercial fleet business case. For various populations of fleets, the LCOC components are shown for four scenarios: baseline, with no DER system or LCFS credit generation; and three ChaaS scenarios which variously include optimal DER system integration and/or the optimal generation of LCFS credits.	38
3.1	Adoption process framework adapted from Bae et al. [30].	43
3.2	Breakdown of sources for papers in the analyzed corpus.	48
3.3	Illustration of perception weight development process.	48
3.4	Average category SSD for the split corpora and the full corpus, $n = 800$	52
3.5	Average variance of normalized saliency between ensembles, $n = 800$	53
3.6	Radar plot of final saliency scores.	54
4.1	Illustration of a populated ABM consisting of five fleets, two utilities, and five OEMs. Fleet agents own assets including vehicles and EV charging plugs (EVSE).	62

4.2	Adoption process followed by fleet agents, adapted from Bae et al. [30].	63
4.3	Visualization of how agents form perceptions on the basis of observations. Observations pertaining to maintenance cost for two vehicle options, an ICEV and a BEV, are shown, where more heavily weighted observations are shown as larger points. Colors identify the means by which observations were made. The overall distribution of observations is shown as a box-and-whisker plot, with the median as a line and the mean as a circle. The vertical axis shows how values map to subjective assessments, on a scale of 0 to 5, with lower maintenance costs mapping to higher subjective assessments.	70
4.4	Illustration of the utility calculation for an ICEV and a BEV. In the lower plots, the widths of bars correspond to the relative weights assigned to vehicle attributes, and the heights correspond to the subjective assessments of those attributes, such that the fraction of area occupied by colored bars is proportional to the vehicle’s perceived utility. This is mapped onto the upper portions of the plots, which show the calculated utility, on a scale of 0 to 5, and the relative contribution of each attribute.	71
4.5	Comparison of up-front costs for a Class 4 delivery truck.	72
4.6	Comparison of long-term costs for a Class 4 delivery truck.	73
4.7	Comparison between the projections made by the TEMPO model, for light-medium duty freight, and our model, for Class 3 delivery vehicles.	74
4.8	Map showing the locations of the simulated fleets. Fleet sizes, in number of vehicles, correspond to the sizes of the circles on the map.	78
4.9	Adoption outcomes from many simulations using different initial random number generator states, but otherwise identical input parameters. The bold trace is from the baseline simulation presented throughout Section 4.6.1.	80
4.10	Illustration of the effect of varying input parameters on EV adoption outcomes. Parameter choices are shown on the horizontal axis. The parameter used in the baseline run is marked by a dotted vertical line. Outcomes for years 2025, 2030, 2035, 2040, 2045, and 2050 are shown, where later years are represented by lighter-colored traces.	81
4.11	EV adoption outcomes when the fleet population is simulated to make adoption decisions with strict consideration of a single attribute. The dashed line is the baseline outcome, in which agents weigh all factors against one another.	85
4.12	Histogram of travel distance across all of the vehicle-days represented in the Fleet DNA dataset.	86
4.13	Electrification outcomes for the population of 10 Class 6-8 fleets in the weight class heterogeneity experiment. Figures on the diagonal illustrate the electrification trajectory, in terms of the percentage of vehicles that are electric, for the two scenarios. The off-diagonal figure (bottom left) shows the subtractive difference between the two scenarios. Thin colored lines are results for individual Class 6-8 fleets. The bold black line aggregates over the Class 6-8 population.	88
4.14	Timing of perception crossover: the first year when each fleet considers one or more EV models to have greater utility than the ICEV model, for the two scenarios. The chronological year of crossover is shown on the left. On the right, the year of crossover is instead expressed relative to the year when each fleet first becomes aware of suitable EV options.	89
4.15	Electrification outcomes for the control group of 14 fleets in the decision preference experiment. The figure is formatted analogously to Figure 4.13.	91

4.16	Perception crossover for the control group of 14 fleets in the decision preference experiment. The figure is formatted analogously to Figure 4.14.	92
4.17	Electrification outcomes for the full population of fleets in the range option experiment. The figure is formatted analogously to Figures 4.13 and 4.15.	94
4.18	Perception crossover for the control group of 14 fleets in the decision preference experiment. The figure is formatted analogously to Figures 4.14 and 4.16.	95
A.1	Comparison of the driving patterns for the vocations modeled. The height of the light-colored area represents the number of vehicle days on which the vehicle was driving at the time of day shown on the horizontal axis. Whereas, for delivery vehicles, driving occurs mostly between 8 a.m. and 4 p.m., passenger transport vehicles operate throughout the afternoon and evening (between noon and midnight).	129
C.1	Agents build perceptions of vehicle options and the categories to which the options belong. When an agent becomes aware of a new option, the prior distribution of perceptions about that option is constructed by inheriting perceptions from other options that occupy the same categories.	150

Chapter 1

Introduction

1.1 Motivation

The transportation system faces pressing sustainability challenges, including the need to reduce the environmental impacts of transport while advancing its increasingly important role in the national and global economy. Medium-duty and heavy-duty (MHD) vehicles account for more than a quarter of U.S. road transportation fuel use and CO₂ emissions, with a projected increase in energy consumption of 11% by 2050 [1]. They also account for 73% of PM_{2.5} emissions globally and distribute exposure of these harmful products inequitably [2]. Plug-in electric vehicles (PEV) are a promising alternative to petroleum-fueled transportation with the potential to eliminate direct MHD emissions. PEVs draw some or all of their power from the electric grid, enabling efficiency and fueling cost improvements and reducing or eliminating tailpipe emissions. However, the ability of PEVs to satisfy operational requirements for some MHD sectors is constrained by technical challenges, including limits to driving range and long recharging times. Furthermore, the cost competitiveness of electric MHD (eMHD) vehicles has been limited by factors including high vehicle production costs and insufficient high-power charging infrastructure [3–5].

Still, the demand for MHD electrification is increasing. On the MD side, Amazon, UPS, FedEx, and DHL have made commitments to purchase substantial numbers of PEVs for their package delivery fleets. On the HD side, recent analysis has suggested that long-haul freight trucks may soon be competitive with conventional trucks based on total cost of ownership (TCO) if fast-charging infrastructure is sufficiently deployed [6–8]. The 2022 Infrastructure Investment and Jobs Act allocates \$7.5 billion to charging infrastructure deployment and acceleration of PEV adoption, \$5 billion to decarbonization of school buses, and tens of billions to grid decarbonization [9]. Other legislation by the federal government has proposed to allocate \$174 billion to vehicle electrification, including electrifying 50,000 diesel transit vehicles, 20% of the school bus fleet, and the

federal vehicle fleet, including that of the USPS, and installing 500,000 PEV chargers by 2030. To actuate ambitious plans such as these requires the guidance of robust and up-to-date modeling and analysis. The timelines under which PEVs will become suitable and cost-competitive in different MHD vocations will differ across the regions of the US. Their readiness for deployment therefore depends on varying sets of factors, for which technical, economic, and environmental assessments can define allocations and trajectories for investments in infrastructure, marketing, and products. No matter how large a role these policies take, the decentralized decision making of fleet purchasers, vehicle and EVSE manufacturers, energy system stakeholders, and local policy-makers will continue to play key roles in defining eMHD availability and uptake. Presently, these stakeholders have incomplete access to information and resources to inform their decision making, which has been identified as a critical obstacle to eMHD market development [10].

1.2 Existing literature and models

Electrification of LD vehicles has drawn substantial research attention, including the development of an array of modeling tools, techno-economic analyses, and adoption/consumer preference studies. Similar research for MHD electrification has lagged somewhat behind, but recent efforts have been made to assess feasibility and cost competitiveness, and estimate adoption projections, for eMHD vehicles.

Much of the eMHD literature concerns the feasibility and cost competitiveness of electrifying the various vocations in the MHD transportation sector, often relying on TEA methods. Consensus on the suitability of eMHD for certain applications remains under debate. For example, research papers published within the last year have reached opposite conclusions on the feasibility of using batteries to power long-haul trucks [8, 11]. This debate is partially in response to key uncertainties, such as the role of future policies and incentives and the pace of technological development. In addition, most studies reduce either the supply or demand side of the market to exogenous variables. Often, studies set the availability and costs (supply side) of fast charging as external parameters to project the ability of PEVs to satisfy existing duty cycles [4, 11, 12]. As a result, these studies

are unable to investigate or represent the dynamics of feedback between supply and demand (of vehicles and charging) that will play a major role in the evolution of the eMHD markets.

Adoption of eMHD vehicles and infrastructure is another emerging focus of the literature. Adoption modeling can be split broadly into two categories, top-down and bottom-up, which terms refer to the primary direction of information aggregation in the model. Prominent top-down adoption modeling approaches include “consumer choice models,” in which consumers are modeled as groups of decision makers responding to changing factors, and diffusion rate models, where the process of acceptance of a new technology is modeled in aggregate, often via regression (e.g. Bass models, Gompertz or others) [13]. Every study of eMHD adoption found in our literature survey uses a top-down approach. LD PHEV adoption, however, has been studied using bottom-up ABM approaches. In these studies, the ABM approach enabled researchers to explore potential avenues for increasing market penetration, including social/behavioral methods such as communicating lifecycle costs to purchasers [14, 15].

In the process of real-world adoption, operators of fleets will largely make final determinations of PEV feasibility and cost competitiveness for themselves, based on their own operations data, rather than based on results reported in research articles. The ability of fleets to conduct full analyses is constrained by access to modeling tools and data, which, if insufficient, may itself serve as a barrier to adoption. A variety of free modeling tools, many developed by DOE, provide partial analysis capabilities to fleets. Prominent examples include AFLEET [16] and EVI-FAST [17]. AFLEET, from Argonne National Laboratories, which can be used to estimate basic economic and environmental metrics for MHD vehicles using a variety of fuels, which enables fleet operators to conduct preliminary assessments for vehicle purchasing. EVI-FAST, from the National Renewable Energy Laboratory, similarly facilitates financial analysis for installation of charging infrastructure. HEVI-Load [18], a tool not yet available to the public from Lawrence Berkeley National Laboratory, employs a top-down approach with exogenous eMHD adoption projections to assess geographic charging infrastructure needs. The development and availability of these and other tools provide immense value to researchers, fleets, utilities, and other stakeholders.

The broader field of mobility-energy system modeling comprises these and many other modeling tools addressing different geographic, temporal, and sectoral resolutions and extents. From this perspective, researchers have recently highlighted the need to integrate models at a system level and have identified a spatiotemporal gap in existing models [19]. Accelerating trends such as electrification, shared mobility, automation, and connectivity are increasing the integration of the transportation and energy systems, resulting in new challenges but also providing new opportunities for optimization. In a critical example of this move toward integration, transportation energy demand and energy supply are becoming increasingly temporally coupled because of electrification. An implication of this coupling is a need for models operating at a high temporal resolution (hourly/sub-hourly) and a wide spatial extent (up to national) as aligned with the characteristics of the electric grid. Energy-mobility system models with these spatiotemporal characteristics have not yet been demonstrated.

1.3 Funding and support

This research was funded by the Department of Energy (DOE, Award #2420-1914). Support was also provided by the Electric Power Research Institute (EPRI), a leading research and development company with decades of experience studying the mobility-energy systems interface, and Lightning eMotors (LeM), an emerging leader in MHD PEV design and production.

1.4 Overview

The work described in this dissertation encompasses the development, integration, and demonstration of two modeling tools. These tools aim to address gaps in transportation/energy system modeling from the perspectives of two key decision-making entities. The first perspective is that of an operator of a MHD fleet, who likely lacks the means to answer critical questions associated with fleet electrification, particularly regarding economics and constraints to operation. The second perspective is that of an entity aiming to influence the trajectory of electrification in MHD fleet settings via regulations, incentives, or other means, whose decisions presently have highly

uncertain outcomes. Although these perspectives have different sets of priorities and needs, the modeling capabilities necessary to address them overlap. Specifically, entities aiming to influence adoption of electrified technologies benefit from understanding the factors that are most critical to the decisions of those agents choosing whether to adopt (that is, fleet operators).

Therefore, a tool facilitating economic and operational analysis of electrified fleets provides value in both perspectives. Our economic analysis tool (Tool #1) operates at a fleet-level spatial scope and a high temporal resolution (sub-hourly). Because sets of outputs can be summed to expand the spatial scope indefinitely, this tool provides the first step toward filling the spatiotemporal gap identified in mobility-energy system modeling. In recognition of this potential, our bottom-up adoption modeling tool (Tool #2) serves to aggregate and analyze sets of outcomes of Tool #1 in the context of interacting sets of agents, resulting in an integrated energy-mobility system model with high spatial scope and temporal resolution. One key capability this enables is to investigate the sensitivity of eMHD adoption to regulatory decisions and other factors.

A visualization of the capabilities of the integrated tool is shown in Figure 1.1. The research questions that were addressed in the development and application of the integrated tool are summarized in Section 1.5.

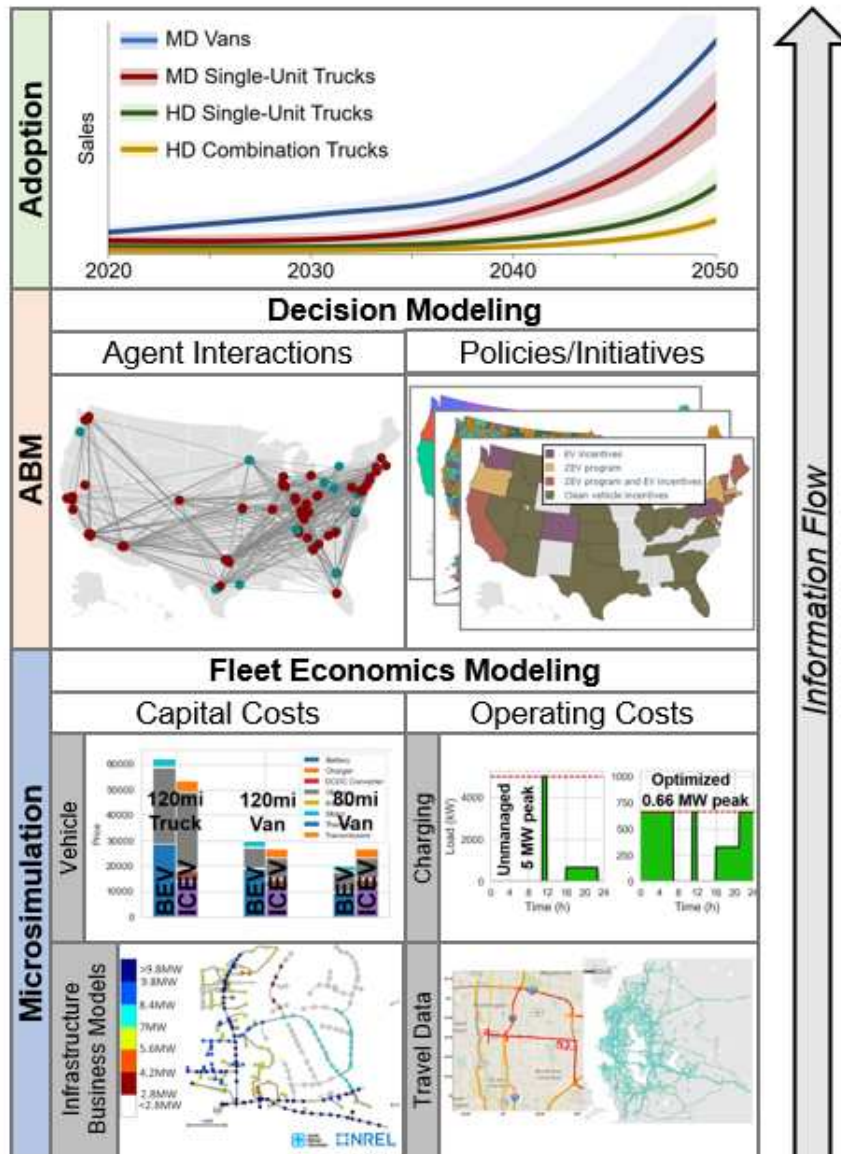


Figure 1.1: Conceptualization of information flow between the modeling tools.

1.5 Research questions

Research Question 1

How do internal and external factors influence the feasibility of electrification for MHD fleets? Internal factors (e.g., fleet vocation, operating schedules) and external factors (e.g., charging infrastructure availability, emissions regulations) together determine the economic and opera-

tional feasibility of adopting PEVs for a fleet. These factors often interact in complex ways, which creates opportunities for electric fleets to optimize aspects of their operation, such as the timing and control of vehicle charging. Consequently, estimating the feasibility of electrification requires performing optimization tasks in addition to evaluation tasks. A modeling tool with these capabilities, tailored to the case of eMHD fleets, facilitates experiments and studies addressing this question. In the broader context of this research, this tool provides decision support to simulated agents in a bottom-up adoption modeling tool developed to answer RQ#2 and #3. We address RQ#1 primarily in Chapter 2.

Research Question 2

How do economic, behavioral, and regulatory factors influence the decisions made by key entities in the MHD transportation system? The work addressing RQ#1 aims to identify and estimate feasibility characteristics that factor strongly in fleet electrification decisions, but research has demonstrated that factors not directly related to economic/operational feasibility, including behavioral/psychological factors, also play a role. Additionally, entities other than fleet operators make decisions that both affect and respond to feasibility outcomes, resulting in a complex set of feedback loops that may limit or encourage MHD PEV adoption. Although human decision making is extremely difficult to model with precision, an improved understanding of these additional factors and decision makers enables the development of useful models of key decision makers in the MHD transportation system. This decision model plays an integral role in our bottom-up adoption modeling tool. We address RQ#2 primarily in Chapter 3.

Research Question 3

How can decisions by stakeholders in the MHD transportation system contribute to desired outcomes? All types of stakeholders in the transportation system, including fleet owners, regulators, and equipment suppliers, have articulated long-term goals for decarbonization and the adoption of MHD PEVs. Suitably designed models can enable stakeholders to explore potential benefits and drawbacks of decisions made in pursuit of these goals, such as design of incentives or

emissions targets. However, existing adoption models, which employ top-down approaches, suffer the disadvantage that they cannot be used to identify or measure causal factors or the means by which they influence adoption, limiting their suitability for decision support. A bottom-up, agent-based modeling (ABM) approach, incorporating the models and insights derived from addressing RQ#1 and RQ#2, enables such decision-support investigations. This culmination of modeling, simulation, and analysis efforts introduces and demonstrates a novel framework for understanding the costs and benefits of electrification, modeling potential long-term transportation system dynamics, and clarifying decisions in pursuit of broad transportation-system goals. We address RQ#3 primarily in Chapter 4.

1.6 Research contributions

The research herein addresses two critical gaps in the field of MHD transportation system modeling:

1. **A lack of models with high temporal resolution and spatial extent.** A high time resolution is necessary because an electrified transportation system is coupled tightly in time with the dynamic modern energy system. A high spatial extent is necessary to be able to model interconnections, feedbacks, travel, and heterogeneity of policy and energy sectors across spatial domains. Successfully addressing this gap requires developing a model that can represent important aspects of the national mobility-energy system at a time resolution smaller than one hour.

This research integrates a high-temporal-resolution (≤ 1 minute) microsimulation model into a scalable-extent, bottom-up, agent-based adoption model, with the result of filling the spatiotemporal modeling gap for the MDHD sector. Employing microsimulation to model MDHD operations enables investigating the feasibility and cost implications of factors that can only be approximated in models with coarser time resolutions. These include optimized charging and scheduling; electricity rate structures that might include coincident peak

demand rates, time-of-use rates, and real-time pricing; detailed dynamics of extreme fast charging; optimized dispatch and integration of DERs; and more.

- 2. A lack of bottom-up vehicle and charging infrastructure adoption models, in which adoption estimates arise from decisions made at an agent level.** A bottom-up approach is required to represent several important factors in the development of an electrified MDHD market. The coevolution of electric vehicle demand and charging infrastructure availability necessitates models that represent the dynamics of, and interactions between, supply and demand of vehicles and charging. Additionally, the heterogeneity of MDHD fleets and their decision-making processes necessitate models that represent decision-making at a high granularity. Successfully addressing this gap requires developing an agent-based model that represents agents including vehicle purchasers, vehicle and infrastructure manufacturers, and utility/EV infrastructure agents, and specific to various MDHD applications and local conditions.

1.7 Organization of chapters

In answer to Research Question #1, we use best practices from economic modeling of transportation system entities, including work by the Department of Energy (DOE), the Electric Power Research Institute (EPRI), the Rocky Mountain Institute (RMI), and others, to develop a comprehensive and scalable model of fleet electrification economics. In Chapter 2, we apply this model to assess the potential of an emerging EV charging business model, Charging-as-a-Service (ChaaS), in a variety of transportation sectors. In answer to Research Question #2, we interpret and build on a conceptual model of technology adoption which incorporates established theories of organizational decision making and which recent scholarship has applied to the specific domain of fleet decarbonization. To build a quantitative model using this framework, we apply automated text mining and machine learning techniques to a corpus of scholarly research characterizing preferences among vehicle and system attributes, which is described and demonstrated in Chapter 3. And in answer to Research Question #3, we leverage best practices in agent-based modeling and

social network modeling to integrate the models and methods, and apply them to the evaluation of scenarios, in Chapter 4. Combining and extending these strands of research is a necessary next step toward achieving urgent transportation and energy system goals. We conclude with an assessment of future research directions in Chapter 5 and a summary conclusion in Chapter 6.

Chapter 2

Electric vehicle and fleet economics modeling

2.1 Abstract

In support of global decarbonization efforts, the adoption of electric vehicles is proceeding rapidly across transportation sectors, population groups, and regions. However, access to electric vehicle charging infrastructure remains sparse and inequitable, and technical, economic, and procedural challenges have encumbered its expansion. Charging-as-a-service, which can mitigate cost and effort risks of charging equipment ownership, adds a means of meeting many of these challenges. As charging-as-a-service has not been analyzed or addressed in the scholarly literature, we synthesize media, marketing, and analogous scholarly publications to categorize and describe the functions asserted by existing charging-as-a-service vendors. We demonstrate and quantify the benefits of many of these functions via four example business cases, sited at a detached residence, a multi-unit dwelling, a consumer-facing business, and a commercial fleet depot. Outcomes show that savings and revenues realized via charging-as-a-service can reduce charging costs, beneficially shape grid loads, and make charging installation economical where it otherwise would not be. Discussion highlights the need to deploy charging-as-a-service in cooperation with initiatives promoting societal benefits, including charging access equity.

This work references a peer-reviewed publication entitled “Transportation and electricity systems integration via electric vehicle charging-as-a-service: a review of technoeconomic and societal benefits”, which has been published in *Renewable and Sustainable Energy Reviews* [20]. See Figure 2.1 for the graphical abstract included in the publication.

2.2 Introduction

Ground transportation is rapidly transitioning toward electric power as a key part of global efforts to meet decarbonization goals [21, 22]. Electric power as a transportation fuel offers the

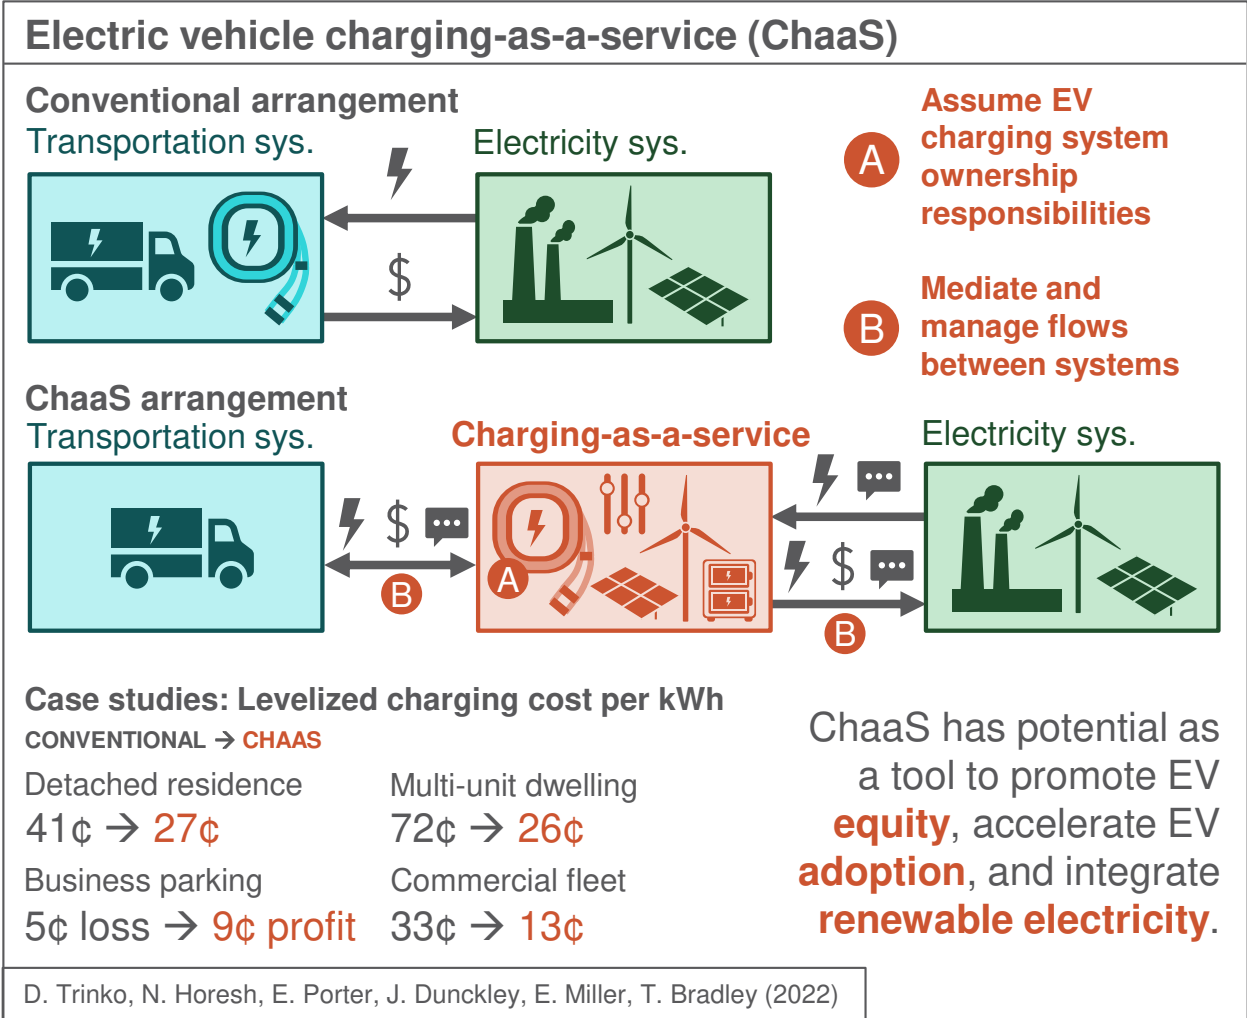


Figure 2.1: Graphical abstract.

benefits of eliminating tailpipe emissions, the potential to be generated from non-emitting sources, lower operating costs, and the potential to leverage the ubiquity of the existing electric grid for recharging [23]. However, access to EV charging infrastructure is currently sparse [24] and inequitable [25], and will require rapid expansion to enable a widespread and just electrification transition.

Early adopters of passenger EVs have recharged primarily at home [23], but it is unclear how ownership models for charging infrastructure will evolve as the vehicle electrification transition broadens across transportation sectors, population groups, and regions. Depending on charging power requirements and local grid characteristics, the costs to procure, install, and use electric vehicle service equipment (EVSE) can range widely. EVSE installed in a residence costs on the order of USD\$1,000¹ per plug, and provides regular charging at residential electricity prices [23]. In contrast, high-power EVSE installations can cost hundreds of thousands of dollars per plug if substantial electricity infrastructure upgrades are required [26–28]. Further, EV charging costs can exceed petroleum fueling costs when high-power charging is subject to expensive peak demand prices or experiences low utilization [29]. To date, these kinds of variability and uncertainty have posed obstacles to assessing and communicating the economics of EV charging, in turn posing an obstacle to the adoption of EVs [30].

Charging infrastructure embodies a set of tight temporal and spatial linkages between the transportation and electricity systems. From the electricity system perspective, utilities must equate energy supply and demand at a high resolution in time and space, a task which is made simpler by steady and predictable usage patterns. However, uncertainties in travel demand and charging behavior, particularly with the use of fast chargers, lead to unsteady and unpredictable charging demand patterns [31]. Optimizing to meet the needs of both systems must begin with a comprehensive understanding of those needs and how they interact. Instead of expecting EV users and electricity system operators to understand and adapt to one another’s comprehensive needs, a third

¹“\$” is used to denote U.S. currency throughout this article.

party can serve as a mediator, performing communication and optimization services for collective benefit.

The set of emerging infrastructure ownership and usage models involving this kind of third party are known as “Charging-as-a-Service” (ChaaS). This concept has not yet been addressed in the scholarly EV literature, despite being offered by many leading charging infrastructure companies. A few scholarly publications reference the term, but few align with the concept of ChaaS we observe in the industry. Many refer to ChaaS as any instance of publicly accessible charging stations [32–34], whereas others have employed the term to describe mobile charging “delivery” by means of large battery packs loaded on vehicles [35, 36]. Singh et al. identify some of the functions we observe in ChaaS in the context of other energy-related services [37].

“X-as-a-service” (XaaS) was originally associated with various functions provided in cloud computing, including communication-as-a-service (CaaS), infrastructure-as-a-service (IaaS), and software-as-a-service (SaaS) [38]. Over time, the XaaS term has been invoked to refer to offerings outside the realm of computing. Examples include lighting-as-a-service, in which a provider covers the up-front cost of energy efficient lighting fixtures that a client pays to use over time, and mobility-as-a-service, which unifies and facilitates access to shared travel resources [39]. We observe the common thread unifying XaaS frameworks to be the offloading of asset ownership responsibilities in exchange for paid access to those assets over time. Accessing these assets “as a service” is typically the direct alternative to owning those assets, and the “as-a-service” modifier indicates a deviation from an ownership model status quo. As with cloud computing, the asset provided “as-a-service” would otherwise be owned by its users in the dominant ownership model.

Most of the ChaaS models discussed in literature (“h” included to avoid confusion with CaaS models such as Containers- [40], Content- [41], and Communication-as-a-Service [42]) describe vehicle charging provision models that do not align with the definition of “as-a-service” we have observed. Publications referring to the general case of public charging stations as ChaaS do not specify a deviation from an ownership model status quo [32–34]. Similarly, “mobile charging delivery,” which has been labeled as ChaaS by some researchers and includes responding to emer-

gencies and providing convenience where charging is unavailable nearby, is unlikely to change vehicle operators' equipment ownership responsibilities [35, 36].

A variety of providers in the EV charging market (including charging network operators [43–46], equipment manufacturers [46, 47], electricity providers [48, 49], and consulting/management entities [50–54]) offer services that fit our ChaaS model. These ChaaS providers purchase and install EVSE at the site of the user and mediate exchanges between the user and the electricity provider. Recent marketing materials for these models emphasize their role in streamlining and speeding the design, installation, and maintenance of EVSE, and reducing up-front costs [43–55].

A recent review of charging infrastructure business models [56] offers a categorization of roles served by public EV charging networks for light-duty transportation, but does not call explicit attention to the role of ChaaS provider. The review highlights the need to assign clear roles among owners, hosts, operators, and other entities in the EV charging system to enable large-scale EV adoption, and raises the timely question of whether infrastructure deployment is a public or private responsibility. In this work, we offer a comprehensive definition of the roles of ChaaS provider and ChaaS customer, each of which can be fulfilled by public and private entities alike. We find ChaaS to be a flexible means for directing public and/or private investments toward the efficient rollout of charging infrastructure.

No comprehensive and non-commercial survey of ChaaS business cases as they appear in the market is currently available to the public. Many potential but yet-unrealized benefits of ChaaS are asserted by commercial providers, but no examples analyzing these asserted benefits are found in the current literature. This article provides an overview of what we understand to be the purpose and functions of ChaaS, economic analyses of four potential use cases for ChaaS, and discussion of the various types of benefits ChaaS can provide to the transportation and energy systems to enable increased adoption and equitability of electric transportation.

2.3 System model of Charging-as-a-Service

The transition to EVs involves establishing a new type and scope of interface between the electricity and transportation systems in the form of vehicle charging. Under dominant EV charging models, this interface is a locus of confusion, miscommunication, and general inefficiency, because it is where the motivations and imperatives of the two previously disconnected systems must be coordinated.

For example, EV drivers benefit from speed and flexibility in charging, whereas electric power providers benefit from predictability and stability in power demand. These incentives can contradict one another, and their reconciliation can tip in favor of one entity or the other. If electricity providers assess electricity prices on an energy basis (\$/kilowatt-hour, kWh) irrespective of time or power demand, which is typical of residential price structures, they incentivize charging flexibility at the risk of distribution system stability (for example, if many EVs charge simultaneously in the same neighborhood and overload the local distribution infrastructure [57]). On the other hand, if electricity providers assess a price for peak demand, which is typical of commercial price structures, they incentivize users to forgo high-power or simultaneous charging, which can restrict the availability of their EVs to provide transport by necessitating long dwell times. The inflexibility of this type of price structure has been identified as a factor in decisions not to adopt EVs in commercial fleets [29, 30].

These and other compromises and incompatibilities at the interface of the two systems may be addressed by an intermediary providing ChaaS. This provider can be understood as a conduit between the supply side (electricity provider) and the demand side (charging customer), facilitating flows of information, electric power, and capital between the two sides. In this framework, which is illustrated in Figure 2.2, ChaaS can serve many functions spanning a range of complexity. A simple example is for a ChaaS provider to pay for the purchase and installation costs for residential EVSE and assess a usage fee, thereby easing customers' capital expenditure burdens. A more complex example is to install and dispatch distributed energy resources (DERs) for supplemental power,

managing vehicle charging and discharging, and providing and selling ancillary services to the electric grid, thereby easing utilities' load integration challenges.

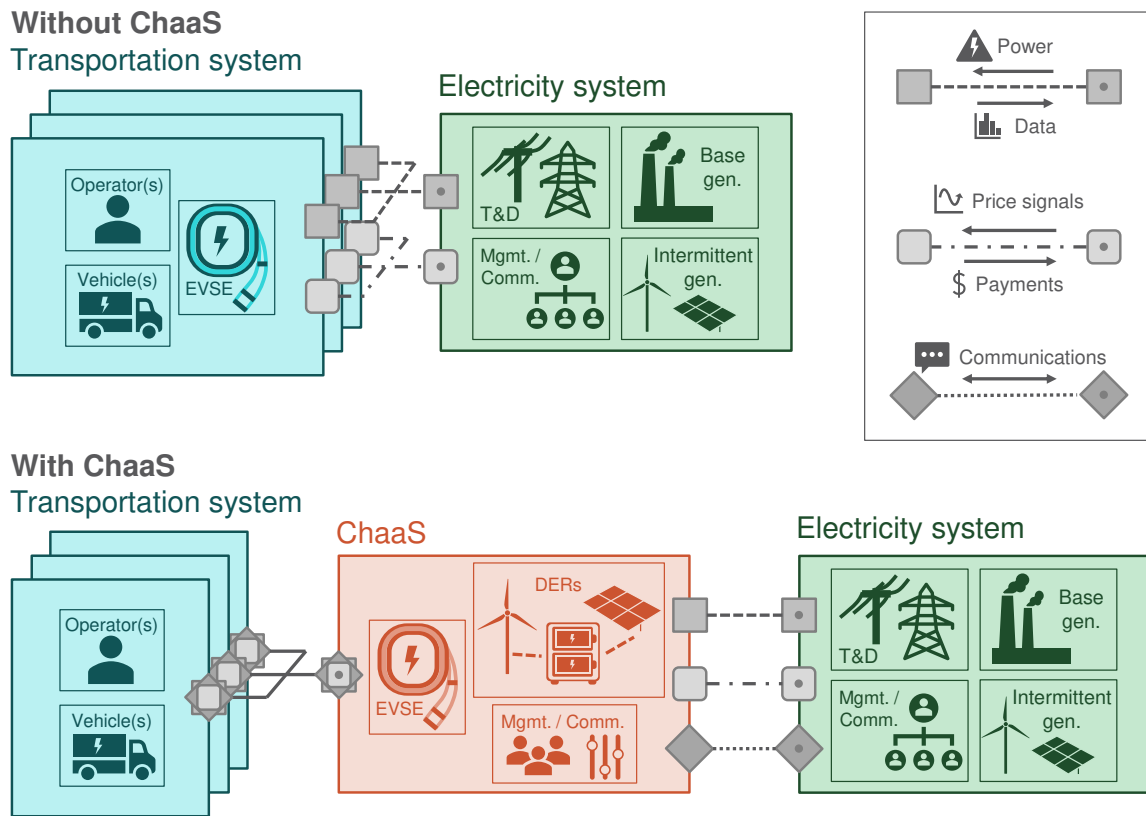


Figure 2.2: System diagram illustrating the roles Charging-as-a-Service can play at the interface between the transportation and electricity systems. ChaaS can serve as a unified conduit for electric, economic, and information flows, and can help manage risks for transportation system and electricity system entities by assuming EVSE ownership responsibilities. Abbreviations: Distributed Energy Resources (DERs); Electric Vehicle Service Equipment (EVSE); Transmission & Distribution (T&D).

ChaaS provides value by mediating flows between systems and/or by offloading risks faced by one or both systems. By mediating flows and communications, ChaaS may address problems such as suboptimal charging system design and operation due to lack of design expertise, communication obstacles between entities, regulatory obstacles (“red tape”), and miscommunications regarding the suitability and functionality of EVs [58–67]. Examples of risks faced by EV customers addressable by ChaaS providers include unexpected electricity costs, equipment upkeep

costs, and obsolescence costs [43, 44, 46, 47, 51, 54, 55]. Examples of risks faced by utilities and/or site owners addressable by ChaaS providers include low equipment utilization and variability in power demand [45, 46, 50–52, 54]. By renegotiating who bears these risks, ChaaS can help to balance the allocation of the costs and benefits of transportation electrification across systems and stakeholders. With this in mind, policymakers and regulators should assess the potential for ChaaS as a tool to advance societal interests.

As with any privately-operated entity providing services to the public, appropriate oversight will be necessary to enable equitable societal benefits and avoid unjust or monopolistic outcomes. For example, protecting ChaaS customers from unjustified price increases may require regulatory intervention, particularly if customers are contractually “locked in” to service from their ChaaS providers. Lessons learned from regulating electric utilities should serve as guidelines for developing a ChaaS regulatory framework [68]. Although it is out of scope for this article, suitable ChaaS regulation is an urgent challenge deserving attention from the research and policymaking communities.

2.4 Functions of charging-as-a-service

In the general business case for ChaaS, a ChaaS provider assesses fees to its customers in exchange for the use of EVSE, and for related services, such as equipment installation, integration, upkeep, monitoring, and management. We assign the potential functions of ChaaS to four broad categories: economic, technical, managerial, and societal. These categories and functions are summarized in Table 2.1. The following sections describe each category in detail.

Table 2.1: Summary of economic, technical, managerial, and social functions of charging-as-a-service.

Category	Subcategory	Function
Economic	Risk mitigation	Mitigating stranded asset risk [43, 44, 46, 47, 55] Mitigating cost variability risk [44, 46, 49, 51, 54, 55]
	Cost reduction	Deferring up-front costs [44, 46, 50, 53, 54] Optimizing system design/operation for cost [45, 47, 48, 51, 52] Obtaining favorable electricity prices [48, 49] Participating in incentive programs [49, 51, 52]
	Revenue generation	Selling charging services [45, 46, 50] Participating in electricity markets [51, 54] Engaging customers for other business [45, 46, 50]
Technical	Expert design	Suitable selection of equipment [44, 47, 48, 51–53]
	Expert operation	Charging management for performance [45, 47, 52, 54] Leveraging on-site generation and storage resources [47, 54] Obtaining and analyzing usage data [45–47, 49, 52, 53]
Managerial	System interfacing	Interfacing with electricity provider [47] Obtaining permits / complying with regulations [47, 52] Simplifying billing [43]
	Life cycle management	Managing design and construction [44, 46, 47, 51] Managing / automating operations [44, 47, 49, 51–53, 55] Managing interactions with secondary customers [43, 45, 46, 53] Accessing dedicated EV customer service [44, 46, 52] Maintaining / updating equipment [43–47, 51–53, 55] Handling system end-of-life [69]
Societal	EV adoption acceleration	Improving economics for EV adopters [49, 54, 55] Streamlining the electrification process [47]
	EV equity improvement	Reducing barriers to low-income electrification [69] Decoupling EVSE economics from utilization [69] Broadening availability of charging resources [55]
	Decarbonization facilitation	Using electricity from renewable sources [48, 51] Participating in decarbonization incentive programs [43, 51]

2.4.1 Economic functions

The economic functions that ChaaS can serve for a customer include mitigating economic risks, reducing costs, and generating revenues, as summarized in Table 2.1.

Risk mitigation is an inherent function of ChaaS because owning an expensive asset, such as EVSE, is riskier than paying to access an asset owned by someone else. ChaaS providers leverage a high risk tolerance and/or access to capital to bear the ownership risks of EVSE, and EVSE users

pay over time to avoid these risks. Ownership risks include the risk that EVSE will unexpectedly lose value (become “stranded”), or, on the other hand, the risk that EVSE will incur unexpected costs. EVSE can lose value for a variety of reasons (repair, abuse, neglect, natural disaster), but a key reason is potential obsolescence. Historically, EVSE connector standards have changed and been discontinued (CHAdeMO [70], AVCON [71], and Magne Charge [72]), and charging rate capabilities have increased rapidly over time [73]. If charging standards continue to change, or if competing EVSE options with improved charging rate capabilities continue to become available, EVSE owners face serious obsolescence risks.

EVSE owners also risk incurring unexpected costs. EVSE has earned a reputation for unreliability [74–76], which, in addition to disrupting operations, can lead to high unplanned maintenance and/or replacement costs. Charging costs can also vary unpredictably, over the short term due to supply-demand dynamics and over the long term due to global politics and other factors. ChaaS providers can regulate operating costs for EVSE users by managing the maintenance and replacement of EVSE, and by creating and enforcing variability limits on charger usage costs. By these and other means, a risk-tolerant ChaaS provider can shield a comparatively risk-averse EV operator from obsolescence and cost risks. However, EV operators participating in ChaaS must accept the risks associated with depending on infrastructure they do not own, including the risk of losing access to charging.

ChaaS can also **reduce or defer costs** for customers. The costs to procure charging infrastructure can be prohibitively high [26–28]. A ChaaS provider can pay or finance the up-front costs of a charging system and assess usage fees. This results in a deferment of up-front equipment costs for the system users that may make EV adoption economically favorable for some customers.

ChaaS providers can also help customers to take advantage of favorable electricity price structures, incentive programs, and policy initiatives that they could otherwise be unaware of or lack the knowledge to navigate optimally. These tasks become especially complex when charging behaviors must be adapted in order to actualize incentives, such as those in time-of-use (TOU) price schedules or demand response programs. A ChaaS provider can offer knowledge and expertise to

actualize all available incentives, and even exert leverage on behalf of their customers to enhance existing opportunities or drive additional opportunities.

In addition to cost savings, ChaaS providers can aid with **revenue generation**. Revenue opportunities include selling charging to secondary customers, selling services to the electric grid, and attracting customers for other business. ChaaS providers can actualize these opportunities if EV operators are unable to because they lack knowledge, time, or scale. To sell charging to secondary customers, a ChaaS provider can provide marketing and/or payment services (for example, if it operates a larger network of EVSE). Participation in electricity markets and provision of grid services are often blocked by procedural obstacles and limited to large power producers. Other markets, such as the market for low-carbon fuel standard (LCFS) credits in California [77], can be similarly complex to navigate. A ChaaS provider can offer technical and procedural assistance to enable fleets to access these and other otherwise unreachable markets.

2.4.2 Technical functions

EV charging systems are best designed with expert knowledge and experience. EVSE system design support documents compiled for consumers [78, 79], residents of multi-unit dwellings (MUDs) [80], businesses [81, 82], and fleets [83] span dozens of pages and topics, which ultimately serve to highlight the unavoidable complexity of designing an EV charging system. Inexperienced system design invites an array of serious risks, beginning with over- or under-specification of equipment, which can increase costs and/or diminish performance. More innovative technical additions to a system, such as automated charging management (smart charging) or the integration of DERs, can drastically reduce operating costs and improve performance (as we demonstrate in Section 2.5.4), but further increase the complexity of design and operation. To address these challenges and enable innovative benefits, ChaaS companies can provide **design and operations expertise** to their customers. ChaaS providers can serve some of these functions using software that they own or license.

2.4.3 Managerial functions

The managerial functions served by ChaaS can be divided into categories of system interfacing and product life cycle management. **System interfacing** tasks include choosing pricing models, identifying applicable programs and incentives, implementing grid connection upgrades, responding to price signals, identifying and following all pertinent regulations, and other responsibilities [79, 81–83]. By offloading these tasks to a ChaaS provider, the user can avoid time-intensive responsibilities and the costs of potential errors in completing them, such as noncompliance penalties. ChaaS providers can also **manage the life cycle** of an EV charging system, which includes design, procurement, construction, operation, maintenance and support, and end-of-life. For many, the capacity to host EVSE may be incumbent on the ability to offload life cycle management responsibilities such as planning, repair, replacement, and disposal, which is a lesson the industry has learned by experience. An early charger-leasing program offered by the Blink network required hosts to manage site preparation and construction, which few were equipped to do, ultimately leading the program to be discontinued [56, 84].

2.4.4 Societal functions

ChaaS can function for societal benefit, including by helping to accelerate the adoption of EVs, addressing equity challenges in electric transportation, and reducing the carbon emissions impact of transportation. Policy incentives or mandates to support these societal aims are in place or are being considered and would enable such societal benefits to be internalized to the ChaaS business model.

Electrification of transportation offers the potential to reduce the environmental impact of the transportation sector, which would benefit human health and mitigate emissions driving climate change [85]. Major barriers to adoption of EVs include unfavorable economics (in terms of cost and risk), technical challenges including lack of access to charging, and managerial challenges such as burdensome EVSE procurement processes [8, 30]. By performing the functions enumerated in the previous sections, ChaaS diminishes these barriers and therefore promotes the uptake of EVs.

Similarly, EV adoption to date has been largely inaccessible to low-income populations due to the high cost of vehicles and EVSE, and due to a lack of charging infrastructure in low-income areas and at MUDs [25, 86]. ChaaS can **improve the equity of electric transportation** by facilitating access to EV charging by low-income and otherwise underserved populations. If a ChaaS provider operates a network of charging infrastructure, they have an incentive to install and maintain infrastructure in areas where standalone infrastructure may not be economical due to low utilization. Whereas a low-utilization charger may not be able to recoup its costs on its own, it can add value to a network by extending the network's reach, including to low-income areas. ChaaS networks thus serve to aggregate benefits of EVSE across users and communities.

EV populations can serve as flexible loads by following the incentives of TOU rates, by participating in demand response programs, or even serving as storage resources, all of which can increase grid integration of renewable and intermittent power sources. ChaaS can reduce the technical complexities of participating in these programs, thereby increasing the utilization of EVs for the purpose of demand management and further **promoting uptake of zero-carbon electricity**.

Summary

The preceding sections list and describe the economic, technical, managerial, and societal functions that ChaaS has been advertised to offer. All of the functions branch out from the fundamental ChaaS function: offloading the responsibilities of owning and operating EVSE from the EV user, and aggregating such responsibilities to a provider that is best positioned to design and implement efficient EVSE systems and operate them effectively. The breadth of functionality indicates potential for ChaaS to provide adaptive value to customers with varying priorities and demands. Customers can range from private entities (consumers, fleets) to representatives of the public interest (governments, nonprofit entities), making ChaaS a versatile tool for expanding, improving, and reducing costs of EV charging.

2.5 Business cases

To quantitatively illustrate some of the economic benefits of ChaaS introduced in Section 2.4.1, we present economic analyses for four example business cases, each representing a different scope and sector. These include (1) detached residences, (2) MUDs, (3) commercial buildings providing parking for customers, and (4) commercial EV fleets. Economic outcomes are sensitive to a variety of specific assumptions, all of which can vary significantly across regions, vehicle models, decision makers, markets, and time periods. Quantitatively analyzing the full range of business cases for ChaaS, which would require an exhaustive sensitivity analysis, is outside the scope of this paper. Because our scenarios and assumptions span only a narrow range, we intend our examples as restricted demonstrations of specific ChaaS functions, not as comprehensive valuations of the ChaaS concept.

Economic analysis methodology

For each business case, we present and compare summary economic metrics for two or more ownership and operation scenarios. Estimating these metrics begins with computing yearly cash flows CF_y , paid by the EVSE user. These cash flows consist of equipment costs, installation costs, electricity costs, ChaaS fees, and system revenues:

$$CF_y = CF_y^{\text{equip}} + CF_y^{\text{inst}} + CF_y^{\text{elec}} + CF_y^{\text{ChaaS}} + CF_y^{\text{rev}} \quad (2.1)$$

where equipment and installation costs CF_y^{equip} and CF_y^{inst} occur in year $y = 0$; ChaaS fees CF_y^{ChaaS} are assessed such that ChaaS providers operate profitably; and revenues $CF_y^{\text{rev}} \leq 0$ are generated via sale of charging (case 3) or LCFS credits (case 4). Electricity costs CF_y^{elec} are a function of yearly driving distance d (km), vehicle energy consumption rate r (kWh/km), electricity prices C (for energy consumed, \$/kWh) and D (for peak demand, \$/kW-month), and the timing and power of charging P (kW, for each hour h of each month m of each year y):

$$\text{CF}_y^{\text{elec}} = \sum_{h \in y} C_h P_h + \sum_{m \in y} D_m \max_{h \in m} P_h \quad (2.2)$$

C and D can vary over time and represent case-appropriate pricing schedules from Southern California Edison (SCE), in the Los Angeles area.² The schedules used in each case are detailed in relevant sections of the Appendix.

Hourly charging power P must meet constraints of vehicle operation, including that all energy consumed during driving is replenished via charging:

$$\sum_{h \in y} P_h = \frac{dr}{\eta} \quad (2.3)$$

where charging efficiency η is assumed to be 0.9. P is further constrained by the timing of driving requirements, battery capacity constraints, and EVSE power constraints. To simulate automated charging management in ChaaS scenarios (and all Case 4 scenarios), P is optimized via linear programming to minimize CF^{elec} subject to these constraints.

Summary metrics are net present cost (NPC) and levelized cost of charging (LCOC):

$$\text{NPC} = \sum_{y=0}^{y_f-1} \frac{\text{CF}_y}{(1+i)^y} \quad (2.4)$$

$$\text{LCOC} = \text{NPC} \times \left(\sum_{y=0}^{y_f-1} \frac{\sum_{h \in y} P_h}{(1+i)^y} \right)^{-1} \quad (2.5)$$

where discount rate i is assessed yearly, and y_f is the charging system lifespan. Other costs pertinent to EVSE ownership and operation, such as labor, maintenance, networking, payment system contracts, and software, are considered out of scope.³

²Many SCE pricing structures include incentives in the form of peak demand prices and TOU price schedules, which we leverage to realize savings. Some utilities do not offer these kinds of incentives, instead supplying energy at a static price (\$/kWh), which may limit their customers' potential for electricity cost savings via ChaaS.

³Although we estimate that additional operating costs would be small relative to the costs in our model, they are nonzero and would increase the fees assessed by the ChaaS provider in ChaaS scenarios.

2.5.1 Detached residence case

Typically, the majority of EV charging at a detached residence (often a single-family home, or SFH) is supplied by Level 2 EVSE⁴ installed to the home electric system, such as in a garage [78, 79]. Residential ChaaS programs are currently offered by many electricity providers and independent entities for SFH [48, 49]. This example demonstrates the potential economic savings from actualizing favorable electricity pricing structures via ChaaS-controlled home charging of a commuter EV. These savings accrue relative to a baseline case, where the EV owner pays up front for EVSE procurement and installation and does not take advantage of any special electricity pricing structures. Key assumptions used in our analysis are shown in Table 2.2.

Table 2.2: Selected assumptions used in the single-family-home business case. A complete version of this table can be found in A.1.1.

<i>Category</i>	<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Economics	Annual discount rate (ChaaS provider)	7%	[87]
	Annual discount rate (resident)	20%	[88]
	System lifespan	10	years
	Electricity pricing (baseline)	Tiered residential rate	A.1.1
	Electricity pricing (ChaaS)	Residential EV rate	A.1.1
Operation	Non-EV electricity consumption	4,000	kWh/year [89]
	Driving distance	18,500	km/year [90]
Vehicle	Vehicle type	Tesla Model 3	[91]
	Electricity consumption rate	161	Wh/km [91]
	Battery capacity (usable)	80.5	kWh [91]
	Battery range	500	km [91]
	EVSE power	7.2	kW

4 cash flows are relevant to this case: CF^{equip} , CF^{inst} , CF^{elec} , and CF^{ChaaS} . In the ChaaS case, the ChaaS provider pays CF^{equip} and CF^{inst} . In exchange, the user pays an access fee CF^{ChaaS} , which the ChaaS provider sets in order to achieve a 7% annual return on investment over 10 years (i.e., break even with a 7% discount rate) [87]. In both scenarios, the EV user pays CF^{elec} , which

⁴Level 2 chargers supply alternating current (AC) and can provide up to 19.2 kW. Higher power charging (50+ kW) is supplied by direct current fast chargers (DCFC).

varies depending on the pricing structure and charging management strategy. In the ChaaS case, the ChaaS provider automates the timing of charging.^{5,6} In the baseline case, charging timing does not affect CF^{elec} because the baseline electricity price does not vary with time.

Results

The savings due to procurement cost deferral (“procurement cost savings”) scale with the difference between the user’s discount rate (20%) and the ChaaS provider’s (7%)⁷ Under our assumptions, the ChaaS provider can assess a \$23.44 monthly usage fee to recoup the initial cost plus a 7% annual return over 10 years. For the end user, 10 years of \$23.44 monthly payments equate to an LCOC of \$0.078/kWh, compared to \$0.122/kWh if they pay up front for procurement and installation. System cost deferral thus yields a procurement cost savings of \$0.044/kWh, as shown in Figure 2.3.

Electricity cost savings can be realized by managing the timing of charging subject to a dedicated TOU rate schedule, in which charging prices increase at high-demand times (4pm-9pm). By avoiding high charging prices, the residence saves \$310 per year compared to a non-EV-specific rate schedule. This translates to an LCOC savings of \$0.096/kWh. Together with the procurement cost savings, ChaaS reduces the LCOC paid by the consumer by \$0.14/kWh (from \$0.41/kWh to \$0.27/kWh). The user must balance these savings against vehicle operation requirements. If charging during high-demand times cannot be avoided, the savings potential diminishes accordingly.

⁵ Optimally managed charging would occur exclusively during the lowest-price time periods, which are in the early morning. However, if this timing were incompatible with EV usage requirements, supplemental higher-cost charging would be necessary.

⁶ Some costs to enable smart charging, including costs for any necessary equipment, are neglected for simplicity. However, equipment compatibility is an important consideration. For example, some ChaaS programs restrict participants’ EVSE choices for controls compatibility [48].

⁷ The procurement cost savings are highly sensitive to the discount rate assumptions. If the resident and ChaaS provider were to assess equivalent discount rates, ChaaS would only be able to offer electricity cost savings.

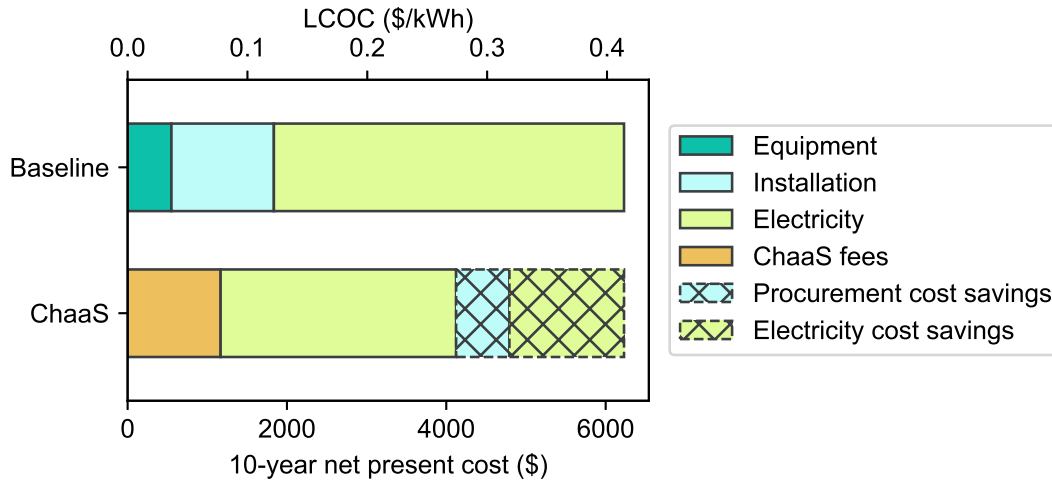


Figure 2.3: Breakdown of costs paid by the customer in the detached residence example, expressed both as a net present cost and as a levelized cost of charging (LCOC).

Summary

In Case #1, two functions of ChaaS (one, deferment of up-front costs, and two, charging management) reduce LCOC for SFH charging by approximately one third. ChaaS provides additional noneconomic sources of value to the EV user. In the sense that refueling infrastructure is conventionally not owned by consumers, this ChaaS arrangement bears some resemblance to the petroleum status quo. Charging infrastructure installation, maintenance, repair, and replacement are novel burdens for new EV adopters, and even if the monetary costs of these responsibilities are small, consumers may perceive their noneconomic costs (inconvenience, risk, unfamiliarity) to be substantial.

2.5.2 Multi-unit dwelling case

Residents of MUDs are much less likely to have access to home charging than SFH residents [92–94]. Unlike for a SFH, there is not an obvious default ownership model for EVSE, which can create barriers for MUD residents wishing to charge at home.⁸ Further, electrical infrastructure is often unavailable near MUD parking, which can necessitate costly line extensions. Finally, if

⁸In many states, “right-to-charge” laws have been proposed or enacted to reduce some of these barriers [95].

departing renters have ownership stake in MUD EVSE, the means and procedures for ownership transfer are not generally defined. Some ChaaS offerings have specifically targeted these challenges [43], functioning to provide home charging access to residents who otherwise would have to charge elsewhere, mitigate or eliminate equipment costs for building owners, and manage the life cycle of the charging system.

Unlike at a SFH with a single EV, EVSE at MUDs can be shared by multiple EVs, either by physically transferring plugs between vehicles or by attaching redundant plugs to each charger and automating power switching between them. In this business case, a ChaaS provider manages EVSE sharing between residents to reduce CF^{equip} and CF^{inst} , and automates charging management to reduce CF^{elec} , in alignment with the usage needs of the residents. EVSE sharing would incur additional, uncertain costs, such as for labor or for redundant plugs, but we assume these costs would be small relative to overall system costs.

The result is contrasted against two different baseline scenarios. In the first baseline scenario, the MUD does not host EVSE, necessitating residents to charge at public EVSE (“public charging” scenario). In the second, the MUD hosts EVSE, but chargers are not shared and residents begin charging their EVs immediately upon arrival at home, without concern for temporal price variation (“suboptimal MUD charging” scenario). In all cases, 15 residents of a MUD are each assumed to operate the same vehicle, driving the same yearly distance, as in the SFH case. Other key assumptions are described in Table 2.3.

Table 2.3: Selected assumptions used in the MUD business case. A complete version of this table can be found in A.1.2.

<i>Category</i>	<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Economics	Annual discount rate (ChaaS provider)	7%	[87]
	Annual discount rate (resident)	20%	[88]
	System lifespan	10	years
	Electricity pricing	Commercial EV rate	Table A.2
	Public charging price	0.31	\$/kWh [96]
Operation	Driving distance per vehicle	18,500	km/year [90]
Vehicle	Vehicle type	Tesla Model 3	[91]
	Electricity consumption rate	161	Wh/km [91]
	Battery capacity (usable)	80.5	kWh [91]
	Battery range	500	km [91]
	EVSE power	7.2	kW
	Vehicle count	15	

Results

Cost outcomes are summarized in Figure 2.4. In the suboptimal MUD charging scenario, the average LCOC for residents is \$0.72/kWh, more than double the public charging cost. In the ChaaS scenario, managing charging to optimize electricity costs reduces the LCOC by up to \$0.22/kWh, and optimal sharing of EVSE reduces LCOC by a further \$0.23/kWh, to an overall LCOC of \$0.26/kWh, for a savings of \$0.05 relative to public charging.




Figure 2.4: Breakdown of costs paid by residents in the MUD example, expressed both as a net present cost for a group of 15 vehicles and as a levelized cost of charging (LCOC).

Summary

Optimal charger sharing and charging management generate savings of more than 60% compared to a suboptimal residential charging scenario, and of about 15% relative to public DCFC. MUD residents unwilling to curtail or share charging time would have less savings potential. However, there are many other reasons for a resident to prefer home charging to public charging, including convenience and reliability. Overcoming the myriad procedural and policy barriers to installing EVSE at MUDs, including the misalignment of incentives for MUD owners/managers, would realize significant noneconomic value.

2.5.3 Business parking case overview

Businesses often install EVSE to supply charging for existing parking spots, whether to attract customers to public businesses (e.g. grocery stores, restaurants, hotels) [81, 96] or to offer charging to employees while they are at work [82]. A ChaaS provider can agree to pay or finance CF^{equip} and CF^{inst} for EVSE hosted by such a business, in exchange for CF^{ChaaS} . A negative net cash flow for the host, equal to the sum of (negative) charging sales revenues CF^{rev} and the host's total costs ($CF^{\text{ChaaS}} + CF^{\text{elec}}$), represents profit for the host. We analyze a scenario in which ChaaS

additionally provides “peak shaving” services to reduce peak demand costs, either by pooling load across a network or by integrating and dispatching energy storage resources. We assume this is done optimally, such that the charging load factor is equal to the charger utilization. In the scenario without ChaaS, we assume all chargers are in simultaneous use at least once during each month with no demand mitigation measures, resulting in maximum peak demand costs. Key assumptions are given in Table 2.4.

Table 2.4: Selected assumptions used in the business parking case. A complete version of this table can be found in A.1.3.

<i>Category</i>	<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Economics	Annual discount rate (ChaaS provider)	7%	[87]
	Annual discount rate (site host)	7%	[88]
	System lifespan	10	years
	Electricity pricing	Commercial TOU rate	Table A.4
	Mean energy price	0.127	\$/kWh
	Peak demand price	11.33	\$/kW-month

Results

As shown in Figure 2.5, profitability is a function of charger utilization, the price set by the host for EV charging, charging power per plug, and the presence or absence of ChaaS peak shaving. As utilization increases, the host can set lower prices for charging and remain profitable, with or without ChaaS. ChaaS peak shaving generally improves profitability, but the benefit diminishes as utilization increases, ending with zero incremental benefit at 100% (24 hours/day) utilization, where peak shaving would have no effect. In such a case, value streams other than peak shaving would be required for ChaaS to produce economic benefits.




Figure 2.5: Charging sales price, in \$/kWh, and utilization, in hours per day, required for a charging system host to achieve a net profit over the system’s lifetime. In the “No ChaaS” scenarios, peak demand is maximized (all chargers are assumed to be used simultaneously at least once per month). In the “ChaaS” scenarios, peak demand is minimized via load pooling, shifting via DERs, or other means, achieving a theoretical lower bound to electricity cost. The savings realized by ChaaS substantially reduce the price that must be assessed to achieve a profit, particularly at low utilization levels.

Figure 2.6 illustrates the levelized revenues (i.e., negative LCOE) if the host sells charging for \$0.31/kWh [96] and the system is utilized 2.7 hours per day, the average public charger utilization level projected for 2025 by one study [97]. At this price and utilization level, peak shaving is necessary to realize a profit at any of the three EVSE power capacities.

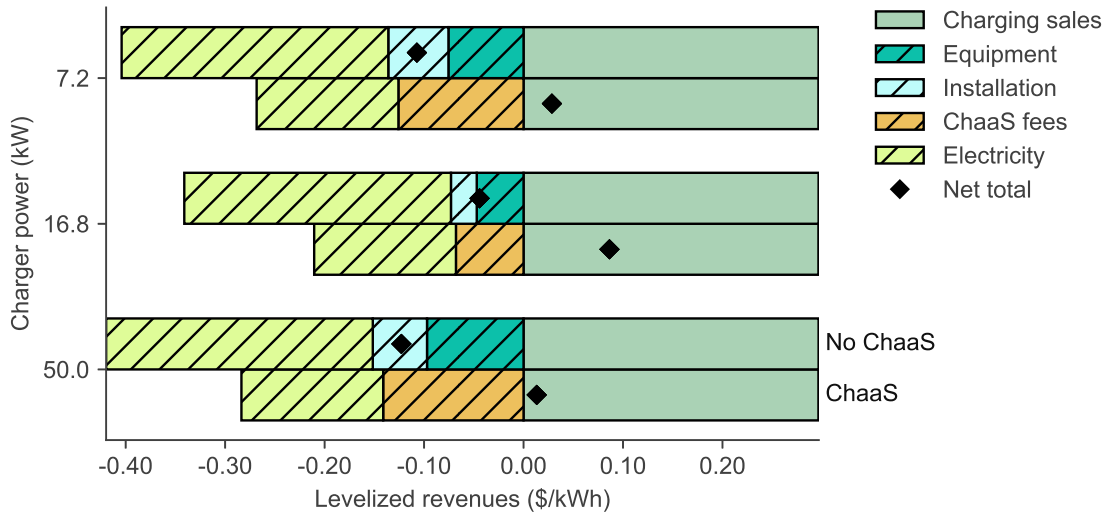


Figure 2.6: Breakdown of value streams paid and accrued by station hosts in the business parking case, expressed as a levelized revenue per kWh sold, when the charging system is utilized 2.7 hours per day and charging is sold at \$0.31/kWh.

Summary

In this example, we demonstrate how profitable charging sales depend on price setting and charger utilization, with peak shaving improving profitability at low utilization. Optimal peak shaving would be difficult to realize even via ChaaS and thus represents a lower bound on electricity costs. Even slight reductions to utilization, or slight increases to electricity costs due to suboptimal peak demand management, diminish profit potential. Further, our analysis does not account for the costs of actualizing peak shaving, which would be substantial if performed using energy storage resources. However, the fact that privately hosted public charging continues to grow, despite questionable profitability [56], suggests that hosts seek value beyond direct profits from charging sales, including attracting customers, amplifying brand image, and expanding charging networks. ChaaS can help to realize these other sources of value by providing and/or managing a platform for finding and utilizing hosted stations, designing and managing the system to optimize utilization, and bearing the financial and non-financial responsibilities of EVSE installation, ownership, and upkeep.

2.5.4 Commercial fleet case overview

For this final business case, concerning commercial EV fleets, we investigate the costs and benefits of applying more innovative cost- and energy-saving techniques via ChaaS. In addition to EVSE sited at fleet depots, the ChaaS provider owns a commercially sized microgrid, equipped with a solar photovoltaic (PV) array to generate power and a storage battery to enable adjusting when power is drawn from the grid. The ChaaS provider also manages the accrual and sale of LCFS credits for EV charging. Because LCFS credits scale with the carbon intensity of the energy source, these two functions are interrelated, including that increased usage of the PV array increases the revenue possible from LCFS. Although fleets are not prevented from generating and selling LCFS credits independently, few electric fleets have done so to date [77]. Recognizing the complexity of applying for, obtaining, and selling LCFS credits, we assume all steps of the LCFS process to be handled by the ChaaS credit accounting provider. We also assume the microgrid can be used to serve multiple separate fleets.

We estimate the costs associated with charging multiple geographically proximal, operationally diverse fleets, all of which are subject to a single electricity pricing model and are served by a single ChaaS provider. In our ChaaS scenario, all EVSE and microgrid resources are jointly owned by the ChaaS provider and are controlled to minimize net operating expenses, including revenues from LCFS. In our baseline scenario, each fleet controls charging to minimize cost, but they do not access microgrid resources or LCFS revenues. We estimate costs for varying numbers of fleets, ranging from 1 to 25, each of which comprises 20 Class 3 electric vans. Each fleet functions for either package delivery or passenger transport, the driving and charging behaviors of which are modeled after those of real electric fleets. Key assumptions are summarized in Table 2.5.

Table 2.5: Selected assumptions used in the commercial fleet business case. A complete version of this table is provided in A.1.4.

<i>Category</i>	<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Economics	Annual discount rate (ChaaS provider)	7%	[87]
	Annual discount rate (commercial fleet)	7%	[87]
	System lifespan	10	years
	Electricity pricing	Commercial EV rate	Table A.2
	LCFS credit price	150	\$/metric ton [98]
DER system	Battery storage power capacity	600	kW [99]
	Battery storage energy capacity	2.4	MWh [99]
	PV nominal power capacity	1	MW [99]
	PV array type	fixed rooftop panels	[99]
	DER system cost	1.97	\$million [99]
Operation	Mean, maximum daily distance (passenger)	91.4, 337.9	km
	Mean, maximum daily distance (cargo)	80.4, 271.0	km
Vehicle	Vehicle type	electric Class 3 van	[100]
	Electricity consumption	327	Wh/km [100]
	Battery capacity (usable)	73.6	kWh [100]
	Electric range	225	km [100]
	Number of vehicles per fleet	20	vehicles

Results

Figure 2.7 illustrates the coordinated dispatch of EVSE and DERs in the ChaaS scenario, in contrast with baseline charging operation, when ChaaS serves 3, 10, and 25 fleets. Whereas baseline charging management strictly avoids high electricity prices in the evening, ChaaS charging management optimizes across several objectives: actualizing and maximizing LCFS revenues based on the time-varying value of LCFS credits; reducing the total energy purchased from the grid by means of DER system resource dispatching; avoiding the purchase of expensive evening power; and meeting the various constraints of the DER system and the energy requirements of the fleets. As it serves more fleets, the DER system’s capacity is more fully utilized, while its relative

savings potential decreases. At the extreme, when serving 25 fleets, the peak charging demand is quadruple the PV generation peak, limiting the DER system's ability to meaningfully reduce grid power demand.

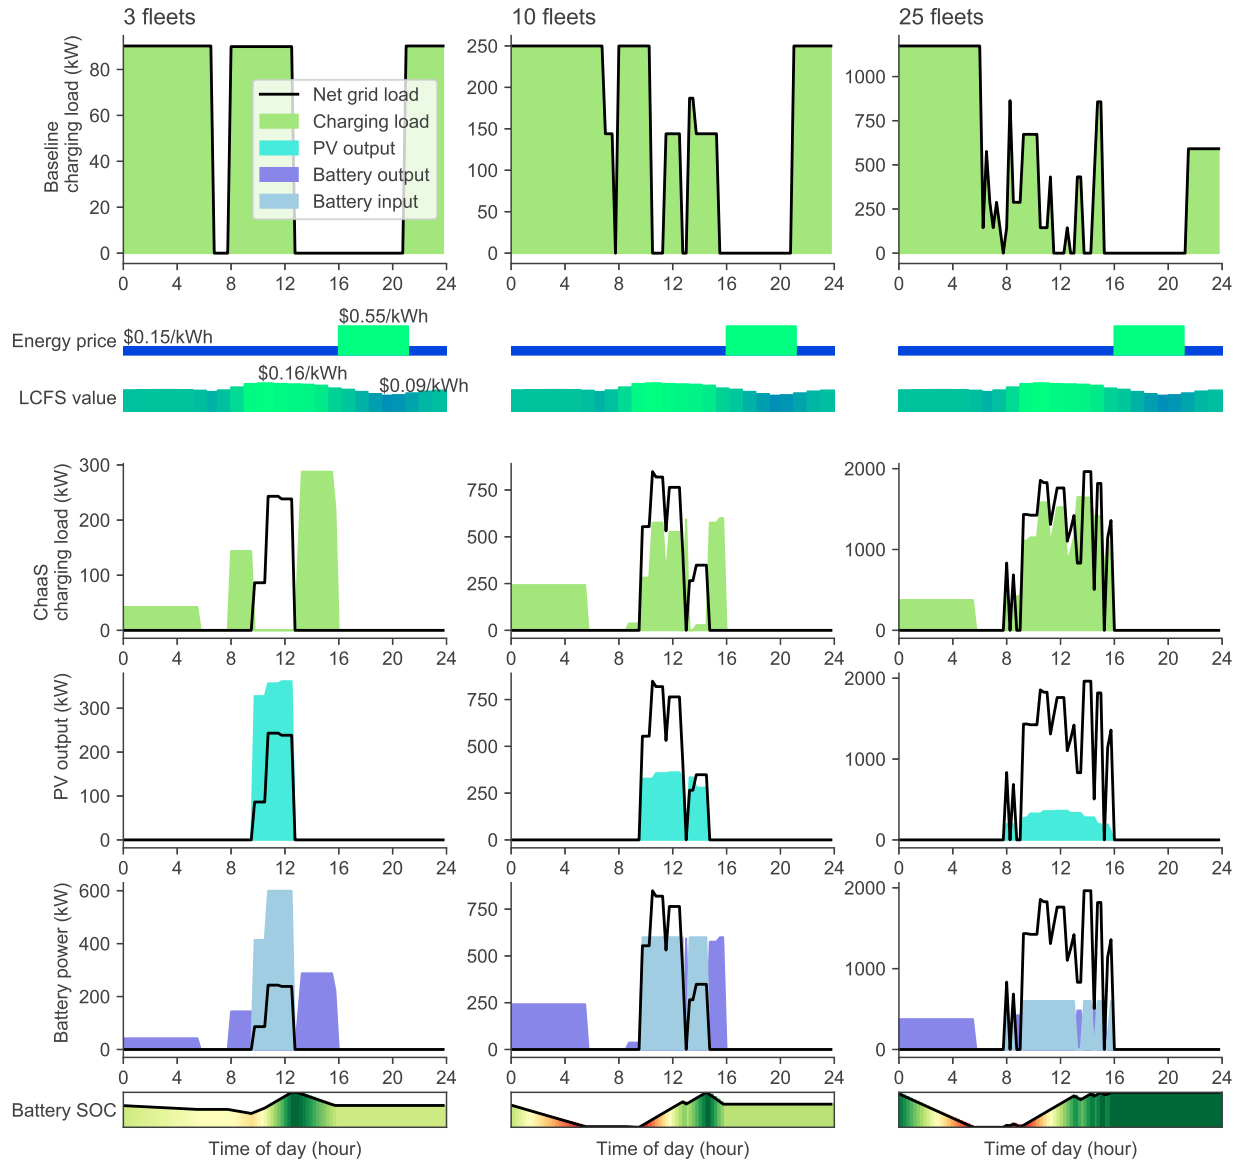


Figure 2.7: Power flows throughout a representative day for the commercial fleet business case, for populations of 3, 10, and 25 fleets. The baseline charging load is shown in the top row. The ChaaS power flows, which include the integration of a PV array and a battery storage system, and the accrual and sale of LCFS credits, are shown in the lower rows.

Figure 2.8 compares LCOC across scenarios for varying numbers of fleets. Serving more fleets, by increasing DER utilization, reduces the fraction of LCOC devoted to repaying up-front costs for the DER system. However, the DER system never realizes a net benefit compared to the non-DER alternative, suggesting that the DER system we reference is too expensive to pay back solely from avoided grid demand and the differential revenues from zero-carbon LCFS credits, even at maximum utilization. In contrast, LCFS revenues, with or without DERs and irrespective of the number of fleets served, actualize substantial LCOC reductions.

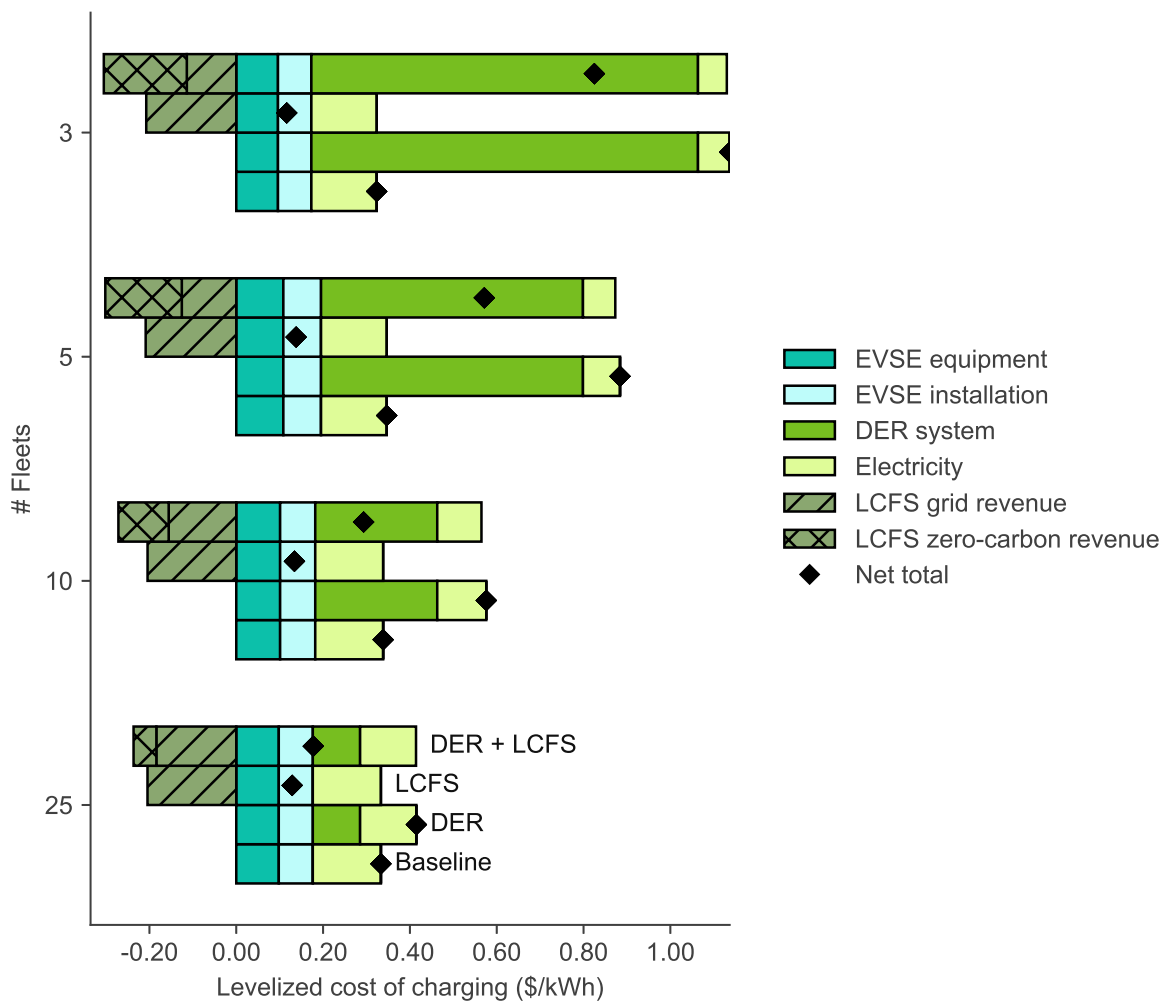


Figure 2.8: Breakdown of value streams paid and accrued by groups of fleets in the commercial fleet business case. For various populations of fleets, the LCOC components are shown for four scenarios: baseline, with no DER system or LCFS credit generation; and three ChaaS scenarios which variously include optimal DER system integration and/or the optimal generation of LCFS credits.

Summary

This example illustrates how scale and diversity improve the savings and revenues realizable via DER dispatch. While the DER system we modeled may not be economically justifiable in the scenario we present, it may become so in a place with higher grid electricity prices, lower DER system costs, or a more carbon-intensive grid. Although we find that LCFS revenues can drastically reduce LCOC at any scale, a large electric fleet – or an aggregation of small ones via ChaaS – may be necessary to justify actualizing LCFS revenues depending on participation costs (application, approval, sales), which we did not model. Both value streams may thus be infeasible for small fleets to access without ChaaS.

Other important aspects of fleet electrification, such as life cycle management, are sufficiently complex that the alternative to ChaaS for some fleets would be to employ a full-time charging system management team. Ensuring grid readiness for fleet electrification is a pressing challenge on its own, often requiring fleets to negotiate with utilities years in advance of EV deployment. This process could be eased in a ChaaS framework, for example, if a utility takes a proactive approach to offering its own ChaaS program. The major ChaaS functions in this example, both of which regulate power demand to the benefit of the grid, provide mutual benefit to fleets and utilities and may serve as means for utilities to meet their own challenges in electrifying fleets.

Here and in the previous examples, unpredictable variation in vehicle availability requirements would impose more constraints to charging management than what we model, such as a need to maintain minimum battery reserves. Such constraints could result in costly demand spikes or disruptions to vehicle availability. If fleets are reluctant to trade operational flexibility for the benefits we describe, the value available from ChaaS, and their electrification potential, would be reduced. Nonetheless, many prominent fleets have set decarbonization goals requiring rapid electrification [101]. For these to be achieved, communication and interest negotiation at the transportation-electricity system boundary, which we identify as the domain of ChaaS, are urgently necessary.

2.6 Conclusion

In this work, we present a summary and systems-level analysis of the purpose and functions of Charging-as-a-Service as it has emerged in the market. We identify functions belonging to four categories – economic, technical, managerial, and societal – across the broad span of potential use cases. We present four business cases to demonstrate some of the key economic value streams ChaaS can actualize across a range of transportation sectors.

ChaaS providers invest in transportation electrification in pursuit of profits. Although our examples consider a steady 7% annual return, the real business case likely assumes significant short- and long-term scaling commensurate with rapid growth expectations of the broader electric transportation system. Considering this, our work may present a rather conservative estimate of the economic benefit sought by ChaaS providers.

To many EV operators, noneconomic benefits of ChaaS may be the primary appeal of ChaaS. Existing ChaaS companies have strongly emphasized convenience functions in their marketing, suggesting they view convenience to be a compelling source of value for customers. In addition, ChaaS bears association with the so-called “subscription economy.” A great deal of research, well outside the scope of this paper, has sought to understand the reasons why many people are so eager to pay over time to avoid owning things [102]. For some customers, psychological factors and other unquantifiable factors likely play a role in ChaaS’s appeal.

ChaaS’s most important contributions, however, could be to the benefit of society. As one example, extending EV charging access to MUD residents, who may be renters and/or have moderate to low incomes [94], is an urgent and unsolved challenge. A failure to improve charging availability at MUDs could hinder the widespread adoption of EVs, with the result that lower-income and disadvantaged communities continue to lack access to EVs and their localized human health benefits [86]. ChaaS providers, potentially aided or spurred by government or nonprofit initiatives, can help to negotiate the mismatched incentives between property owners and potential EV adopters and assuage the capital cost burdens of MUD charging installations. If sufficient to ease the MUD charging access problem, the societal benefits could vastly exceed individual benefits.

In the broader context of this dissertation, this chapter primarily addresses Research Question #1, which asks how factors internal to and external to MHD fleets influence the feasibility of electrification. The technical and economic analysis methods introduced and demonstrated in this chapter illustrate, across a range of use cases, the influence a variety of factors exert on the costs, benefits, and overall feasibility of vehicle electrification. These methods are attentive to temporal and spatial constraints defined, collectively, by the actions of disparate agents across the electricity and transportation systems. Case #4 applies these methods to a set of MHD fleets, showcasing the technical and economic analysis methods available to fleet agents in our broader model (Chapter 4). The constraints in Case #4 include fleets' operating requirements (defined by fleets), technical equipment characteristics (defined by fleets, equipment manufacturers, and/or ChaaS providers), electricity prices (defined by utilities), the value of grid incentives (defined by utilities, governing bodies, and the market), and the availability of power (defined by utilities and the environment). Many of these constraints are novel to an electrified transportation system, adding complexity to the system while generating new opportunities to optimize the allocation of costs and benefits across stakeholders. These modeling capabilities can thus serve dual purposes: first, decision and design *support* to individual fleets; and second, decision and design *simulation* in our agent-based model.

Chapter 3

Decision and preference modeling

3.1 Abstract

The transportation sector must rapidly decarbonize to meet its emissions reduction targets. Medium- and heavy-duty decarbonization is lagging the light-duty sector due to technical and operational challenges and the choices made by medium- and heavy-duty fleet operators. Research investigating the procurement considerations of fleets has relied heavily on interviews and surveys, but many of these studies suffer low rates of participation and are difficult to generalize. To model fleet operators' decision-making priorities, we apply a robust text analysis approach based on latent Dirichlet allocation to a broad corpus of fleet adoption studies. We find operational compatibility to be the most salient factor, followed by long-term economics, technological familiarity, and perceived reliability.

3.2 Introduction

Given the aggressive greenhouse gas emission reduction goals set out on every level from international goals ("Paris Agreement," 2016) to federal [103] and state [104] legislation, the fleet transportation landscape is poised to enable the adoption of a new generation of electrified vehicles. Battery electric vehicles (BEV) are currently the most used zero emissions vehicles [105].

While the use of BEVs for personal transportation is rapidly expanding [106], and the number of medium- and heavy-duty (MDHD) BEV fleets is growing [107], the latter are experiencing a slower transition. The reasons for the delay have received little concentrated effort in academic research. So far, researchers have relied on fleet manager interviews and surveys using a variety of techniques and with varying success to address this question. For example, ref. [30] synthesizes theories of organizational adoption in the context of interviews conducted with 20 fleet managers in order to characterize the adoption process framework typically followed by MDHD fleet man-

agers (see Figure 3.1). They identify a list of factors that play an important role in the adoption decision (the middle box in Figure 3.1) for each of 20 fleet managers interviewed, including relative advantage characteristics (cost reductions, emissions benefits), operational compatibility and complexity, and sources of uncertainty such as fuel price risks.

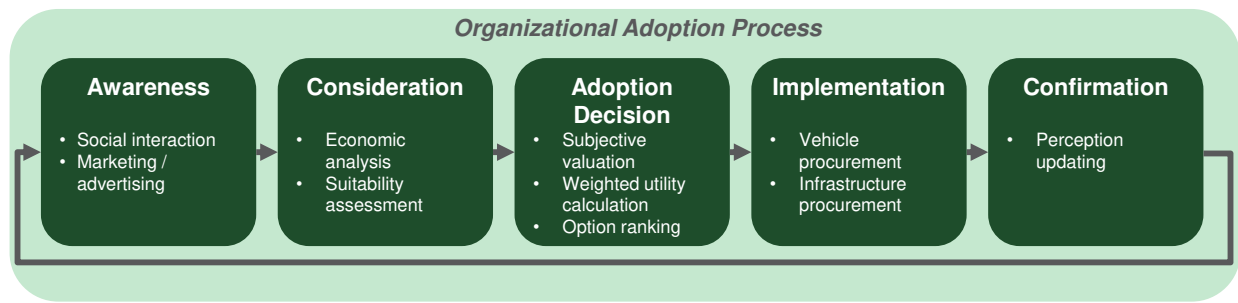


Figure 3.1: Adoption process framework adapted from Bae et al. [30].

In addition to the work of [30], other efforts [61, 64, 66, 108] attempt to construct a quantifiable model through ranking statements from interviews using area expertise as the basis for these rankings. Each of these studies is based on a small number of fleets that cannot possibly represent the wide vocational, organizational, and geographic diversity of the MDHD transportation sector. Such efforts are also vulnerable to known methodological errors including interviewer/interviewee bias, selective nonresponse, and other non-sampling errors.

3.2.1 Contributions

To address the shortcomings of interview-based research methods for modelling of fleet adoption, this study employs latent Dirichlet allocation (LDA), a commonly used unsupervised machine learning algorithm, to analyze a corpus, or collection of written material about a particular subject. In our case, we define the corpus to be a collection of existing academic papers consisting of alternative-fuel fleet vehicle adoption studies that have been indexed in Google Scholar over the past three decades. LDA is part of a family of text-mining approaches, sometimes referred to as natural language processing (NLP), that yield a topic model, or statistical characterization of the

topics, keywords, or phrases that are common to a corpus. The present research advances the understanding of MDHD EV fleet adoption by employing this novel approach to assess and rank the saliency of agent perceptions. In our case, the agent perceptions are used to model the views, opinions, assessments, insights, or beliefs of academic papers concerning the electric fleet-vehicle purchasing process.

Employing topic models to generate models of adoption considerations is an important and growing area of investigation in several disciplines. For example, Kumar and Ng [109] look at success and growth factors of renewable energy projects with LDA and non-negative matrix factorization (NMF). Debnath et al. [110] employ social media data to build a corpus for the political, economic, social, technological, legal, and environmental aspects of consumer EV adoption. Additionally, Ruan & Lv [111] analyze Reddit data with LDA sentiment analysis to model the public perception of EVs.

LDA is widely used in fields other than transportation and energy, but in other contexts it typically relies on coherency scores to justify topics (see Kapadia [112] for a broader discussion of topics, topic models, coherency, and topic coherence). For the technical reasons described below, we rely on saliency rather than coherence as our metric of model quality and performance, and we contend that saliency scores provide a more direct and consistent way to rank the various attributes that can be used to profile the appetite of fleet managers for EV adoption.

The remainder of this paper is organized as follows. Section 3.3 briefly discusses the concept of text mining. Section 3.4 covers the methods used to train, fine-tune, and validate our LDA model. Section 3.6 presents our results and Section 3.7 discusses their impacts. Section 3.8 outlines planned future extensions of our work and Section 3.9 draws relevant conclusions.

3.3 Brief overview of text mining

Both text mining and natural language processing (NLP) share their theoretical origins in the 1950s. These ideas remained in the realm of theory due to the lack of available computing power until the 1990s. The onset of the Big Data era in combination with rapid development of computing

technology allowed text mining practices to enter the mainstream [96]. It is now commonplace in several industrial sectors such as business analytics, engineering, public health, and the physical and social sciences [113]. Text mining is the process of wrangling and analyzing unstructured text data. Tasks in this field include content extraction, sentiment analysis, pattern recognition, topic modeling, dimensionality reduction and many more. The connecting theme between these tasks is using term frequencies and semantic relationships in tandem with clustering and classification methods [114]. The model employed in this work heavily relies on the progression of the text mining field over the past 70 years.

Latent Dirichlet allocation [115] is a product of some of the most cited authorities in data science, including David Blei, Andrew Ng and Michael Jordan, who collaborated to eliminate issues in overfitting experienced when using its predecessor, probabilistic latent semantic analysis (pLSA) [116]. pLSA is one of many alternatives to LDA. Other methods include latent semantic analysis (LSA) [117] and non-negative matrix factorization (NMF) [109]. LDA holds the capability, unique among these options, of assigning multiple topics to a single document. LDA also fulfills our applications' requirements of scalability, accuracy, and explainability. This allows LDA to return a more accurate representation of important terms in the corpus. NMF does show promising topic modeling results in [109] and it reduces the amount of topic guessing that occurs in the training process, but LDA was selected over NMF because NMF is usually reserved for shorter length texts such as social media posts and web chat captions since an NMF model may generate topics that are less interpretable [118].

LDA models are applied in everything from social media fake news detection [119] to movie reviews [120], to medical journal database organization [121]. Our usage of LDA slightly differs from this common implementation in that we are evaluating the global importance of terms instead of limiting our view to their contribution to generated topics. Although it is well suited for this purpose, we have not found similar examples in the literature.

LDA relies on the major assumption that every document is a random combination of words where word order does not matter. The set of words (w) are then placed into a set of N topics (z)

by defining the likelihoods of how many topics a given document will contain (α) and how many words are associated with each topic (β). These hyperparameters, in tandem with the marginal topic distribution (θ) are used to build the resulting joint distribution of a topic mixture, as specified in Equation 3.1:

$$p(\theta, z, w \mid \alpha, \beta) = p(\theta \mid \alpha) \prod_{n=1}^N p(z_n \mid \theta) p(w_n \mid z_n, \beta) \quad (3.1)$$

Integrating over this distribution and summing over the topics, we arrive at the marginal distribution of topics in a document, as specified in Equation 3.2. Such distributions provide the information on topic probability needed to calculate saliency, the final scoring metric (defined below):

$$p(w \mid \alpha, \beta) = \int p(\theta \mid \alpha) \left(\prod_{n=1}^N \sum_{z_n} p(z_n \mid \theta) p(w_n \mid z_n, \beta) \right) d\theta \quad (3.2)$$

With the framework of distributions created, they need to be optimized. LDA completes this task by using Gibbs sampling. Depending heavily on Monte Carlo Markov Chains and Bayes Theorem, this algorithm pulls the word under consideration out of its initial randomly assigned topic and infers the best fitting (e.g., most closely aligned) topic based on the number of times that the document containing the word uses a topic and the number of times a topic uses the word. One LDA training iteration consists of this process for each word in the corpus. Tomar [122] provides a deeper explanation of the execution of Gibbs sampling.

3.3.1 Performance metrics

The key output of a typical LDA model is a “coherency score”. However, the methods proposed here measure “saliency” as the basis for an attribute’s ranking in the fleet decision process. In layperson’s terms, saliency is simply the extent to which an item is distinct or pronounced.

Saliency is a metric originally developed for TERMite, a visualization package produced by Chuang et al. [123] that shows the overall importance of a word in a corpus. Using the distributions

derived in Equation 3.1 and 3.2, saliency ranks each word based on its frequency scaled by its distinctiveness as defined in 3.3:

$$\text{saliency} = P(w) \times \text{distinctiveness}(w) \quad (3.3)$$

where w is a word in the corpus and distinctiveness is defined in Equation 3.4:

$$\text{distinctiveness}(w) = \sum_T P(T | w) \log \frac{P(T | w)}{P(T)} \quad (3.4)$$

where T is topic. Given a specific word, distinctiveness is the sum of each topic's conditional probability multiplied by the Kullback-Leibler divergence [124] of the topic and word under consideration.

This step is implemented to help eliminate irrelevant terms by looking for term importance within a document, as well as ensuring it is not a common term across the corpus. To ensure important words were placed into relevant categories, several preliminary models were trained to find the terms with the highest saliency. These terms were then placed in the categories described in Section 3.6.

3.4 Methods

The present research includes the development and implementation of LDA with a novel corpus and process. To present this process, we first assemble and pre-process a corpus to make a digestible input for the LDA model. Next, we validate both corpus and model consistency. We then aggregate results to produce a single saliency score per category. Figure 3.2 summarizes the sources of papers in the corpus; Figure 3.3 depicts the overall process.

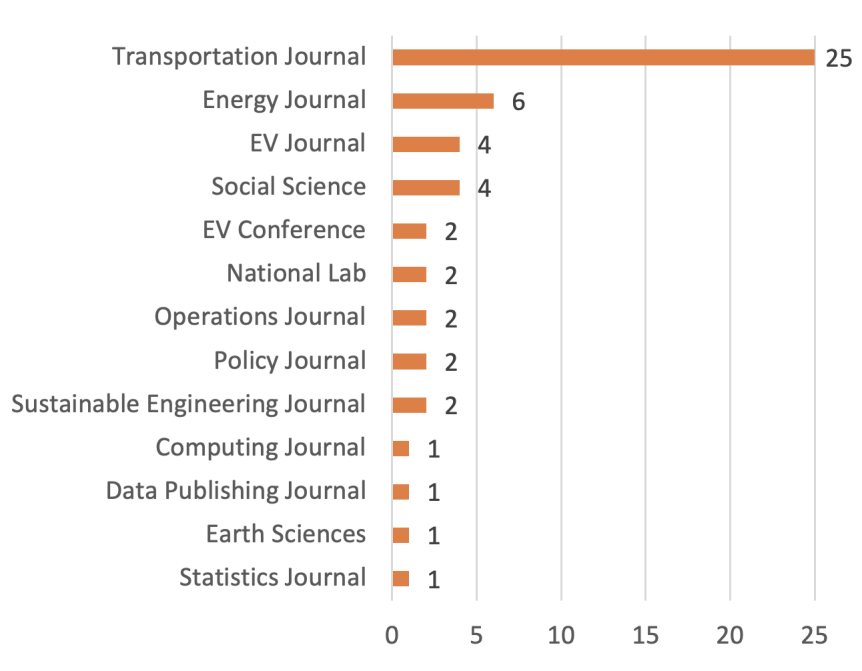


Figure 3.2: Breakdown of sources for papers in the analyzed corpus.

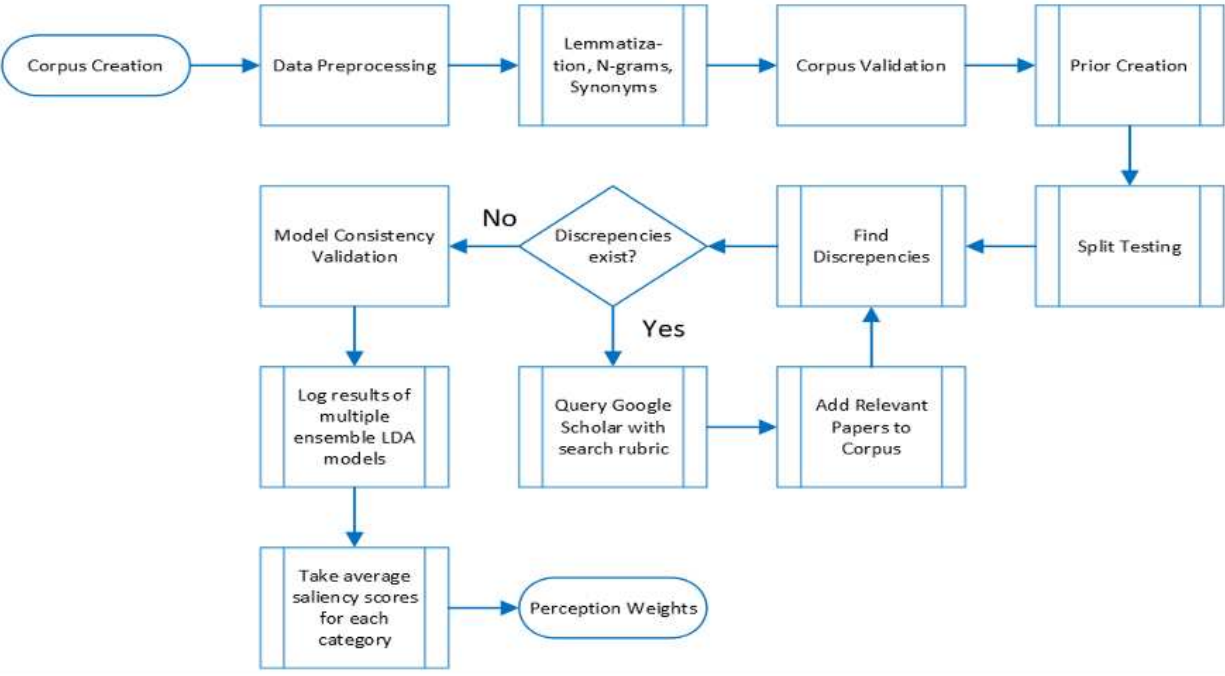


Figure 3.3: Illustration of perception weight development process.

3.4.1 Corpus description

We compiled our original corpus consisting of 39 papers from a series of literature searches conducted prior to conceptualizing this work, all of which concern the alternative-fuel (AFV) and electric vehicle adoption behaviors of commercial fleets. It was necessary to consider both AFVs and BEVs since there are not enough papers on strictly BEV fleet adoption to effectively train an LDA model. During the process of running split-corpus validation, our corpus grew to 53 papers. Augmentation of the corpus is explained in Section 3.8.3. Figure 3.2 shows the makeup of paper sources, with Transportation Research proving to be the prevailing source for relevant papers on MDHD EV adoption. The corpus paper selection process is a potential source of bias, since many of these papers focus on fleet manager and fleet user perceptions derived from interview processes, while others report survey results from fleet adoption attitude research projects.

3.4.2 Data wrangling/preprocessing

To be suitable for input to a topic modeling engine, a corpus must undergo several pre-processing steps. These include elimination of stop words, creating n-grams, lemmatization, and synonym grouping. We briefly describe each of these steps below.

“Stop words” are common words such as “a” and “the” that do not contribute to coherent topics. They may appear frequently enough to alter topic distribution results and must be filtered out of the corpus. 173 application-specific words (e.g. “afv,” ”bev,”and ”literature”) were added to the default stop words list in the Natural Language Processing Tool Kit (NLTK) [125] for removal.

N-grams are combinations of two or more neighboring words, such as “battery electric vehicle” (a 3-gram, or trigram). Gensim, the package used to implement LDA [126], supports bigrams, two-word phrases, and trigrams, three-word phrases. Each combination of neighboring words is tested in the corpus and terms with a score higher than the pre-established threshold of 100 are added to the corpus [127], as defined in Equation 3.5:

$$\text{score} = \frac{(\text{n-gram count} - \text{min count}) \times \text{length}}{\text{word}_a \text{ count} \times \text{word}_b \text{ count}} \quad (3.5)$$

where

- n-gram count = number of times the n-gram is present in corpus
- min count = minimum collocation count threshold
- length = size of vocabulary
- word_a count = number of times the first word occurs, and,
- word_b count = number of times the second word occurs

Lemmatization allows for all words with a similar root to be represented by a single word. For instance, “truck”, “trucks”, and “trucking” would all be represented by the word “truck”. We chose to use lemmatization rather than similar operations such as “stemming” because (1) the root words are guaranteed to be English words, therefore increasing the interpretability, and (2) processing time was not a constraint [128].

A synonym group is a group of highly salient words that all refer to the same idea. Most adoption and perception categories can be referenced using any of several synonyms. For example, “environment” and “green” can be categorized under “sustainability.” While manually choosing which words represent certain topics does introduce a certain amount of bias, this process aggregates the saliency of individual terms into a more intuitive single concept. The perception categories that emerge from our model are brand image, familiarity/knowledge, long-term cost, norms/attitudes, policy/incentives, reliability/uncertainty, capital cost, suitability/compatibility, and sustainability. The synonyms that fall in each of these groups are shown in Appendix B.

3.4.3 LDA implementation

Referencing Equation 3.1 and 3.2, the values of α , β , and k have a significant impact on the results of the model. Initially, we implemented a three-dimensional grid search to optimize parameter selection. After noticing inconsistent results between model runs, we searched for a solution using the ensemble LDA approach explained in Section 3.5.2.

3.5 Verification

A key step in the development of this tool was the validation process. Two separate components required validation. First, the corpus itself was balanced with split-corpus validation testing (Section 3.5.1) to show that the resulting saliency scores were not a factor of a small portion of the papers. Next, the consistency of model results on the corrected corpus were verified.

3.5.1 Split-corpus testing

To check if there was similar influence from all members of the corpus, it was split into two different validation corpora. These were used to build separate LDA models. If there was a discrepancy in saliency scores for a term between the two models, this term was incorporated into a Google Scholar query with a uniform format. From here, the earliest results that are relevant to electric vehicle fleet adoption were added in equal parts to the two validation corpora and the process was repeated until the average difference in saliency scores for the targeted terms reached the pre-determined threshold. This is sometimes referred to as corpus augmentation.

3.5.2 Model repeatability

Typical LDA implementations sample from several random variables during Gibbs sampling (Section 3.3), and therefore suffer from a lack of repeatability. To address this issue, we used the ensemble LDA (eLDA) [129] routine, as implemented in Gensim. eLDA uses density-based spatial clustering of applications with noise [130] to categorize the number of topics and beta Dirichlet distributions that create stable results. The process returns a single LDA model with the α , β , and k values shown to produce convergent results. To ensure consistency, 100 ensembles of 8 LDA models each were trained. Section 3.7.1 shows the results from the validation efforts.

3.6 Results

3.6.1 Validation results

The split-corpus validation resulted in two rounds of corpus augmentation before saliency scores were within 1.2% of each other on average. Figure 3.4 shows the average sum squared difference (SSD) between the validation sets for each category. Corpus augmentation 2 shows an average SSD difference of just under 0.004 which translates into 1.2%.

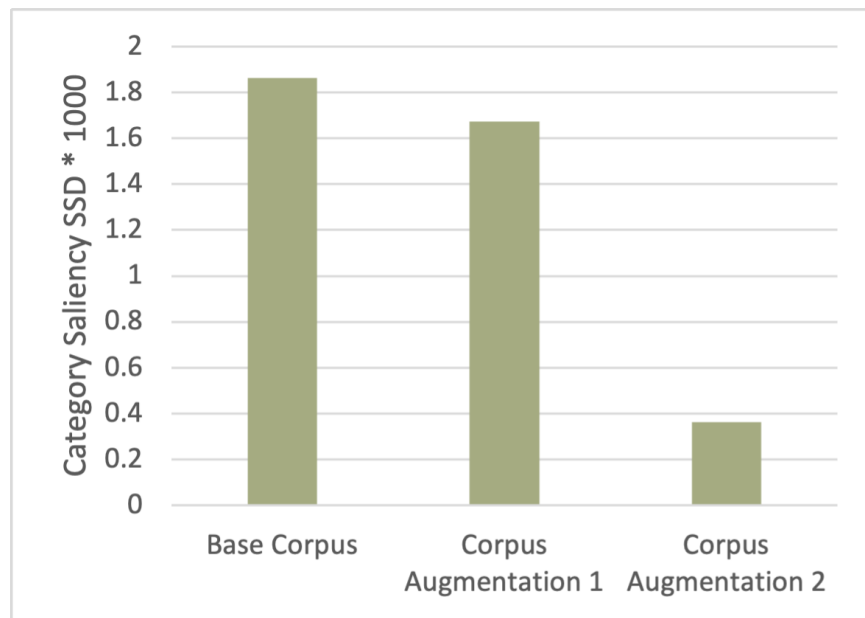


Figure 3.4: Average category SSD for the split corpora and the full corpus, $n = 800$.

Figure 3.5 shows the difference in term saliency as a percent between ensembles. The average difference in saliency scores between ensembles was 1.49% over 100 independent models, suggesting that our implementation of LDA is highly repeatable. The norms/attitudes category produced the outlier at just above 4.5%.




Figure 3.5: Average variance of normalized saliency between ensembles, $n = 800$.

3.6.2 Final saliency scores

Figure 3.6 presents the final normalized saliency scores from the validation process. The scores were normalized by dividing by the maximum value for visualization purposes. These scores are the average term saliency from the 100 generated ensembles.

Suitability/compatibility shows significant dominance over the other categories. Long term cost, familiarity/knowledge, and reliability/uncertainty are considered equally salient since the level of uncertainty shown in Section 3.6.1 results in the possibility of these categories swapping rankings. This situation extends to the relationship between sustainability and brand image which is covered in Section 3.7.3 5.3.

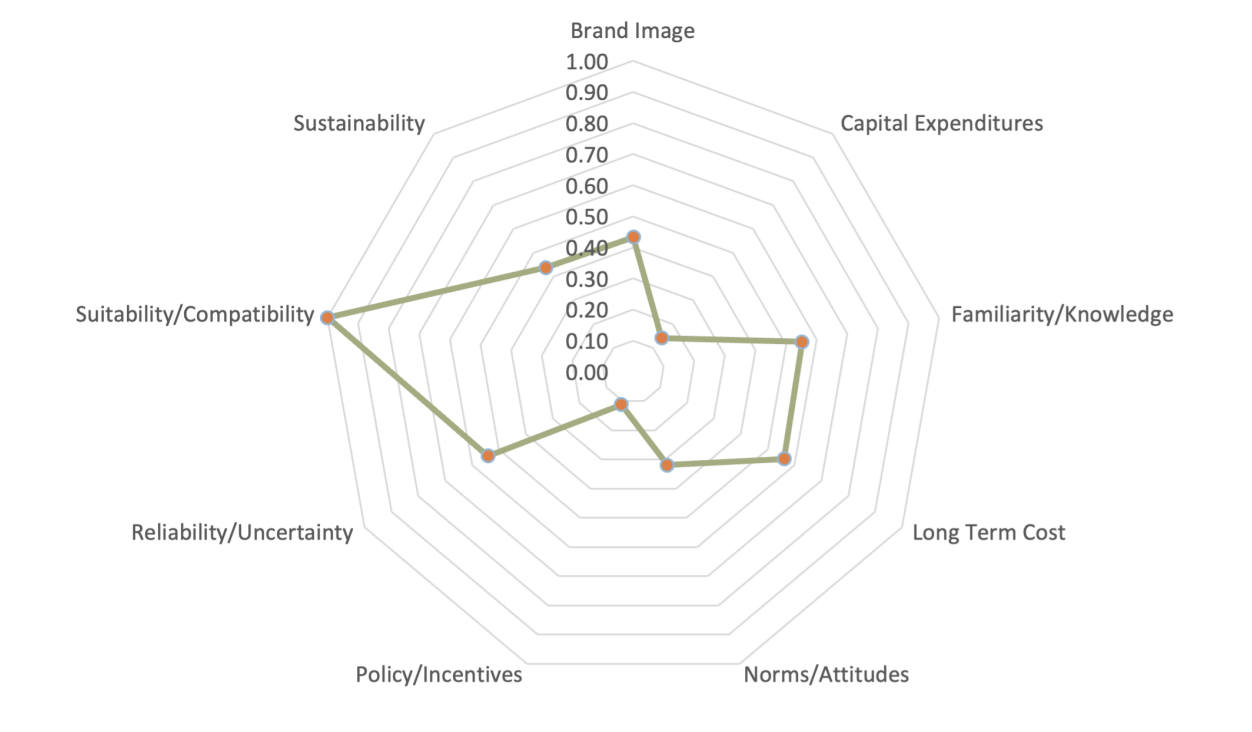


Figure 3.6: Radar plot of final saliency scores.

3.7 Discussion

Results show conclusions that agree with past work in some areas while contributing a different perspective in other areas. Suitability/compatibility, sustainability, and long-term cost saliency scores reflect trends brought to light in more subjective approaches. Capital cost stands out with a low saliency score contradicting intuition.

3.7.1 Suitability/compatibility

We consider a vehicle’s suitability/compatibility to be a fleet operator’s perception of the ability to use it to fulfill business requirements. This definition, which we used to guide the selection of synonyms for the “suitability/compatibility” group, is adapted from the definition of compatibility offered in Bae et al. [30]. For many MDHD vocations, such as long-haul freight, the future suitability/compatibility of BEVs remains disputed [8, 11]. If a vehicle indeed cannot satisfy the requirements for which it was purchased, any other potential benefits to the fleet, including

cost or emissions reductions, are rendered irrelevant. With this in mind, the dominant salience of suitability/compatibility makes sense.

Because BEV options have not yet been demonstrated in many MDHD vocations, fleet operators' ability to assess BEV suitability is often limited, which leads to a perception of incompatibility. This is somewhat at odds with decarbonization goals, some of which exhibit an assumption of universal BEV suitability. For instance, Colorado plans for a 100% zero-emission transit fleet by 2050 [104]. While important for meeting emissions goals, the actual usability of a 100% electric fleet must first be verified. As a counter example, Li [131] mentions the use of the word "feasible" in the California 2045 mandate where it states that all MDHD vehicles must be electric by 2045, if feasible. With some municipalities considering options that may not be entirely possible, and others using ideas of suitability in the wording of the policy itself, it is logical that suitability is an important topic when addressing fleet adoption.

3.7.2 Capital cost

While it seems like capital cost should be an important factor when considering a large purchase [10], several other works agree with our results that it is not a high-ranking priority. Anosike et al. [132] dive deeper into explaining the relationship between capital costs and fleet adoption while Knupfer et al. [133] suggests that cost barriers are high due to small and niche battery suppliers (which is known to be changing). Lithium prices forecasted by Sun et al. [134] suggest that the industry is experiencing a temporary price spike and that battery materials will return to a lower stable price. From a fleet-specific perspective, the idea of capital cost may not be as important to managers, in general, since fleet vehicles tend to have higher annual mileage so that the initial cost is spread out over more miles when accounting for the purchase. Lower maintenance and fuel costs also help balance the high up-front cost [135].

3.7.3 Sustainability and long-term cost

From the economic, environmental, and social perspectives [24], comprehensive sustainability can be represented, in part, by long-term cost, including corpus terms such as "economy" and

“investment.” The environmental component is captured in the sustainability category with words such as “green” and “environment.” Concepts of social sustainability fall under the norms-attitudes category with words such as “social” and “behavior.”

While the saliency of social sustainability (norms/attitudes) is not high in our results, long-term cost and environmental sustainability have above average scores. This is in line with findings reported elsewhere. Kumar & Alok [24] indicate that \$38 million can be saved by setting up electric school buses in a vehicle-to-grid system, suggesting that economic sustainability does impact adoption decisions. Based on the results of a small study involving 11 interviews, Sierzchula [136] concludes that environmental impact is a less important factor to fleet managers. On the other hand, Globisch et al. [137] claim to have surveyed a more relevant pool than Sierzchula [136] and conclude that climate protection is the second most important factor in EV fleet adoption. With the wide differences in results presented previously, we posit that environmental sustainability is of above average importance but trails behind ideas such as suitability, reliability, and familiarity. One point of agreement from other research is the relationship between brand image and sustainability. Sierzchula [136] and Globisch et al. [137] agree that environmental sustainability is motivated by brand image. In our results there is a one percentage point difference in the sustainability and brand image average saliency scores.

3.8 Future work

The immediate results of this work are used to populate a model of the fleet procurement decision, but there remains room to continue our topic modeling efforts. Opportunities include the regularization of the synonym process, comparison to other topic modeling algorithms, further academic corpus augmentation, and comparison to a collection of industry reports on the same topic.

3.8.1 Synonym normalization

While the synonym process is important for returning interpretable results, we acknowledge that our process for creating the perception categories introduces bias into the overall system. One way to mitigate this would be to develop a method for changing the weight of a category based on the number of terms linked to that category. Eliminating this factor does have disadvantages as well. The fact that some categories have more terms relating to them can also be taken as a sign of importance of that category.

3.8.2 Model comparison

The focus on LDA resulted from it fulfilling all our criteria for the research. To get a better idea of performance with less emphasis on explainability, other algorithms should be investigated. For example, Gui et al. [138] implement reinforcement learning to optimize an autoencoder for topic modeling purposes, and, in a more parallel study, Kumar and Ng [109] compare the results of LDA and non-negative matrix factorization (NNMF) for trends in renewable energy adoption. Comparing LDA, NNMF and reinforcement learning optimized autoencoders on the same corpus would give a more global perspective of the importance of terms in a corpus.

3.8.3 Corpus amendments

During the corpus augmentation process, we adhered to a specific search criterion to reduce the possibility of selection bias. This included the use of the “allintitle” command on Google Scholar. This command reduces search results to works that contain all search terms within the title. If the “allintitle” requirement was to be lifted, several more papers related to fleet electrification could be added to the corpus. This path does involve zooming out from MDHD fleets, in general, to focus on a broader collection of EV fleets (i.e., light-duty vehicles), since most papers on MDHD EV adoption are already included in our original corpus. With this said, including any academic papers on this topic of which we are not aware would obviously be a useful contribution.

3.8.4 Industry corpus

The lack of discussion in the corpus on terms such as capital costs raises the question of the relevance of academic papers to the actions of industry professionals operating under different influences. To gain a different perspective on MDHD EV adoption, a corpus of solely industry-sourced articles should be compiled. The resulting saliency scores for each category should be compared to those of the academic corpus to find discrepancies in perception between industry and academia. Social media represents another potential source of industry perceptions, similar to its use in Debnath et al. [110] relative to light-duty EVs.

3.9 Conclusion

We have presented a set of methods to characterize and score the factors considered during MDHD EV adoption decisions. By integrating the results of studies spanning a range of vocations and research methods, our approach mitigates sources of error inherent to surveys and interviews. Some outcomes match intuitive expectations, including that operational compatibility is a key factor in adoption decisions. In contrast, unexpected results, such as that purchase price is a low-importance factor, highlight gaps in the corpus or reveal errors in our intuitive preconceptions. Our method can be repeated using augmented and alternative corpora, to characterize changes to the academic understanding of adoption decisions over time, and to reveal insights that may be missing from the current literature. Finally, these methods are not limited to our topic or discipline. We hope researchers in other disciplines, particularly those similarly challenged by subjectivity error, find use for them.

The outcomes of this chapter pertain to Research Question #2, which asks how economic, behavioral, and regulatory factors influence the decisions made by key entities in the MHD transportation system. Many researchers have offered answers to restricted versions of this general question, for example, via surveys applied to samples of particular types of MHD fleets within restricted geographic scopes. In this chapter, we synthesized the text of many such studies in pursuit of a more comprehensive answer to RQ#2. The outcome, a quantification and ranking of the

adoption decision factors discussed throughout our assembled corpus, offers a new understanding of the existing literature, as well as an approximate model that we can instrumentalize to simulate the adoption decisions of fleets in our agent-based model (Chapter 4).

Chapter 4

Agent-based model development and application

4.1 Abstract

Reaching transportation decarbonization targets will require rapid uptake of zero-emission vehicles, including electric vehicles. This chapter presents a bottom-up, agent-based model of electric vehicle uptake at a fleet-level resolution based on the latest theories of new technology adoption by commercial fleets. The high granularity of the model permits detailed analysis of the dynamics of adoption among fleets, including social, economic, environmental, and technical factors. We present the design, verification, and validation of the model to demonstrate its suitability for application and further development. Preliminary simulation results demonstrate the model's capabilities relative to other adoption modeling approaches. Benefits include results at a high spatiotemporal resolution and the representation of nonlinear dynamics of choice. Challenges include high computational cost and lack of data availability.

4.2 Introduction

Technology adoption models are often developed and applied to make market predictions, even though human decision-making is often irrational [139] and even highly respected adoption models are without predictive power, particularly when making predictions early in a technology's adoption life cycle [13, 140]. As an alternative, adoption models can be designed as experimental tools. The effects of costly interventions aiming to influence technology adoption, such as investments in policy, research, and marketing, usually cannot be evaluated until after significant costs have been committed. If adoption models can serve as testbeds for these kinds of interventions, they can provide valuable decision support, even if predicted outcomes are only accurate relative to each other. When used primarily to experiment with cause-effect dynamics, qualitative fidelity becomes more

important than quantitative predictive accuracy, shifting the requirements the model must satisfy and how it should be conceptualized.

Adoption models commonly conceptualize potential adopters in aggregate, which can limit their suitability as experimental tools. Two examples of aggregated adoption modeling approaches are consumer choice models and diffusion rate models. Consumer choice models group consumers into a small number of discrete segments, each exhibiting uniform sets of preferences and behaviors. Diffusion rate models apply regression-like methods (e.g. Bass models, Gompertz or others), sometimes to an entire consumer population at once [13]. If we categorize these aggregative models, which begin from a high-level, low-detail concept of the consumer population, as “top-down” adoption models, we can contrast them with “bottom-up” models that begin from disaggregated sets of decision makers, individuating their needs and preferences. Causality in adoption decisions can thus be represented more directly than in top-down models, at the cost of higher complexity. Bottom-up models might be especially well applied to understand the dynamics of eMHD fleet adoption, because each single decision-making entity, i.e., each fleet vehicle procurement manager, can be responsible for deciding and actuating large vehicle procurements. eMHD adoption studies published to date have exclusively applied models with a top-down architecture. Examples from the eLD adoption modeling literature illustrate how bottom-up approaches can directly model dynamics that top-down approaches cannot, including social interactions and the evolution of perceived behavioral norms [14, 15].

In agent-based modeling (ABM), the bottom-up approach we employ in this work, discrete agents make decisions and perform actions individually. Classic examples of ABM demonstrate how the collective behavior of individuals can drive non-intuitive outcomes that top-down approaches cannot capture. The Schelling segregation model, for example, demonstrates how a small amount of in-group preference (every agent preferring at least 1/3 of neighbors to be from the same group as them) can result in disproportionately high segregation between groups [141]. In the adoption of new technologies, decision makers are also likely to respond to the choices of their neighbors and competitors as well as to their own unique requirements. If these and other non-

aggregable factors affect adoption decisions, top-down models will fail to represent their effects and, further, the effects of interventions that target them.

In this work, we design and apply an ABM comprising fleets, electric utilities, and automakers (or alternatively, original equipment manufacturers, OEMs). We describe the model’s structure and the adoption process model in Section 4.3, and detail how modeled agents learn and interact in Section 4.4. We present the results of verification and validation procedures in Section 4.5, before presenting and discussing simulation results and sensitivities in Sections 4.6 and 4.7.

4.3 Adoption process model

The model’s structure is illustrated in Figure 4.1. Agents are organized in networks, one each for fleets, utilities, and OEMs. Fleet agents own and control vehicles and charging infrastructure, about which they form and share opinions. Utility agents control the costs of installing and operating electric infrastructure. OEM agents control the attributes and availability of vehicles, including costs.

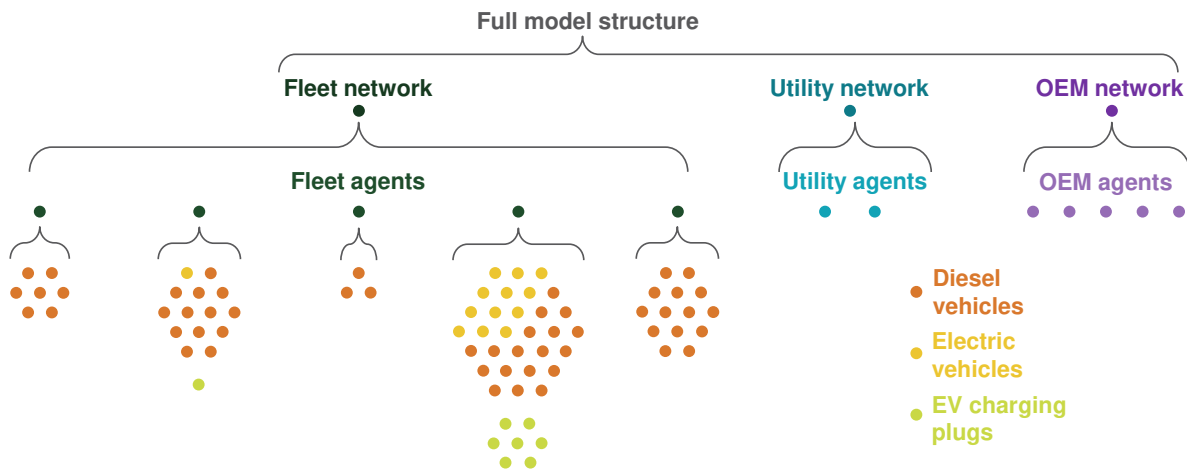


Figure 4.1: Illustration of a populated ABM consisting of five fleets, two utilities, and five OEMs. Fleet agents own assets including vehicles and EV charging plugs (EVSE).

To this ABM population, we apply a conceptual framework for alternative-fuel vehicle adoption by commercial fleets [30]. Bae et al. developed this framework by integrating three reference

theories and reinforced it based on interviews with HD fleet managers. The theories underpinning their framework account for (1) how innovations spread across social networks (Diffusion of Innovations theory [142]), (2) factors particular to organizational adoption (Technology-Organization-Environment framework [143]), and (3) the roles played by individuals within the organization, including decision makers and those directly operating the new technologies [144]. The framework they present consists of an adoption process and, within that process, a characterization of the factors influencing the decision to adopt. The framework is adapted and presented in Figure 4.2. The following sections describe how each step in the process is modeled.

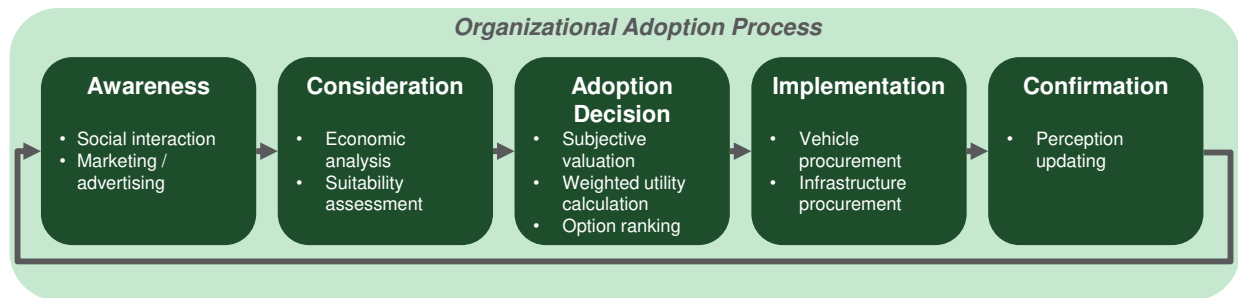


Figure 4.2: Adoption process followed by fleet agents, adapted from Bae et al. [30].

Awareness

Agents' state of awareness is binary, that is, they are either unaware or aware of each vehicle available to them. They can gain awareness either through the advertising and outreach efforts of OEMs and others, or through interactions with other fleets. Advertising is modeled as a stochastic process, where all agents that are unaware of a particular OEM have a probability of discovering it each year. Interactions between agents are modeled in detail (see Section 4.4).

Consideration

Fleets will consider adopting a vehicle if they perceive it to be capable of satisfying their operational requirements. These requirements are represented in terms of range and cargo/passenger capacity. A vehicle will not be considered for adoption if it has less range than the longest distance

traveled in a day by existing vehicles, or if it has less cargo/passenger capacity than the greatest capacity used by existing vehicles. Vehicles perceived to be operationally suitable are also subjected to economic analysis, details of which can be found in Chapter 2 and Appendix A. Finally, agents assess the relative utility of all the vehicles they are aware of, using the process described in 4.4.

Adoption decision

Having evaluated vehicles for suitability and subjectively assessed them, agents decide whether to replace each existing fleet vehicle, and what to replace it with. (Vehicles that were purchased within the previous 8 years are never replaced.) Every existing vehicle is paired with the highest-utility option suitable for replacing it and entered into Equation 4.1 to find the replacement probability p_{replace} :

$$p_{\text{replace}} = \begin{cases} 1 & u_1 \geq u_0 \\ 10^{k(u_1 - u_0)} & u_1 < u_0 \end{cases} \quad (4.1)$$

where u_0 is the utility of the existing option, u_1 is the utility of the potential replacement, and k is a constant. Thus, if the potential replacement is perceived to be equal or superior to the existing vehicle, it is always adopted. Otherwise, it may still be adopted, but with a probability that decreases sharply with its perceived inferiority. This enables fleets to adopt vehicles as pilot projects, to gain information and evaluate feasibility. (Each vehicle option can only be “piloted” once.)

Implementation and confirmation

Finally, fleets procure and operate new vehicles and charging infrastructure. In the implementation step, fleets sell vehicles that are being replaced, procure new vehicles, and design a minimum-cost EV charging system, accounting for existing charging assets, to enable any and all BEVs to complete their drive cycles. (All BEV charging is assumed to be performed at the fleet’s home base.) For the confirmation step, fleet agents record observations of all attributes of all vehicles they own, old and new (see “direct experience” in Section 4.4).

4.4 Learning and opinion formation

Each fleet agent's state of awareness of the vehicle options available on the market is measured and tracked. Agents form perceptions over time about the options of which they are aware, based on direct experience, economic analysis, social knowledge sharing, and observational inference. These *perceptions* arise from *observations*, which are modeled as data points corresponding to each attribute (Table 4.1) of each vehicle option.⁹ Agents continually make observations about all the vehicles they are aware of throughout the simulation. Observations are assigned weights depending on the means of observing them, with direct experience most heavily weighted.

Fleet agents are currently the only type of agent subject to this learning model. In future work, we will evaluate this model's suitability for utility and OEM agents. In several senses, however, decisions made by utilities (tariff and incentive design, resource allocation) and OEMs (vehicle design, price setting, sales) are distinct from fleet purchase decisions. We have conceptualized the vehicle purchase decision – choosing one option out of a limited number of options – as a discrete optimization problem, whereas many utility and OEM decisions are best conceptualized as continuous optimization. Adapting the model presented here into a continuous learning and decision-making model for utility and OEM agents is a next step.¹⁰

⁹We assume fleet agents are able to directly observe all of the attributes in Table 4.1 for vehicles they own (i.e., via “direct experience”). During economic analysis, agents can only estimate attributes in the Economics category. Observations of all attributes can be communicated via social inference and social interaction.

¹⁰Fleets also make non-discrete decisions, such as trip and charge scheduling. In the example of charge scheduling, we assume fleets optimize for cost, constrained by static trip schedules, enabling us to collapse those decisions into the discrete vehicle procurement decision. This simplification comes at the cost of flexibility in trip and charge scheduling decisions. A continuous decision model for utility and OEM agents could be adapted to expand the decision-making power, and flexibility, of fleet agents.

Table 4.1: List of vehicle and technology attributes that fleet agents incorporate in their adoption decisions.

Category	Attribute
Economics	Total cost of ownership
	Purchase price
	Infrastructure cost
	Fueling cost
	Maintenance cost
	Resale value
Performance	Fuel economy
	Range suitability
	Reliability
	Refueling convenience
Risk / uncertainty	Price stability
	Fuel security
	Familiarity
Strategy / other	Driver acceptance
	Brand image
	Environmental sustainability

Agents abstract the observations they make to the categories to which vehicles belong. This enables agents to assess new options on the basis of their similarity to existing options. Weighting observations enables observations of more similar options to exert more influence. For example, an agent’s initial assessment of a new Class 3 electric van draws from observations the agent has made about other electric vehicles, with the greatest weight for vehicles that are Class 3, are vans, or both. This observation abstraction process is further detailed in Appendix C. The different methods of observation available to agents are described in the sections below.

Direct experience

Observations made by experience, i.e., through owning a vehicle, are weighted more heavily than other observations because they are the most direct means of observation. At each simulation step, fleet agents measure the attributes of all the vehicles they own and record the measurements as observations. The models for the various attributes are described in Appendix C.

Economic analysis

Agents also assess vehicle options by conducting economic analyses of potential reconfigurations of their fleets. This method only pertains to economic attributes, which include TCO, purchase price, maintenance cost, charging infrastructure cost, and fueling cost. These observations, particularly of TCO and fueling costs, depend on fleet configuration details and are calculated with the assumption of optimal fleet and charging system design. An example of this optimization and accompanying analysis can be found in Section 2.5.4. We provide more detail about the economic analysis procedure in Appendix A.

Social interaction

Each pair of fleets is either socially connected or socially unconnected. If two fleets are connected, they communicate with each other and can mutually share their observations. Our model for this sharing of information is inspired by deGroot learning [145, 146]. In DeGroot learning, nodes in a directed social network share “opinions” by a simple mechanism. At each of a sequence of time steps, each node’s opinion is updated as the average of its own opinion and the opinions of its neighbors. This average can be weighted based on the relative trust a node assigns to each of its neighbors (and to itself).

In our model, rather than recomputing its opinion as a weighted average of a set of opinions, an agent learns from its neighbors by sampling from their observations and adding those samples to its own set of observations. For simplicity, we assume that all of an agent’s social connections are equally weighted.

Social inference

Whether socially connected or not, it is possible for one fleet to learn about another fleet's procurement decision, particularly if it is a large or high-visibility procurement decision. When observing other fleets' purchases, fleet agents can use inference to make observations about vehicle attributes. Intuitively, large and/or highly visible purchase orders have positive implications for characteristics of the vehicles being purchased.

We conceptualize this process as a Bayesian inference, where an agent's prior knowledge (its previous observations of a vehicle's characteristics) are assessed for compatibility with the observed evidence (another fleet buying that type of vehicle). If, given the agent's prior state of knowledge, the observed evidence would be unlikely, the agent updates its prior knowledge to account for the new observation. In our model, Bayesian inference would be computationally infeasible to model directly. As an alternative, we use Approximate Bayesian Computation (ABC) [147], which offers a fast approximation of Bayesian inference that is compatible with our observation-perception model. We provide details of our ABC model in Appendix C.4.

Initial bias

In our model, agents do not abstract observations between vehicles powered by different fuels. As a result, when agents are introduced to vehicles powered by a new type of fuel, none of their prior observations apply to it. Instead, we specifically initialize a set of observations about that fuel type, from which all vehicles powered by that fuel can inherit. These initial observations assume the new fueling option to be inferior, by a fixed proportion (5% by default), to existing options. This is to reflect an expectation that fleets would begin from a place of skepticism regarding new fueling options. (We test alternative values of this proportion in the sensitivity analysis in Section 4.6.1.)

Perceptions

Whereas observations consist of a continually growing collection of discrete data points, making procurement decisions requires summarizing the mass of observed data into simpler, compara-

ble quantities. We model this summarization procedure in two steps: (1) subjective valuation and (2) utility calculation.

Subjective Valuation

In this step, a distribution of observations, pertaining to a single attribute of a single vehicle, is collapsed into a single value, ranking it on a subjective scale of 0 to 5 relative to all known vehicles. If o_a^p is the p^{th} percentile of all observations an agent has made pertaining to attribute a , and o_{avi} is the i^{th} observation of attribute a for vehicle option v , about which N observations have been made, then the subjective value s is:

$$s_{av} = 5 \times \frac{1}{N} \sum_{i=1}^N \frac{o_{avi} - o_a^5}{o_a^{95} - o_a^5} \quad (4.2)$$

Intuitively, s is the weighted mean of the observations *pertaining to one vehicle*, mapped linearly between the 5th and 95th percentile of the observation distribution *for all known vehicles*. As the agent makes more observations, o_a^5 might decrease and o_a^{95} might increase, altering the support of the valuation function. As a result, for example, becoming familiar with a new, unprecedentedly superior vehicle will depress the subjective values of all previously known vehicles, without the underlying observations changing. Figure 4.3 illustrates some of these concepts.

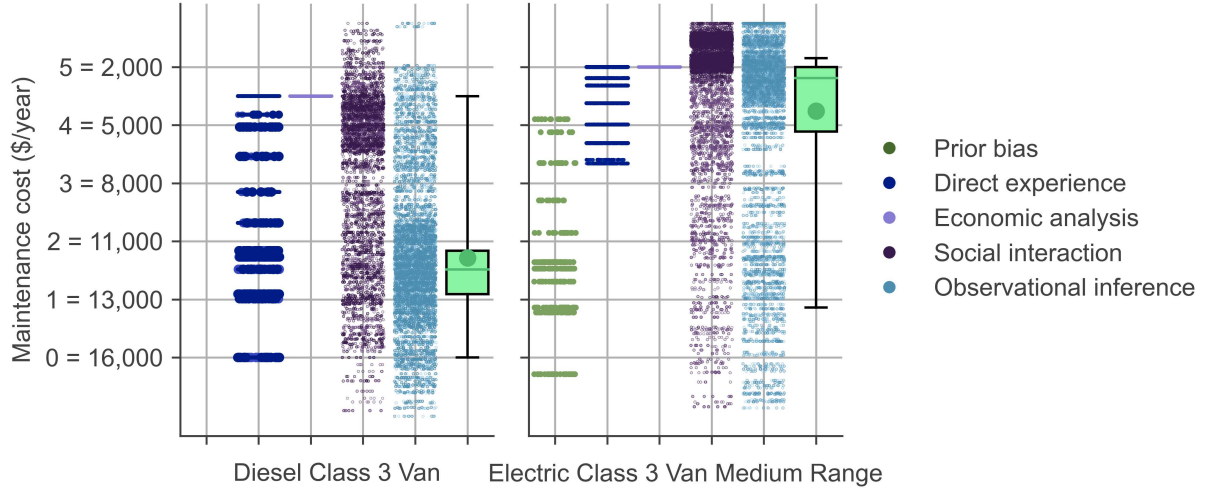


Figure 4.3: Visualization of how agents form perceptions on the basis of observations. Observations pertaining to maintenance cost for two vehicle options, an ICEV and a BEV, are shown, where more heavily weighted observations are shown as larger points. Colors identify the means by which observations were made. The overall distribution of observations is shown as a box-and-whisker plot, with the median as a line and the mean as a circle. The vertical axis shows how values map to subjective assessments, on a scale of 0 to 5, with lower maintenance costs mapping to higher subjective assessments.

Utility Calculation

With subjective valuations s_{av} for all attributes a of all known vehicles v , the final step in assessing a vehicle is to evaluate its holistic value, or “overall utility,” (u_v). In Chapter 3, we investigated and evaluated how fleets prioritize the various vehicle attributes in their purchasing decisions. Using that outcome to define the priority weight w assigned to each attribute a , the utility is simply calculated as the weighted sum of the subjective valuations over all attributes A :

$$u_v = \sum_{a \in A} w_a \times s_{av} \quad (4.3)$$

The calculation of u_v for two vehicle options is illustrated in Figure 4.4.




Figure 4.4: Illustration of the utility calculation for an ICEV and a BEV. In the lower plots, the widths of bars correspond to the relative weights assigned to vehicle attributes, and the heights correspond to the subjective assessments of those attributes, such that the fraction of area occupied by colored bars is proportional to the vehicle’s perceived utility. This is mapped onto the upper portions of the plots, which show the calculated utility, on a scale of 0 to 5, and the relative contribution of each attribute.

4.5 Verification and validation

Validating the predictive accuracy of the model would ideally involve comparisons with historical data of eMHD adoption, but this was not possible due to the immaturity of the eMHD market and the inaccessibility of data. However, because the purpose of this model is for qualitative experimentation rather than quantitative prediction, face validity is a more relevant requirement than predictive accuracy. To demonstrate the model’s validity, we present the results of four standard verification and validation procedures: submodule verification, comparison with another model, peer review, and sensitivity analysis (in the next section) [148].

4.5.1 Submodule verification

Economics model verification and validation

Our economics model synthesizes standard engineering economics methods, methods and assumptions from industry and U.S. government [149], and up-to-date component cost data from a real eMHD OEM. In this section, we present up-front and long-term cost results for representative vehicles and compare them to results from a leading economics model [149].

Up-front cost model

We construct up-front costs from models of the costs of individual high-value components of vehicles. Although the comparison model also builds from component costs, their component cost assumptions are not publicly available. Figure 4.5 shows the comparison model’s up-front cost estimates for a Class 4 delivery vehicle, alongside ours, for 2020 and 2025. Our estimates match closely with the “pessimistic” case from [149], perhaps suggesting that the battery costs paid by LeM are nearer to pessimistic projections than optimistic projections.

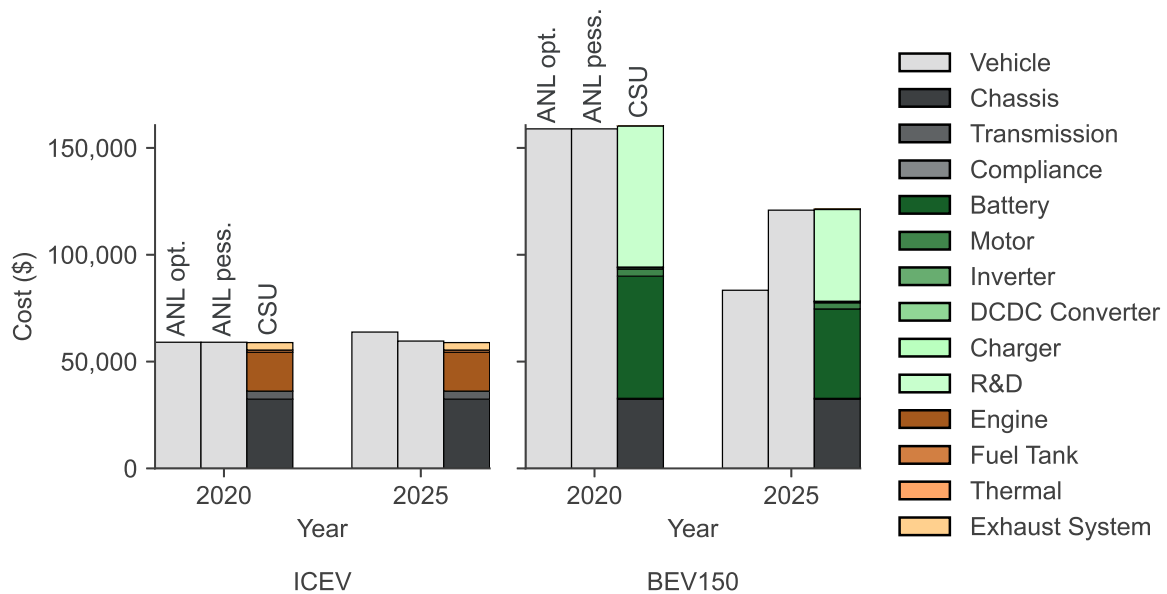


Figure 4.5: Comparison of up-front costs for a Class 4 delivery truck.

Long-term cost model

We use many of the same methods as published in [149] for the various components of long-term cost. For purchase and financing, [149] use their optimistic estimate of up-front costs, which is substantially lower than ours, accounting for the difference between long-term cost estimates for the BEV150. Our model also uses a more detailed model of electricity costs. Whereas they assume a non-varying energy price of \$0.11/kWh, our model takes utility rates, which can vary by time of day and season, and can include demand prices. For this comparison, shown in Figure 4.6,

we apply the same electricity pricing assumptions as the authors of [149]. (For all other results in this paper, we apply a comprehensive utility rate.)

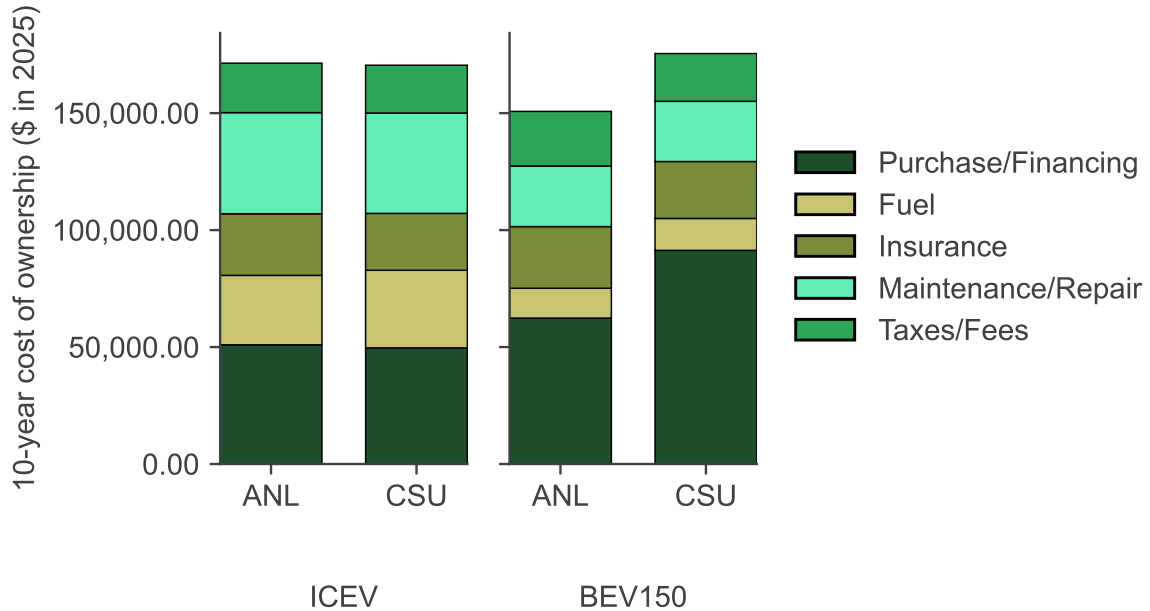


Figure 4.6: Comparison of long-term costs for a Class 4 delivery truck.

4.5.2 Comparison with other projections

Whereas a variety of LD EV adoption models have been demonstrated [150], few MHD EV adoption models or projections exist yet to compare to ours. The EIA AEO 2022, for example, projects zero EV sales for freight vocations between now and 2050 [151], which would be trivial to compare our results to. One eMHD adoption model that is not yet publicly available is the Transportation Electrification Mobility Pathway Options (TEMPO) model, from which some early eMHD sales projections have been published [152]. TEMPO estimates adoption, energy consumption, and other metrics by simulating mode and technology choice in discrete segments of the passenger and freight transportation sectors. The MHD sector is divided into segments described by cargo weight and distance (ton-miles). For each segment, TEMPO simulates the choice of mode (light/medium/heavy truck, rail, ship, air) and technology (ICEV, BEV, etc.). Although its model of MHD transportation only includes freight, not vocational transport, and encompasses

the entire freight transportation system rather than a small group of fleets, we can make qualitative comparisons between vehicle sales in our model and in TEMPO's, keeping in mind the differences in scope and assumptions. Figure 4.7 shows the vehicle sales over time for the two models. In the case of TEMPO, the results are for "light-medium freight," as defined in the AEO, and ours are for the mix of vocations described in 4.6.

For the results shown in Figure 4.7, we simulated only delivery (regional and last-mile) fleets, to match TEMPO's focus on freight vehicles. As shown in the figure, our model displays a qualitatively similar trend of rapidly decreasing ICEV sales, ending around 2035, which are then replaced by BEV sales. Note that, in TEMPO, the deployed vehicle stock increases over time, whereas the total number of vehicles deployed in our model remains constant throughout.

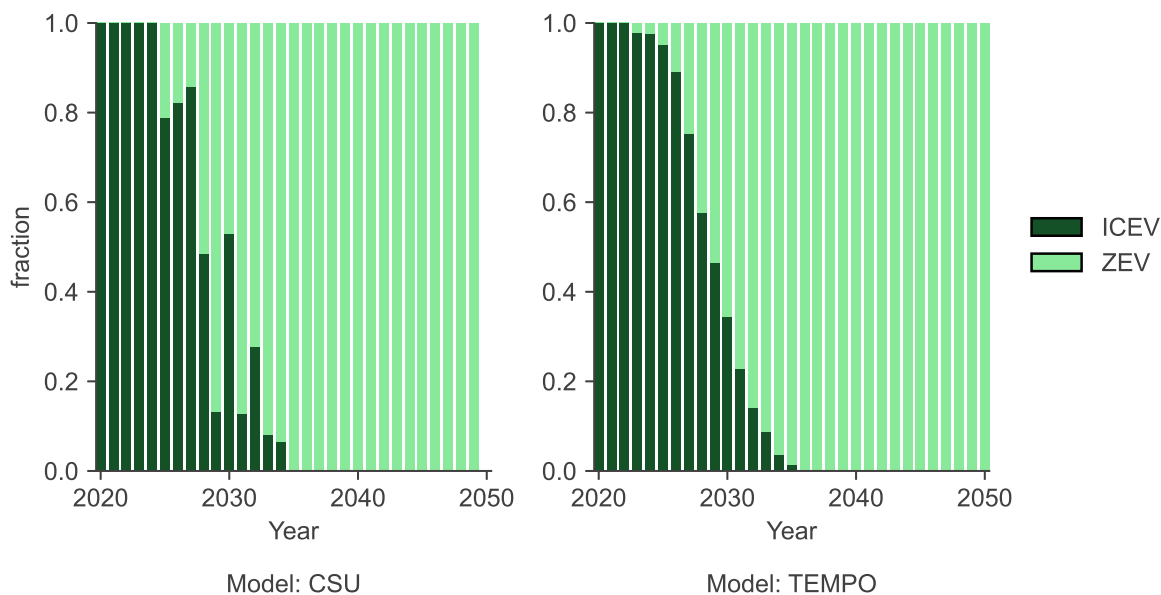


Figure 4.7: Comparison between the projections made by the TEMPO model, for light-medium duty freight, and our model, for Class 3 delivery vehicles.

4.5.3 Peer review

The final piece of verification and validation of our model was to invite regular input and criticism from peers and experts, . While developing the model, we delivered progress presentations to

project sponsors at the U.S. Department of Energy, project partners at EPRI and Lightning eMotors, and a range of experts from CSU and the National Renewable Energy Laboratory (NREL). These progress presentations were valuable opportunities for others to assess our assumptions and methods. Every meeting yielded actionable input that has been integrated in the current model, such as directing us to new sources of data and brainstorming alternative approaches to decision weighting. These efforts will continue as we continue to develop and apply our model.

Table 4.2: List of major peer-review instances and outcomes. Acronyms: Colorado State University (CSU, author’s PhD committee), U.S. Department of Energy Vehicle Technologies Office (DOE VTO), Electric Power Research Institute (EPRI), Lightning eMotors (LeM), National Renewable Energy Laboratory (NREL), Renewable and Sustainable Energy Reviews journal (RSER), Sacramento Municipal Utility District (SMUD).

Date	Party	Outcome
11/18/2020	EPRI, SMUD	Economics model scoping (spatial, temporal)
12/18/2020	LeM	Data scoping and procurement plan
03/02/2021	EPRI	Feedback on economics model (suggested sources for validation)
04/07/2021	DOE VTO	Submission of funding application (approved July 2021)
11/01/2021	CSU	Feedback on overall plan and scoping
01/18/2022	NREL	Feedback on decision model (early assessment of face validity)
04/25/2022	DOE VTO	Validation of economics model (request comparison with other DOE models), feedback on decision model (early assessment of face validity)
05/23/2022	RSER	Submission of Charging-as-a-Service manuscript (second revision in review as of 12/12/2022)
07/27/2022	LeM	Refined data procurement plan
08/12/2022	NREL	Feedback on decision model and agent-based model design (midway assessment of face validity)
09/19/2022	DOE VTO	Feedback on decision model and agent-based model design (midway assessment of face validity)
12/09/2022	NREL	Assessment of validation/verification results and future work suggestions
01/24/2023	DOE VTO	Feedback on agent-based model validation/verification and heterogeneity experiments

4.6 Demonstration and sensitivity assessment

We design and test a set of input data and parameters to simulate a group of 30 fleets that operate Class 3 (MD) vehicles to serve a variety of vocations. We use drive cycle data from real fleets, which was provided by Lightning eMotors on behalf of its customers. The drive cycles represent delivery, airport shuttle, and ride-share vocations. We site our simulated fleets in the metropolitan area around Los Angeles, USA, and subject them to a commercial utility pricing structure from Southern California Edison (SCE) (Table A.2 in Appendix A). Procurement decisions are made by fleets every simulation year, from 2020 to 2050. We apply battery price and diesel price projections from the EIA AEO [151]. All BEVs are assumed to qualify for the IRA subsidy [153] throughout the simulation (although it is scheduled to end in 2032). All fleets weight vehicle attributes identically in their decision-making, as defined in Chapter 3.

For simplicity, we simulate only two OEMs: one which produces a single class 3 van powered by diesel, and one which produces a class 3 electric van, which can be equipped with a short-range battery (80 miles) or a medium-range battery (140 miles). Other attributes of the vehicle options are provided in the appendix.

Fleet specification

Fleet sites were chosen from five searches in Google Maps in the Los Angeles area. Queries and the corresponding vocations are shown in Table 4.3. For each query, a random subset of locations was chosen from the search results, and a fleet was specified from the LeM data for each location. Fleet specifications include number and type of vehicles, and drive cycle for each vehicle. All fleets were assumed to operate Class 3 diesel vans. The locations are shown in Figure 4.8. Drive cycles are drawn from real-world data for five vocations: regional delivery, last-mile delivery, airport shuttle, ride share (passenger transport), and microtransit (passenger and cargo transport).

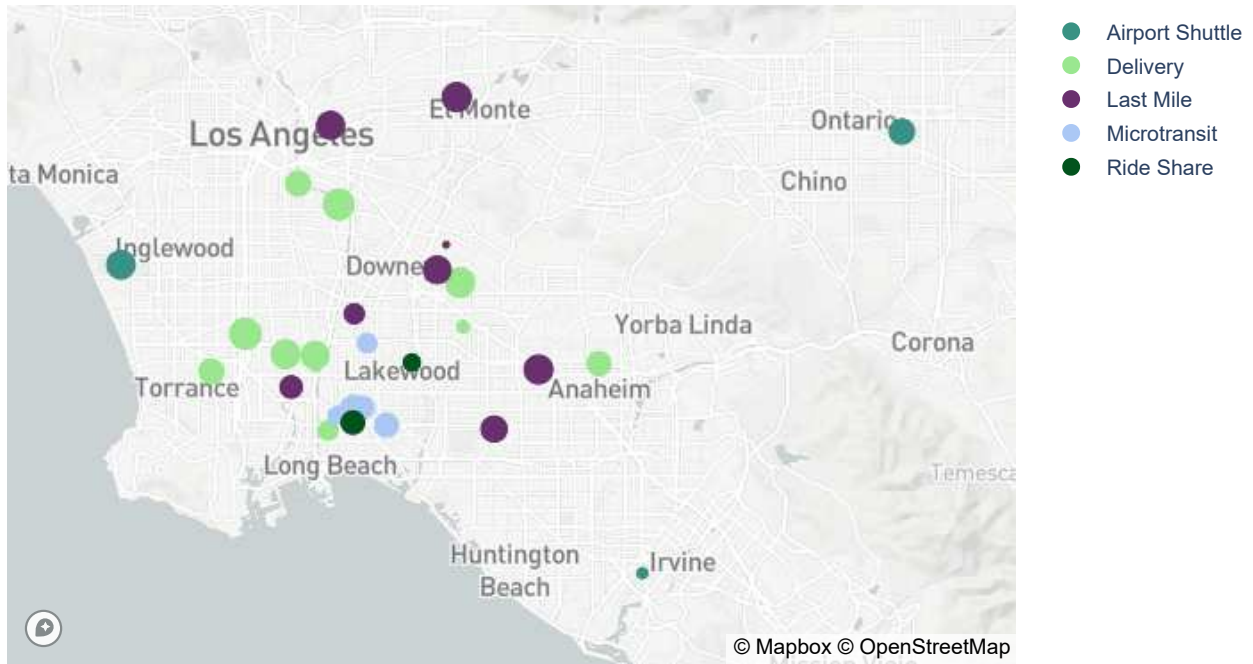


Figure 4.8: Map showing the locations of the simulated fleets. Fleet sizes, in number of vehicles, correspond to the sizes of the circles on the map.

Table 4.3: Google Maps queries used to find sites for simulated fleets, of five different vocations, in the Los Angeles metropolitan area.

Query	Count	Vocation(s)
airport	3	Airport shuttle
amazon warehouse	8	Last-mile delivery and regional delivery
car rental	7	Ride-share and microtransit
distribution center	8	Last-mile delivery and regional delivery
fedex ground	4	Last-mile delivery and regional delivery

4.6.1 Sensitivity analysis

In a sensitivity analysis, we test the effect of changes to input parameters on simulated outcomes. The design and outcomes of the baseline scenario were described previously (Section

4.6); here, we present the parameters we varied and what effect those variations have on adoption outcomes.

Note on computational costs

The ABM and sensitivity analysis computation approach is costly, particularly in terms of computation time and storage. The baseline scenario presented in Section 4.6 takes over 2 hours to simulate and generates over 2 GB of output data (most of which are records of agents' observations). For the sensitivity analysis, we simulated a subset of the fleets from the baseline scenario (7 fleets out of 30), reducing the simulation time to 45 minutes and enabling the full set of sensitivity scenarios to be completed in a few days, rather than a week or two. Reducing the computational cost of the model will be necessary to make simulating larger populations tractable. (Future work could investigate how results change as the simulated population increases. Experiments on larger populations may not require one-to-one simulations.)

Stochasticity

Some of the model's dynamics require generating pseudo-random numbers to draw from probability distributions. Examples include how fleet agents decide to adopt pilot vehicles, how they discover previously unknown OEMs, and which observations fleets share with each other during social interaction (Section 4.4). Repeatable outcomes are possible by specifying the initial state of the pseudo-random number generator (RNG). Figure 4.9 shows the relative stock of BEVs in the simulated population for several simulations, each of which has a different initial RNG state. The bold trace is from the baseline run presented elsewhere in this section. The effect of stochasticity is most significant in the middle years of the simulation. In all cases, adoption is near zero until approximately 2025, when it begins to increase, before reaching 100% around 2040. In the intervening years, different RNG states can cause the overall adoption curve to lag or lead by one or two years.




Figure 4.9: Adoption outcomes from many simulations using different initial random number generator states, but otherwise identical input parameters. The bold trace is from the baseline simulation presented throughout Section 4.6.1.

Parameter variation

The sensitivity of adoption outcomes to variations in select model parameters is shown in Figure 4.10 and described in the sections below.

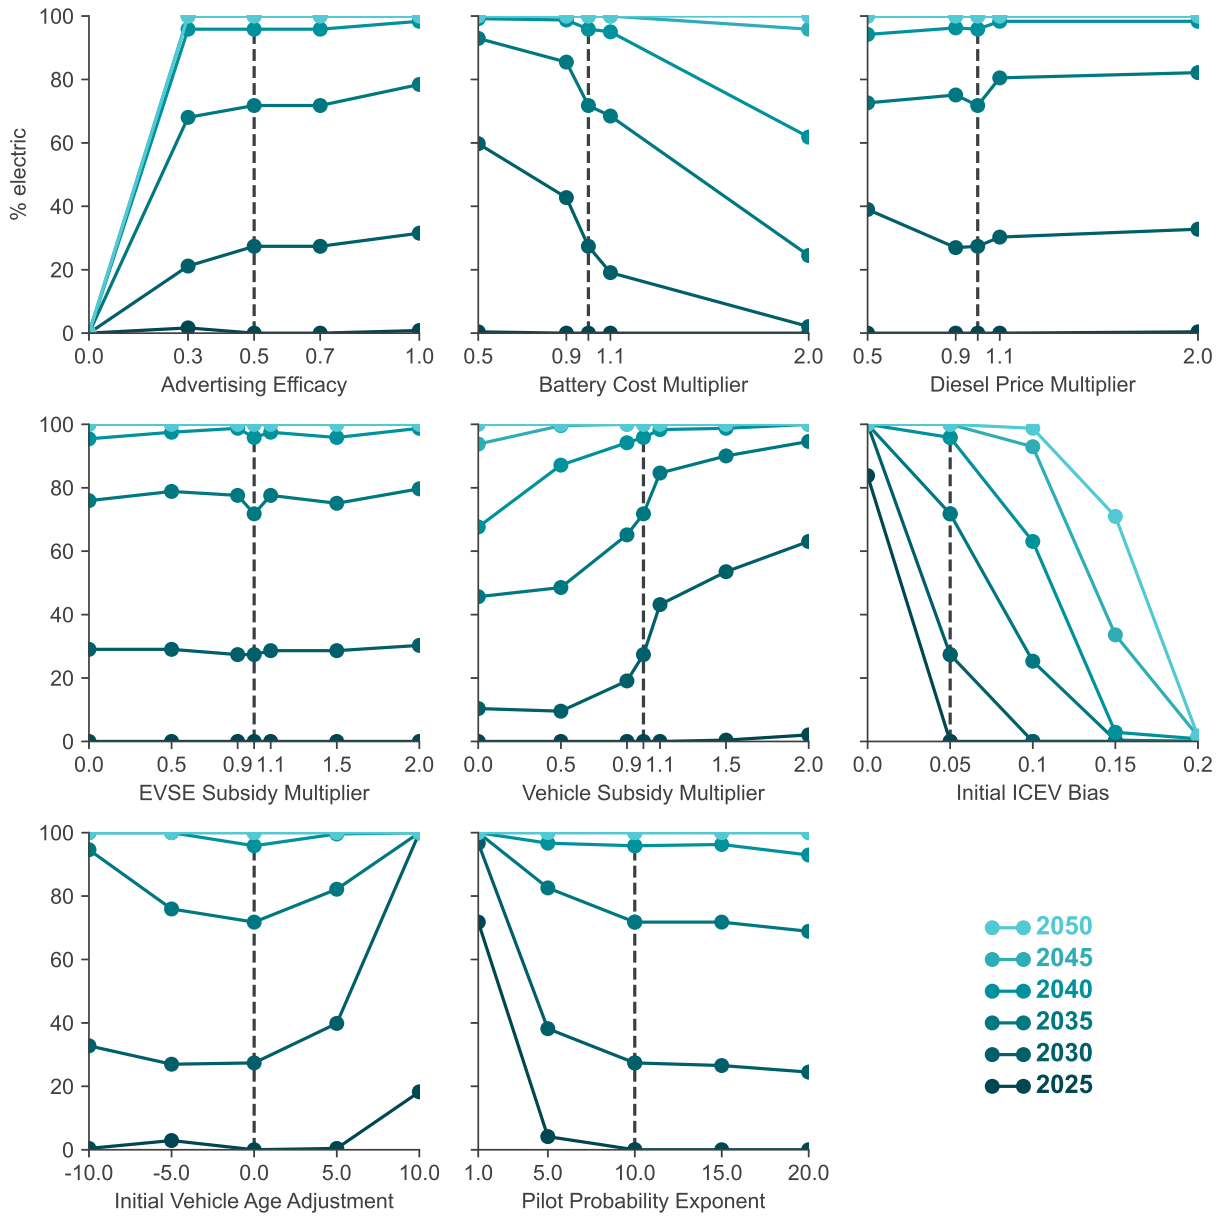


Figure 4.10: Illustration of the effect of varying input parameters on EV adoption outcomes. Parameter choices are shown on the horizontal axis. The parameter used in the baseline run is marked by a dotted vertical line. Outcomes for years 2025, 2030, 2035, 2040, 2045, and 2050 are shown, where later years are represented by lighter-colored traces.

Advertising efficacy

Advertising efficacy is the probability that at least one fleet agent will become aware of a new OEM (i.e., the electric OEM) via advertising during each simulated year. Because no fleets have

discovered the electric OEM at initialization, this is the only means by which the population can be first introduced to it. In the trivial case, where the probability is zero, adoption of EVs remains at zero throughout. Otherwise, early in the simulation, there is a slight positive relationship between EV adoption and advertising efficacy. Later on, EV adoption approaches 100% in every nontrivial case.

Cost multipliers

Results are highly sensitive to battery costs, which are controlled by a cost multiplier, where high costs significantly delay EV adoption and low costs accelerate it. The effect of diesel prices is smaller, but generally, higher diesel prices speed EV adoption. In a few instances, notably 2025 in the 0.5x case, this trend does not hold true, which may be an example of nonlinear and stochastic processes in the model causing unintuitive outcomes.

Subsidy multipliers

By default, fleet agents take advantage of the EV and EVSE subsidies enacted in the Inflation Reduction Act [153]. Adoption outcomes are less sensitive to EVSE subsidies, likely because the EVSE required for the simulated vehicle types are inexpensive in relation to the vehicles themselves. In contrast, the EV subsidy appears to be a strong driver of adoption, particularly early on. In 2025, the model predicts the IRA subsidy to result in twice as many EVs in 2025 as the case without it, and if the subsidy were doubled by doubling the subsidy multiplier, the result would be twice again as many. In later years, adoption reaches significant levels irrespective of the subsidy amount.

Initial ICEV bias

In the baseline scenario, fleets initially assume BEVs to have 5% lower utility than ICEVs. Figure 4.10 shows that model outcomes can be quite sensitive to this assumption. If the initial bias is zero instead of 5%, EV uptake is significant and almost immediate; if the initial bias is 20%, EV uptake remains at zero throughout. This highlights how our model of learning has an implicit inertia, causing fleets' knowledge to lag the modeled ground truth. If the initial bias is too pessimistic, a great deal of new information is needed to overcome it enough for any fleets to consider

adoption. Moreover, we model all fleets to begin with the same bias. In reality, different fleets have different levels of bias and states of knowledge, which would change the overall dynamics of adoption. A more realistic scenario might include fleets that are highly EV-pessimistic alongside fleets that are highly EV-optimistic.

Initial vehicle age adjustment

Vehicles owned by fleets are initialized with a uniform distribution of age between 0 and 15 years. The initial vehicle age adjustment tests how the outcome changes if the initial ages of vehicles are increased or decreased (clipped at 0 and 15). At the low extreme (-10), fleets begin with mostly new vehicles, and at the high extreme (+10), fleets begin with vehicles nearing retirement. Both extremes result in faster EV adoption than the baseline case. If most vehicles are new at the start of simulation, initial uptake of EVs is delayed (recall that vehicles cannot be replaced before they are 8 years old). Due to declining battery costs and rising diesel costs, among other factors, the utility of EVs increases over time. The longer fleets are prevented from forming impressions about EVs, the fewer poor observations (e.g., of high costs or poor reliability) they can make as a population. Once they begin experiencing EVs in the late 2020s and early 2030s, then, their initial impressions are improved on average relative to the baseline case, accelerating adoption in later simulation years. At the other extreme, EVs experience highly fortunate timing: entire fleets of vehicles are retiring just as EVs are becoming feasible replacements. Because the vehicles must be replaced anyway, many fleets take a chance on EVs early and quickly discover that they offer superior holistic value to ICEVs.

Pilot probability exponent

The pilot probability exponent is the constant k in Equation 4.1, which scales the difference in utility, measured on a scale of 0-5, between vehicle options. Recall that smaller values of k increase the probability of piloting a vehicle. If k increases by 1, the probability of piloting decreases by a factor of $10^{u_1 - u_0}$. As shown in the figure, $k = 1$, which simulates an overwhelming eagerness to pilot new kinds of vehicles, results in rapid uptake of EVs. For $k > 1$, there is a slight trend of slower uptake with increasing k .

Isolated decision-making attributes

In another sensitivity analysis approach, we compare the baseline scenario results with outcomes when fleet agents' adoption decisions consider each vehicle attribute (Table 4.1) in isolation, rather than a holistic assessment. As shown in Figure 4.11, the results vary drastically between different attributes. For attributes for which BEVs are indisputably superior to ICEVs from the beginning (TCO, brand image, environmental sustainability, fuel cost, and maintenance cost), uptake begins early and reaches 100% quickly. Because ICEVs have the opposite advantage with range suitability, refueling convenience, and infrastructure cost, they remain dominant throughout in those simulations. In between these extremes, there are attributes for which neither fuel type has an overwhelming advantage (reliability, fuel security), or for which BEVs begin inferior, but become superior over time (purchase price). These simulation results, particularly when considering only purchase price, somewhat resemble the results of holistic assessment.

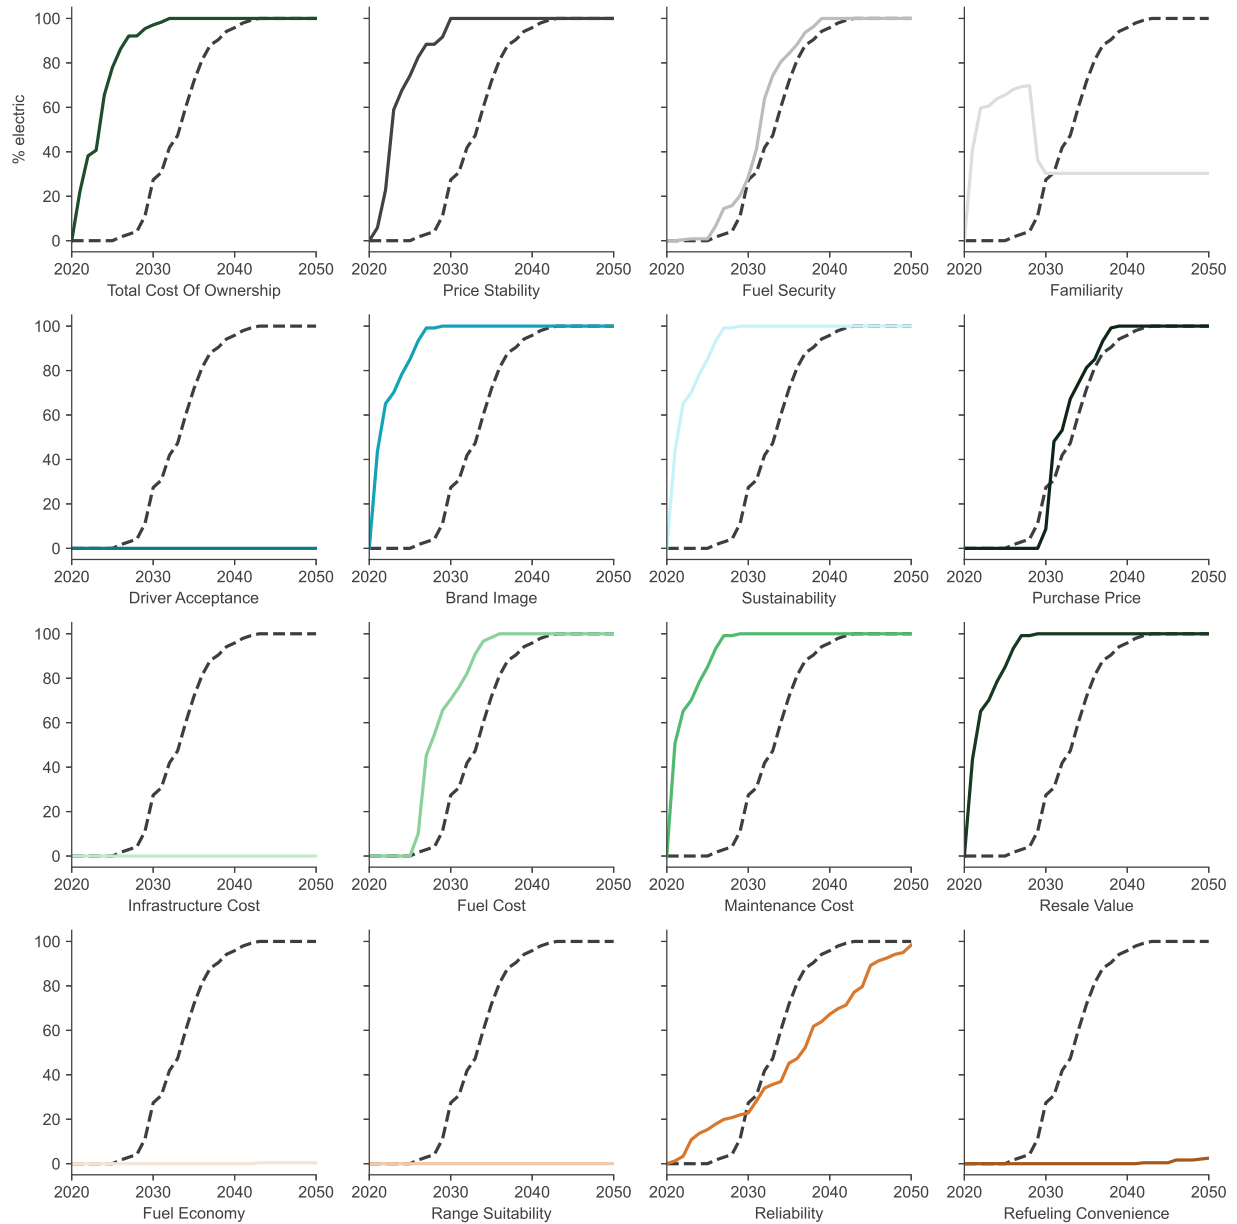


Figure 4.11: EV adoption outcomes when the fleet population is simulated to make adoption decisions with strict consideration of a single attribute. The dashed line is the baseline outcome, in which agents weigh all factors against one another.

4.7 Simulating more heterogeneous populations

In the results presented so far, the simulated population has been heterogeneous in several respects: each fleet operates vehicles on different drive cycles, forms a unique set of perceptions,

and owns a unique set of vehicle and EVSE assets. However, for simplicity during validation and verification, we restricted some other kinds of heterogeneity that the model is designed to simulate. For example, fleets only operated Class 3 vehicles (homogeneous weight classes), all agents made decisions based on the same set of priorities (homogeneous decision weights), and EVs were only available with two electric-range options (nearly homogeneous range options). In this section, we present experiments to test how introducing and expanding these additional kinds of heterogeneity can affect simulation outcomes.

As input to these experiments, we use an alternative set of drive cycle data representing a greater variety of drive cycles and weight classes. The Fleet DNA database [154], from NREL, includes operating data for commercial vehicles ranging from class 2 (up to 10,000 lbs) to class 8 (maximum allowable weight vehicles), representing eight vehicle types: delivery vans, delivery trucks, school buses, transit buses, bucket trucks, service vans, tractors, and refuse trucks. The distribution of daily travel distance for vehicles in the database, shown in Figure 4.12, skews low, with a mode of about 40 miles per day (75% of travel days have a distance of 85 miles or less; 95% are less than 180 miles). These data enable our experiments to include a greater variety of vocations, vehicle types, and weight classes, but still have limited variety in daily travel distance.




Figure 4.12: Histogram of travel distance across all of the vehicle-days represented in the Fleet DNA dataset.

Three experiments test the effect of greater heterogeneity in three attributes. In each of these experiments, we compare the adoption outcomes between full simulations that include different levels of heterogeneity. Every simulation considers 20 fleets completing drive cycles drawn from Fleet DNA. The experiments address the following questions:

1. **Weight class heterogeneity:** How does electrification of lower weight classes, which is generally more economically and operationally feasible, affect the electrification trajectory of fleets operating higher weight classes? *Hypothesis: the adoption trajectory of an HDV population will accelerate if the population can observe the actions of a MDV population.*
2. **Preference heterogeneity:** How do differences in the range and variety of decision-making preferences in a population affect electrification outcomes? *Hypothesis: a population with average decision preferences adopt EVs more quickly if the population can observe a population with especially EV-positive preferences. Oppositely, it will adopt EVs more slowly if observing a population with especially EV-negative preferences.*
3. **Range option heterogeneity:** How does the set of available electric range options affect the trajectory of EV uptake? *Hypothesis: a population will adopt EVs more quickly if a greater span and number of electric range options are available.*

4.7.1 Weight class heterogeneity

For this experiment, the population comprises 10 “heavier” fleets operating Class 6-8 vehicles (>19,500 lbs) and 10 “lighter” fleets operating Class 3-5 vehicles (14,000-19,500 lbs). We compare two scenarios. In (a), there is no communication between the heavier and lighter populations (i.e., no social interaction and no social observation). In (b), we impose no such division between populations, so heavier and lighter fleets can interact with and observe one another. The specifications of the fleets (drive cycles, number and type of vehicles) are identical between scenarios. Looking only at adoption outcomes for the heavier population of fleets, we investigate whether and how interacting with lighter fleets affects the timing and magnitude of electrification of the heavier fleets. Figure 4.13 illustrates and compares adoption outcomes for the two scenarios. Figure 4.14

illustrates when fleets first consider one or more EV model to have superior utility to the ICEV model they have conventionally used.

Table 4.4: Summary of scenarios considered in the weight class experiment.

Scenario	Description
a	Heavier/lighter populations divided
b	Heavier/lighter populations not divided

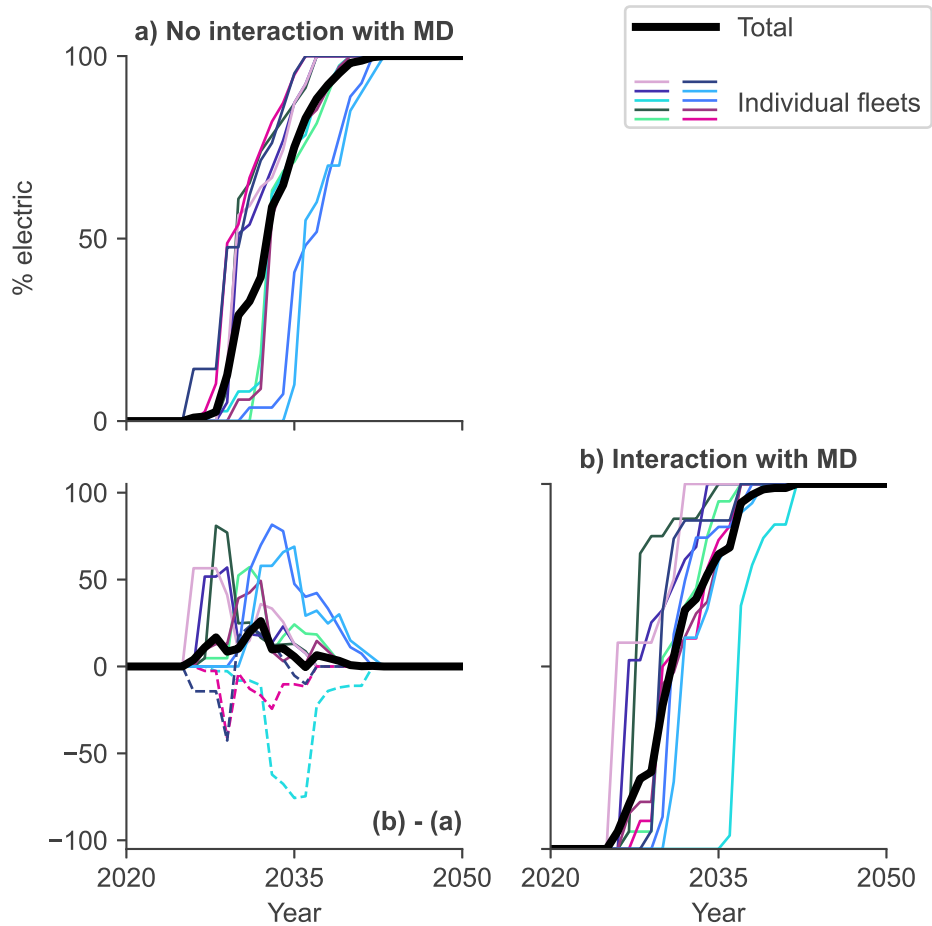


Figure 4.13: Electrification outcomes for the population of 10 Class 6-8 fleets in the weight class heterogeneity experiment. Figures on the diagonal illustrate the electrification trajectory, in terms of the percentage of vehicles that are electric, for the two scenarios. The off-diagonal figure (bottom left) shows the subtractive difference between the two scenarios. Thin colored lines are results for individual Class 6-8 fleets. The bold black line aggregates over the Class 6-8 population.

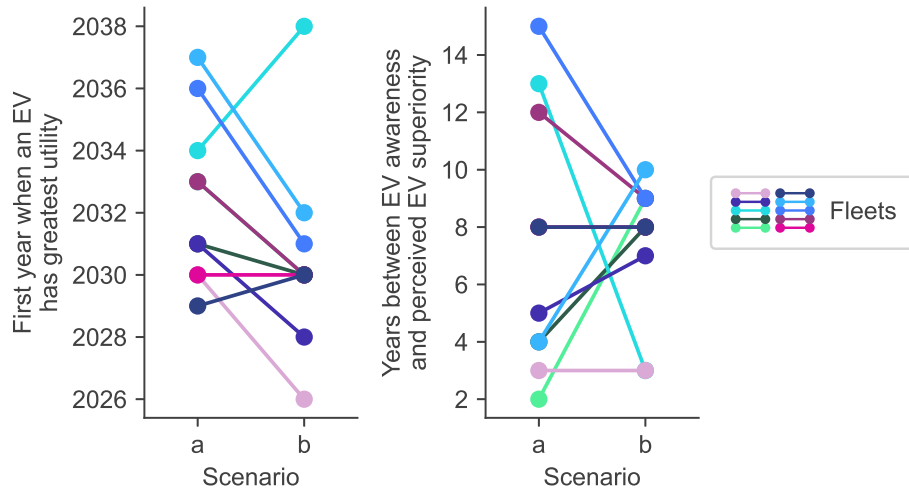


Figure 4.14: Timing of perception crossover: the first year when each fleet considers one or more EV models to have greater utility than the ICEV model, for the two scenarios. The chronological year of crossover is shown on the left. On the right, the year of crossover is instead expressed relative to the year when each fleet first becomes aware of suitable EV options.

As shown in Figure 4.13, interacting with lighter fleets indeed appears to accelerate EV adoption among heavier fleets. For 7 of 10 fleets, the difference in adoption between (b) and (a) is positive throughout the simulation. In aggregate, overall adoption leads in Scenario (b) by up to approximately 25% (around 2030). The results in Figure 4.14 suggest that this acceleration is at least partially caused by fleets becoming aware of EVs earlier. The crossover point for EVs arrives several years sooner in (b), although, in many cases, the time from awareness to crossover is increased in (b). Put another way, when they can interact with lighter fleets, heavier fleets begin considering EVs earlier, during the periods in which suitable EVs are still too costly for them. Despite being introduced prematurely, this gives the fleets a head start to gather observations overcoming the initial ICEV bias, enabling them to adopt quickly once EV costs are acceptable to them.

For the 3 fleets that lag in (b), the lag is due to a delay in awareness of suitable options, a stochastic dynamic. This indicates that subtle network effects which express in aggregate, such as the EV-normalizing influence of lighter fleets in a population demonstrated here, can be overridden by stochasticity at more localized scales. So, although we showed in Figure 4.10 that varying

advertising efficacy (the parameter governing stochastic awareness) has a small global effect, this experiment shows that it can have a significant effect on individual fleets.

4.7.2 Preference heterogeneity

In this second experiment, we test how agents with different decision-making preferences influence one another. Our hypothesis is that a population that includes agents with preferences tailored to near-term EV adoption would serve as a normalizing example, thereby speeding adoption by the rest of the population. Oppositely, we test to see whether the opposite is true: if a population includes agents with preferences tailored to the status quo, will adoption by the rest of the population be slowed?

The simulated population is divided into a control group of 14 Class 3-5 fleets with uniform decision preferences, for which results are analyzed, and an “influencing” group of 6 Class 3-5 fleets. We modify the decision preferences of the influencing group and examine whether and how this affects EV adoption in the control group. Analogous to the previous section, Table 4.5 summarizes the design for the simulated scenarios, Figure 4.15 illustrates and compares adoption outcomes for the four scenarios, and Figure 4.16 illustrates the timing of perception crossover.

Table 4.5: Summary of scenarios considered in the decision preference experiment.

Scenario	Description
a	Uniform preferences throughout
b	Influencing population heavily prioritizes brand image
c	Influencing population heavily prioritizes purchase price
d	Influencing population: half prioritize brand image, half prioritize purchase price

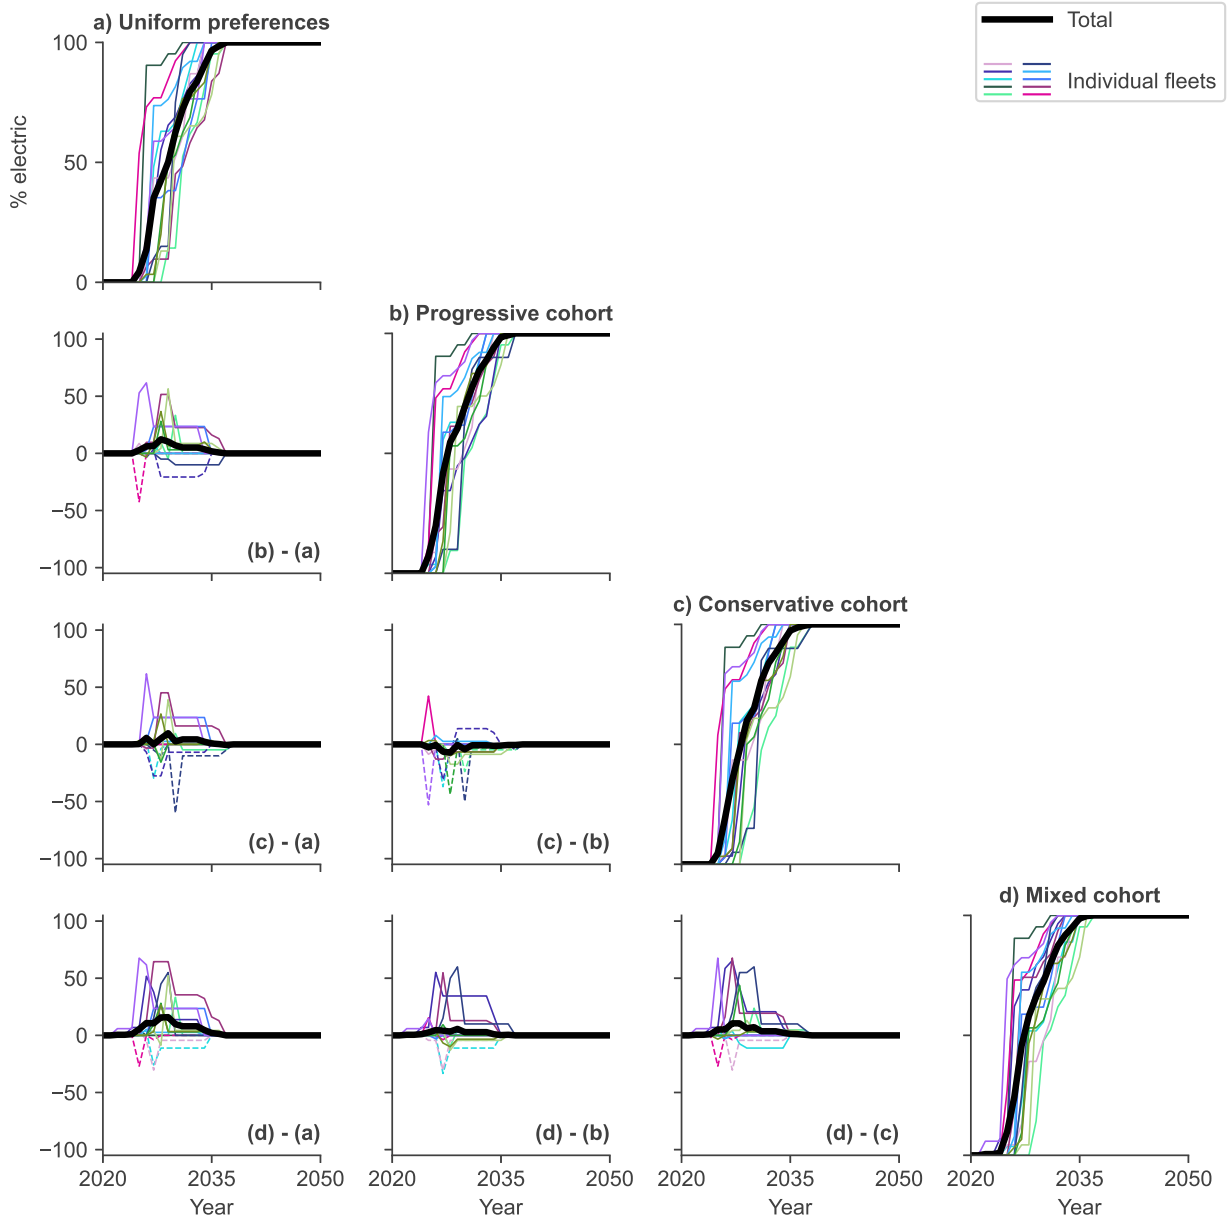


Figure 4.15: Electrification outcomes for the control group of 14 fleets in the decision preference experiment. The figure is formatted analogously to Figure 4.13.

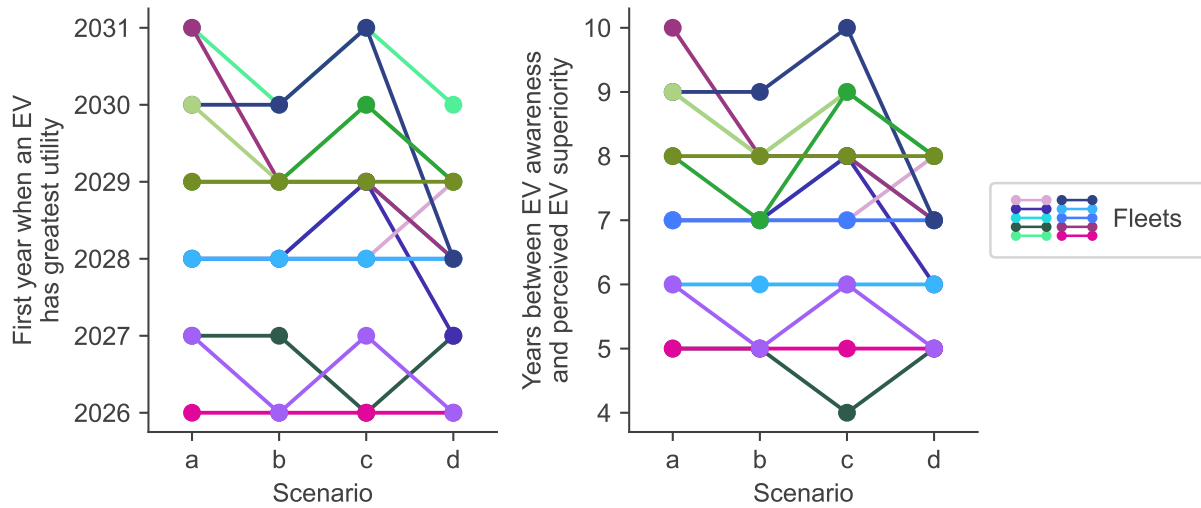


Figure 4.16: Perception crossover for the control group of 14 fleets in the decision preference experiment. The figure is formatted analogously to Figure 4.14.

Similar to the previous experiment, results support the hypothesis that early adopters can normalize a technology for others. Compared to the baseline simulation in (a), adoption outcomes are accelerated in the control group if there are “progressive”, i.e. brand image prioritizing, fleets to observe and interact with. For (c), one might expect equal and opposite results to (b) – delayed adoption due to a “conservative”, up-front cost prioritizing, influence – but this is not the case. This effect occurs because fleets are biased against EVs by default. Progressive fleets influence the control group to revise their perceptions of EVs via their surprising (in a statistical sense, informative) adoption of EVs. In contrast, conservative fleets’ behavior of continuing to buy ICEVs is uninformative to the control group, because it is expected behavior. As a result, (c) has only a small effect on adoption outcomes, mostly due to stochastic effects. Similarly, although (d) includes fewer progressive fleets than (b) and includes some conservative fleets, adoption is accelerated again relative to all other scenarios. The fact that (d) slightly leads (b) in aggregate could suggest that more diversity in a population, including both detractors and promoters of a technology, improves adoption outcomes of the technology, perhaps due to efficient exchange of information. However, the difference is small and may only be due to stochastic effects. Further experiments are

needed to understand the detailed effects of preference diversity on technology adoption, although the evidence from these tests supports the hypothesis that early adopters accelerate EV adoption population-wide.

4.7.3 Range option heterogeneity

In this final experiment, we test the effects of varying the set of vehicle options available for purchase. We hypothesize that a greater variety of electric range options, i.e. smaller differences between the ranges available for purchase, would result in faster adoption. Intuitively, a fleet manager should choose the shortest-range vehicle that satisfies the fleet’s requirements. Any additional range would superfluously add cost to the vehicle. A finer resolution of range availability would thus be more likely to include an optimally suitable option for a given fleet, enabling them to adopt EVs sooner. For every scenario, we simulate and analyze a population of 20 Class 3-5 fleets. Scenarios differ only in which electric range options are available to the fleets. As with the other experiments, Table 4.6 summarizes the design for the simulated scenarios, Figure 4.17 illustrates and compares adoption outcomes for the four scenarios, and Figure 4.18 illustrates the timing of perception crossover.

Table 4.6: Summary of scenarios considered in the range option experiment.

Scenario	Range options available
a	50 mile
b	250 mile
c	(50, 150, 250) mile
d	(50, 100, 150, 200, 250) mile
e	(50, 75, 100, 125, 150, 175, 200, 225, 250) mile

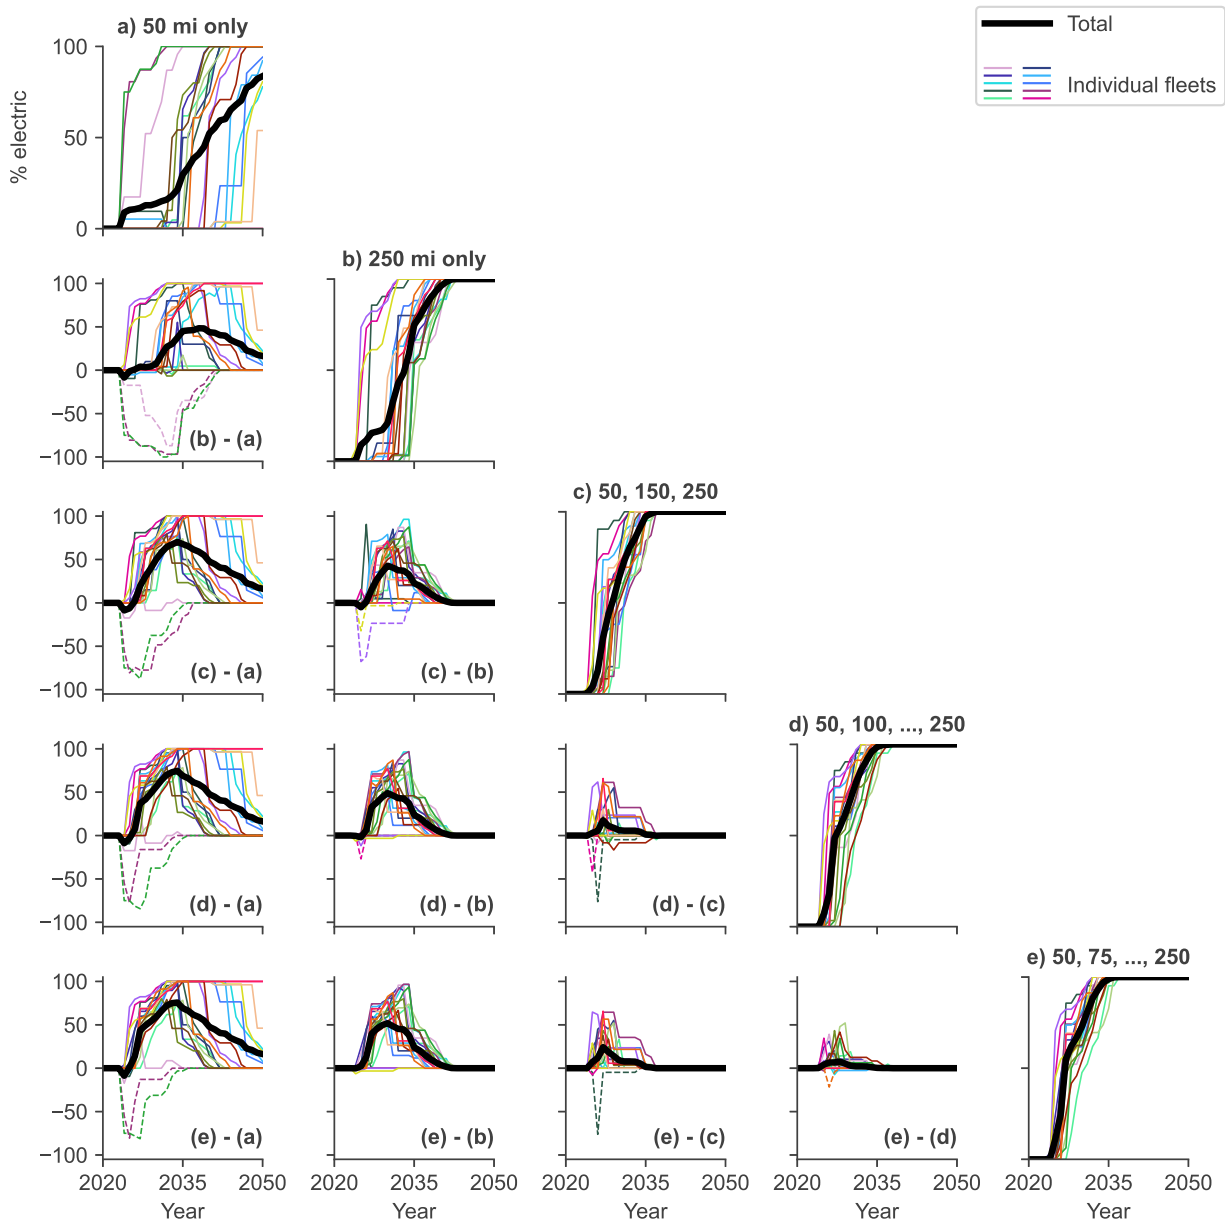


Figure 4.17: Electrification outcomes for the full population of fleets in the range option experiment. The figure is formatted analogously to Figures 4.13 and 4.15.




Figure 4.18: Perception crossover for the control group of 14 fleets in the decision preference experiment. The figure is formatted analogously to Figures 4.14 and 4.16.

As expected, more range options generally result in faster EV uptake. The successive improvement between scenarios can be seen in the first set of off-diagonal subplots in Figure 4.17 (b-a, c-b, d-c, and e-d). Each successive comparison exhibits a diminishing improvement over the scenario with the next-fewest range options: whereas (c) has greatly accelerated uptake compared to (a) and (b), the benefit of (e) compared to (d) is small.

A less expected result is evident in the leftmost column of Figure 4.17. For 3 of the 20 fleets, the fastest adoption occurs in (a). For these fleets, an EV with 50 miles of range is the obvious choice early in the simulation of (a) because of its affordability, and its suitability for their short-range operations. In (b), a vehicle with 250 miles of range is superfluous and costly, and, unsurprisingly, they do not buy any until nearly halfway through the simulation. However, in (c)-(e), despite having 50 miles of range as an option, they are slower to adopt than in (a). Recall that fleets make observations of all available options, and their perceptions of different models are partially shared as a function of similarity between models. Thus, these fleets’ observations of expensive, high-range models “shade” their perception of the 50-mile model sufficiently to delay their first purchases of it. We cannot conclude whether this dynamic translates to the real world, but note that it bears some resemblance to the phenomenon of “choice paralysis” – the “negative effects

on consumer choice behavior due to high product variety” [155]. At minimum, it serves as an instructive example of the kinds of nonintuitive outcomes that can arise from an ABM. Given that our ABM is a highly simplified model of the real-world system, then, this outcome is a reminder that vastly more nonintuitive outcomes, which no adoption model can anticipate with precision, are certain to occur in reality. Bottom-up approaches such as ABM, and the general study of complex systems, are increasingly necessary tools for exploring these kinds of dynamics, even if predicting them is impractical.

4.8 Conclusion

In this work, we presented the motivation, design, validation, and demonstration of an agent-based model for testing the effects of interventions on uptake of electric medium- and heavy-duty vehicles. By combining a detailed model of EV economics and operation with a Bayes-inspired model of social learning and opinion formation, we achieve results consistent with other models. Having tested the model for coherency and validity, we justify its suitability for further development and experimentation.

We focused the majority of this chapter on demonstrating the validity and potential of our approach. In addition to its purpose as a validation step, our sensitivity analysis includes examples of experiments that are simple to run using our approach but would not be compatible with top-down approaches. For example, the sensitivity analyses of advertising efficacy, pilot probability, and initial vehicle age demonstrate how the dynamics of fleets’ initial exposure to EVs might influence short- and long-term adoption outcomes. Results from simulations where fleet agents consider single vehicle attributes in isolation, which outcomes vary drastically for different attributes, illustrate the importance of understanding fleets’ priorities.

Section 4.7 demonstrates more types of experiments made possible by bottom-up adoption modeling. Factors that are not considered in conventional adoption models, particularly the interactions between decision makers, can substantially alter adoption outcomes. As we show in two of the experiments, outcomes for individual fleets are sensitive to the requirements and decision

priorities of other fleets with which they can interact and which they can observe. In addition to serving as examples of bottom-up adoption modeling experiments, these results illustrate the unavoidable complexity of understanding, let alone modifying, population-level choice dynamics.

Because our model attempts to represent complex dynamics that are absent from preceding models, experimenting with our model requires specifying parameters that cannot be directly measured and, instead, must be estimated, complicating the predictive utility of the results. (An alternative to parameter estimation is model fitting, which is not currently feasible due to a lack of data and a high computational cost.) By externalizing parameters, top-down models can avoid measuring or estimating them. But, if a parameter has significant real-world effects, externalizing it is equivalent to hard-coding an assumed value for it. This works if the assumed value is correct and remains correct over time. However, if any externalized parameters deviate from their implicitly assumed values, for example, as a result of a policy intervention, top-down models are unable to adapt. This trade-off between simplicity and capability accounts for each of the key costs and benefits of increasingly bottom-up modeling.

Research Question #3 asks how decisions by stakeholders in the MHD transportation system can contribute to desired outcomes – for example, equitable decarbonization and improvements to air quality, human mobility, and freight transportation efficiency. The model we have developed internalizes many important parameters that are typically externalized in adoption modeling, which enables experiments to test the potential effects of interventions on EV adoption, how it is distributed, and how the system must adapt to enable it. Our initial experiments illustrate some of the ways in which individual actions, taken separately and in aggregate, can cause down- and cross-network effects that cannot be represented in dominant modeling approaches, but that make a real difference in the electric transportation market.

Chapter 5

Future work

Supply-side agent enhancements

No dynamics of interaction between supply-side agents and demand-side agents have been explicitly illustrated thus far. This is because supply-side agents (utilities and OEMs) have initially served simplified roles to facilitate fleet agent development and verification. However, the actions of utilities and OEMs in the real world can place significant and complicated constraints on fleets' electrification potential. Long lead times for electricity infrastructure upgrades, prohibitive pricing structures, monopsonization of entities along the EV supply chain, and other supply-side-driven dynamics have already constrained the adoption and use of eMHD vehicles. Understanding and integrating the actions driving these dynamics is a critical next step in characterizing the development of an eMHD market.

As mentioned in Chapter 4, the observation-perception framework applied for fleet decisions will not be trivial to adapt to modeling utility and OEM decisions. Supply-side agents not only pursue distinct sets of goals from fleets, but those decisions may not be discretizable in the same way as the vehicle adoption decision, or at all. Decision-making in a more fundamentally continuous space, for example, the decision by a utility of how to allocate resources across competing and overlapping demands, may call for a decision model that is distinct from the one applied for fleet agents. These and other questions – does a sufficient literature exist characterizing utility operator decisions for us to apply the process from Chapter 3? Do sufficient commonalities even exist between different utility companies to enable using a common conceptualization for all of them? – make this a challenging, but necessary, next step.

Leveraging techniques from the field of reinforcement learning

The observation-perception model employed by fleet agents can be characterized, in a sense, as an instance of reinforcement learning (RL). However, we did not explicitly employ methods from the field of RL. It may be possible to enhance or extend our model by doing so. Briefly, an RL simulation consists of agents, each with a dynamic state, the ability to perform actions, and the ability to observe the output of a reward function based on its actions and the surrounding environment. In this concept, the agents in our model apply a determinate and limited reward function, that is, the perceived utility of each of the vehicles the agent owns. It might be beneficial to reconceptualize this reward function more broadly. Owning and operating good vehicles is a good thing, but it is only a specific means to more general ends – increasing profits, for example. However, would reconceptualizing the reward function as a representation of profit require modeling the purchase behaviors of fleets’ clients, a tremendously greater undertaking? Referencing the RL field more explicitly might lead to a middle solution to tractably enhance current and future agent models.

Scope expansion

The experiments in Chapter 4 involved tens of fleets, but characterizing the real market may involve samples consisting of hundreds, thousands, or even more. Doing so will be challenging in terms of data procurement and potentially infeasible in terms of computing resources. As we have not prioritized computing efficiency so far, efficiency improvements may be sufficient to make significantly greater-scope runs feasible. However, bottom-up modeling is inherently costly. Instead of indefinitely increasing the scope of simulations in pursuit of one-to-one representation, it may be more sensible to represent larger markets via effective design of small sample populations.

Modeling more, and more varied, populations might even push the novelty of this approach to its limits. All of the experiments we have presented were designed from a validation standpoint: testing simple hypotheses using relatively simple simulations of limited populations, with results that can be easily understood and, more or less, comprehensively explained. As the scope increases, however, results will likely become increasingly difficult to explain, just like the dynamics

of real markets. The key difference is that we can observe all the states of all the agents at any time, unlike in real markets. Even so, will this approach eventually produce results that are not fully explainable, not because of a lack of observability, but because the effort to trace causes and effects becomes intractable – despite the model consisting of easily-understood building blocks? What utility, if any, do modeling results have if we cannot fully understand them?

Chapter 6

Conclusions

In this work, we apply bottom-up modeling techniques to simulate electric vehicle (EV) adoption in medium- and heavy-duty (MHD) fleets. Our approach begins with comprehensive technoeconomic analysis of fleets' procurement and operating costs, including the spatial and temporal constraints that arise with increasing electrification. In Chapter 2, we use this approach to demonstrate the potential for an emerging business model, charging-as-a-service, to facilitate the increasing integration of the transportation and electricity systems. Next, we use text mining and analysis tools to characterize a fleet vehicle adoption decision model, which recognizes that technical and economic attributes are only a subset of the attributes influencing adoption decisions. We analyze a corpus drawn from the scholarly literature and introduce its implications for our decision model in Chapter 3. Finally, in Chapter 4, we integrate our technoeconomic analysis approach and adoption decision model in an agent-based model (ABM) of the MHD transportation system. We use this model to simulate EV adoption dynamics at a fleet-level resolution and demonstrate some of the novel experiments made possible by our bottom-up approach. These include simulating a heterogeneous network of fleets, for which aggregate outcomes can vary as a function of sectoral and attitudinal diversity within the network.

Our model arrives as part of a multi-discipline and multi-domain push to build and apply transportation system-level models with greater flexibility and scope [19]. Bottom-up, or microeconomic, conceptualizations of behavior enable practical experimental capabilities that top-down, macroeconomic approaches do not. Top-down models can often make useful predictions, but they have limited capacity to investigate questions about cause and effect or to model counterfactual scenarios. Bottom-up models open up these capacities by internalizing more attributes of the modeled population, including its members' fundamentally heterogeneous requirements, preferences, and access to information.

These novel capabilities come at a cost. Bottom-up models are more burdensome to design, verify, validate, and apply than typical top-down models. Internalizing more attributes increases the amount and quality of data required for design and validation, or, alternatively, the number of unmeasured assumptions that must be estimated. Many MHD transportation and electricity system entities choose not to make the data they gather openly available, often for reasons that are unlikely to change (protecting customer security and/or market positioning), meaning that these challenges will likely persist. Furthermore, due both to data constraints and computational cost constraints (memory, processing, and storage), a one-to-one simulation of the transportation and electricity systems, rather than a representative sample of them, may never be feasible. Researchers applying models like ours will thus continue to bear the burden of designing useful experiments and drawing useful conclusions from them, tolerating some hard-to-verify assumptions and representing only a subset of the real-world systems.

Nonetheless, addressing these challenges, rather than avoiding them, enables new and important means of inquiry. Inexact answers to critical questions, urgently relevant to policymaking, investment allocation, and planning by all members of the system, are better than none at all. So long as researchers take care to communicate their findings with precision, characterizing the uncertainties inherent to their models, these kinds of research outputs can and should fill the gaps left by traditional models. Further, untraditional means of measuring and gathering new kinds of data, enabled by machine learning and artificial intelligence, are rapidly evolving. (Our corpus analysis approach in Chapter 3 is one example.) These methods and capacities will help to mitigate the data sourcing challenge – requiring careful attention, of course, to defining and understanding the ethics and limitations of their use.

Chapter 4 offers early hints of the kinds of hypothesis testing enabled by our model. In place of questions such as “In what year will EVs make up the majority in sector X ?”, for which top-down models are better suited, we envision answering more qualitative, exploratory questions, pertaining to dynamics rather than to end results. For example, our model is well suited to simulate the potential roll-off effects of ending an incentive such as the IRA EV purchase subsidy. How might

demand for EVs differ between the final year the subsidy is active and the subsequent year? How do outcomes differ if the subsidy is retired gradually instead of all at once, and what kinds of intuitive or nonintuitive dynamics occur with different retirement schedules? Pointing to another direction for future work, one of the experiments in Chapter 4 only begins to assess the effects of attitudinal heterogeneity in a population of fleets. As evidenced by the research in the corpus we analyzed to define the decision model, the particular trade-offs weighed by decision makers for fleets are critical to their adoption dynamics, are heterogeneous between fleets within and across vocations and sectors, and are formidably difficult to characterize with precision, especially at large scale. The tools we have presented here are poised to adapt to new findings about the decision models of fleets as they become available, and to assess the consequences to individual actors as well as the broader network.

Finally, the broad uptake of electric transportation hinges on effective interfacing between supply (of vehicles, electrical equipment, and power) and demand, as we detail in Chapter 2. Some of the most pressing challenges arise from problems at this interface: long lead times and high costs for utilities to procure and install the electricity infrastructure needed for electric fleets; poor reliability of charging equipment and vehicles; and insufficient and unequal access to suitable MHD EVs. Although our ABM represents agents on the supply side (automakers and utilities) in addition to fleets on the demand side, our supply-side agents currently have little decision-making power. To represent the consequences of the challenges at the supply-demand interface, the supply-side agents in our model must be able to impose more constraints, by means of greater simulated autonomy and a more complex agent model. By expanding the scope and heterogeneity of our ABM this way, we expect to find and characterize more unexpected and unintuitive adoption and usage dynamics – along with an additional wealth of meaningful questions about the MHD system and the people whose decisions guide its evolution. With the foundation presented here, we will also have the means to begin answering them.

Bibliography

- [1] 21st Century Truck Partnership. *21st Century Truck Partnership Research Blueprint*. Tech. rep. Feb. 2019. URL: https://www.energy.gov/sites/prod/files/2019/02/f59/21CTPResearchBlueprint2019%5C_FINAL.pdf.
- [2] Drew Kodjak. *Policies to Reduce Fuel Consumption, Air Pollution, and Carbon Emissions from Vehicles in G20 Nations*. Tech. rep. May 2015. URL: https://theicct.org/sites/default/files/publications/ICCT%5C_G20-briefing-paper%5C_Jun2015%5C_updated.pdf.
- [3] Angel Alejandro Juan et al. “Electric Vehicles in Logistics and Transportation: A Survey on Emerging Environmental, Strategic, and Operational Challenges”. In: *Energies* 9.2 (2016), p. 86. DOI: 10.3390/en9020086.
- [4] Emir Cabukoglu et al. “Battery electric propulsion: An option for heavy-duty vehicles? Results from a Swiss case-study”. In: *Transportation Research Part C: Emerging Technologies* 88 (2018), pp. 107–123. ISSN: 0968-090X. DOI: 10.1016/j.trc.2018.01.013.
- [5] Miguel Jaller, Daniel Rivera, and John Harvey. *Spatio-Temporal Analysis of Freight Patterns in Southern California*. Tech. rep. Nov. 2020. URL: <https://escholarship.org/uc/item/1259f9s1>.
- [6] Bernd Heid et al. *What’s sparking electric-vehicle adoption in the truck industry?* Tech. rep. McKinsey & Company, Sept. 2017. URL: <https://www.mckinsey.com/industries/automotive-and-assembly/our-insights/whats-sparking-electric-vehicle-adoption-in-the-truck-industry>.
- [7] Nawa Raj Baral et al. “Biomass Feedstock Transport Using Fuel Cell and Battery Electric Trucks Improves Lifecycle Metrics of Biofuel Sustainability and Economy”. In: *Journal of Cleaner Production* 279 (2020), p. 123593. ISSN: 0959-6526. DOI: 10.1016/j.jclepro.20.123593.

- [8] Amol Phadke et al. *Why Regional and Long-Haul Trucks are Primed for Electrification Now*. Tech. rep. Lawrence Berkeley National Laboratory, Mar. 1, 2021. DOI: 10.2172/1834571. URL: <https://www.osti.gov/biblio/1834571>.
- [9] *Infrastructure Investment and Jobs Act*. Jan. 2021. URL: <https://www.congress.gov/117/bills/hr3684/BILLS-117hr3684enr.pdf>.
- [10] Mehmet Efe Biresselioglu, Melike Demirbag Kaplan, and Barbara Katharina Yilmaz. “Electric mobility in Europe: A comprehensive review of motivators and barriers in decision making processes”. In: *Transportation Research Part A: Policy and Practice* 109 (Mar. 1, 2018), pp. 1–13. ISSN: 0965-8564. DOI: 10.1016/j.tra.2018.01.017. URL: <https://www.sciencedirect.com/science/article/pii/S0965856417311771>.
- [11] Kate Forrest et al. “Estimating the technical feasibility of fuel cell and battery electric vehicles for the medium and heavy duty sectors in California”. In: *Applied Energy* 276 (Oct. 15, 2020), p. 115439. ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2020.115439. URL: <https://www.sciencedirect.com/science/article/pii/S030626192030951X>.
- [12] Cornelius Moll et al. “Are Battery-Electric Trucks for 24-Hour Delivery the Future of City Logistics?—A German Case Study”. In: *World Electric Vehicle Journal* 11.1 (2020), p. 16. DOI: 10.3390/wevj11010016.
- [13] Baha M. Al-Alawi and Thomas H. Bradley. “Total cost of ownership, payback, and consumer preference modeling of plug-in hybrid electric vehicles”. In: *Applied Energy* 103 (2013), pp. 488–506. ISSN: 0306-2619. DOI: 10.1016/j.apenergy.2012.10.009.
- [14] Margaret J. Eppstein et al. “An agent-based model to study market penetration of plug-in hybrid electric vehicles”. In: *Energy Policy* 39.6 (2011), pp. 3789–3802. ISSN: 0301-4215. DOI: 10.1016/j.enpol.2011.04.007.
- [15] Margaret J. Eppstein et al. “Using National Survey Respondents as Consumers in an Agent-Based Model of Plug-In Hybrid Vehicle Adoption”. In: *IEEE Access* 3 (2015), pp. 457–468. ISSN: 2169-3536. DOI: 10.1109/access.2015.2427252.

- [16] Argonne National Laboratory. *AFLEET*. URL: <https://afleet.es.anl.gov/home/>.
- [17] National Renewable Energy Laboratory. *EVI-FAST*. URL: <https://www.nrel.gov/transportation/evi-fast.html>.
- [18] Ning Wang, Qingxin Yang, and Hengjun Zhang. “Modeling and analysis of dynamic wireless charging for electric vehicles under different working scenarios”. In: *The International Journal of Electrical Engineering & Education* (2020), p. 002072092092854. ISSN: 0020-7209. DOI: 10.1177/0020720920928541.
- [19] Matteo Muratori et al. “Future integrated mobility-energy systems: A modeling perspective”. In: *Renewable and Sustainable Energy Reviews* 119 (2020), p. 109541. ISSN: 1364-0321. DOI: 10.1016/j.rser.2019.109541.
- [20] David Trinko et al. “Transportation and electricity systems integration via electric vehicle charging-as-a-service: A review of techno-economic and societal benefits”. In: *Renewable and Sustainable Energy Reviews* 175 (2023), p. 113180. ISSN: 1364-0321. DOI: 10.1016/j.rser.2023.113180.
- [21] United Nations Framework Convention on Climate Change. *Nationally determined contributions under the Paris Agreement*. Tech. rep. Nov. 2021.
- [22] Intergovernmental Panel on Climate Change. *Climate Change 2022: Mitigation of Climate Change*. Tech. rep. Apr. 2022.
- [23] Brennan Borlaug et al. “Levelized Cost of Charging Electric Vehicles in the United States”. In: *Joule* 4.7 (2020), pp. 1470–1485. ISSN: 2542-4351. DOI: 10.1016/j.joule.2020.05.013.
- [24] Rajeev Ranjan Kumar and Alok Kumar. “Adoption of electric vehicle: a literature review and prospects for sustainability”. In: *Journal of Cleaner Production* 253 (Apr. 20, 2020), p. 119911. ISSN: 0959-6526. DOI: 10.1016/j.jclepro.2019.119911. URL: <https://www.sciencedirect.com/science/article/pii/S095965261934781X>.

- [25] Kathryn Canepa, Scott Hardman, and Gil Tal. “An early look at plug-in electric vehicle adoption in disadvantaged communities in California”. In: *Transport Policy* 78 (2019), pp. 19–30. ISSN: 0967-070X. DOI: 10.1016/j.tranpol.2019.03.009.
- [26] Marcus Alexander. *Electric Vehicle Supply Equipment Installed Cost Analysis*. Tech. rep. Dec. 2013.
- [27] U.S. Department of Energy. “Costs Associated With Non-Residential Electric Vehicle Supply Equipment”. In: (Nov. 2015).
- [28] Sebastian Rivera et al. “Electric Vehicle Charging Infrastructure”. In: *IEEE Industrial Electronics Magazine* 15.2 (2021), pp. 37–51. ISSN: 1932-4529. DOI: 10.1109/mie.2020.3039039.
- [29] Jean-Baptiste Gallo, Ted Bloch-Rubin, and Jasna Tomic. *Peak Demand Charges and Electric Transit Buses*. Tech. rep. CALSTART, Oct. 2014.
- [30] Youngeun Bae et al. “Factors influencing alternative fuel adoption decisions in heavy-duty vehicle fleets”. In: *Transportation Research Part D: Transport and Environment* 102 (Jan. 1, 2022), p. 103150. ISSN: 1361-9209. DOI: 10.1016/j.trd.2021.103150. URL: <https://www.sciencedirect.com/science/article/pii/S1361920921004454>.
- [31] E Akhavan-Rezai et al. “Uncoordinated Charging Impacts of Electric Vehicles on Electric Distribution Grids: Normal and Fast Charging Comparison”. In: *2012 IEEE Power and Energy Society General Meeting* 1 (2012), pp. 1–7. DOI: 10.1109/pesgm.2012.6345583.
- [32] Shuocheng Guo, Xinwu Qian, and Jun Liu. “Charging-as-a-Service: On-demand battery delivery for light-duty electric vehicles for mobility service”. In: *arXiv* (2020). eprint: 2011.10665.
- [33] Shota Miyahara et al. “Charging Prioritization of Electric Vehicles Under Peak Demand in Commercial Facility: Destination Charging as a Service”. In: *2020 International Conference on Smart Grids and Energy Systems (SGES)* 00 (2020), pp. 226–231. DOI: 10.1109/sges51519.2020.00047.

- [34] Dominik Danner and Hermann de Meer. “Quality of service and fairness for electric vehicle charging as a service”. In: *Energy Informatics* 4.Suppl 3 (2021), p. 16. DOI: 10.1186/s42162-021-00175-3.
- [35] Xu Zhang et al. “Mobile Charging as a Service: A Reservation-Based Approach”. In: *IEEE Transactions on Automation Science and Engineering* 17.4 (2020), pp. 1976–1988. ISSN: 1545-5955. DOI: 10.1109/tase.2020.2983819.
- [36] Jiahua Qiu and Lili Du. “A Charging-as-a-Service Platform for Charging Electric Vehicles On the Move: New Vehicle Routing Model and Solution”. In: *arXiv* (2021). eprint: 2104.00730.
- [37] Mahendra Singh et al. “Emergence of digital and X-as-a-service platforms in German energy sectors”. In: *1st International Association for Energy Economics Online Conference*. June 2021.
- [38] Henry E. Schaffer. “X as a Service, Cloud Computing, and the Need for Good Judgment”. In: *IT Professional* 2.9 (2009), pp. 4–5. DOI: 10.1109/mitp.2009.112.
- [39] Warwick Goodall et al. “The rise of mobility as a service”. In: *Deloitte Review* 20 (Jan. 2017), pp. 114–129.
- [40] *What is CaaS?* Jan. 2020. URL: <https://www.redhat.com/en/topics/cloud-computing/what-is-caas>.
- [41] *Using content infrastructure as CaaS: Content as a service*. URL: <https://www.contentful.com/r/knowledgebase/content-as-a-service/>.
- [42] *Communications as a Service (CaaS)*. URL: <https://www.gartner.com/en/information-technology/glossary/communications-as-a-service-caas>.
- [43] Bethany Villarreal. *Charging as a Service: A New Amenity for Electrified Living*. July 2020. URL: <https://semaconnect.com/blog/charging-as-a-service-a-new-amenity-for-electrified-living/>.

- [44] *Shell Recharge Solutions*. URL: <https://shellrecharge.com/en-us/solutions/product/charging-as-a-service>.
- [45] *EV Charging Business As A Service: What you need to know*. July 2021. URL: <https://www.virta.global/blog/the-demise-of-the-gasmobile-ev-charging-as-a-service>.
- [46] *ChargePoint as a Service*. URL: <https://www.chargepoint.com/products/cpaas>.
- [47] *Lightning eMotors: Charging As A Service*. URL: <https://lightningemotors.com/charging-as-a-service/>.
- [48] “Xcel Energy Home Charging Advisor”. In: (). URL: <https://ev.xcelenergy.com/home-charging-advisor>.
- [49] Robert Walton. *Duke may offer some EV customers 'all you can charge' for just \$19.99/month (restrictions apply)*. Feb. 2022. URL: <https://www.utilitydive.com/news/duke-may-offer-some-ev-customers-all-you-can-charge-for-just-1999month/619210/>.
- [50] Vanessa Peng. *WattLogic: What is EV charging as a service?* Oct. 2021. URL: <https://wattlogic.com/blog/ev-charging-as-a-service/>.
- [51] *Amplify Power Charging-as-a-Service*. URL: <https://amplypower.com/charging-as-a-service/>.
- [52] *EV Connect Powering What's Next*. URL: https://www.evconnect.com/?utm%5C_source=release%5C&utm%5C_medium=email%5C&utm%5C_campaign=evcaas.
- [53] *EO Charging: Charging-as-a-service*. URL: <https://static1.squarespace.com/static/5b96995789c172d9106f1574/t/5ebc2147ee7a6919fe41715a/1589387593050/Charging-as-a-Service+%5C%28EO+Charging%5C%29+2020.pdf>.
- [54] *AECOM: Charging as a Service*. URL: <https://publications.aecom.com/transportation-electrification/article/charging-as-a-service>.

- [55] Steve Tengler. 'Charging As A Service' For Electric Vehicles Growing As A Market Offering. Feb. 2022. URL: <https://www.forbes.com/sites/stevetengler/2022/02/16/charging-as-a-service-for-electric-vehicles-growing-as-a-market-offering/?sh=6fe886ff5939>.
- [56] Sarah LaMonaca and Lisa Ryan. "The state of play in electric vehicle charging services – A review of infrastructure provision, players, and policies". In: *Renewable and Sustainable Energy Reviews* 154 (2022), p. 111733. ISSN: 1364-0321. DOI: 10.1016/j.rser.2021.111733.
- [57] Matteo Muratori. "Impact of uncoordinated plug-in electric vehicle charging on residential power demand". In: *Nature Energy* 3.3 (2018), pp. 193–201. DOI: 10.1038/s41560-017-0074-z.
- [58] Kevin Nesbitt and Daniel Sperling. "Myths regarding alternative fuel vehicle demand by light-duty vehicle fleets". In: *Transportation Research Part D: Transport and Environment* 3.4 (1998), pp. 259–269. ISSN: 1361-9209. DOI: 10.1016/s1361-9209(98)00006-6.
- [59] Layla Mohammed, Eva Niesten, and Dimitri Gagliardi. "Adoption of alternative fuel vehicle fleets – A theoretical framework of barriers and enablers". In: *Transportation Research Part D: Transport and Environment* 88 (2020), p. 102558. ISSN: 1361-9209. DOI: 10.1016/j.trd.2020.102558.
- [60] Kevin Nesbitt and Daniel Sperling. "Fleet purchase behavior: decision processes and implications for new vehicle technologies and fuels". In: *Transportation Research Part C: Emerging Technologies* 9.5 (2001), pp. 297–318. ISSN: 0968-090X. DOI: 10.1016/s0968-090x(00)00035-8.
- [61] Heather Klemick et al. "Heavy-duty trucking and the energy efficiency paradox: Evidence from focus groups and interviews". In: *Transportation Research Part A: Policy and Practice* 77 (July 1, 2015), pp. 154–166. ISSN: 0965-8564. DOI: 10.1016/j.tra.2015.04.004. URL: <https://www.sciencedirect.com/science/article/pii/S0965856415000816>.

- [62] Roger Bennett. “Fleet vehicle buyers’ intentions to purchase electric vehicles: antecedents and possible consequences”. In: *International Journal of Electric and Hybrid Vehicles* 7.4 (2015), p. 362. ISSN: 1751-4088. DOI: 10.1504/ijehv.2015.074677.
- [63] Sigal Kaplan et al. “Intentions to introduce electric vehicles in the commercial sector: A model based on the theory of planned behaviour”. In: *Research in Transportation Economics* 55 (2016), pp. 12–19. ISSN: 0739-8859. DOI: 10.1016/j.retrec.2016.04.006.
- [64] Marshall Miller, Qian Wang, and Lew Fulton. *Truck Choice Modeling: Understanding California’s Transition to Zero-Emission Vehicle Trucks Taking into Account Truck Technologies, Costs, and Fleet Decision Behavior*. Tech. rep. Nov. 1, 2017. URL: <https://escholarship.org/uc/item/1xt3k10x>.
- [65] Xavier Demeulenaere. “The use of automotive fleets to support the diffusion of Alternative Fuel Vehicles: A Rapid Evidence Assessment of barriers and decision mechanisms”. In: *Research in Transportation Economics* 76 (2019), p. 100738. ISSN: 0739-8859. DOI: 10.1016/j.retrec.2019.100738.
- [66] Benedikt Anderhofstadt and Stefan Spinler. “Factors affecting the purchasing decision and operation of alternative fuel-powered heavy-duty trucks in Germany – A Delphi study”. In: *Transportation Research Part D: Transport and Environment* 73 (Aug. 1, 2019), pp. 87–107. ISSN: 1361-9209. DOI: 10.1016/j.trd.2019.06.003. URL: <https://www.sciencedirect.com/science/article/pii/S1361920918310599>.
- [67] Miguel Jaller, Leticia Pineda, and Yasar Gueldas. *Fostering the Use of Zero and Near Zero Emission Vehicles in Freight Operations*. Tech. rep. July 2020. URL: <https://escholarship.org/uc/item/64k579cv>.
- [68] Sonya Ziaja and Mohit Chhabra. “Climate Adaption for Energy Utilities: Lessons Learned From California’s Pioneering Regulatory Actions”. In: *Frontiers in Climate* 3 (2021), p. 718472. DOI: 10.3389/fclim.2021.718472.
- [69] Mark Duvall. personal communication. Feb. 17, 2022.

- [70] Electrify America. *National ZEV Investment Plan: Cycle 3*. Tech. rep. June 2021.
- [71] *My Model 3 Charged Off a Piece of History: AVCON*. Oct. 2020. URL: <https://callasev.com/my-model-3-charged-off-a-piece-of-history-avcon/>.
- [72] Chevrolet/Geo Communications. *INDUCTIVE TECHNOLOGY TO CHARGE GM'S NEW ELECTRIC VEHICLES*. Sept. 1996. URL: https://en.wikipedia.org/wiki/Magne%5C_Charge.
- [73] Abby Brown et al. "Evolution of Plug-In Electric Vehicle Charging Infrastructure in the United States". In: June 2020.
- [74] Thomas Caldwell. *EV Charger Reliability is Critical*. Feb. 2022. URL: <https://www.forbes.com/sites/forbestechcouncil/2022/02/14/ev-charger-reliability-is-critical/?sh=69d515882d40>.
- [75] Jennifer Sensiba. *Charger Reliability is the Next Challenge for the EV Industry*. Apr. 2021. URL: <https://cleantechnica.com/2021/04/21/charger-reliability-is-the-next-challenge-for-the-ev-industry/>.
- [76] David Rempel et al. "Reliability of Open Public Electric Vehicle Direct Current Fast Chargers". In: *Arxiv preprint* (Mar. 2022).
- [77] California Air Resources Board. *Low Carbon Fuel Standard (LCFS) Guidance 20-04 Requesting EER-Adjusted Carbon Intensity Using a Tier 2 Pathway Application*. Tech. rep. Apr. 2020.
- [78] EPRI. *Consumer Guide to Electric Vehicle Charging*. Oct. 2019.
- [79] National Renewable Energy Laboratory. "Plug-In Electric Vehicle Handbook for Consumers (Brochure)". In: (). DOI: 10.2172/1026556.
- [80] California Plug-in Electric Vehicle Collaborative. *Plug-in Electric Vehicle Charging Guide for Property Owners, Managers and Homeowner Associations of Multi-unit Dwellings*. Nov. 2013.

- [81] National Renewable Energy Laboratory. “Plug-In Electric Vehicle Handbook for Public Charging Station Hosts (Brochure)”. In: (2012). DOI: 10.2172/1042497.
- [82] National Renewable Energy Laboratory. “Plug-In Electric Vehicle Handbook for Workplace Charging Hosts (Brochure)”. In: (2012).
- [83] National Renewable Energy Laboratory. “Plug-In Electric Vehicle Handbook for Fleet Managers (Brochure)”. In: (2012). DOI: 10.2172/1039084.
- [84] Dahlia Garas, Gustavo O. Collantes, and Michael A. Nicholas. *City of Vancouver EV Infrastructure Strategy Report*. Tech. rep. Dec. 2016.
- [85] Andrew Burnham et al. “Regional Emissions Analysis of Light-Duty Battery Electric Vehicles”. In: *Atmosphere* 12.11 (2021), p. 1482. DOI: 10.3390/atmos12111482.
- [86] Chih-Wei Hsu and Kevin Fingerma. “Public electric vehicle charger access disparities across race and income in California”. In: *Transport Policy* 100 (2021), pp. 59–67. ISSN: 0967-070X. DOI: 10.1016/j.tranpol.2020.10.003.
- [87] Budget, U.S. Office of Management and. *Circular A-94: Guidelines and Discount Rates for Benefit-Cost Analysis of Federal Programs*. Nov. 2015. URL: <https://obamawhitehouse.archives.gov/sites/default/files/omb/assets/a94/a094.pdf>.
- [88] Sofie E Miller. “One Discount Rate Fits All? The Regressive Effects of DOE’s Energy Efficiency Rule”. In: *Policy Perspectives* 22 (2015), pp. 40–54. ISSN: 1085-7087. DOI: 10.4079/pp.v22i0.15110.
- [89] Dennis Silverman. *Southern California Household Energy Savings*. Tech. rep. Oct. 2007. URL: <https://www.physics.uci.edu/%5Ctextbackslashtextasciitildesilverma/actions/HouseholdEnergy.html>.
- [90] US Environmental Protection Agency. *Greenhouse Gas Emissions from a Typical Passenger Vehicle*. URL: <https://www.epa.gov/greenvehicles/greenhouse-gas-emissions-typical-passenger-vehicle>.

- [91] *Tesla Model 3*. URL: <https://www.tesla.com/model3>.
- [92] David Peterson. *Addressing Challenges to Electric Vehicle Charging in Multifamily Residential Buildings*. Tech. rep. UCLA Luskin Center for Innovation, June 2011.
- [93] Sara Baldwin, Amanda Myers, and Michael O’Boyle. *Increasing Electric Vehicle Charging Access at Multi-Unit Dwellings: Workshop Summary Report*. Tech. rep. Energy Innovation Policy & Technology LLC, Sept. 2020. URL: https://energyinnovation.org/wp-content/uploads/2020/09/Increasing-Electric-Vehicle-Charging-at-Multi-Unit-Dwellings%5C_FINAL3.pdf.
- [94] Jae Hyun Lee et al. “Exploring electric vehicle charging patterns: Mixed usage of charging infrastructure”. In: *Transportation Research Part D: Transport and Environment* 79 (2020), p. 102249. ISSN: 1361-9209. DOI: 10.1016/j.trd.2020.102249.
- [95] NESCAUM. *Right To Charge Laws*. Oct. 2019.
- [96] David Trinko et al. “Combining Ad Hoc Text Mining and Descriptive Analytics to Investigate Public EV Charging Prices in the United States”. In: *Energies* 14.17 (Jan. 2021), p. 5240. ISSN: 1996-1073. DOI: 10.3390/en14175240. URL: <https://www.mdpi.com/1996-1073/14/17/5240>.
- [97] Gordon Bauer et al. *Charging Up America: Assessing the Growing Need for U.S. Charging Infrastructure Through 2030*. Tech. rep. July 2021.
- [98] Pacific Gas & Electric. *LCFS Calculator*. URL: https://fleets.pge.com/lcfs/?kwhDisplacedForLcfs=230000%5C&gridIntensityInGramsCO2ePerMJ=0%5C&vehicleSets=%5C%5B%5C%5D%5C&bevRateName=Business+High+Use+EV%5C&bevBlocks=5%5C&dieselPrice=3.1%5C&gasolinePrice=3.3%5C&lpgPrice=3.5%5C&yearsOfOperation=10%5C&btmConstructionCost=10000%5C&btmEvseCost=10000%5C&customVehicles=%5C%5B%5C%5D%5C&customChargers=%5C%5B%5C%5D%5C&isTaxExempt=Yes%5C&isDAC=No%5C&isPublic=No%5C&replacesIceVehicle=Yes%5C&organizationType=private%5C_school%5C_districts%5C&zipCode=94108.

- [99] Vignesh Ramasamy et al. *U.S. Solar Photovoltaic System and Energy Storage Cost Benchmarks: Q1 2021*. Tech. rep. 1. Delft, Nov. 2021, pp. 319–340. eprint: 2011.10665. URL: <https://doi.org/10.1016/bs.hescom.2018.02.003>.
- [100] *Lightning Electric Zero Emission Transit Cargo Van*. URL: <https://lightningemotors.com/lightningelectric-ford-transit-cargo/>.
- [101] BOOSTER. *The Drive Towards Decarbonization*. Tech. rep. Aug. 2021.
- [102] *The power of subscription models*. June 2020. URL: <https://applied.economist.com/articles/a-notion-worth-subscribing-to>.
- [103] Dan Lashof et al. “After Inflation Reduction Act: 5 Ways to Hit the US Emissions Goal”. In: (Aug. 12, 2022). URL: <https://www.wri.org/insights/inflation-reduction-act-emissions-gap>.
- [104] CDOT. *Planning Initiatives*. Colorado Department of Transportation. 2018. URL: <https://www.codot.gov/programs/innovativemobility/electrification/planning-initiatives-and-documents>.
- [105] Nadia Adnan et al. “A comprehensive review on theoretical framework-based electric vehicle consumer adoption research”. In: *International Journal of Energy Research* 41.3 (2017), pp. 317–335. ISSN: 1099-114X. DOI: 10.1002/er.3640. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/er.3640>.
- [106] Fatemeh Nazari, Ehsan Rahimi, and Abolfazl (Kouros) Mohammadian. “Simultaneous estimation of battery electric vehicle adoption with endogenous willingness to pay”. In: *eTransportation* 1 (Aug. 1, 2019), p. 100008. ISSN: 2590-1168. DOI: 10.1016/j.etrans.2019.100008. URL: <https://www.sciencedirect.com/science/article/pii/S2590116819300086>.
- [107] Brennan Borlaug et al. “Heavy-duty truck electrification and the impacts of depot charging on electricity distribution systems”. In: *Nature Energy* 6.6 (June 2021), pp. 673–682. ISSN: 2058-7546. DOI: 10.1038/s41560-021-00855-0. URL: <https://www.nature.com/articles/s41560-021-00855-0>.

- [108] Roger Bennett, Rita Kottasz, and Stephen Shaw. “Factors potentially affecting the successful promotion of electric vehicles”. In: *Journal of Social Marketing* 6.1 (Jan. 1, 2016), pp. 62–82. ISSN: 2042-6763. DOI: 10.1108/JSOCM-08-2015-0059. URL: <https://doi.org/10.1108/JSOCM-08-2015-0059>.
- [109] Munish Kumar and Janine Ng. “Using text mining and topic modelling to understand success and growth factors in Global Renewable Energy projects”. In: *Renewable Energy Focus* 42 (Sept. 2022), pp. 211–220. ISSN: 17550084. DOI: 10.1016/j.ref.2022.06.010. URL: <https://linkinghub.elsevier.com/retrieve/pii/S1755008422000527>.
- [110] Ramit Debnath et al. “Political, economic, social, technological, legal and environmental dimensions of electric vehicle adoption in the United States: A social-media interaction analysis”. In: *Renewable and Sustainable Energy Reviews* 152 (Dec. 1, 2021), p. 111707. ISSN: 1364-0321. DOI: 10.1016/j.rser.2021.111707. URL: <https://www.sciencedirect.com/science/article/pii/S1364032121009813>.
- [111] Tao Ruan and Qin Lv. “Public perception of electric vehicles on reddit over the past decade”. In: *Communications in Transportation Research* 2 (Dec. 1, 2022), p. 100070. ISSN: 2772-4247. DOI: 10.1016/j.commtr.2022.100070. URL: <https://www.sciencedirect.com/science/article/pii/S2772424722000208>.
- [112] Shashank Kapadia. *Evaluate Topic Models: Latent Dirichlet Allocation (LDA)*. Medium. Dec. 29, 2020. URL: <https://towardsdatascience.com/evaluate-topic-model-in-python-late-nt-dirichlet-allocation-lda-7d57484bb5d0>.
- [113] Murugan Anandarajan, Chelsey Hill, and Thomas Nolan. *Practical Text Analytics: Maximizing the Value of Text Data*. Vol. 2. Advances in Analytics and Data Science. Cham: Springer International Publishing, 2019. ISBN: 978-3-319-95662-6 978-3-319-95663-3. DOI: 10.1007/978-3-319-95663-3. URL: <http://link.springer.com/10.1007/978-3-319-95663-3>.

- [114] Angel R. Martinez. “4 - Data Mining of Text Files”. In: *Handbook of Statistics*. Ed. by C. R. Rao, E. J. Wegman, and J. L. Solka. Vol. 24. Data Mining and Data Visualization. Elsevier, Jan. 1, 2005, pp. 109–131. DOI: 10.1016/S0169-7161(04)24004-4. URL: <https://www.sciencedirect.com/science/article/pii/S0169716104240044>.
- [115] David M Blei, Michael Jordan, and Andrew Ng. “Latent Dirichlet Allocation”. In: (2003), p. 30.
- [116] Thomas Hofmann. *Probabilistic Latent Semantic Analysis*. Jan. 23, 2013. DOI: 10.48550/arXiv.1301.6705. arXiv: 1301.6705[cs,stat]. URL: <http://arxiv.org/abs/1301.6705>.
- [117] Rob Churchill and Lisa Singh. “The Evolution of Topic Modeling”. In: *ACM Computing Surveys* (Jan. 12, 2022), p. 3507900. ISSN: 0360-0300, 1557-7341. DOI: 10.1145/3507900. URL: <https://dl.acm.org/doi/10.1145/3507900>.
- [118] Pranav Suri and Nihar Ranjan Roy. “Comparison between LDA & NMF for event-detection from large text stream data”. In: *2017 3rd International Conference on Computational Intelligence & Communication Technology (CICT)*. 2017 3rd International Conference on Computational Intelligence & Communication Technology (CICT). Feb. 2017, pp. 1–5. DOI: 10.1109/CICT.2017.7977281.
- [119] João Souza et al. “A systematic mapping on automatic classification of fake news in social media”. In: *Social Network Analysis and Mining* 10 (Dec. 1, 2020). DOI: 10.1007/s13278-020-00659-2.
- [120] Jiguang Liang et al. “Sentiment Classification Based on AS-LDA Model”. In: *Procedia Computer Science*. 2nd International Conference on Information Technology and Quantitative Management, ITQM 2014 31 (Jan. 1, 2014), pp. 511–516. ISSN: 1877-0509. DOI: 10.1016/j.procs.2014.05.296. URL: <https://www.sciencedirect.com/science/article/pii/S1877050914004736>.
- [121] M. Selvi et al. “Classification of Medical Dataset Along with Topic Modeling Using LDA”. In: *Nanoelectronics, Circuits and Communication Systems*. Ed. by Vijay Nath and Jyotsna

- Kumar Mandal. *Lecture Notes in Electrical Engineering*. Singapore: Springer, 2019, pp. 1–11. ISBN: 9789811307768. DOI: 10.1007/978-981-13-0776-8_1.
- [122] Ankur Tomar. *Topic modeling using Latent Dirichlet Allocation(LDA) and Gibbs Sampling explained!* Analytics Vidhya. July 25, 2019. URL: <https://medium.com/analytics-vidhya/topic-modeling-using-lda-and-gibbs-sampling-explained-49d49b3d1045>.
- [123] Jason Chuang, Christopher D. Manning, and Jeffrey Heer. “Termite: visualization techniques for assessing textual topic models”. In: *Proceedings of the International Working Conference on Advanced Visual Interfaces - AVI '12*. the International Working Conference. Capri Island, Italy: ACM Press, 2012, p. 74. ISBN: 978-1-4503-1287-5. DOI: 10.1145/2254556.2254572. URL: <http://dl.acm.org/citation.cfm?doid=2254556.2254572>.
- [124] James M. Joyce. “Kullback-Leibler Divergence”. In: *International Encyclopedia of Statistical Science*. Ed. by Miodrag Lovric. Berlin, Heidelberg: Springer, 2011, pp. 720–722. ISBN: 978-3-642-04898-2. DOI: 10.1007/978-3-642-04898-2327. URL: https://doi.org/10.1007/978-3-642-04898-2_327.
- [125] Steven Bird, Ewan Klein, and Edward Loper. *Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit*. 2009. URL: <https://github.com/nltk/nltk>.
- [126] Radim Řehůřek and Petr Sojka. *Software Framework for Topic Modelling with Large Corpora*. Version 4.1.0. Valetta, MT, May 2010. URL: <http://is.muni.cz/publication/884893/en>.
- [127] Tomas Mikolov et al. *Distributed Representations of Words and Phrases and their Compositionality*. Oct. 16, 2013. DOI: 10.48550/arXiv.1310.4546. arXiv: 1310.4546[cs,stat]. URL: <http://arxiv.org/abs/1310.4546>.
- [128] Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütz. “Stemming and lemmatization”. In: *Introduction to Information Retrieval (online version)*. Cambridge University Press, 2008. URL: <https://nlp.stanford.edu/IR-book/html/htmledition/stemming-and-lemmatization-1.html>.

- [129] Tobias Brigl. “Extracting Reliable Topics using Ensemble Latent Dirichlet Allocation”. PhD thesis. Munich: Technische Hochschule Ingolstadt, 2019. URL: https://www.sezanceb.de/machine_learning/ensemble_LDA/EnsembleLDA.pdf.
- [130] Martin Ester, Hans-Peter Kriegel, and Xiaowei Xu. “A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise”. In: (1996), p. 6.
- [131] Vy Thuy Li. “Identification of Potential Charging Infrastructure Locations for the Electrification of Medium to Heavy-Duty Vehicles, In the State of California”. In: (2021), p. 54.
- [132] Anthony Anosike et al. “Exploring the challenges of electric vehicle adoption in final mile parcel delivery”. In: *International Journal of Logistics Research and Applications* 0.0 (Sept. 23, 2021), pp. 1–25. ISSN: 1367-5567. DOI: 10.1080/13675567.2021.1978409. URL: <https://doi.org/10.1080/13675567.2021.1978409>.
- [133] Markus Knupfer et al. “Cross impact analysis of Vehicle-to-Grid technologies in the context of 2030”. In: Oct. 3, 2016, pp. 1–5. DOI: 10.1109/ICPDS.2016.7756682.
- [134] Xin Sun, Minggao Ouyang, and Han Hao. “Surging lithium price will not impede the electric vehicle boom”. In: *Joule* 6.8 (Aug. 17, 2022), pp. 1738–1742. ISSN: 2542-4785, 2542-4351. DOI: 10.1016/j.joule.2022.06.028. URL: [https://www.cell.com/joule/abstract/S2542-4351\(22\)00302-6](https://www.cell.com/joule/abstract/S2542-4351(22)00302-6).
- [135] Maximilian Schücking et al. “Charging strategies for economic operations of electric vehicles in commercial applications”. In: *Transportation Research Part D: Transport and Environment* 51 (Mar. 1, 2017), pp. 173–189. ISSN: 1361-9209. DOI: 10.1016/j.trd.2016.11.032. URL: <https://www.sciencedirect.com/science/article/pii/S1361920915301358>.
- [136] William Sierzchula. “Factors influencing fleet manager adoption of electric vehicles”. In: *Transportation Research Part D: Transport and Environment* 31 (Aug. 1, 2014), pp. 126–134. ISSN: 1361-9209. DOI: 10.1016/j.trd.2014.05.022. URL: <https://www.sciencedirect.com/science/article/pii/S1361920914000571>.

- [137] Joachim Globisch, Elisabeth Dütschke, and Martin Wietschel. “Adoption of electric vehicles in commercial fleets: Why do car pool managers campaign for BEV procurement?” In: *Transportation Research Part D: Transport and Environment*. The contribution of electric vehicles to environmental challenges in transport. WCTRS conference in summer 64 (Oct. 1, 2018), pp. 122–133. ISSN: 1361-9209. DOI: 10.1016/j.trd.2017.10.010. URL: <https://www.sciencedirect.com/science/article/pii/S136192091630774X>.
- [138] Lin Gui et al. “Neural Topic Model with Reinforcement Learning”. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. EMNLP-IJCNLP 2019. Hong Kong, China: Association for Computational Linguistics, Nov. 2019, pp. 3478–3483. DOI: 10.18653/v1/D19-1350. URL: <https://aclanthology.org/D19-1350>.
- [139] Daniel Kahneman et al. *Judgment under uncertainty: Heuristics and biases*. Cambridge university press, 1982.
- [140] U.S. Energy Information Administration. *Annual Energy Outlook 2022 Retrospective*. Tech. rep. Sept. 2022. URL: <https://www.eia.gov/outlooks/aeo/retrospective/pdf/retrospective.pdf>.
- [141] Thomas C. Schelling. “Dynamic models of segregation”. In: *The Journal of Mathematical Sociology* 1.2 (1971), pp. 143–186. ISSN: 0022-250X. DOI: 10.1080/0022250x.1971.9989794. eprint: <https://doi.org/10.1080/0022250X.1971.9989794>. URL: <https://doi.org/10.1080/0022250X.1971.9989794>.
- [142] Everett M Rogers. *Diffusion of innovations*. New York: Free Press, 1983.
- [143] Louis G Tornatzky, Mitchell Fleischer, and Alok K Chakrabarti. *Processes of technological innovation*. Lexington books, 1990.
- [144] Ruud T Frambach and Niels Schillewaert. “Organizational innovation adoption: a multi-level framework of determinants and opportunities for future research”. In: *Journal of Busi-*

- ness Research* 55.2 (2002), pp. 163–176. ISSN: 0148-2963. DOI: 10.1016/s0148-2963(00)00152-1.
- [145] Morris H DeGroot. “Reaching a consensus”. In: *Journal of the American Statistical association* 69.345 (1974), pp. 118–121.
- [146] Ali Jadbabaie et al. “Non-Bayesian social learning”. In: *Games and Economic Behavior* 76.1 (2012), pp. 210–225. ISSN: 0899-8256. DOI: 10.1016/j.geb.2012.06.001.
- [147] Mikael Sunnåker et al. “Approximate bayesian computation”. In: *PLoS computational biology* 9.1 (2013), e1002803.
- [148] Robert G. Sargent. “Verification and Validation of Simulation Models”. In: *Proceedings of the 2003 Winter Simulation Conference, 2003* 1 (2003), pp. 37–48. DOI: 10.1109/wsc.2003.1261406.
- [149] Andrew Burnham et al. “Comprehensive Total Cost of Ownership Quantification for Vehicles with Different Size Classes and Powertrains”. In: (Jan. 4). DOI: 10.2172/1780970.
- [150] Lucy Maybury, Pdraig Corcoran, and Liana Cipcigan. “Mathematical modelling of electric vehicle adoption: A systematic literature review”. In: *Transportation Research Part D: Transport and Environment* 107 (2022), p. 103278. ISSN: 1361-9209. DOI: 10.1016/j.trd.2022.103278.
- [151] U.S. Energy Information Administration. *Annual Energy Outlook 2022*. Tech. rep. Mar. 2022. URL: https://www.eia.gov/outlooks/aeo/pdf/AEO2022%5C_Narrative.pdf.
- [152] Catherine Ledna et al. *Decarbonizing Medium- & Heavy-Duty On-Road Vehicles: Zero-Emission Vehicles Cost Analysis*. Mar. 2022. URL: <https://www.nrel.gov/docs/fy22osti/82081.pdf>.
- [153] John Yarmuth. *Inflation Reduction Act of 2022*. Aug. 2022. URL: <https://www.congress.gov/bill/117th-congress/house-bill/5376/text>.
- [154] National Renewable Energy Laboratory. *Fleet DNA: commercial fleet vehicle operating data*. Tech. rep. URL: <https://www.nrel.gov/transportation/fleettest-fleet-dna.html>.

- [155] Frank Huber et al. “Dazing diversity: Investigating the determinants and consequences of decision paralysis”. In: *Psychology & Marketing* 29.6 (2012), pp. 467–478.
- [156] Southern California Edison. *Residential Rate Plans*. URL: <https://www.sce.com/residential/rates/Standard-Residential-Rate-Plan>.
- [157] Southern California Edison. *Residential Rate Plans: Standard Tiered*. URL: <https://www.sce.com/residential/rates/Standard-Residential-Rate-Plan>.
- [158] Southern California Edison. *Commercial Rate Plans: Electric Vehicle Rates*. URL: <https://www.sce.com/business/rates/electric-car-business-rates/business/rates/electric-car-business-rates>.
- [159] Southern California Edison. *TOU-EV-9 Rate Fact Sheet*. URL: [https://www.sce.com/sites/default/files/inline-files/TOU-EV-7_8_9%5C%20Rate%5C%20Fact%5C%20Sheet_WCAG%5C%20\(2\).pdf](https://www.sce.com/sites/default/files/inline-files/TOU-EV-7_8_9%5C%20Rate%5C%20Fact%5C%20Sheet_WCAG%5C%20(2).pdf).
- [160] Southern California Edison. *General Service Time-of-Use Rates*. URL: <https://www.sce.com/business/rates/time-of-use>.
- [161] Southern California Edison. *Rates & Pricing Choices*. URL: <https://www.sce.com/regulatory/tariff-books/rates-pricing-choices>.
- [162] A Dobos. “PVWatts Version 5 Manual”. In: (Sept. 2014). DOI: 10.2172/1158421.

Appendix A

Electric vehicle and fleet economics modeling

A.1 Details of business cases

A.1.1 Detached residence

Electricity pricing details

For the ChaaS case, we simulate a TOU rate schedule offered specifically for customers who charge EVs at home, enabling customers to access low electricity prices at low-demand times of day [156]. (The peak price window is 4pm-9pm; prices are off-peak at all other times.) During the summer, the low-demand price is \$0.22/kWh (compared to up to \$0.54/kWh at high-demand times), and during the winter it is \$0.21/kWh (compared to \$0.50/kWh at high-demand times). For the baseline case, we simulate a tiered rate structure [157], in which the energy price (\$0.28/kWh) does not vary by time of day but increases if monthly usage thresholds are exceeded. These thresholds vary by climate zone; we model the residence to be in zone 9, which is categorized by SCE as a “moderate” climate zone.

Table of assumptions

Economics assumptions

<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Annual discount rate (ChaaS provider)	7%	[87]
Annual discount rate (resident)	20%	
System lifespan	10	years

Electricity pricing assumptions

<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Electricity pricing structures	Tiered residential rates	[157]
	EV time-of-use rates	[156]
California climate zone	9 (moderate)	[157]

Duty cycle assumptions

<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Baseline household energy	4,000	kWh/year [89]
Driving distance	18,500	km/year [90]

Vehicle assumptions

<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Vehicle type	Tesla Model 3	[91]
Electricity consumption	161	Wh/km [91]
Battery capacity (usable)	80.5	kWh [91]
Electric range	500	km [91]

EVSE assumptions

<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Equipment cost	550	\$/plug [23]
Installation cost	1,286	\$/plug [23]

Table A.1: Assumptions used in the single-family-home business case.

A.1.2 Multi-unit dwelling

Electricity pricing details

For this and the commercial fleet case, we used the commercial EV-specific rate TOU-EV-9 from Southern California Edison [158, 159]. The summer rates apply during the months June-

September, and the winter rates apply otherwise. The details of this rate structure are given in Table A.2.

Season	Day of week	Time window	Price (\$/kWh)
Winter	Any	8am-4pm	0.080
		4pm-9pm	0.280
		9pm-midnight, midnight-8am	0.121
Summer	Weekends	9pm-midnight, midnight-4pm	0.117
		4pm-9pm	0.246
	Weekdays	9pm-midnight, midnight-4pm	0.117
		4pm-9pm	0.450

Table A.2: TOU-EV-9 rate schedule [159] used in cases 2 and 4.

Table of assumptions

Economics assumptions

<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Annual discount rate (ChaaS provider)	7%	[87]
Annual discount rate (customer)	20%	
System lifespan	10	years

Electricity pricing assumptions

<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Electricity pricing structure	Commercial EV rate	[158]
Public DCFC charging price	0.31	\$/kWh [96]

Duty cycle assumptions

<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Driving distance	18,500	km/year [90]
Weekday departure time (morning)	9:00 a.m.	
Departure time standard deviation	1	hour
Weekday return time (evening)	5:00 p.m.	

Vehicle assumptions

<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Vehicle type	long-range electric sedan	[91]
Electricity consumption	161	Wh/km [91]
Battery capacity (usable)	80.5	kWh [91]
Electric range	500	km [91]
Number of vehicles	15	vehicles

EVSE assumptions

<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Charging power	7.2	kW AC
Hardware costs	3,541	\$/plug
Installation costs	2,835	\$/plug [99]

Table A.3: Assumptions used in the multi-unit dwelling business case.

A.1.3 Business parking

Electricity pricing details

For this case, we used a general service commercial rate from Southern California Edison [160]. The summer rates apply during the months June-September, and the winter rates apply otherwise. The details of this rate structure are given in Table A.4. In addition to the time-of-use pricing, a peak demand charge of \$11.33/kW/month is assessed.

Season	Day of week	Time window	Price (\$/kWh)
Winter	Weekends	Any	0.087
	Weekdays	8am-9pm	0.130
		9pm-midnight, midnight-8am	0.087
Summer	Weekends	Any	0.104
	Weekdays	8am-noon	0.185
		noon-6pm	0.382
		6pm-11pm	0.185
		11pm-midnight, midnight-8am	0.104

Table A.4: General Service TOU Option R rate schedule [160] used in case 3.

Table of assumptions

Economics assumptions

<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Annual discount rate (ChaaS provider)	7%	[87]
Annual discount rate (host)	10%	
System lifespan	10	years

Electricity pricing assumptions

<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Electricity pricing structure	General Service TOU Option R	[161]
Mean energy price	0.127	\$/kWh
Peak demand price	11.33	\$/kW-mo

EVSE assumptions

<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Power options	7.2; 16.8; 50	kW
Hardware costs	3,541; 5,150; 31,600	\$/plug
Installation costs	2,835; 2,835; 17,756	\$/plug [28]
Annual maintenance cost	1%	% hardware cost [23]

Table A.5: Assumptions used in the business parking business case.

A.1.4 Commercial fleet

Vehicle and EVSE model details

We model a Class 3 electric van and use specifications claimed by the manufacturer [100]. Each fleet is assumed to have 20 of these vehicles, and to have installed one 7.2 kW EVSE unit per vehicle. EVSE units are assumed to be curtailable to any power between 0 and 7.2 kW. Charging efficiency (input and output) is assumed to be 90%.

EVSE hardware cost assumptions for every business case except for Case 1 are the mean per-plug hardware sales prices across a variety of manufacturers. These costs were obtained by the authors from representatives of Lightning eMotors.

Duty cycle details

Duty cycles are modeled from measured real-world driving of 15 electric Class 3 vans. We obtained driving data for 10 delivery vans, encompassing 217 vehicle-days, and 5 passenger-transport vans, encompassing 83 vehicle-days. Simulated fleets are randomly assigned as delivery or passenger-transport fleets (with equal probability of either), and the vehicles in each fleet are modeled to perform (and, for simplicity, repeat) a travel day randomly selected from the duty cycle data corresponding to the fleet type.

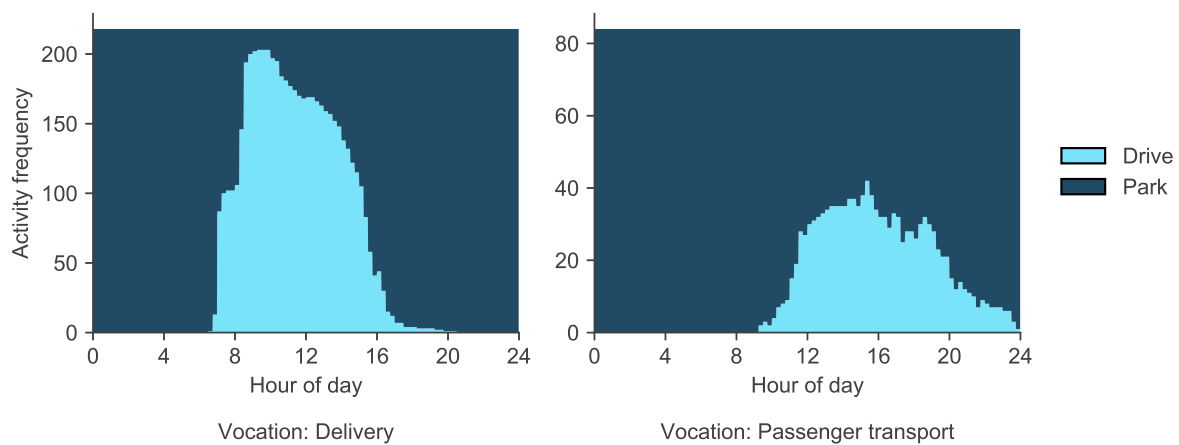


Figure A.1: Comparison of the driving patterns for the vocations modeled. The height of the light-colored area represents the number of vehicle days on which the vehicle was driving at the time of day shown on the horizontal axis. Whereas, for delivery vehicles, driving occurs mostly between 8 a.m. and 4 p.m., passenger transport vehicles operate throughout the afternoon and evening (between noon and midnight).

Control and optimization details

We use a linear programming (LP) formulation to co-optimize charging, storage battery dispatch, and PV curtailment schedules with the objective of globally minimizing electricity costs. Decision variables include charging power for each vehicle in each fleet during each time interval of a day and power input or output from the storage battery during each time interval of a day. Solutions are constrained in order to fulfill driving requirements, replenish all driving energy that is consumed, replenish all storage battery energy that is depleted, adhere to the power and state-of-charge constraints of the battery, and adhere to the power constraints of the EVSE. The

cost variable (annual operating cost) is computed according to the electricity rate schedule, the time value of LCFS credits throughout the day, and the availability of PV power generation. The 24-hour simulation is discretized to 15-minute intervals.

Low-carbon fuel standard details

A Low-Carbon Fuel Standard establishes a market for carbon displacement credits with the goal of reducing the overall carbon intensity of energy consumption in a state or region’s transportation sector [77]. To date, it has only been instituted in California. Subject to the standard, entities whose energy use has a greater carbon intensity than the government-set benchmark must purchase LCFS credits to offset their emissions, whereas entities whose energy use has a lower carbon intensity than the benchmark earn credits that they can sell. The price of LCFS credits is set by the market and can fluctuate. For transportation fuels including electricity, one LCFS “credit” is earned when one metric ton of carbon dioxide emissions is avoided as a result of using a low-carbon fuel. If the low-carbon fuel used is electricity from the grid, the effective carbon intensity varies with time of day, because the mix of generation resources fluctuates (e.g., more solar at midday). Electricity directly obtained from a DER, such as a wind turbine or a solar array, has a carbon intensity of zero. Thus, powering a fleet strictly with DER-derived solar power generates credits equal to the amount of carbon that a conventional fleet would emit to complete the same driving requirements.

PV modeling details

The power generated by the PV array is obtained from the PVWatts API [162]. Array sizing (1 MW), specifications (rooftop fixed panels), and location (latitude/longitude in the Los Angeles metropolitan area) are provided to the API, which returns a year-long timeseries of power generation. For the LP optimization and Figure 2.7, we use the daily timeseries from the median day of the year (computed in terms of total energy production, in kWh).

Table of assumptions

Economics assumptions

<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Annual discount rate (ChaaS provider)	7%	[87]
Annual discount rate (commercial fleet)	7%	[87]
System lifespan	10	years

DER system assumptions

<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Battery storage power capacity	600	kW [99]
Battery storage energy capacity	2.4	MWh [99]
PV nominal power capacity	1	MW [99]
PV array type	fixed rooftop panels	[99]
DER system cost	1.97	\$million [99]

Electricity pricing assumptions

<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Electricity pricing structure	Commercial EV rate	[158]
Energy economy ratio (EER, medium-duty)	3.4	[77]
Low-carbon fuel standard (LCFS) credit price	150	\$/metric ton [98]
Grid carbon intensity	variable by hour	[77]

Vehicle assumptions

<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Vehicle type	electric Class 3 van	[100]
Electricity consumption	327	Wh/km [100]
Battery capacity (usable)	73.6	kWh [100]
Electric range	225	km [100]
Number of vehicles per fleet	20	vehicles

EVSE assumptions

<i>Parameter</i>	<i>Value</i>	<i>Units, source</i>
Charging power	7.2	kW AC
Equipment cost	3,541	\$/plug
Installation cost	2,303	\$/plug

Table A.6: Assumptions used in the commercial fleet business case.

Appendix B

Decision and preference modeling

Word	Category
cost	long term cost
market	long term cost
diesel	long term cost
value	long term cost
infrastructure	long term cost
year	long term cost
total	long term cost
future	long term cost
investment	long term cost
economic	long term cost
expect	long term cost
interest	long term cost
utility	long term cost
economy	long term cost
saving	long term cost
plan	long term cost
purchase	short term cost
battery	short term cost
price	short term cost
cell	short term cost
present	short term cost
procurement	short term cost

pay	short term cost
buy	short term cost
short	short term cost
policy	policy incentives
public	policy incentives
state	policy incentives
government	policy incentives
incentive	policy incentives
agency	policy incentives
city	policy incentives
program	policy incentives
regulation	policy incentives
tax	policy incentives
national	policy incentives
country	policy incentives
fuel	suitability compatibility
time	suitability compatibility
charge	suitability compatibility
transportation	suitability compatibility
range	suitability compatibility
operation	suitability compatibility
need	suitability compatibility
business	suitability compatibility
demand	suitability compatibility
operational	suitability compatibility
change	suitability compatibility
performance	suitability compatibility

available	suitability compatibility
operate	suitability compatibility
customer	suitability compatibility
require	suitability compatibility
plug	suitability compatibility
conventional	suitability compatibility
power	suitability compatibility
compare	suitability compatibility
current	suitability compatibility
requirement	suitability compatibility
technical	suitability compatibility
station	suitability compatibility
mobility	suitability compatibility
recharge	suitability compatibility
driving range	suitability compatibility
research	familiarity knowledge
study	familiarity knowledge
analysis	familiarity knowledge
make	familiarity knowledge
information	familiarity knowledge
potential	familiarity knowledge
relate	familiarity knowledge
paper	familiarity knowledge
report	familiarity knowledge
theory	familiarity knowledge
understand	familiarity knowledge
knowledge	familiarity knowledge

literature	familiarity knowledge
estimate	familiarity knowledge
initial	familiarity knowledge
conduct	familiarity knowledge
network	familiarity knowledge
review	familiarity knowledge
manufacturer	familiarity knowledge
know	familiarity knowledge
fact	familiarity knowledge
discuss	familiarity knowledge
similar	familiarity knowledge
training	familiarity knowledge
insight	familiarity knowledge
learn	familiarity knowledge
introduce	familiarity knowledge
expert	familiarity knowledge
decision	norms attitudes
manager	norms attitudes
consider	norms attitudes
private	norms attitudes
barrier	norms attitudes
firm	norms attitudes
perceive	norms attitudes
see	norms attitudes
preference	norms attitudes
management	norms attitudes
behavior	norms attitudes

influence	norms attitudes
attitude	norms attitudes
influence	norms attitudes
intention	norms attitudes
indicate	norms attitudes
suggest	norms attitudes
positive	norms attitudes
social	norms attitudes
decision make	norms attitudes
consideration	norms attitudes
perspective	norms attitudes
standard	norms attitudes
willingness	norms attitudes
negative	norms attitudes
perception	norms attitudes
behaviour	norms attitudes
seem	norms attitudes
alternative	brand image
organization	brand image
innovation	brand image
design	brand image
improve	brand image
strategy	brand image
first	brand image
organizational	brand image
org	brand image
early	brand image

corporate	brand image
introduction	brand image
lead	brand image
early adopter	brand image
image	brand image
promote	brand image
energy	sustainability
evs	sustainability
environmental	sustainability
afvs	sustainability
hybrid	sustainability
efv	sustainability
afv	sustainability
reduce	sustainability
impact	sustainability
emission	sustainability
green	sustainability
lng	sustainability
reduction	sustainability
bev	sustainability
hydrogen	sustainability
electricity	sustainability
sustainable	sustainability
environment	sustainability
bevs	sustainability
alternative fuel	sustainability
clean	sustainability

cng	sustainability
gas	sustainability
renewable	sustainability
variable	reliability uncertainty
service	reliability uncertainty
risk	reliability uncertainty
work	reliability uncertainty
problem	reliability uncertainty
likely	reliability uncertainty
issue	reliability uncertainty
maintenance	reliability uncertainty
concern	reliability uncertainty
availability	reliability uncertainty
lack	reliability uncertainty
access	reliability uncertainty
uncertainty	reliability uncertainty
limited	reliability uncertainty
reliability	reliability uncertainty
safety	reliability uncertainty
user	driver acceptance
driver	driver acceptance
support	driver acceptance
employee	driver acceptance
acceptance	driver acceptance
experience	driver acceptance
individual	driver acceptance
share	driver acceptance

consumer	driver acceptance
operator	driver acceptance
usage	driver acceptance
personal	driver acceptance
ease	driver acceptance
implementation	driver acceptance

Table B.1: Index of keywords (left column) and the categories to which they were assigned.

Appendix C

Agent-based model development and application

C.1 Vehicle attributes and calculation methods

C.1.1 Range Suitability

From Chapter 3, range suitability is a dominant factor in fleets' decision making. Intuitively, a vehicle with insufficient range to complete its drive cycle is not suitable; a vehicle with plenty of range has maximum suitability; and in between, there is some finite amount of extra range buffer, beyond what is nominally needed for the drive cycle, up to which range suitability increases. For example, a 120-mile EV is more suitable for a 100-mile drive cycle than a 105-mile EV would be, but a 150-mile EV may not be much more valuable for the use case. We use a logistic function to model this intuition:

$$RS = \frac{1}{1 - e^{-k(r-r_0)}} \quad (\text{C.1})$$

where r is the excess range of a vehicle, above what is required by its drive cycle, and k and r_0 are parameters. k and r_0 are selected so that the value of RS is 0.95 when r equals 20 miles, and is 0.05 when r equals zero miles.

C.1.2 Total cost of ownership

Our concept of total cost of ownership (TCO) is equivalent to Net Present Cost (NPC) in Chapter 2 (Equation 2.4). The TCO of an individual vehicle is the NPC of owning and operating it. Cash flows that are specific to that vehicle (e.g. vehicle purchase price, fueling costs not due to peak demand, insurance, tax, depreciation) are straightforward to include in the TCO. However, some cash flows must be shared between vehicles, for example, the costs to purchase, install, and maintain charging equipment, and peak electricity demand costs. These shared cash flows are divided between vehicles according to their "relative contribution," as directly as possible.

For instance, the cash flows for an EVSE unit are split according to the energy drawn from each vehicle during the lifetime of the EVSE. This apportioning method avoids “double charging” EVSE costs, and also seeks to avoid overburdening the TCOs of the vehicles that happen to be purchased at the same time as the EVSE (if future vehicles can take advantage of the EVSE too). However, apportionment becomes more complicated as a fleet makes additional purchases over time. If new fast chargers are procured for a fleet that was previously making do with slow, inexpensive chargers – the new chargers being needed to support a procurement of vehicles that need faster charging – should existing vehicles bear those costs, if they make use of the new EVSE, even though their needs did not precipitate the purchase? If so, the TCO for those vehicles suddenly increases, and the TCO for the new vehicles is accordingly lessened – potentially unfairly. In summary, there are multiple defensible ways to apportion some of the cash flows contributing to TCO, and choosing between them can have a non-negligible effect on adoption outcomes.

C.1.3 Purchase price

“Purchase price” includes the price to procure a vehicle *and any charging equipment it requires*, minus any subsidies / incentives. EVSE purchase prices are apportioned between the vehicles that use them, as discussed and subject to the same ambiguities as in the TCO section above.

C.1.4 Maintenance cost

Maintenance cost observations by agents are expressed as the cost to maintain and repair a vehicle during the most recent year. These are estimated using methods drawn from Ref. [149]. Because maintenance costs tend to increase throughout the life of a vehicle, the observations for individual vehicles change over time. We assume agents do not incorporate future predictions of maintenance cost to their observations.

C.1.5 Fuel economy

Fuel economy observations for electric vehicles are expressed as miles per equivalent gallon of gasoline or diesel (whichever is the prevailing fuel for an equivalent ICEV):

$$FE_{\text{elec}}[\text{mpgde}, \text{mpgge}] = \frac{k[\text{kWh/gde}, \text{gge}]}{C[\text{kWh/mi}]}$$

where C is the nominal electric energy consumption for a vehicle, and k is 33.7 for gasoline and 37.9 for diesel.

C.1.6 Resale value

The resale value of a vehicle is a function of its purchase price, the rate of depreciation, and the year of purchase. Our depreciation model is adapted from Ref. [149].

C.1.7 Familiarity

An agent's familiarity with a vehicle is a fundamentally different concept than most other observations, because it increases with *the number of observations the agent has made pertaining to a vehicle*, and *the confidence with which those observations were made*. Thus, instead of tracking familiarity as a distribution of data points, familiarity is a determinate function of the sum of the weights of the agent's observations. We model this as a linear function with a ceiling:

$$F = \begin{cases} \sum w & \sum w < w_{max} \\ w_{max} & \sum w \geq w_{max} \end{cases}$$

where w is a vector of the observation weights, and w_{max} is a parameter.

C.1.8 Driver acceptance

We model driver acceptance as a random variable. The driver acceptance value for each vehicle is drawn from a distribution set by its OEM. There have been conflicting research findings about the driver acceptability of commercial EVs compared to ICEVs, which is likely a function of ease of use, reliability, personal preferences, and other factors. For simplicity, all OEMs were assumed to have the same distribution of driver acceptance:

$$DA \sim \text{Beta}(\alpha, \beta)$$

where we have selected $\alpha = 20$ and $\beta = 1$, for a beta distribution skewing somewhat sharply toward a mode of 1.

C.1.9 Fueling price stability

Fueling price stability is modeled as a dimensionless function of fuel type and year. Currently, nominal values are selected from normal distributions irrespective of year:

$$PS \sim \begin{cases} N(80, 10) & \text{electricity} \\ N(65, 10) & \text{gasoline and diesel} \end{cases}$$

However, agents have the ability to track and save fuel prices over time, enabling a direct model of perceived price volatility/stability. In future work, price stability should be a function of price histories.

C.1.10 Fuel security

Similar to fueling price stability, fuel security is a function of fuel type and year, but nominal values are selected from normal distributions:

$$FS \sim \begin{cases} N(80, 10) & \text{electricity} \\ N(77, 10) & \text{gasoline and diesel} \end{cases}$$

C.1.11 Refueling convenience

Refueling convenience is represented by the time that must be added to a vehicle's daily schedule to accommodate refueling/recharging, where smaller values are preferred. For ICEVs, the value is always assumed to be 0, because ICEV schedules are considered the baseline. For BEV schedules where all charging occurs at the depot during normal dwell hours, the value is also 0. Positive

values only occur when charging must be completed during a scheduled route, in which case the time spent charging is a function of the vehicle's energy needs, the power used for charging, and the time spent deviating from the route to get to a charger (assumed 5 minutes). This model can be made more detailed in the future by tracking charging stations geographically, modeling queues, and so on.

C.1.12 Fuel cost

Our fueling cost model from Chapter 2 is converted to \$/mile for agent observations.

C.1.13 Reliability

As with driver acceptance, we model reliability as a random variable set by the OEM. Again, for simplicity in early experiments, all OEMs are assumed to have the same distribution of reliability:

$$R \sim \text{Beta}(\alpha, \beta)$$

where we have selected $\alpha = 50$ and $\beta = 1$, for a beta distribution skewing quite sharply toward a mode of 1.

C.1.14 Brand image and environmental sustainability

Brand image and environmental sustainability are currently modeled as static, dimensionless functions of fuel type, where electricity has a significant advantage over gasoline and diesel. (We give BEVs 2x better brand image than ICEVs, and 10x better environmental sustainability.) In future work, these models could be made to incorporate the greenhouse gas intensity of grid electricity, tailpipe emissions regulations and benchmarks, and other social factors.

C.2 Adoption process

C.2.1 Awareness

For a fleet agent to consider adopting a particular vehicle, it must first be aware that the vehicle exists. Fleet agents can become aware of new vehicles via social interaction, social inference, or as a result of marketing and outreach efforts. By default, fleet agents are initially unaware of all eMHDV options.

Marketing and outreach

During each simulation step, there is a probability that each agent will become aware of each OEM it was not previously aware of, which is meant to emulate marketing and outreach activities by OEMs and other stakeholders. This probability is a modifiable parameter that we vary in sensitivity analysis (Section 4.6.1). By default, each OEM becomes known to at least one new agent with probability 0.5, each year. (Before at least one agent is aware of an OEM, this is the only means by which the fleet agent population can learn of it.)

Social interaction

Agents share observations and make inferences through socialization, which is detailed in Section 4.4. These activities also function to disseminate awareness of OEMs. When one agent is learning from another and knowledge is shared about vehicle options/OEMs that are new to the first agent, this causes the first agent to become aware of them and to initialize a set of observations about them.

C.2.2 Consideration

Once aware of a vehicle option, fleet agents first evaluate it for operational suitability. If it can fulfill the operational needs of the fleet, they analyze the economics of scenarios in which they adopt it, and estimate its overall utility according to the fleet's needs and priorities.

Operational suitability evaluation

Operational suitability is a broad term that we interpret here to be the answer to the question, “Can this vehicle do what our current fleet vehicles do?” For the answer to be “yes,” two conditions must be satisfied: (1) the vehicle must have a range greater than or equal to the greatest distance traveled in a day by the vehicle it would replace, and (2) the vehicle must have a cargo or passenger capacity greater than or equal to that of the vehicle it would replace. In some ways, these are quite conservative conditions. Fleets can consider replacing an existing vehicle with multiple smaller vehicles, for example, to satisfy the range and/or capacity conditions. The range requirements can also be reduced with en-route charging. Such choices could be feasible if associated costs (labor, charging) are less than the cost of a large replacement vehicle, or if other factors (brand image, environmental sustainability) outweigh them. In future work, these conditions will be relaxed, but they are used here for simplicity.

Finally, vehicles will not be replaced if they were purchased in the previous 8 years.

Exploratory economic analysis

Next, fleets evaluate the economics of replacing existing vehicles with each suitable vehicle option. They design scenarios, run short- and long-term economic analyses of each scenario, and save the results as new observations. Potential replacement options are determined as a pairing of a vehicle and the drive cycle it runs. For each pairing, a fleet agent follows a procedure:

1. Determine which of the existing vehicles could be replaced by the vehicle under consideration if it runs the drive cycle under consideration.
2. Design a scenario wherein all existing vehicles that could be replaced by this vehicle are replaced by this vehicle.
3. Run the techno-economic analysis tool, which does the following:
 - Specifies a minimum-cost EV charging system to support the new fleet’s operations, accounting for any existing EV charging equipment.
 - Designs a minimum-cost charging schedule, using the minimum-cost charging system, to enable the fleet to complete its drive cycles.

- Estimates the charging costs (for EVs) and fueling costs (for ICEVs) according to the charging schedule and drive cycles.
 - Predicts and amortizes total capital and operational cash flows, including vehicle procurement, maintenance, financing, taxes, etc., across a specified time horizon (by default, 10 years).
4. Extract the cash flows and net present costs incurred as a result of making the vehicle replacements.
 5. Divide these extracted costs (TCO, purchase price, maintenance cost, charging infrastructure cost, fueling cost) by the number of new vehicles that were added as replacements, and save those average costs as observations to the corresponding vehicle / drive cycle record.

In addition to replacement scenarios, the above procedure is applied for the “baseline” fleet. In this baseline scenario, the only procurements are to replace vehicles that have become inoperable due to age with new versions of the same model. This enables estimating the costs *not* to replace each vehicle of the fleet.

Utility evaluation

The final step before deciding whether to buy new vehicles is to assess the overall utility of each option known to the fleet, based on all the observations that have been made to characterize it. This procedure is described in Section 4.4.

C.2.3 Decision

In the previous step, vehicles were (a) culled for suitability, (b) evaluated economically, and (c) assessed holistically (evaluated for overall utility). In the Decision step, existing vehicles are marked for replacement with new vehicles using a semi-stochastic process. For each existing vehicle and its best potential replacement, the replacement is exercised with probability p_{adopt} , which is a function of the overall utilities for the two options (Equation 4.1). If the best potential replacement has an equal or greater utility (u_1) than the existing vehicle (u_0), the replacement is made with probability 1. In the alternative case, some fleet operators choose to adopt a new vehicle

before they are confident that it is a superior option, to gain information and evaluate feasibility (as a “pilot project”). In our model, the probability of piloting a vehicle is a function of the utility difference, where vehicles with smaller utility disadvantages are more likely to be piloted:

where u_0 is the perceived utility of an existing vehicle, u_1 is the perceived utility of the new option that would replace it, and k is a factor that scales the utility difference. For $(u_1 - u_0) \geq 0$, $p_{\text{pilot}} = 1$. Otherwise, for every additional $\frac{1}{k}$ disadvantage in utility, the probability of pilot adoption is reduced by a factor of 10. For example, if $k = 5$, a vehicle with a utility disadvantage of $1/5 = 0.2$ will be piloted with probability 0.1; a vehicle with utility disadvantage $2/5 = 0.4$ will be piloted with probability 0.01; and so on. We chose $k = 10$ for the experiments we present. The sensitivity of model outcomes to choice of k is investigated in Section 4.6.1. Equation 4.1 is evaluated for every vehicle a fleet is aware of but has not yet owned, every time step.

C.2.4 Implementation

To ensure the new fleet is operable, the techno-economic analysis tool is run once more with the new specifications, one outcome of which is a list of any EV charging equipment that may need to be procured and installed.

Having specified a new fleet and charging system, the fleet agent procures new equipment and sells existing equipment. Fleet agents procure equipment by transacting with OEM agents, which set prices and control production volume, and consulting with and/or transacting with electric utility agents to enable installation of charging equipment. Equipment sales value is a function of initial cost and depreciation, which is tracked as part of the economics model.

At this stage, we do not model the OEM and electric utility agents to place constraints on fleets’ implementation capabilities. However, vehicle procurement and grid readiness lead times are well-known challenges. Future experiments with this model will investigate how such constraints can influence EV adoption outcomes.

C.2.5 Confirmation

The final step of the process feeds forward into future iterations. Owned vehicles are evaluated continually via direct experience (Section ??), which updates the agents' observations and perceptions of those vehicles. For newly adopted options, the fleet agent's knowledge prior to adoption was strictly from initial assumptions, socially obtained impressions, and economic analysis. Moving forward, the fleet agent can now accrue observations of these options via direct experience, which may confirm or contradict the observations that led the agent to adopt them. For example, a fleet might adopt an unfamiliar EV with the expectation of high reliability, but quickly find out that it has unsatisfactory reliability. These new observations would begin to skew the agent's perceptions, potentially enough that the agent replaces it again with the vehicle it had been purchased to replace. Alternatively, if perceptions are confirmed or exceeded, the agent is more likely to make additional, similar replacements in the future.

C.3 Observation abstractions

Further, agents generalize their observations to assess new, previously unknown vehicle options. When an agent observes something about a vehicle it owns (e.g. its reliability), that observation is recorded as a data point associated with that specific vehicle. To generalize it, the observation also becomes associated with the drive cycle, body type, weight class, and fuel type categories to which the vehicle belongs. When the agent becomes aware of a new vehicle option, its prior state of observation "inherits" the observations associated with each of the new option's categories. These inherited perceptions are given less weight as their specificity decreases. As a result, the initial distribution of observations for a new type of Class 3 electric van would be drawn heavily from observations about another Class 3 electric van; somewhat less from observations of a Class 3 electric cargo truck; and only marginally from observations of a Class 8 electric tractor. Arranged from most general to most specific, the categories are fuel type, weight class, vehicle body type, vehicle make/model, vehicle make/model/drive cycle, and individual vehicle. This arrangement is illustrated in Figure C.1.

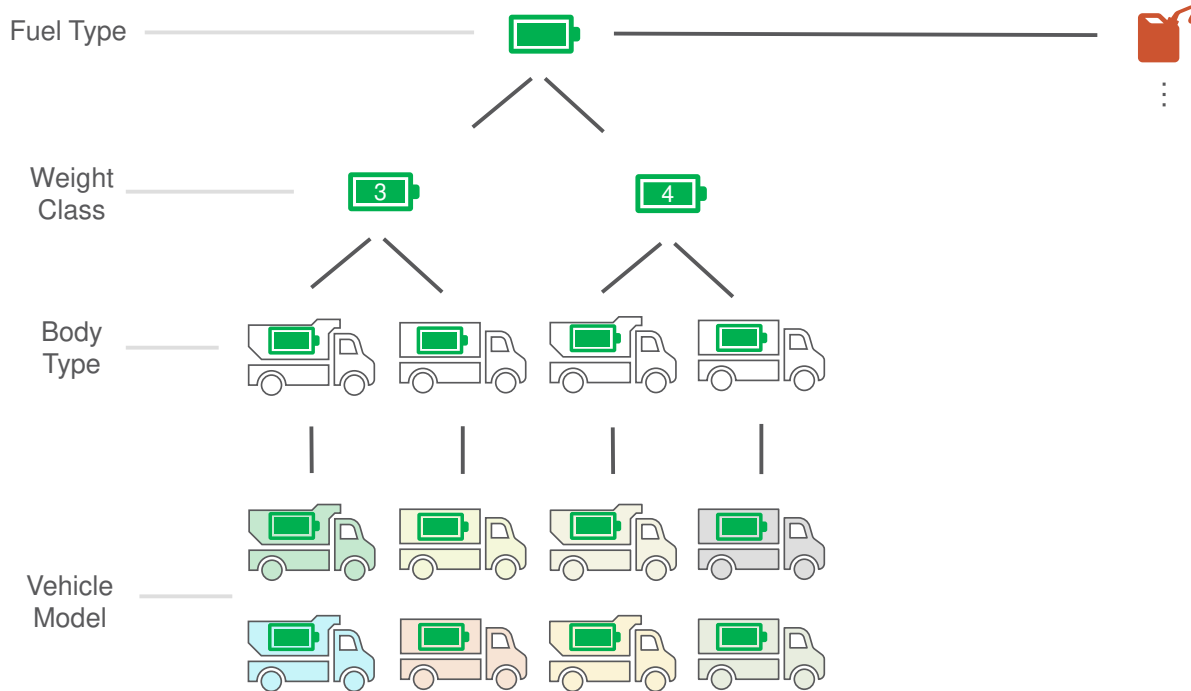


Figure C.1: Agents build perceptions of vehicle options and the categories to which the options belong. When an agent becomes aware of a new option, the prior distribution of perceptions about that option is constructed by inheriting perceptions from other options that occupy the same categories.

The methods agents use to make observations, summarized in Table C.1, are detailed in the sections below.

Table C.1: Methods by which fleet agents make observations, and the relative importance weight assigned to those observations

Method	Relative weight per year
Direct experience	5
Economic analysis	3
Social interaction	1
Social inference	1
Preliminary assumptions	5 (once, at initialization)

C.4 Observational inference process details

Just as in true Bayesian inference, in ABC an observer is modeled to update its prior knowledge to account for new information. True Bayesian inference requires directly computing a posterior distribution, which is often computationally infeasible, or estimating it via methods such as Markov Chain Monte Carlo, which can still be computationally costly. In contrast, ABC makes a simpler estimate of the posterior distribution, trading precision for tractability. The observer draws a large number of samples from its prior distribution, evaluates each sample for consistency with the observed evidence, and discards any samples that are insufficiently consistent with the evidence according to a threshold. The posterior distribution is thus constituted from the non-discarded samples.

We use ABC to generate new observations *in terms of vehicle attributes* on the basis of evidence *of other fleets' procurement choices*. In isolation, the ABC method would involve replacing an agent's observations with newly sampled observations (the estimated posterior). Instead, we sample from that estimated posterior to generate new observations, without discarding the old observations.

We illustrate the procedure by example. Fleet A observes that Fleet B has purchased Vehicle X as replacement for Vehicle Y. This implies that Fleet B judged Vehicle X to be superior to Vehicle Y, but Fleet A cannot directly know the basis for this judgment. Fleet A takes a large number of samples (10,000) from its existing observations about Vehicles X and Y, and for each sample, assesses which vehicle the sample suggests is superior (according to Fleet A's priorities, because we assume Fleet A does not know Fleet B's priorities). Observations for which Vehicle X is superior to Vehicle Y form the posterior distribution. To preserve memory and storage, Fleet A draws a relatively small number of samples (20) from this posterior distribution, and records these samples as new observations.