

KONSTANTINA VALOGIANNI

Sustainable Electric Vehicle Management using Coordinated Machine Learning



SUSTAINABLE ELECTRIC VEHICLE
MANAGEMENT USING COORDINATED
MACHINE LEARNING

Sustainable Electric Vehicle Management using Coordinated Machine Learning

Duurzaam management van elektrische voertuigen met behulp van
gecoördineerde machine learning

Thesis

to obtain the degree of Doctor from the
Erasmus University Rotterdam
by command of the
rector magnificus

Prof.dr. H.A.P. Pols

and in accordance with the decision of the Doctorate Board

The public defense shall be held on

Thursday the 30th of June 2016 at 15:30 hrs by

KONSTANTINA VALOGIANNI
born in Larisa, Greece

Erasmus University Rotterdam

The Erasmus University logo, featuring the word "Erasmus" in a stylized, cursive script.

Doctoral Committee

Promoters: Prof.dr.ir. H.W.G.M. van Heck
Prof.dr. W. Ketter

Other members: Prof.dr. A. Gupta
Prof.dr.ir. M.B.M. de Koster
Prof.dr. R. Unland

Erasmus Research Institute of Management - ERIM

The joint research institute of the Rotterdam School of Management (RSM)
and the Erasmus School of Economics (ESE) at the Erasmus University Rotterdam
Internet: <http://www.erim.eur.nl>

ERIM Electronic Series Portal: <http://repub.eur.nl/pub>

ERIM PhD Series in Research in Management, 387

ERIM reference number: EPS-2016-387-LIS

ISBN 978-90-5892-450-6

©2016, Konstantina Valogianni

Design: B&T Ontwerp en advies www.b-en-t.nl

This publication (cover and interior) is printed by haveka.nl on recycled paper, Revive®.

The ink used is produced from renewable resources and alcohol free fountain solution.

Certifications for the paper and the printing production process: Recycle, EU Flower, FSC, ISO14001.

More info: <http://www.haveka.nl/greening>

All rights reserved. No part of this publication may be reproduced or transmitted in any form or by any means electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without permission in writing from the author.



To my parents, Athanasia and Thanasis

Acknowledgments

Conducting a PhD has been one of the most exciting journeys in my life. First and foremost, it allowed me to pave the way to my academic career and develop my personality toward multiple directions. Second, I was honoured to meet and collaborate with great people each of whom I feel the need to thank separately.

I am indebted to my supervisors Prof. Eric van Heck and Prof. Wolfgang Ketter who gave me the opportunity to conduct this PhD project under their valuable guidance. Eric, thank you for your curious questions, which made me a better researcher and for all this positive attitude, which always kept me motivated. Wolf, thank you for all the support and trust throughout my PhD trajectory and showing me the way to succeed. Your energy taught me that I should never give up, while brainstorming with you was one of the most fun parts of my PhD. Thank you both!

Besides my two supervisors, I would like to express my gratitude to my co-authors who have been great collaborators during this journey. Prof. Alok Gupta invited me to visit the department of Decision and Information Sciences at University of Minnesota. Alok besides being an exceptional academic host, provided me with amazing guidance and taught me how to continuously improve and become a better academic. Additionally, I would like to thank Alok for being a member of my PhD inner committee and flying all the way to Rotterdam to be part of my defence. Dr. John Collins was always there since the beginning of my PhD to guide me in the world of energy systems. Furthermore, I am grateful to John's diligence to meet me whenever I was in Minneapolis, to brainstorm and come up with innovative ideas. I extremely appreciate John's willingness to fly from Minneapolis to Rotterdam to be a member of my PhD plenary committee. Prof. Gedas Adomavicius co-authored one of the articles presented in this dissertation (Chapter 5). I am thankful to Gedas' excellent guidance and the time he spent on our mathematical discussions. Dr. Soumya Sen from University of Minnesota has been a co-author of the article presented in Chapter 4 of this dissertation. Soumya was always there during my

research visit in Minnesota to discuss about ideas and help me to put them on paper. Dr. Dmitry Zhdanov has been a co-author of two papers of this dissertation (Chapter 2 and 3). I am thankful to his feedback to improve the papers.

In addition, I would like to thank the members of my PhD inner committee: Prof. René de Koster and Prof. Rainer Unland. I am grateful to René's thorough feedback on my dissertation and his willingness to serve as the secretary of my PhD inner committee. I am thankful to Rainer Unland who served as a member of my PhD inner committee and provided feedback on my dissertation. Furthermore, I would like to express my gratitude to the members of my PhD plenary committee: Prof. Maria Gini, Dr. Andreas Symeonidis and Dr. Jan van Dalen. I am thankful to Maria who welcomed me in her group meetings in the Computer Science Department of University of Minnesota during my time there, and was always willing to provide feedback on my work. I am incredibly grateful to Andreas who, as my MSc thesis supervisor in January 2011, played a pivotal role in my decision to pursue a PhD at Rotterdam School of Management. Many thanks go to Jan van Dalen who was willing to serve in my PhD committee and be part of my defence.

I am extremely grateful to have worked in such a great place like the Department of Technology & Operations Management. The collegial environment since the first day I joined has played an important role in my academic and personal development. Even more important are the friendships I made. Irina, Nick, Panos, Sarita, Wouter, Clint, Paul thanks for all the great moments in "Department 1". Derck, Markus, Micha, Yixin thanks for being awesome officemates. Furthermore, Derck and Micha thanks for being my paranymphs. Christina, thanks for all the interesting discussions. Furthermore, I would like to thank the rest of the "BIM crew": Arthur, Dimitris, JoHee, Ksenia, Mark, Mo, Otto, Rodrigo, Thomas, Ting, Yashar for giving feedback on my work during department seminars, the scientific developers: Govert and Erik for their valuable help whenever needed and the best support staff one could wish for: Cheryl, Ingrid and Carmen.

I would also like to thank the Erasmus Research Institute of Management (ERIM) for the financial support throughout these years to attend conferences, workshops, external courses and build a strong academic portfolio. Special thanks to the colleagues in the ERIM office: Natalija Gersak, Kim Harte, Miho Iizuka, Tineke van der Vhee. In addition, I would like to acknowledge that part of my PhD was funded by the EU FP7 project Cassandra and thank all the Cassandra consortium members for the lively meetings and

the interesting discussions.

Last but not least, a special thanks goes to my family and friends for their unconditional love and support. To my mom, Athanasia and my dad, Thanasis who raised me with the aspiration to make the world a better place and never give up on my dreams (Στη μητέρα μου Αθανασία και στον πατέρα μου Θανάση που με μεγάλωσαν με την φιλοδοξία να κάνω τον κόσμο καλύτερο και ποτέ να μην εγκαταλείπω τα όνειρά μου). To my brother, George, who was always there to understand my “crazy mathematical” questions and have a meaningful conversation about constellations. To my friends Vivian, Vasia, Athina, Thodoris, Fryni, Andreas, Eri, Vasilis, Thanasis, George, Christina, Jan Willem, Yannis, Konstantinos, Elizabeth, Ioannis, Natassa, Anna, Sonia, Maria, Agapi, who proved to me that no matter which part of the world you are in, your friends will always be there for you. Thank you all! Ευχαριστώ πολύ!

Konstantina Valogianni
Rotterdam, April 2016

Table of Contents

Acknowledgments	vii
1 Introduction	1
1.1 Research Question	6
1.2 Research Methodology	7
1.3 Dissertation Structure	8
1.4 Declaration of Contribution	8
2 Sustainable Demand Side Management for Electric Vehicles using Personalized Learning Agents	11
2.1 Introduction	11
2.2 Background and Related Work	13
2.2.1 Smart Electricity Grid and Electric Mobility	13
2.2.2 Sustainability and Green IS	16
2.3 Adaptive Management of EV Storage (AMEVS)	17
2.3.1 Model Assumptions	18
2.3.2 Input Module	20
2.3.3 Learning Module	22
2.3.4 Optimization Module	24
2.3.5 AMEVS Variations based on Energy Consumption Utility	26
2.4 Experimental Evaluation	28
2.4.1 Simulation Environment and Data Description	28
2.4.2 Benchmarks	31
2.4.3 Numerical Results	35
2.4.4 Sensitivity Analysis	45
2.5 Policy Recommendations	48
2.6 Conclusions & Future Work	50

3	Sustainable Electric Vehicle Charging: A Data-driven Approach	53
3.1	Introduction	53
3.2	Mobility Integrated Energy Management Model	55
3.2.1	Photovoltaic generation forecasting	57
3.2.2	Electric Vehicle Charging Scheduling	59
3.3	Experimental Evaluation	60
3.3.1	Data Description	61
3.4	Numerical Results	64
3.4.1	PV forecasting	65
3.4.2	Peak Demand Reduction	67
3.5	Conclusion & Future Work	68
4	Maximizing Social Welfare in Grid Resource Allocation for Electric Vehicle Charging	71
4.1	Introduction	71
4.2	Related Work	73
4.3	Model Formulation and Structural Analysis	74
4.3.1	Assumptions	75
4.3.2	Multiple Vickrey Auction	76
4.4	Results	78
4.4.1	Theoretical Properties	78
4.4.2	Data Description	86
4.4.3	Empirical Evaluation	87
4.4.4	Robustness Check	92
4.5	Conclusions & Future Work	93
4.6	Acknowledgements	94
5	A Dynamic Pricing Mechanism to Coordinate Electric Vehicle Charging	97
5.1	Introduction	97
5.2	Related Work	99
5.3	Hybrid Coordination	100
5.3.1	EV Driver's Agent Decision Problem	100
5.3.2	Smart Grid Manager's Decision Problem	102
5.3.3	Learning Component	105
5.4	Multiagent Simulation	105
5.4.1	Scenarios & Assumptions	107
5.4.2	Data Description	108

5.4.3	Benchmarks	110
5.5	Numerical Results	112
5.5.1	Impact on Power Demand	113
5.5.2	Shaping Aggregate Power Demand	115
5.5.3	Sensitivity Analysis	116
5.6	Conclusions & Future Work	116
5.7	Acknowledgements	117
6	Conclusion	119
6.1	Synopsis of Main Findings	120
6.1.1	Chapter 2	120
6.1.2	Chapter 3	121
6.1.3	Chapter 4	121
6.1.4	Chapter 5	122
6.2	Combination of IS artifacts	123
6.3	Generalizability and Methodology Discussion	123
6.4	Theoretical Contribution and Practical Impact	125
6.5	Limitations and Directions for Future Research	126
	References	129
	Summary	139
	Nederlandse Samenvatting (Summary in Dutch)	141
	Περίληψη στα Ελληνικά (Summary in Greek)	143
	Curriculum Vitae	147
	Portfolio	149

Chapter 1

Introduction

The advent of new technologies has resulted in a major business transformation in modern societies (Bharadwaj et al., 2013). One important and quickly transforming aspect is the way that electricity is produced, delivered, and consumed. In the past, electricity was generated mainly using fossil fuels, distributed in a purely centralized way and the end customers were passively consuming at a regulated price. Currently, we live in a transformation period, in which the electricity grid has been equipped with advanced information and communication technology (ICT) at every stage of the electricity supply chain. The ICT-enabled electricity grid is known as a *smart grid* (Amin and Wollenberg, 2005) and has revolutionized the structure and the capabilities of the traditional electricity grids.

The modern smart grid, besides its ability to incorporate technologically advanced equipment and process large amounts of data, has in its core a significant sustainability component. A main characteristic of the smart grid is the large scale integration of renewable sources (wind turbines, solar panels, electric vehicles) that reduce society's dependence on fossil fuels, mitigate carbon emissions and increase sustainability levels.

This dissertation investigates how intelligent algorithms can support electricity customers within the electricity grid. Specifically, it focuses on electricity customers that own electric vehicles (EVs). Electric vehicles require particular attention for two reasons. First, they consume more electricity than an average household; for example, a US household might consume on average 30 kWh¹ a day and an EV battery needs from 24kWh to 80kWh to charge fully (a full charge might cover the driving needs of 1-1.5 days). Second, with rising EV adoption, new challenges arise for sustainable societies. All these new EVs added to the electricity grid will put its infrastructure under critical strain, since extra demand will be added to the grid. This demand was previously covered by gasoline, but

¹<http://www.eia.gov/tools/faqs/faq.cfm?id=97&t=3> [Date Accessed: March 22nd, 2016]

now the electricity grid operators need to be able to cater for this new customer segment, namely EV owners. In order for the grid to be reliable and able to service all EV customers, new infrastructure needs to be installed. The new infrastructure must be able to accommodate this extra electricity demand and prevent blackouts. A typical example is the California grid, which must now accommodate demand from more than 100,000 EVs. For this purpose, the California grid operator has decided to expand the capacity of the grid². However, this solution is unsustainable since more raw materials (such as copper, required for cable installation) need to be consumed and the investment will be excessive.

The goal of this dissertation is to design algorithms that will coordinate EV charging, making the best use of the available information. Coordinating EV charging will increase the levels of sustainability on the smart grid since less infrastructure will be needed to cover peaks. A coordinated EV charging electricity demand is expected to be less volatile with increased capacity utilization. This means that the “peaks” and “valleys” in the demand curve will be reduced, increasing the grid’s stability. Coordination mechanisms can be either decentralized (*bottom-up*) or centralized (*top-down*) (Dias et al., 2006), both of which have advantages and disadvantages depending on the market they are applied to.

On the one hand, decentralized approaches require no formal coordinating entity and assume that each individual electricity customer communicates with the electricity grid via pricing signals. These signals have the ultimate goal of incentivizing consumers to charge the car when demand is low (low price period) and provide counter-incentives for charging when there is peak in the electricity demand (high price period). Bottom-up approaches have as their main benefit that the customers have the freedom to schedule their power consumption based on individual preferences. However, the main disadvantage is that since the same price signals are provided to all the customers, the power consumption schedules coincide, leading to *herding* behavior (Gottwalt et al., 2011). Specifically, since agents tend to be cost minimizers, they are inclined to shift power demand to the cheaper time periods, creating new peaks in the demand.

On the other hand, centralized mechanisms assume an external coordinator, who is usually the grid operator. This actor is responsible for maintaining the stability and reliability on the grid and usually prevents electricity consumption during periods when electricity demand is peaking. The benefits of the top-down coordination mechanisms are that they easily satisfy the constraints imposed by the coordinator (e.g. smart grid manager), leading to a reliably balanced system. The most important challenge is that

² <http://www.greentechmedia.com/articles/read/california-steps-up-again-on-electric-vehicles> [Date Accessed: March 22nd, 2016]

the coordinator must intervene and exogenously control the EV battery, violating the EV driver's preferences.

This dissertation is divided into four separate studies as follows (see Figure 1.1). Each study proposes a different coordination mechanism dependent on the involved stakeholders (grid operators, electricity providers, EV users) and the objective (peak demand reduction, sustainability increase). The benefits for the electricity market and its environmental sustainability are investigated and discussed.

In each study we use different types of machine learning to learn from the available data (customer preferences, price information, market conditions, grid capacity utilization) and optimize stakeholders' decisions. However, each of the stakeholders involved (grid operators, energy providers and EV owners) has objectives that might be conflicting. For example, the grid operators strive for grid stability and low demand volatility, whereas the EV customers desire maximum comfort and fully charged batteries. Similarly, the energy providers are profit maximizers, which might be in conflict with the previous objectives. Therefore, the machine learning algorithms used at each stakeholder's side need to be coordinated so that in the end the outcome is beneficial for as many stakeholders involved as possible. In a multi-actor system, such as the electricity grid, it is possible that all the conflicting objectives might lead to outcomes that are not beneficial either for the market or for the customers. For example, in an electricity market where the grid operator broadcasts variable prices to the customer portfolio, it is very likely that all customers will strive to minimize their costs, charging their EVs when prices are low, creating new peaks during these low price periods. This phenomenon is known in the literature as herding (Gottwalt et al., 2011), and it is one of the situations when two conflicting objectives might lead to a worse outcome for both sides involved.

In the presented studies the machine learning algorithms are implemented in the form of an *IS artifact*. According to Hevner et al. (2004) an IS artifact is a construct which provides a solution to a common problem. An artifact demonstrates feasibility of the solution and facilitates the comparison with other similar artifacts designed to address the same problem. IS artifacts allow for thorough evaluation of the proposed solutions and can be used for scenario analysis.

Chapter 2 – Decentralized EV Coordination: We propose a bottom-up EV charging coordination mechanism. This study designs IS artifacts that represent EV owners. These IS artifacts learn their owners' preferences and schedule EV charging so that these preferences are satisfied and the benefits for each individual EV owner are maximized. Each EV owner is represented by an IS artifact which is implemented through an intelligent software agent (Wooldridge and Jennings, 1995; Adomavicius

et al., 2009). We demonstrate that adopting the proposed IS artifact yields financial benefits for each individual customer as well as the grid. Specifically, we observe that heterogeneity of customer preferences gives rise to an emergent charging coordination. By emergent coordination we mean a bottom-up coordination which emerges from each individual maximizing his/her own utility function.

Chapter 3 – Decentralized EV Coordination, smart home perspective: An extension of Chapter 2 is the integration of the IS artifact in a smart home combined with an EV. In this study, we implemented a mechanism which accounts for matching the EV charging with a solar panel generation. Using supervised learning, the proposed algorithm estimates the solar panel output and it tries to charge the EV with as much renewable energy as possible. The outcome is two-fold: the smart home has an electricity cost reduction, since the solar panel belongs to the smart home and the smart grid has increased sustainability levels, since significant amount of conventional energy is substituted with renewable energy.

Chapter 4 – Centralized EV Coordination: We take the stand-point of the smart grid manager (or grid operator) who is represented by an IS artifact that schedules EV charging using an auction mechanism. Its main objective is to service as many EV owners as possible without overloading the grid. We allocate the available charging capacity and payments to the electricity customers under the social welfare maximization objective and demonstrate the outcomes for all stakeholders involved (EV users, grid operator, energy providers). We observe that the grid remains stabilized and the social welfare is maximized. We prove the properties that ensure maximum social welfare and provide managerial insights to the grid operators and energy providers. This study is proposing a theoretical auction-based framework for EV charging scheduling and payment allocation, and is going to be evaluated empirically using a mobile app experiment.

Chapter 5 – Hybrid EV Coordination: We propose a hybrid coordination mechanism that combines benefits both from the decentralized and centralized approaches. We show that this mechanism is capable of reducing peak demand and satisfying individual preferences. What is more important, this mechanism can mitigate *herding* behavior that is present in bottom-up mechanisms and overcomes the practical implementation barriers of the top-down approaches. It worths mentioning that this mechanism is also fair, since it broadcasts the same price function to all customers in the market.

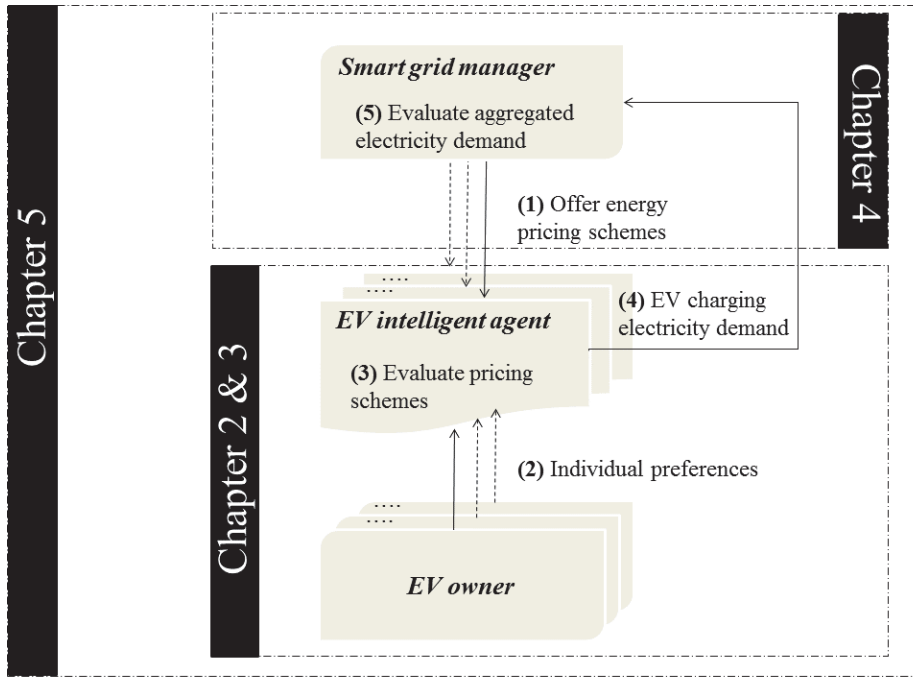


Figure 1.1: Areas of contribution for the four studies that comprise the dissertation project

Scientifically, this work contributes to the implementation of IS artifacts that assist electricity customers in the complex smart grid’s environment (Adomavicius et al., 2009; Bichler et al., 2010), and to the nascent IS research stream of Energy Informatics (Watson et al., 2010). Following the categorization proposed by Melville (2010), the proposed IS artifacts answer the research question: “What design approaches are effective for developing information systems that influence human actions about the natural environment?”. With the proposed IS artifacts we are providing incentives to EV owners so that their charging behavior creates benefits for them in terms of electricity cost savings and for the electricity grid in terms of peak demand and volatility reduction.

This dissertation’s managerial and societal relevance lie in its contributions to assisting EV customers in their complex decisions and eventually coordinating EV charging. Firstly, providing assistance to the EV owner’s complex decision problems overcomes bounded rationality issues (Simon, 1996) and provides them with a broader decision spectrum. This means that EV owners can arrive at better decisions benefiting from information available. Secondly, coordinating EV charging, so that peak demand is mit-

igated, is important for the electricity grid since capacity expansion investments can be reduced (Strbac, 2008). This means that by employing IS artifacts, sustainable societies can reduce peaks and volatility in the electricity demand, and the grid’s stability will be increased. A more stable and reliable grid requires less capacity expansion investments which is both more cost effective and more sustainable. Consequently, the proposed IS artifacts can be used by two different stakeholders: by energy customers to support them in the electricity markets and by the grid managers to mitigate peak demand and increase sustainability levels.

1.1 Research Question

EV owners have to take real time charging decisions, making the best use of information available. Typically, they have to estimate how much energy to charge so that their driving needs are covered without violating their preferences. In variable pricing regimes they have to account for the price variation as well, so that their EV charging costs are minimized. One can understand that for a human these type of decisions are not easy to make on the spot. Therefore, most humans just plug their car for charging once they return home from work and unplug it the next morning. This leads to excessive charging that overloads the electricity grid and decreases its sustainability levels (a full charge might give more electricity than the EV owner actually needs to cover her driving needs for the coming day ³).

To tackle this challenge, this dissertation proposes intelligent IS artifacts that manage EV charging, support individual EV owners, and ultimately coordinate EV charging to mitigate peak demand. The main research question of this dissertation is:

How should IS artifacts be designed to achieve effective EV charging coordination under various market conditions, in a sustainable manner?

We define an effective coordination mechanism as an approach that schedules EV charging so that peaks in the electricity demand are reduced and benefits for the individual EV owners are created. In other words, we aim at creating mechanisms that achieve win-win situations (Kambil and van Heck, 2002) for both parties involved in the EV charging process. We aim to answer this question through a **design science** approach (Hevner et al., 2004; Gregor and Jones, 2007; Walls et al., 2004). Our artifacts offer new solutions

³https://www.teslamotors.com/nl_NL/blog/driving-range-model-s-family [Date Accessed: March 22nd, 2016]

to the existing problem of managing EV charging in an effective way. Therefore, from the design science point of view, they are categorized under the solution group “Improvement: New Solutions for Known Problems” (Gregor and Hevner, 2013). These new solutions strive to yield the maximum possible benefits to the individual EV owners using the information available. We evaluate our IS artifacts with respect to validity, utility, quality and efficacy (Hevner et al., 2004; Gregor and Hevner, 2013) in simulations built with real-world data and real-world experiments using a mobile application (Koroleva et al., 2014).

We define the various market conditions in the presented studies, as different pricing regimes that might create various incentives for EV customers. We show how the grid is influenced by variable pricing or time-of-use regimes, as well as pricing schemes dependent on the charging speed. We present results examining all these scenarios and provide insights to energy policymakers.

1.2 Research Methodology

We adopt different methods to answer this dissertation’s research question. Employing multiple methods has the advantage of revealing different aspects of the problem and the respective subproblems, which might not be apparent when employing a single-method research approach. We now give an overview of the methods implemented, while the detailed outcomes are presented in each chapter separately.

Mathematical Modeling - Optimization. Mathematical modeling is the essence of representing real-world decision processes with mathematical concepts and language (Bender, 2012). It can be a powerful tool, when abstraction is needed, so that a problem can be solved mathematically and provide intuition to decision makers. In this dissertation we employ mathematical models to optimize certain stakeholder decisions (e.g. EV customer charging decisions) based on certain objectives. We calibrate the parameters of our models with real-world data, so that we achieve realistic results and provide the right guidance to stakeholders.

Statistical Machine Learning. Statistical machine learning can leverage the power of data, so that intelligent recommendations can be offered to users. In this dissertation we employ both supervised and unsupervised learning to *learn* from data and offer personalized recommendations to electricity customers. For instance, we use reinforcement learning (Sutton and Barto, 1998) to learn from electricity consumption data for each individual customer and offer personalized EV charging recommendations. Furthermore, we implement a supervised learning approach to estimate a photovoltaic panel’s output from weather data and couple this output with EV charging.

Specifically in this dissertation, we assume that both the grid operators and the EV customers have a machine learning algorithm which assists them in their decisions. Each of the involved actors (EV customers and grid operators) have different and potentially conflicting objectives. Therefore, the learning mechanisms might make the electricity markets converge to sub-optimal situations. For example, the grid operator might set prices that will lead to shifting electricity demand to low price periods, creating herding in EV charging and new peaks in the demand curve. For this purpose, we are using *coordinated* machine learning in the sense that we are setting the incentives on both sides of the electricity market (grid operators and EV customers) so that we prevent the market from converging to suboptimal outcomes (such as the herding in EV charging).

Simulations. Since goal of this dissertation is to demonstrate how the proposed intelligent algorithms are able to assist electricity customers, we conduct large scale simulations with various objectives. Each chapter of this dissertation assumes different market conditions and stakeholder objectives, therefore simulation-based experiments are necessary to provide evidence that the designed algorithms are effective. Simulation-based experiments are important when examining complex problems (such as smart grid related problems) because they are able to reveal unexpected outcomes that might not be apparent in a purely mathematical analysis. Our simulations use assumptions from Power TAC (Ketter et al., 2016) ⁴.

1.3 Dissertation Structure

This dissertation is structured as follows (Figure 1.2). The current chapter provides an introduction to the research context. Chapters 2, 3, 4 and 5 present the detailed studies and their outcomes. Chapter 6 summarizes the results of the separate studies and provides directions for future research.

1.4 Declaration of Contribution

Here, I declare my contribution to the chapters of this thesis as well as co-authors' contributions.

Chapter 1: The main part of the work was done independently by the author of this thesis. The promoters and inner committee gave feedback, which was implemented accordingly by the author.

⁴Power TAC is a smart grid simulation platform where energy brokers compete to attract electricity consumers and make profits.

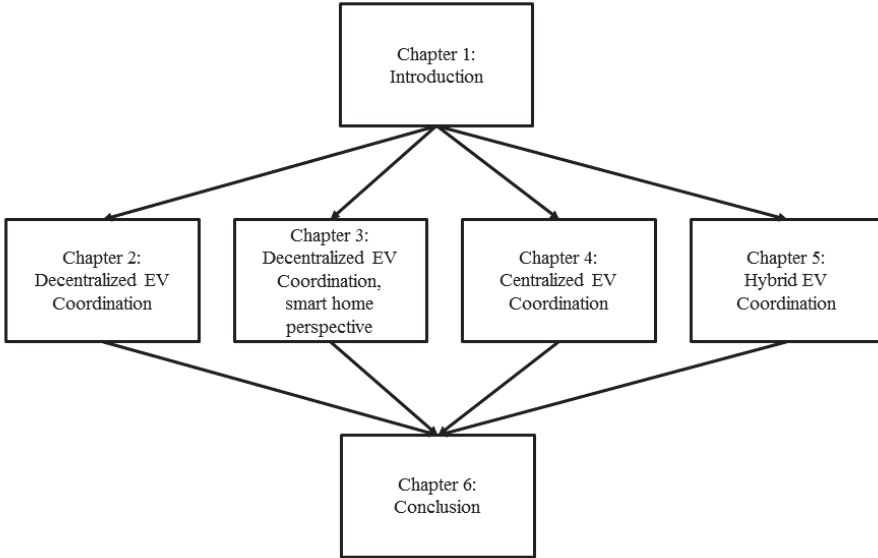


Figure 1.2: Structure of the dissertation

Chapter 2: The author of this thesis conducted the majority of the work for this chapter. Specifically, the research question was formulated by the author after consulting with the co-authors. The algorithm implementation was done by the author independently. The data was obtained by the author in consultation with the co-authors. The results and the discussion were analyzed and written by the author independently. Feedback was given by the co-authors and was implemented accordingly by the author. This chapter has appeared in peer reviewed artificial intelligence and information systems conference proceedings. Furthermore, this chapter is now under review at a top information systems journal and the author of this dissertation is the first author.

Chapter 3: The author of this thesis conducted the majority of the work for this chapter. Specifically, the research question was formulated by the author after consulting with the co-authors. The algorithm implementation was done by the author independently. The data was obtained by the author in collaboration with the co-authors. The results and the discussion were analyzed and written by the author independently. Feedback was given by the co-authors and was implemented accordingly by the author. This chapter has appeared in peer reviewed information systems conferences and the author of this thesis is the first author.

Chapter 4: The author of this thesis conducted the majority of the work for this chapter. Specifically, the research question was formulated by the author in collaboration with the co-authors. The algorithm implementation was done by the author independently. The data was obtained by the author independently. The results and the discussion were analyzed and written by the author independently. Feedback was given by the co-authors and was implemented accordingly by the author. This chapter has appeared in peer reviewed information systems conferences and the author of this thesis is the first author.

Chapter 5: The author of this thesis conducted the majority of the work for this chapter. Specifically, the research question was formulated by the author in collaboration with the co-authors. The algorithm implementation was done by the author independently. The data was obtained by the author independently. The results and the discussion were analyzed and written by the author independently. Feedback was given by the co-authors and was implemented accordingly by the author. This chapter has appeared in peer reviewed information systems and artificial intelligence conferences and the author of this thesis is the first author.

Chapter 6: The main part of the work was done independently by the author of this thesis. The promoters and inner committee gave feedback, which was implemented accordingly by the author.

Chapter 2

Sustainable Demand Side Management for Electric Vehicles using Personalized Learning Agents¹

2.1 Introduction

Electric vehicles (EVs) are potentially a major factor in creating a sustainable future (European Commission et al. (2011), US Department of Energy²). EV adoption is heavily incentivized by governments and policymakers as EVs lower local carbon footprint, reduce noise pollution, and have much higher engine efficiency. At an individual level, many commuters tend to adopt EVs as part of their effort to become more environmentally

¹Parts of this chapter have appeared in the following peer reviewed conference proceedings:

Valogianni, K., Ketter, W., Collins, J. & Zhdanov, D.(2014). Enabling Sustainable Smart Homes: An Intelligent Agent Approach. In Proceedings of the International Conference on Information Systems, (ICIS 2014), Auckland, New Zealand, Association for Information Systems (AIS).

Valogianni, K., Ketter, W., Collins, J. & Zhdanov, D.(2014). Adaptive Learning Agents for Electric Vehicle Customer Decision Support. In Conference on Information Systems and Technology, (CIST 2014), San Francisco, California.

Valogianni, K., Ketter, W., Collins, J. & Zhdanov, D.(2014). Effective Management of EV Storage using Smart Charging. In Proceedings of the Twenty-Eighth Conference of Advancement on Artificial Intelligent (AAAI 2014), pp. 472-478, July, 27-31, Quebec, Canada.

Valogianni, K., Ketter, W., & Collins, J. (2014). Learning to Schedule Electric Vehicle Charging given Individual Customer Preferences. In Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), Lomuscio, Scerri, Bazzan, Huhns (eds.) pp.1591-1592, May 5-9, 2014, Paris, France.

Valogianni, K., Ketter, W., & Collins, J. (2013). Smart Charging of Electric Vehicles. In 2013 AAAI Workshop - Trading Agents Analysis and Design (TADA 2013), Bellevue, Washington.

This paper is currently under review at a top-ranked information systems journal and the author of this dissertation is the first author.

²<http://energy.gov/articles/history-electric-car> [Date Accessed: March 22nd, 2016]

aware and sustainable (Sovacool and Hirsh, 2009). Consequently, psychological motives play a role in EV adoption in addition to economic and policy incentives.

Besides the numerous benefits they offer to individuals and society, widespread EV adoption threatens the stability of the power grid. This is because they consume large amounts of electricity to charge their batteries; for example a US household might consume on average 30 kWh³ a day and an EV battery needs from 24kWh to 80kWh to charge fully. Therefore, EVs will become significant electricity consumers. In the absence of a scheme for managing this additional load, large EV penetrations have the potential to destabilize segments of distribution grids and even the higher-voltage transmission grids. This could occur if, for example, all residents of a neighborhood own an EV, and they decide to start charging them around the same time in the evening, which is already a time of peak consumption in many areas of the world. The traditional solution to this challenge would be to build extra electricity infrastructure with higher capacity (MW) that would be capable of handling such peak loads. However, this solution would be costly as well as highly unsustainable (Strbac, 2008; Watson et al., 2010).

Taking a closer look at the smart grid, we see that the expectation of peak demand (maximum electricity demand over some time horizon) is the primary determinant of grid capacity planning (Kirschen, 2003). In other words, if there is peak demand of 3MW for only one hour during the day, the electricity grid needs to have capacity for at least 3MW, despite that the demand seldom exceeds 1MW during the rest of the day. Therefore, unanticipated changes in peak demand will be an important issue as increasing numbers of EVs are introduced. Because of the time required to charge an EV, owners may fear that their car batteries will run out at inconvenient times and places, and therefore desire to charge them up to full capacity at high charging speeds. This creates high peak loads that might last short periods of time, but put the infrastructure under strain. This feeling of the customer is known as *range anxiety* (Franke et al., 2011) and influences charging behavior since the EV is often the owner's main means of transportation.

We aim to examine the effect of large scale EV integration in the electricity grid via an IS artifact that represents an individual EV owner and personalizes EV charging based on his/her own preferences (driving habits, household occupancy, car characteristics, etc.). It supports individual EV owners in their charging decisions, overcoming the barrier that many commuters have of not knowing how and when to charge their car batteries. The core of our artifact is an adaptive smart charging algorithm that benefits from information about individual preferences and market conditions, and reallocates EV charging to time intervals with lower demand.

³<http://www.eia.gov/tools/faqs/faq.cfm?id=97&t=3> [Date Accessed: March 22nd, 2016]

In order for the IS artifact to be effective, it must be implemented through a charging controller on the EV. The controller should identify when the car is plugged in and should be capable of receiving signals from the grid infrastructure. Tesla motors currently provides an online portal⁴, where the drivers can enter their expected driving distance, the prices they pay and will then receive information about the charging duration based on their charging speed/charging outlet they use. We extend this logic by implementing our artifact along the same lines but using machine learning and optimization to learn and schedule the EV charging based on personalized preference inputs.

Our method produces a win-win outcome, providing individual savings as well as increase in societal welfare through grid-level peak reductions. We show that our approach benefits from heterogeneity of preferences among EV owners, bringing cost savings to individuals that consume when demand is low and use electricity from their batteries when demand is high. Furthermore, we observe benefits for the distribution network by reducing the strain of peak demand related to EV charging. Consequently, by satisfying individual objectives we observe benefits for the grid at the same time. This can be interpreted as an emergent charging coordination across a population of EV owners, stemming from their heterogeneous preferences. We quantify the benefits to the grid using peak-to-average-ratio (PAR) and load factor (LF) metrics.

The rest of this chapter is structured as follows. First, we outline the background of electric mobility and the role of EVs in a modern electricity grid (smart grid). Second, we describe the IS artifact and evaluate its output, calibrating it with real world data. We benchmark it against existing charging strategies, and we present a sensitivity analysis of our results. Furthermore, we present selected policy recommendations based on the simulated scenarios. Summarizing, we discuss our main findings and outline streams for future work.

2.2 Background and Related Work

Our work builds on research related to smart grid, electric mobility, energy informatics and green IS.

2.2.1 Smart Electricity Grid and Electric Mobility

The smart electricity grid is the evolution of the traditional electricity delivery infrastructure, with technological advancements playing increasingly crucial roles in generation,

⁴<http://www.teslamotors.com/charging/calculator> [Date Accessed: March 22nd, 2016]

transmission, distribution, and consumption. It is becoming a smart market (Bichler et al., 2010) for electricity; decisions can be facilitated by intelligent software agents that can act on behalf of people or organizations (Ketter et al., 2012). Blumsack and Fernandez (2012) identify three aspects that make the smart grid powerful in serving customer electricity needs: real-time monitoring at the transmission level, automation of various aspects of distribution systems, and smart-metering for electricity customers (Chrysopoulos et al., 2014). The transmission level consists of high capacity lines bringing electric power from large-scale generation facilities to local distribution networks. Power in the transmission grid is largely produced by generation companies (GenCos) that own single generation plants or a portfolio of power generation sources (Kirschen and Strbac, 2005). Local distribution networks are the second level of the smart grid where electricity is delivered (at lower voltage than the transmission level) to end customers located around a power substation. These two levels describe the physical infrastructure that reliably delivers electricity to end customers.

Above the physical infrastructure lie the economic mechanisms that allow for financial exchanges between energy customers and energy providers (in the literature also known as aggregators or energy utility companies). Energy utility companies offer energy tariffs (contracts) to consumers and aim to make profits through transactions with them. EV owners rely on electricity to cover their mobility needs, charging their cars so that their EV batteries contain sufficient energy for their driving needs. EV owners are billed for their mobility services based on electricity prices rather than oil prices as conventional car owners are.

A large scale EV integration in the electricity grid will bring demand increase which might threaten the grid's stability. The effect of EV integration on the physical infrastructure is outlined by Peças Lopes et al. (2010). On the positive side, using available information to make better EV charging decisions shows promising results for the grid and the EV owners. Vandael et al. (2013) aggregate EV customer profiles with the objective to coordinate their charging. This charging coordination is performed by an aggregator (EV fleet operator). This centralized mechanism achieves peak power demand reduction without accounting for individual preferences and behavioral patterns. Gerding et al. (2013) propose a two-sided market approach to allocate charging timeslots among EV customers and to avoid charging congestion. Gerding et al. (2011) present an online auction mechanism where EV owners state their timeslots available for charging and also bid for power. Finally, Stein et al. (2012) describe an online mechanism with pre-commitment for coordinating the EV charging.

At a macroscopic view, Brandt et al. (2012) outline EV integration towards the new era of sustainable societies and relate it to new business models, whereas Wagner et al. (2013) present a specific business model targeting EV integration that will allow for a smoother transition to electrified mobility. Tomic and Kempton (2007) and Kahlen and Ketter (2015) describe business models for EV fleet owners that try to benefit from using EVs' capacities to trade energy in the wholesale market, while Kempton et al. (2013) map the area of creating new business models for supporting EV integration in the electricity grid. At an EV fleet level, Almuhtady et al. (2014) introduce a new battery-swapping policy model so that fleet owners can optimize their decisions about maintenance costs and make significant savings, while at the same time maintaining the *green character* of their fleet. From an infrastructure standpoint Avci et al. (2014) and Mak et al. (2013), describe optimal placement plans for charging poles so that the EV charging is facilitated properly to serve EV owners.

The previous solutions aim at coordinating the EV charging from a grid operators' point of view, so that the grid becomes balanced. We contribute to this discussion by trying to understand what will happen to the electricity grid, if every EV customer is represented by an *intelligent software agent* (March et al., 2000) that is responsible for charging the EV. Hence, we adopt a bottom-up approach and examine the overall effect of EV charging on the electricity grid. Specifically, we focus on individual preferences and customers' valuation of electricity consumption that are important factors for understanding individual EV charging behaviors. We propose the *Adaptive Management of EV Storage - AMEVS*, an IS artifact, implemented through a software agent, that takes the standpoint of the individual, and schedules charging and discharging so that preferences are satisfied. Our end goal is to examine how the adoption of such an artifact affects the electricity grid and provide energy policy recommendations. Therefore we aim to answer the following research question:

How will the electricity grid be affected by the adoption of IS artifacts (individual intelligent software agents) aiming at smart EV charging, given certain individual preferences?

Regarding individual preferences, we base our models on microeconomic theory focused on electricity consumption. Samadi et al. (2010) and Avci et al. (2014) provide the basic principles based on which they derive quadratic preferences for electricity consumption. Li et al. (2011) and Shao et al. (2011) use some generic forms of utility functions in similar environments, whereas Houthakker (1951) uses a linear transformation of electricity and other fuels (gas in this case) to express the utility consumers get from consuming energy. We draw from this literature and experiment with various types of utility functions which represent different preferences and consumption behavior.

By embedding these preferences into AMEVS, we can schedule EV charging in order to be beneficial for each individual EV owner while taking grid congestion into account. We evaluate our artifact in heterogeneous and homogenous preference populations. We show that if each individual adopts AMEVS rather than uncontrolled charging and discharging, not only is the customer better off, but the grid also benefits in terms of lower *peak-to-average ratio* (PAR), therefore lower costs. Furthermore, while evaluating AMEVS we observe that the impact of individual preferences and customer's valuation towards energy consumption is significant as the EV adoption rates grow. Therefore, these preferences should be accounted for in designing incentives for EV adoption and energy policy in general. We examine this impact in specific scenarios and provide respective policy recommendations for EV adoption.

2.2.2 Sustainability and Green IS

We design AMEVS artifact so that it supports environmental sustainability and is in accordance with the Green IS principles (Watson et al., 2010; Dao et al., 2011). We follow the *design science* research approach as presented by Hevner et al. (2004); Gregor and Hevner (2013); Walls et al. (1992). More specifically, with the presented IS artifact we aim to satisfy individual energy customers' preferences while at the same time we strive to reduce the negative impact of energy consumption to the electricity grid and the environment. The designed IS artifact is implemented in the form of an intelligent agent (March et al., 2000) that, by gathering information from its environment, tries to overcome the information overload that energy customers have to face in the smart grid (Amin and Woltenberg, 2005) of the future. Tilson et al. (2010) argue that an important and overlooked type of an IS artifact is a digital infrastructure. We are contributing to this literature by presenting the design of an IS infrastructure component for smart EV charging, which can have a transformational impact on consumer behavior and smart grid operations.

Energy informatics is a young and very important research area (Dedrick, 2010) with the focus on smart grid being one of the key concerns (Watson et al., 2010). We are building on the nascent Green IS literature (vom Brocke et al., 2013) and trying to boost the development of the Green IS field overall and energy informatics in particular. We focus on two of the questions identified by Melville (2010): creating effective design for developing information systems that influence human actions about the environment; and studying the association between information systems and organizational and sustainability performance. In similar terms, from the elements of an energy informatics framework

presented by Kossahl et al. (2012), we focus on construction of IT architecture, control systems and standards.

It terms of recent research, our work relates to the issue of carbon management systems (Corbett, 2013), as conceptually a carbon management system is a form of Green IT. By extension, systems for improving performance of smart grid are also valid examples of both organizational and personal Green IT artifacts. Some IS researchers explicitly consider EVs as a critical element of a smart grid: for example, Brandt et al. (2013) take a household view on the vehicle-to-grid capability.

According to Malhotra et al. (2013) employing Green IS to support environmental sustainability is expected to create significant societal impact. Specifically, they propose the following steps in creating effective green IS artifacts: *conceptualize, analyze, design, create impact*. Adopting these steps we conceptualize the problem of support for EV users in scheduling their EV charging, analyze the mobility and consumption data presented in section 2.4.1, design our Adaptive Management of EV Storage artifact (section 2.3) and observe its impact on the electricity grid (Section 2.4).

Since our end goal is to examine the effect of AMEVS adoption on the grid, we simulate numerous scenarios representing different electricity populations. Therefore, we build a simulation of the energy market and investigate the performance of AMEVS artifact. Following Kane and Alavi (2007) directives for building effective simulations, we specify our criteria for an effective simulation. For our simulation environment these are the high accuracy in learning the household consumption behavior, the ability to minimize individual costs, the prevention of herding behavior in energy consumption. All the aforementioned factors, if satisfied, will ensure an effective simulation environment for our evaluation.

2.3 Adaptive Management of EV Storage (AMEVS)

The Adaptive Management of EV Storage (AMEVS) is an IS artifact that acts on behalf of each individual who owns an EV, satisfying his/her preferences and driving profile. It is implemented in an intelligent agent (Wooldridge and Jennings, 1995) that consists of the following components (modules): an *input module*, a *learning module*, and an *optimization module* (Figure 2.1). The *input module* receives inputs related to the individual customer's behavior and creates a driving profile for that particular customer. The *learning module* is responsible for learning the household consumption profile using reinforcement learning (RL) (Sutton and Barto, 1998) trained on past consumption behavior (*offline*) and incorporating observed consumption behavior (*online*). This profile will serve as a basis

on top of which the EV charging will be scheduled. The proposed artifact will strive to fill the valleys and shave the peaks of the household profile by scheduling EV charging accordingly. This is in the interest of individuals, since they can reduce consumption (shave peaks) when prices are high, and increase consumption (fill valleys) when prices are low. Finally, the *optimization module* takes input from the other two modules and optimizes charging and discharging, ensuring customer mobility service whenever the customer needs the EV. In Figure 2.1 we present an overview of AMEVS which will later on be calibrated with specific data. Furthermore, in Figure 2.2 we present the sequence of activities taken by AMEVS in order to present the user with the optimal EV charging profile. In the Appendix we present a table of notation and a table of abbreviations, which explain all the parameters and variables of AMEVS.

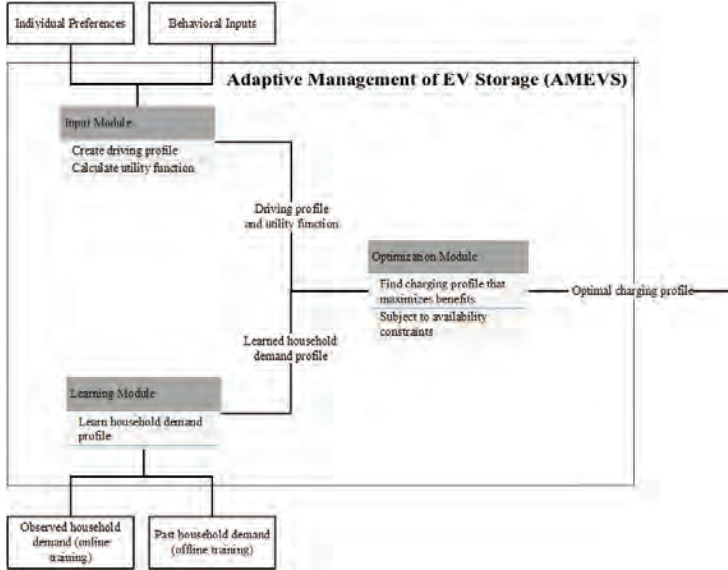


Figure 2.1: AMEVS Overview.

2.3.1 Model Assumptions

Below we list the assumptions of our artifact drawing the boundaries and the communication with its environment.

- We examine EV charging in combination with household consumption since both of these components constitute a customer's total electricity demand. Therefore,

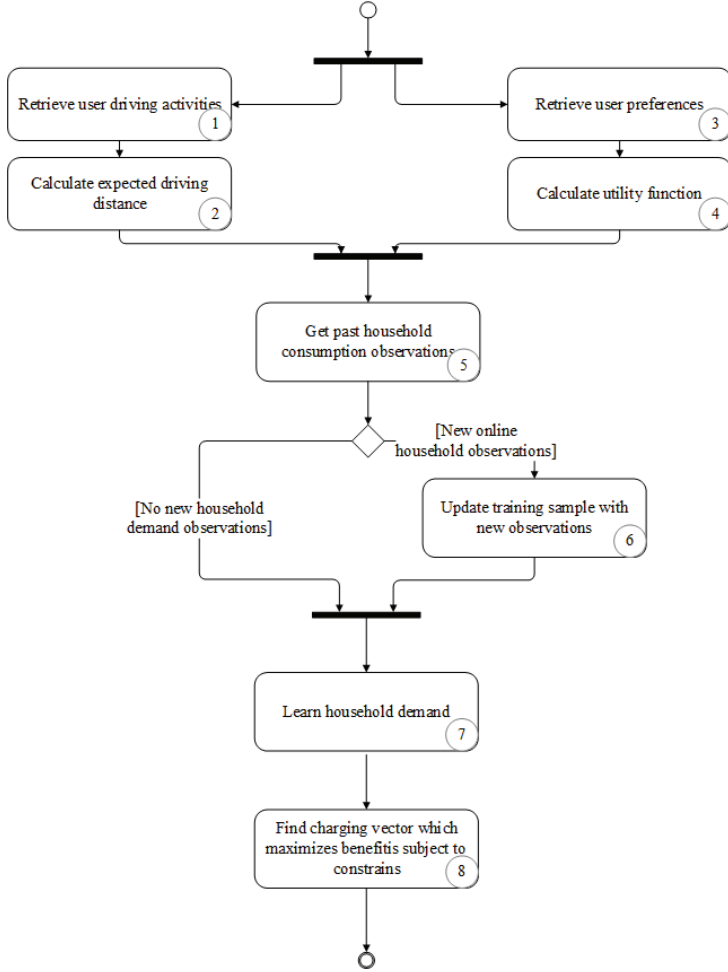


Figure 2.2: AMEVS activity diagram.

we assume the total electricity consumption to be the summation of EV charging demand and household electricity consumption.

- Each household owns one purely electric car. With this assumption, we focus our results on electrically powered cars and avoid situations where a household owns a second conventional car. This double car ownership would influence our outcome since household owners might substitute the conventional car for the electric one.

- Each EV owner can charge her car either at home or at her workplace, assuming that modern businesses allow for EV charging on their premises. However, the EV owner and not the company is billed for this charging.
- Since the consumers can charge at home or work, we assume Level 1 charging (3.3-3.6 kW charging speed).
- EV owners are exposed to variable prices that reflect availability of energy in the grid. Therefore, the prices serve as a signal to the EV owner for charging or not.
- Each EV owner has her own energy consumption valuation and this can be expressed through calibration of one of the utility functions presented in Section 2.3.5.
- All EVs drive within a region that has consistent energy prices.
- We assume Vehicle to Grid functionality to be in place (V2G). This allows customers to sell surplus energy back to the grid. For this process, we assume 3.6% distribution line and DC (EV battery) to AC (electricity grid) conversion losses as suggested by Reichert (2010). Since they buy and sell at the same price, taxes and VAT are included for both processes.
- The price for EV charging or discharging during a time period t is the same (the users can buy and sell electricity at the same price over a time period t).

2.3.2 Input Module

The *input module* gathers information about each EV owner, including driving behavior and individual preferences. All these inputs are described by the input set $\mathbf{I} = \{I_1, \dots, I_N\}$, which in our case represents all driving activities of an EV owner per day ($n \in \{1, N\}$ is the number of driving activities per day and per EV owner). I_n denotes the various driving activities per person. Each driving activity, I_n , is tied to a departure time t_d^n and an arrival time t_a^n . When $t \in \{t_d^n, t_a^n\} \ \forall n \in \{1, N\}$ then the EV customer is not available for charging. This is the step ① in Figure 2.2. From these activities, the expected driving distance will be calculated and a representative driving profile of each particular EV owner will be produced. This profile expresses the expected driving distance for this particular EV owner for each time slot t , $\mathbb{E}[Dist_t]$. In real-life these inputs can be recorded via a mobile application. Currently, we bootstrap the process with driving activities from our data sets (Section 2.4.1), but in future we plan to update them with recorded data from a mobile application experiment. So, the *input module* uses the following mapping

function $F(\cdot)$ to create the output $\mathbb{E}[Dist_t]$, which is the expected driving demand needed (per EV owner) for time slot t , step ② in Figure 2.2 (more details about deriving F in this particular data set are presented in Section 2.4.1):

$$\mathbb{E}[Dist_t] = F(I_1, \dots, I_N) \quad (2.1)$$

where the number of inputs, N , might change depending on the input data set. The expected driving distance $\mathbb{E}[Dist_t]$ for time t provides an estimation of the minimum amount needed to be charged, so that the customer has enough energy to drive.

We assume that customers own purely electric cars such as Nissan Leaf and Tesla S, and that they can charge them when they are at home and at work (“standard” charging with direct billing to the customer), which has been implemented by large businesses to encourage their employees to drive “green.” The model’s output is the EV charging demand at each point in time (*timeslot*) according to the inputs given. Charging must take place within the *charging envelope* shown in Figure 2.3. It displays the feasible charging region bounded by the minimum and maximum state of charge. Charging should end when the customer needs to use the car. Therefore, if the EV charges at nominal charging rate, the battery will fill up to a certain capacity lower or equal to the maximum state of charge ($C_{t,max}$) (depending on the start and end time). The slope of the max charge rate indicates the fastest charging speed that can be achieved while the car is plugged in.

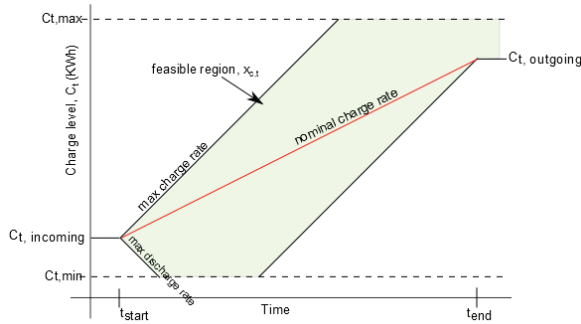


Figure 2.3: Electric vehicle *charging envelope*

Individual Utility and Consumer Benefit

Besides driving activities ($\mathbf{I} = \{I_1, \dots, I_N\}$), individual preferences (the way that EV owners value electricity) comprise the second input. Using these preferences, the input module

determines the individual utility function and consequently the utility the customer gets from energy consumption. Assuming that the total electricity consumption consists of two components: $x_{h,t}$, household demand (kWh) and $x_{c,t}$ demand from EV charging (kWh), we have individual utility from household and consumption and EV charging $U(x_{h,t}, x_{c,t})$. In Section 2.3.5, we present AMEVS after getting as inputs some common utility function families from the literature. According to consumer theory (Mas-Colell et al., 1995), the individual consumer benefit is defined as in Equation (2.2). This equation summarizes steps ③ and ④ in Figure 2.2.

$$W(x_{h,t}, x_{c,t}) = U(x_{h,t}, x_{c,t}) - (x_{h,t} + x_{c,t}) \cdot \hat{P}_t \quad (2.2)$$

where \hat{P}_t is the estimate of price per energy consumption unit (€/kWh). In this chapter, we assume that EV owners receive the demand they request from the grid, so $x_{h,t}$, and $x_{c,t}$ are only constrained by the capacity they have in their households and their EV batteries.

2.3.3 Learning Module

The artifact’s *learning module* uses reinforcement learning (RL) to learn the customers’ household consumption profile. RL is based on a *reward mechanism* that provides the algorithm with positive and negative rewards for preferred or non-preferred decisions, respectively. In our particular problem, the learning module explores a two-dimensional space with one dimension being the time $\mathbf{T} = \{t\}$ over which the household pattern is learned and the second dimension being the possible household consumption levels $\mathbf{L} = \{l\}$, shown in Figure 2.4. Each transition from a state to another is associated with a reward. The rewards are the highest for actions that lead to states that form the target policy (household consumption behavior). Since we want our learning module to learn the household consumption behavior of the particular EV owner it represents, we provide negative rewards for all the “wrong” actions and give a positive reward to the action we want our algorithm to choose. Consequently, the learning module searches the space for actions with the highest rewards. We randomly permute this reward matrix to prevent the algorithm from overfitting. More formally, the customer agent’s decision problem can

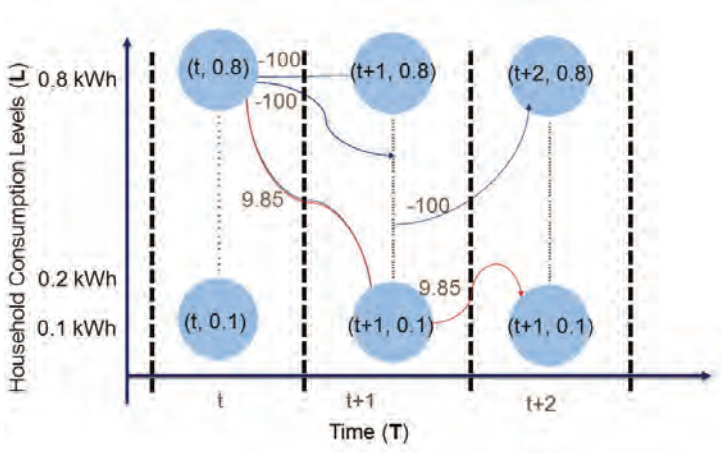


Figure 2.4: Learning space.

be modeled as the following Markov Decision Process (MDP) (Puterman, 1994):

$$\text{State Space } \mathbf{S} = \mathbf{T} \times \mathbf{L} = \{(t, l) | t \in \mathbf{T}, l \in \mathbf{L}\}$$

$$\text{Action Space } \mathbf{A} = \{A_i\}$$

$$\text{Rewards } R(t, l, A_i) = \begin{cases} W(l) & \text{if } l \text{ the desired household consumption for time } t \\ -\Theta & \text{else} \end{cases} \quad (2.3)$$

$\mathbf{T} = \{t\}$ is the set with the time intervals where $t \in \{1, \dots, T\}$ and $\mathbf{L} = \{l\}$ with $l \in \{1, \dots, L\}$ is the set with the consumption levels per hour, with L being the maximum household consumption level. The time slot we use is 1h, however this can be tailored to the practical needs of the artifact user without influencing the algorithm. The set \mathbf{L} is discretized at the level of 1 kWh. The learning rewards are the welfare of the household consumption level l , in case that this is the one that the algorithm needs to learn. Otherwise, the learning module receives a negative reward $(-\Theta)$, which is a sufficiently large number to discourage the module from taking this action again. The Action Space, $\mathbf{A} = \{A_i\}$ includes all transitions ($|L \times L|$ in total) from (t, l) to $(t+1, l) \forall t \in \{1, \dots, T\}$. The index i indicates each transition in the examined time horizon T .

The training of the learning algorithm is done on past household consumption observations (*offline training*, step ⑤ in Figure 2.2) and throughout the course of the simulation new data points are added (*online training*, step ⑥ in Figure 2.2). The customer agent has to learn the individual household consumption pattern through rewards $R(t, l, A_i)$ that are offered to it for each state (t, l) . The learning component strives to maximize

the total rewards accumulated, aiming for the desired household consumption levels. The optimal evaluation of the states gives the learned household consumption ($\hat{\mathbf{x}}_{\mathbf{h}}$ is a vector over the temporal dimension) (step ⑦ in Figure 2.2):

$$\hat{\mathbf{x}}_{\mathbf{h}} = \underset{\mathbf{x}_{\mathbf{h}}}{\operatorname{argmax}} \mathbb{E} \left\{ \sum_{l=1}^L \sum_{t=1}^T \gamma^{|l \times t|} R(t, l, A_i) \right\} \quad (2.4)$$

where $\gamma \in [0, 1]$ is the discount factor and practically expresses the weight of the previous state rewards. In Figure 2.4 we present a stylized example of how RL iterates over all possible states in our example, in order to find the optimal path (learned household consumption pattern). For this example we assume $\Theta = 100$ and $W(l, 0) = U(l, 0) - (l+0) \cdot \hat{P}_t$ (we examine only the household so $x_{c,t} = 0$ for this example), where $U(l, 0) = 100 \cdot l - 5 \cdot l^2$ (this is an example of the quadratic function family described in Section 2.3.5, the reader may refer to this section for more details) and $\hat{P}_t = 1$ monetary units per consumption l for each time t of this example.

For the *learning module* we selected RL over forecasting methods (ARIMA models, exponential smoothing, etc.) because RL is more flexible and can adapt easily to exogenous shocks in the household demand profile. With the reward function we can *steer* the artifact towards learning the household profile that represents the household, rather than being influenced by the most recent observations that many time-series models would do. Specifically, we compared RL with an ARMA model ($\text{ARMA}(1, 1) \times (1, 1)_{21}$) created by the same training set (it is non stationary, so no integrative components were required), an exponential smoother and a 6th degree polynomial regression estimator. RL showed an error of 6.8% compared to 55.8%, 18.5% and 39.6% of the ARMA, exponential smoother and polynomial regression, respectively.

2.3.4 Optimization Module

Once the individual household consumption estimate $\hat{x}_{h,t}$ and the individual driving profile $\mathbb{E}[\text{Dist}_t]$ are known, the *optimization module* schedules the EV charging so that the individual benefit is maximized. For this process, the household consumption pattern ($\hat{\mathbf{x}}_{\mathbf{h}}$) is required, Equation (2.4), from the learning module so that the optimization module schedules the EV charging on top of that. It aims to ensure minimal energy cost for the individual customer accounting both for household consumption ($\hat{\mathbf{x}}_{\mathbf{h}}$) and EV charging ($\mathbf{x}_{\mathbf{c}}^*$) costs. Therefore, the IS artifact needs to have already learned the household consumption, so that the EV charging can be optimized on top of the learned household consumption. The variable $\hat{\mathbf{x}}_{\mathbf{h}}$ contains stochasticity since it is dependent on past obser-

vations of household demand. For each planning horizon T , the customer agent calculates the charging vector \mathbf{x}_c^* based on (2.5). Since the model depends on exogenous stochastic inputs such as the learned household profile ($\hat{\mathbf{x}}_h$) and the estimated driving distance per hour ($\mathbb{E}[Dist_t]$), it optimizes the expected individual benefit $W(\cdot)$, obtained by the agent over time T (step ⑧ in Figure 2.2):

$$\mathbf{x}_c^* = \underset{\mathbf{x}_c}{\operatorname{argmax}} \mathbb{E} \left\{ \sum_{t=1}^T W(\hat{x}_{h,t}, x_{c,t}) \right\} \quad (2.5)$$

subject to the constraints (2.6), (2.7), (2.8):

$$-X_{max,t} \leq x_{c,t} \leq X_{max,t} \quad \forall t \in \{1, \dots, T\} \quad (2.6)$$

Constraint (2.6) ensures that the charging power is within the range allowed by the grid and the charger. The upper bound $X_{max,t}$ represents the maximum power that the customer agent can charge from the network per time slot t and is the same absolute value as the discharging power. The negative sign ($-X_{max,t}$) indicates discharging back to the grid (V2G). This represents the main network constraint and is dependent on the characteristics of the residential connection. Therefore, the charge and discharge rates are equal since they depend on the electricity connection characteristics (assuming that EV owners charge and discharge on electricity connections with same characteristics).

$$x_{c,t} = C_t - C_{t-1} + \mathbb{E}[Dist_t] \cdot \rho \quad \forall t \in \{1, \dots, T\} \quad (2.7)$$

$$C_0 = SoC_{min} \quad (2.8)$$

where C_t is the state of charge at timeslot t , ρ is the capacity/distance rate given by specifications of the automotive industry, and SoC_{min} is the minimum allowed state of charge that does not damage the battery's lifetime.

The prices \hat{P}_t at each hour are predicted by the intelligent agent using a moving window of the previous seven days, averaged over each hour respectively. This method maintains the price variation over a daily horizon, being influenced by the previous week's observations which have higher correlation with the current week's observations. Table 2.1 presents the general formulation of AMEVS algorithm in pseudo-code.

Table 2.1: Adaptive Management of EV Storage - AMEVS

AMEVS pseudo-code	
1	Initialization
2	Calculate $\mathbb{E}[Dist_t] = F(I_1, \dots, I_N)$ from the input set $\mathbf{I} = \{I_1, \dots, I_N\}$
3	Receive preference input $U(\cdot)$
4	Calculate $W(x_{h,t}, x_{c,t}) = U(x_{h,t}, x_{c,t}) - (x_{h,t} + x_{c,t}) \cdot \hat{P}_t$
5	for each $(t, l, A_i) \in \mathbf{T} \times \mathbf{L} \times \mathbf{A}$
6	Calculate reward as: $R(t, l, A_i) = \begin{cases} W(l) & \text{if } l \text{ the desired household consumption for time } t \\ -\Theta & \text{else} \end{cases}$
7	end for
8	Calculate household demand profile as: $\hat{\mathbf{x}}_h = \operatorname{argmax}_{\mathbf{x}_h} \mathbb{E}\{\sum_{l=1}^L \sum_{t=1}^T \gamma^{ l \times t } R(t, l, A_i)\}$
9	For horizon T calculate optimal charging vector as: $\mathbf{x}_c^* = \operatorname{argmax}_{\mathbf{x}_c} \mathbb{E}\{\sum_{t=1}^T W(\hat{x}_{h,t}, x_{c,t})\}$
10	Subject to constraints:
11	(1) $-Xmax_{x,t} \leq x_{c,t} \leq Xmax_{x,t} \quad \forall t \in \{1, \dots, T\}$ AND
12	(2) $x_{c,t} = C_t - C_{t-1} + \mathbb{E}[Dist_t] \cdot \rho \quad \forall t \in \{1, \dots, T\}$ AND
13	(3) $C_0 = SoC_{min}$
14	return \mathbf{x}_c^*

2.3.5 AMEVS Variations based on Energy Consumption Utility

Due to lack of empirical studies on the measurement of energy preferences in the EV context, we experiment with some of the most commonly used utility functions in the literature and compare the results. These utility functions are approximations to help us derive concrete results.

LinAMEVS

Firstly, following the functional form proposed by Houthakker (1951), we assume that the customer's utility towards energy consumption ($x_{h,t}$) and EV charging consumption ($x_{c,t}$) is linear, both getting different weights depending on the customer but also depending on time (Equation (2.9)).

$$U(x_{h,t}, x_{c,t}) = \beta_t \cdot x_{h,t} + \eta_t \cdot x_{c,t} \quad (2.9)$$

where $\beta_t + \eta_t = 1$. This utility factor represents customers that suffer from higher range anxiety since their marginal utility is constant until they reach the desired amount of electricity. They get equal satisfaction (constant marginal utility) for each extra consumption unit until they reach the desired consumption target. Using (2.2), (2.5), (2.5) and $\beta_t = 1 - \eta_t$ the charging vector becomes:

$$\mathbf{x}_c^* = \operatorname{argmax}_{\mathbf{x}_c} \mathbb{E}\left\{\sum_{t=1}^T \{(1 - \eta_t - \hat{P}_t) \cdot \hat{x}_{h,t} + (\eta_t - \hat{P}_t) \cdot x_{c,t}\}\right\} \quad (2.10)$$

subject to constraints (2.6), (2.7), (2.8). We call this variation of the algorithm LinAMEVS and it represents the higher range anxiety customers whose marginal utility is constant until they reach the desired consumption amount.

QuadAMEVS

Secondly, instead of a linear relationship between utility and energy consumption, we assume quadratic utility along the lines of the utility functions proposed by Samadi et al. (2010), Fahrioglu and Alvarado (2000) and Avci et al. (2014):

$$U(x_{h,t}, x_{c,t}) = \begin{cases} \omega \cdot (x_{h,t} + x_{c,t}) - \frac{\alpha}{2} \cdot (x_{h,t} + x_{c,t})^2 & , 0 \leq (x_{h,t} + x_{c,t}) \leq \frac{\omega}{\alpha} \\ \frac{\omega^2}{2\alpha} & , (x_{h,t} + x_{c,t}) > \frac{\omega}{\alpha} \end{cases} \quad (2.11)$$

where ω , stands for the *level of satisfaction obtained by a user as a function of his/her energy consumption* and varies among customers. The variable α (assuming ω is constant) indicates for how much consumption $x_{h,t} + x_{c,t}$ the utility function will be saturated. In other words, it shows the boundary after which no utility will be gained and it is inversely analogous to the total consumption needed to saturate the utility function. In our analysis, as in (Fahrioglu and Alvarado, 2000), we also use the value 0.5 since for the ω spectrum that we will investigate in the sensitivity analysis it gives relatively high saturation bounds (from 10kWh to 100kWh). We want to have high saturation bounds so that we get the full solution area of the optimization algorithm. Low saturation bounds (i.e., high α) will lead to losing part of the feasible area of the solution and losing part of the optimal charging paths.

Customers with quadratic valuations towards energy consumption have lower range anxiety since their marginal utility decreases until they reach the desired goal (i.e., desired amount of electricity they need to drive for the coming day). And the charging vector using (2.2), (2.5) and (2.11) now becomes:

$$\mathbf{x}_c^* = \underset{\mathbf{x}_c}{\operatorname{argmax}} \mathbb{E} \left\{ \sum_{t=1}^T \left\{ (\omega - \alpha \cdot \hat{x}_{h,t} - \hat{P}_t) \cdot x_{c,t} - \frac{\alpha}{2} \cdot x_{c,t}^2 \right\} \right\} \quad (2.12)$$

We call this variation of the algorithm QuadAMEVS. For our simulation experiments, we create populations with all possible parameter combinations in Equations (2.9) and (2.12) so that all possible cases are represented in our sample. Both utility functions employed are displayed in Figure 2.5. The quadratic function reaches the maximum goal, with

decreasing marginal valuation, whereas the linear utility function reaches the maximum with constant marginal valuation.




Figure 2.5: Quadratic and linear utility function examples

2.4 Experimental Evaluation

We evaluate AMEVS in different population types and examine its effect on the individual and on the aggregate demand curve. We see that the adoption of AMEVS by all customers leads to peak demand and price reduction in the market, just by satisfying each individual EV owner’s objective. This means that AMEVS achieves an implicit coordination of charging without the presence of an actual coordinator. Furthermore, we examine how AMEVS influences the EV charging landscape as a function of the EV ownership penetration.

2.4.1 Simulation Environment and Data Description

Evaluation is a critical step in the design of an IS artifact (Hevner et al., 2004), therefore we build a simulation environment that approximates the conditions of an energy market, where EV owners have to purchase electricity to cover their household consumption needs and to charge their EVs. We create our simulation based on Power Trading Agent Competition (Power TAC) (Ketter et al., 2015) software platform since it is a well proven smart grid simulation.

Household Consumption

Our experimental setting consists of diverse EV customer populations (see Numerical Results) whose household consumption comes from data provided by a European Energy Utility. The dataset includes 15-minute household consumption information from

the Netherlands, aggregated in hourly intervals to comply with the granularity of our analysis which is $1h$. The histogram of our household consumption data set is displayed in Figure 2.6. We observe that the maximum power consumption a household can have in this data is $3kW$ with quite low frequency of observations. The peak demand might even be higher than $3kW$ for short periods of seconds, and the value $3kW$ stands for the hourly average of this peak. However, despite this low frequency, the value of peak demand determines the capacity specifications of the distribution grid. Therefore, energy policy makers should try to mitigate this peak demand in order to reduce capacity investments and increase sustainability.




Figure 2.6: Histogram of daily household consumption data (source: European energy utility company)

Figure 2.7 displays the average daily profiles of some typical households of our data set (anonymized because of privacy). We see that the peak demand value varies across households, while in most households peak hours seem to be similar. Therefore, reallocating charging demand to different time intervals will be promising for flattening the demand curve.




Figure 2.7: Examples of average daily energy demand

Driving Profiles

To calibrate AMEVS with driving profiles we use data provided by the Dutch Bureau of Statistics⁵ (CBS). This particular data set provides us with average driving distances of individual drivers based on their driving activities, depending on the customer segment they belong to. Therefore, the input module receives the input set $\mathbf{I} = \{I_1, \dots, I_N\}$ with the various activities the driver performs per day (work, shopping, business trips, visits, leisure activities, school). A driver might be a part-time employee, full-time employee, student, unemployed, retired and depending on the segment she belongs to, the different driving activities are associated with different driving distances.

Driving Episodes Using the outputs of Equation (2.1) the input module simulates driving episodes which relate to EV owners driving as a result of particular activities (work, leisure, etc.). From the driving episodes AMEVS can estimate probabilities of departure $p_{t_d}^{I_n}$ and arrival $p_{t_a}^{I_n}$ for each corresponding activity I_n . These probabilities accompanied with the estimated driving distance $\mathbb{E}[Dist_t]$ comprise the tuple $(\mathbb{E}[Dist_t], p_{t_d}^{I_n}, p_{t_a}^{I_n})$ which we bootstrap our simulation with, to create realistic driving conditions. Figure 2.8 shows an aggregation of 1000 driving episodes averaged over a day. We observe a morning peak until 10.00 and an afternoon peak starting from 13.00 and decreasing until 20.00 when most commuters have returned home.




Figure 2.8: Driving demand data (source: Central Bureau of Statistics, Netherlands)

Pricing Schemes

We use two pricing schemes to calibrate AMEVS. First we use the EPEX SPOT prices (Figure 2.9)⁶ since we want to create a variable pricing scheme, where the price difference indicates the energy availability. However, all models can be trained on different data sets

⁵www.cbs.nl [Date Accessed: March 22nd, 2016]

⁶<http://www.epexspot.com/en/market-data/intraday/intraday-table/> [Date Accessed: March 22nd, 2016]

(e.g., US mobility, pricing data) to examine effects on different populations. Our basic assumption is that the EV customers interact with the energy market through an energy provider (broker (Peters et al., 2013)) and buy energy from the market to cover both their household and their EV charging needs. Figure 2.9 depicts a typical weekly price curve. From EPEX data we derive a demand-price relationship ($P = f(D)$) that we use to see how changes in the demand influence prices.




Figure 2.9: Typical weekly price curve (source: EPEX SPOT)

Secondly, we use a time-of-use (TOU) pricing scheme which is a three-part tariff currently in use in California⁷. This is one of the first multi-tier tariff designed specifically for retail use and in particular for EVs. Therefore, we are using this as an exemplary pricing scheme in our simulations. We have converted the currency to reflect European prices. We select the median tariff: €0.38 for Peak [14:00-21:00], €0.20 for Part Peak [7:00-14:00] and [21:00-23:00], and €0.10 for Off-peak [23:00-7:00]. This pricing scheme is displayed in Figure 2.10 and represents a more predictable three part energy tariff. AMEVS facing this scheme does not require estimate of prices, since this is a scheme announced in advance to consumers. Therefore, in this case we have $\hat{P}_t = P_t$. Comparing these two pricing schemes we can see the effect of using AMEVS in dynamic and static pricing regimes.

2.4.2 Benchmarks

In order to derive insights about AMEVS performance in real world scenarios, we compare it with commonly used benchmarks. Firstly, we evaluate AMEVS compared to observed charging data derived from the Netherlands during 2013, capturing the charging behavior of all the EV owners in the country. Furthermore, we compare our algorithm to charging

⁷<http://www.pge.com/en/myhome/environment/whatyoucando/electricdrivevehicles/rateoptions/index.page>
[Date Accessed: March 22nd, 2016]




Figure 2.10: Time-of-Use pricing scheme (source: Pacific Gas and Electric Company)

algorithms representing various decision criteria (heuristic, cost minimizing, etc.). In all these charging algorithms we make the same assumptions as in AMEVS regarding the charging speed and the electricity prices for buying and selling. Namely, the consumers can charge at home and at work, so we assume Level 1 charging. Furthermore, we assume the same price for buying and selling energy, since it will not be sustainable to have a business model that always favors selling energy to the grid (as presented in the modeling assumptions).

Observed Charging Data

To quantify the observed real-world charging, we collected data from the Netherlands during 2013. This data set includes EV charging transactions starting from January 11th, 2013 to December 31st, 2013. In total, it represents 1500 EV owners and 231,995 charging transactions with the grid. The mean and standard deviation of the daily charging demand is shown in the boxplot diagram of Figure 2.11. The charging varies from low speed charging (3kW) to fast charging (25kW). We see that most of the customers charge their cars from 9:00 to 14:00-15:00 which indicates charging at work. Another peak occurs between 19:00 and 20:00 when they have returned home. We suspect that this charging behavior includes redundant charging attributed to the customers *range anxiety* (Franke et al., 2011) since there is no control related to how much and when customers charge. This benchmark serves as our baseline, reflecting the current EV charging situation.

Naive Charging

This benchmark reflects the naive charging behavior EV owners have at the moment: they recharge the EV whenever possible (mostly at night when back home from work) so that they have enough battery to drive. We call this Naive charging (Table 2.2) and we assume the charging is conducted by an intelligent agent representing the customer. The




Figure 2.11: Observed charging: average demand and variability

agent acts based on the logic that when the EV battery is not full and the EV owner is available for charging, the EV battery should be charged.

The basic difference between this benchmark and the observed charging is that the observed data come from the actual behavior that the EV owners currently have. At the moment EV owners mostly charge during the day since it is more convenient for the them to charge at work. On the other hand, naive charging assumes also a lot of charging during the night, since the EV charging is conducted by the software agent. Therefore, it is important to have both benchmarks, since we want to examine the effect on the electricity grid, when EV customers start to charge in different locations and not only at their work premises. The naive charging is expected to differ from the observed data during the night period. As we will see in the numerical results section, this hypothesis is confirmed.

Table 2.2: Charging Benchmark 2

Naive Charging	
1	Initialization
2	for each $t \in \{1, ..T\}$
3	Calculate $CA_t, \mathbb{E}[Dist_t]$
4	if $CA_t == TRUE$ & $C_t < \mathbb{E}[Dist_t] \cdot \rho$
5	$D_t = x_{h,t} + x_{c,t}$
6	endif
7	endfor
8	return D

Here, CA_t is the charging availability vector ($\forall t \in \{1, ..T\}$), $\mathbb{E}[Dist_t] \cdot \rho$ is the expected capacity needed for driving up to timeslot t , $\mathbf{D} = \{D_t\}$ is the total demand vector and $\mathbf{x}_h, \mathbf{x}_c$ are the household and charging demand vectors over time, respectively.

Heuristic Charging

We use a Heuristic Charging approach as our third benchmark. Here, the agent predicts the prices over a time horizon T , using a simple moving average model. Assuming that \hat{P}_t stands for the energy price estimate per kWh and $x_{c,t}$ for the charging demand at the timeslot t , the customer agent acts based on the following heuristic: if $\hat{P}_t \leq \hat{P}_{t+1}$ charge the pre-scheduled amount, resulting from the behavioral model, otherwise split the charging demand (i.e. the respective charging time) evenly to the time horizon T . The variable $X_{max,t}$ stands for the maximum amount that can be charged from the grid per time slot t . The Heuristic Benchmark is described in Table 2.3. More specifically, here we use the myopic approach of this heuristic with $T = 2$ to compare AMEVS which uses weekly planning $T = 168$ with a totally myopic benchmark.

Table 2.3: Charging Benchmark 3

Heuristic Charging - HC	
1	Initialization
2	for each $t \in \{1, ..T\}$
3	Calculate $CA_t, \mathbb{E}[Dist_t]$
4	if $\hat{P}_t \leq \hat{P}_{t+1}$ & $C_t < \mathbb{E}[Dist_t] \cdot \rho$
5	$x_{c,t} = X_{max,t}$
6	else
7	$x_{c,t} = \frac{X_{max,t}}{T}$ and $x_{c,t+T-1} = \frac{X_{max,t}}{T} + X_{max,t+T-1}$
8	endif
9	$D_t = x_{h,t} + x_{c,t}$
10	endfor
11	return D

Cost Minimization Charging

The cost minimization approach has been used in the literature (He et al., 2012; Halvgaard et al., 2012) to solve the EV charging problem. We use it here to benchmark our approach and to measure the outcomes of the comparison. This charging results from the customer agent's goal to minimize costs without accounting for individual preferences. To make a fair comparison we assume the same constraints as in AMEVS and the following optimization mechanism as a decision rule:

$$\mathbf{x}_c^* = \underset{\mathbf{x}_c}{\operatorname{argmin}} \mathbb{E} \left\{ \sum_{t=1}^T \hat{P}_t \cdot (\hat{x}_{h,t} + x_{c,t}) \right\} \quad (2.13)$$

with the input $\hat{x}_{h,t}$ coming from the learning module and subject to the following constraints:

$$-X_{max,t} \leq x_{c,t} \leq X_{max,t} \quad \forall t \in \{1, ..T\} \quad (2.14)$$

Similar to AMEVS, here constraint (2.14) ensures that the charging demand per hour (charging speed) is within the range allowed by the grid. The upper bound $X_{max,t}$ represents the maximum power that the customer agent can charge from the network per timeslot t and is the same absolute value as the discharging power. The negative sign ($-X_{max,t}$) indicates discharging back to the grid (V2G). This represents the main network constraint and is dependent on the characteristics of the residential connection.

$$x_{c,t} = C_t - C_{t-1} + \mathbb{E}[Dist_t] \cdot \rho \quad \forall t \in \{1, ..T\} \quad (2.15)$$

$$C_0 = SoC_{min} \quad (2.16)$$

where C_t is the state of charge at timeslot t , and ρ is the capacity/distance rate given by specifications of the automotive industry and SoC_{min} is the minimum allowed state of charge that does not destroy the battery's lifetime.

2.4.3 Numerical Results

We examine the performance of AMEVS compared to the other charging benchmarks with respect to the ability to reduce energy peaks and energy prices (peak and average). Our goal is to observe how AMEVS reduces peak demand by satisfying the individual objectives. We start our experiments with homogeneous populations, for example, populations with customers that all have high range anxiety (LinAMEVS) or low range anxiety (QuadAMEVS). In Section 2.5, we simulate scenarios that reflect heterogeneous populations, where EV owners represent particular penetration rates in the total population (low, medium, and high). For demonstration purposes, we use the parametrization of the algorithm displayed in Table 2.4. We use values that are used in the literature or are those that demonstrate insightful results of the algorithm. More detailed sensitivity analysis is presented at the end of this section. For AMEVS planning horizon T , we assumed a weekly planning ($T = 168h$) since a person's driving activities may vary in a week, but the driving behavior is usually repeated across weeks. For LinAMEVS, we initially experiment with equal weights of EV charging and household demand on the utility function ($\beta_t = \eta_t = 0.5$) and later in the sensitivity analysis, we show how the change of weights influences the outcome. For QuadAMEVS parameter α , we use a sufficiently low value that does not restrict our feasible solution area and we do not lose solutions (as explained

in Section 2.3.5). For QuadAMEVS parameter ω , we noticed significant changes as ω changes, and therefore we simulated a large population where values in the spectrum of $\omega \in [0.1 - 100]$ are represented.

Table 2.4: AMEVS Parametrization

Parameter	Short Description	Value
T	AMEVS planning horizon	168
β_t	weight of household consumption in the linear utility function	0.5
η_t	weight of EV charging in the linear utility function	0.5
α	coefficient of the quadratic utility function	0.5
ω	coefficient of the quadratic utility function	[0.1 - 100]

Peak Demand

Simulation experiments show that AMEVS can reshape an individual demand curve. Therefore, we examine its effect on peak demand in a population of 10^7 customer agents (all EV owners). We chose a large number of EV owners to see how large customer populations behave with the adoption of our artifact. The average individual demand curve is reshaped as shown in Figure 2.12. We observe that the customers adopting QuadAMEVS tend to consume more when prices are low (e.g., 00:00-05:00) and sell energy back to the grid when prices are high (e.g., 16:00-20:00) (assuming the same price for buying and selling energy back to the grid). As a result, the individual curve is reshaped and has fewer peaks and lower volatility. Furthermore, we observe that using QuadAMEVS brings peak demand reduction not only compared with the naive charging (Figure 2.12) but also compared with the case where no EVs exist on the smart grid. This is attributed to the storage features of EVs that buffer part of the demand to lower demand periods, flattening the demand curve. This creates a strong incentive for EV adoption. In a population where customers are conventional car owners, policymakers should incentivize the EV adoption against a conventional car, if they identify low range anxiety among the population. These incentives will make the demand curve smoother, yielding benefits for both individual customers and the distribution grid.

In contrast, although LinAMEVS shifts demand peaks to earlier slots, it creates the same volatility. Consequently, after adopting LinAMEVS for some time, the energy and price peaks are shifted to earlier time periods. At first glance, this seems to have a negative effect. However, in a heterogeneous population where both customers with linear preferences exist (LinAMEVS) together with other customers who might use naive charging or charge without any control, the presence of LinAMEVS will offset peaks, yielding a




Figure 2.12: Left Pane: Individual demand curve of no EV charging, AMEVS and Naive Charging. Right Pane: Individual demand curve of no EV charging, AMEVS and Observed Charging

less volatile aggregate demand curve. This effect of heterogeneous populations is exemplified in Figure 2.13 and analyzed in Section 2.5. In Figure 2.13 we see that if we combine an EV owner adopting LinAMEVS artifact and an EV owner using Naive charging, we get an average across the combined population which is significantly smoother and, thus, highly beneficial for the smart grid.

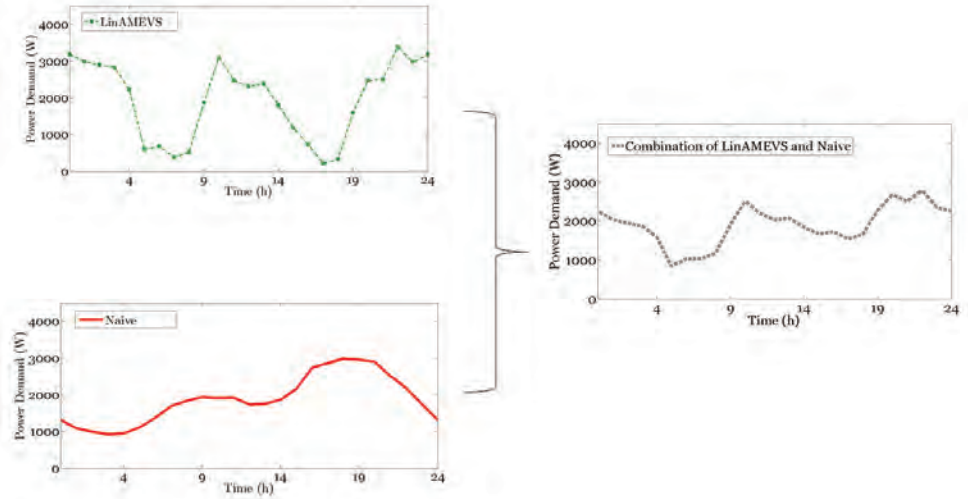


Figure 2.13: Combination of LinAMEVS and Naive Charging: Less volatile average combined demand curve

Finally, MixedAMEVS assumes both customers with linear and quadratic utility with regards to energy. It shows less volatility compared to LinAMEVS, reducing the peaks by a small amount, but compared to QuadAMEVS it still performs worse. In Figure 2.14, we compare QuadAMEVS with the naive charging not only based on the average steady

state result, but also based on the worst case scenario (outlined by the upper bound of the error bars in Figure 2.14). We observe that even the worst case volatility scenario of QuadAMEVS flattens the demand curve compared to Naive Charging.




Figure 2.14: Individual demand curve variability: QuadAMEVS and Naive Charging

In Figures 2.15-2.16, we illustrate the average individual power demand curve resulting from applying AMEVS in Time-of-Use (TOU) pricing regimes. We observe that, in general, the individual demand curve is more volatile, which is an immediate result of the energy prices in Figure 2.9. Since the prices show some variability but also some plateaus, EV owners trying to benefit from the price changes, show more volatility in their power demand during the change periods. In particular, customers with lower range anxiety (QuadAMEVS) have a more volatile demand curve under TOU pricing than under Real-Time-Pricing (RTP). Similarly, the demand curve for customers with high range anxiety (LinAMEVS) is more volatile when they are exposed to TOU than to RTP regimes. Overall, we see that the demand curve is more volatile compared to the scenario when they face RTP.




Figure 2.15: Left Pane: Individual demand curve QuadAMEVS exposed to RTP and TOU. Right Pane: LinAMEVS exposed to RTP and TOU

From Figures 2.15-2.16 we also conclude that TOU pricing yields no peak demand reduction compared to the scenario where no EVs exist in the market. This means that




Figure 2.16: Individual demand curve: MixedAMEVS exposed to RTP and TOU

energy policymakers have to apply RTP if they want to create higher incentives for EV adoption and also higher benefits for the distribution grid.

To provide a more complete comparison (Table 2.5), we use the peak-to-average ratio (PAR) reduction ($PAR = \frac{x_{peak}}{x_{rms}} = \frac{x_{peak}}{\sqrt{\frac{1}{T} \sum_{t=1}^T x_t^2}}$) and the peak demand reduction. PAR is also known as “crest factor” and indicates how extreme the peaks are in a waveform. PAR reduction is important because much of the cost of energy supply is driven by peak demand. This metric, besides academic literature, is used by the US Energy Information Administration (EIA)⁸ to measure the effect of peaks on power demand. In order to have higher sustainability in the electricity grid, we need a lower PAR.

Table 2.5: Energy peak reduction - AMEVS RTP and AMEVS TOU

			AMEVS RTP		AMEVS TOU	
			PAR red.(%)	Peak red.(%)	PAR red.(%)	Peak red.(%)
LinAMEVS	vs.	Naive	−1.6	−13.8	−16.7	−48.3
QuadAMEVS	vs.	Naive	15.4	24.9	14.2	27.2
MixedAMEVS	vs.	Naive	8.7	11.6	−3.9	2.6
LinAMEVS	vs.	Observed	−0.3	−4.6	−17.4	−35.0
QuadAMEVS	vs.	Observed	16.5	30.9	15.3	33.8
MixedAMEVS	vs.	Observed	9.9	18.7	−2.5	10.4
LinAMEVS	vs.	Cost Min.	6.7	15.1	26.9	31.1
QuadAMEVS	vs.	Cost Min.	22.3	44.0	21.1	45.7
MixedAMEVS	vs.	Cost Min.	16.1	34.1	2.4	20.9
LinAMEVS	vs.	no EVs	2.2	−34.0	−11.8	−58.8
QuadAMEVS	vs.	no EVs	18.5	11.5	17.3	14.2
MixedAMEVS	vs.	no EVs	12.3	−4.1	−0.1	−14.8

In Table 2.5, we observe that the linear variation of the algorithm, which expresses customers with high range anxiety, does not perform well compared to naive and observed charging. This assumes a population in which everybody adopts LinAMEVS (homogeneous population). Later on we will examine the effect of LinAMEVS in lower adoption rates. It is interesting to mention that LinAMEVS performs significantly better than

⁸<http://www.eia.gov/todayinenergy/detail.cfm?id=15051> [Date Accessed: March 22nd, 2016]

the cost minimization benchmark, which could be considered as a suitable mechanism to solve the EV charging problem. Cost minimization brings less peak reduction compared to AMEVS because it is solely driven by price changes, creating new peaks in periods when prices are lower. This effect is found in the literature as “avalanche effect” or herding behavior (Gottwalt et al., 2011). This term captures the concurrent shifting of behavior to the same time periods (low price periods) creating new peaks in the demand, just during other time intervals. We can assume that our algorithm creates less herding compared to standard cost minimization approaches that are driven primarily by differences in energy prices. Our algorithm shifts behavior of individuals to different time intervals based on their individual preferences, mitigating the chance that they will coincide and create new peaks. Comparing the left and right pane of Table 2.5 (where AMEVS is evaluated with dynamic and static time-of-use prices), we see that there are more negative PAR reductions (i.e., PAR increases) in the TOU regime. By applying TOU pricing the power demand curve becomes less smooth and increases PAR. This is not desirable for the distribution grid either, because it increases volatility and necessitates the use of expensive power reserves in order to cover this volatile demand without any disruptions.

Load Factor

Load factor (LF), as introduced by Watkins (1915), is a metric that allows us to measure the volatility of the power demand curve as produced by AMEVS in each scenario. It is defined by Equation (2.17) as the fraction of the mean power demand over a time period (time horizon T in our simulation) divided by the peak demand over this period. Therefore, if LF is closer to 1, it means that the power demand is not very volatile since the peak and the average are not significantly different. On a broader scale, high LF means higher capacity utilization in the electricity grid which indicates higher levels of sustainability (Strbac, 2008).

$$LF = \frac{\frac{1}{T} \sum_{t=1}^T x_t}{x_{peak}} \quad (2.17)$$

LF receives values in the interval $LF \in [0, 1]$ and the higher it gets, the less volatile the power demand is. Therefore, a high load factor is more desirable because it shows a relatively constant demand without extreme peaks. Low load factor indicates that there is capacity that remains idle for long periods and gets started only in short time intervals to cover peaks that might occur. This creates extra expenses because power sources that can be ramped up quickly are mostly powered by expensive fuels.

In Table 2.6 we show the effect of AMEVS to the LF both under dynamic (RTP) and static (TOU) pricing regimes. We see that the effect of AMEVS is mostly beneficial for the

Table 2.6: Load Factor Increase AMEVS RTP and AMEVS TOU

			AMEVS RTP	AMEVS TOU
			LF inc.(%)	LF inc.(%)
LinAMEVS	vs.	Naive	-8.2	-30.9
QuadAMEVS	vs.	Naive	24.0	21.9
MixedAMEVS	vs.	Naive	11.7	-6.4
LinAMEVS	vs.	Observed	-7.4	-30.3
QuadAMEVS	vs.	Observed	25.2	23.1
MixedAMEVS	vs.	Observed	12.8	-5.8
LinAMEVS	vs.	Cost Min.	58.6	14.8
QuadAMEVS	vs.	Cost Min.	114.3	110.7
MixedAMEVS	vs.	Cost Min.	93.1	61.3
LinAMEVS	vs.	no EVs	-3.1	-27.4
QuadAMEVS	vs.	no EVs	31.0	28.9
MixedAMEVS	vs.	no EVs	18.0	-1.4

grid's stability since in most scenarios it increases load factor and thus makes the power demand less volatile. However, when populations with high range anxiety (LinAMEVS) dominate the demand, they make it more volatile. Therefore, we have evidence that EV adoption above certain percentages, should be discouraged in populations with high range anxiety, so that the high range anxiety EV owners are sufficient to offset peaks created by other EV owners, but without dominating the market and creating extra peaks. It is notable that cost minimization charging leads to lower LF compared to AMEVS in all situations. This results from the highly sensitive nature of the cost-minimizing algorithm to price changes. Therefore, even though cost minimization might yield financial benefits for individuals (mostly short-term), it is not beneficial for supporting grid's stability or for promoting sustainability. In a TOU pricing regime, we see (AMEVS TOU column) more negative effects for the distribution grid, since the power demand curves become more volatile and consequently more unstable. As a consequence, expensive power reserves need to be used, leading to higher electricity costs and usually to lower sustainability (power reserves are mostly fueled by oil).

Energy Price

An immediate result of the previous figures is that besides peak demand, the average energy price also decreases. Consequently, this price reduction is diffused in the market because of the demand shift and peak reduction. To calculate the exact numbers we use the price demand relationship derived by the EPEX data. In Figure 2.17, we show this reduction for a scenario with values in the whole spectrum of ω , showing a maximum of 38% at 100% QuadAMEVS adoption (against Naive). We also show that a penetra-

tion above 40% of LinAMEVS has a negative effect on the grid. Therefore, assuming a population of consumers with a linear utility function, penetrations higher than 40% should be discouraged, since they would necessitate constructing additional infrastructure to accommodate the new increased peaks.




Figure 2.17: Left Pane: average price reduction–AMEVS vs. Naive Charging. Right Pane: Average price reduction–AMEVS vs. Observed charging

In contrast, on the right pane of Figure 2.17 (comparison with observed charging), we see that the price reduction can reach up to 48% in a mixed population. This significant price reduction indicates that EV owners in the Netherlands (where this data comes from) are currently charging excessive amounts of power without actually needing all of it for their commuting needs. Therefore, by using AMEVS to optimize charging and to reduce redundant amounts, we see a reduction in prices and an increase in the overall welfare. The pattern of price reduction is similar in populations where only a type of customer preferences is present (either linear or quadratic). However, currently this is not an issue since the EV adoption rates are still below 20% (Daziano and Achtnicht, 2013) in an area with full-density charging infrastructure.

We now consider peak price reduction as shown in Table 2.7. From Tables 2.5 and 2.7, we conclude that LinAMEVS does not yield any benefits to individuals or to the market when high range anxiety dominates the market (homogeneous population of high range anxiety customers). This results from its linear behavior which is more driven by price changes. On the other hand, QuadAMEVS reshapes the demand curve, reducing the peaks and the prices. This results from the customer's decreasing utility for each extra unit of energy he/she consumes. Intuitively, quadratic behavior is more realistic since naturally customers consume until one saturation point above which they get no extra utility (Avci et al., 2014; Samadi et al., 2010; Fahrioglu and Alvarado, 2000; Hall and Mishkin, 1982). However, there are customers in the market that get equal marginal utility for each extra power unit they consume until they reach the desired goal (i.e., have a fully charged battery or have enough power to cook their dinner for 2 hours). In EV charging terms, these customers have higher range anxiety and gain maximum

utility once their desired goal is achieved (LinAMEVS). When we implement AMEVS

Table 2.7: Energy price reduction–AMEVS RTP

			Avg. price red.(%)	Peak price red.(%)
LinAMEVS	vs.	Naive	-99.5	-96.1
QuadAMEVS	vs.	Naive	37.8	57.6
MixedAMEVS	vs.	Naive	14.3	30.1
LinAMEVS	vs.	Observed	1.7	-17.3
QuadAMEVS	vs.	Observed	54.9	66.2
MixedAMEVS	vs.	Observed	37.8	44.9
LinAMEVS	vs.	Cost Min.	29.9	38.8
QuadAMEVS	vs.	Cost Min.	70.9	82.4
MixedAMEVS	vs.	Cost Min.	59.9	71.3
LinAMEVS	vs.	no EVs	-160.1	-140.4
QuadAMEVS	vs.	no EVs	-9.5	30.7
MixedAMEVS	vs.	no EVs	-51.1	-12.9

in a TOU pricing environment, we see different effect on the prices (Figure 2.18). Here, customer populations with high range anxiety (LinAMEVS) have a maximum average price reduction of 8% at 20% adoption rate compared to naive charging populations. In a mixed range anxiety population, the best adoption rate for maximum price reduction is 60%.




Figure 2.18: Average price reduction: AMEVS–TOU vs. Naive Charging. Average price reduction: AMEVS–TOU vs. Observed charging

Comparing the current situation as reflected by the observed charging data with scenarios where AMEVS is used in a TOU pricing regime, we see that the maximum price reduction is around 40% at 100% adoption and is created in a scenario where all customers have low range anxiety. Table 2.8 summarizes the results shown in Figure 2.18, adding the comparisons with the other benchmarks. Overall, looking at Tables 2.7 and 2.8, we conclude that TOU pricing is less beneficial for large populations than RTP schemes. These conclusions support findings in the literature (Strbac et al., 1996; Palensky and Dietrich, 2011).

Table 2.8: Energy price reduction - AMEVS TOU

			Avg. price red.(%)	Peak price red.(%)
LinAMEVS	vs.	Naive	-138.4	-225.9
QuadAMEVS	vs.	Naive	45.5	61.4
MixedAMEVS	vs.	Naive	15.0	7.5
LinAMEVS	vs.	Observed	-19.5	-95.1
QuadAMEVS	vs.	Observed	57.3	70.0
MixedAMEVS	vs.	Observed	33.5	28.1
LinAMEVS	vs.	Cost Min.	-12.2	-32.6
QuadAMEVS	vs.	Cost Min.	59.6	82.3
MixedAMEVS	vs.	Cost Min.	37.1	57.6
LinAMEVS	vs.	no EVs	-215.4	-290.1
QuadAMEVS	vs.	no EVs	3.9	36.9
MixedAMEVS	vs.	no EVs	-49.8	-51.1

The suggested adoption rates of EVs and AMEVS more specifically, are summarized in Tables 2.9 and 2.10, providing specific percentages for our case study (Netherlands 2013). The Most Beneficial Adoption Rate is the maximum adoption rate that yields the highest energy price reduction. The Maximum Adoption Rate is the maximum adoption rate that brings no negative effects to the customers (i.e., no price increase). Similar results can be extracted by applying our algorithm to other data sets and case studies. Therefore, AMEVS apart from scheduling EV charging for individual commuters AMEVS can analyze EV fleets and can yield specific aggregate results with regards to pricing and electric mobility incentives.

Table 2.9: EV and AMEVS adoption rates in various populations under RTP

Population Type			Most Beneficial Adoption Rate (%)	Maximum Adoption Rate (%)
Low Range Anxiety (QuadAMEVS)	vs.	naive charging	100	100
High Range Anxiety (LinAMEVS)	vs.	naive charging	20	40
Mixed Range Anxiety (MixedAMEVS)	vs.	naive charging	55	100
Low Range Anxiety (QuadAMEVS)	vs.	current situation	100	100
High Range Anxiety (LinAMEVS)	vs.	current situation	50	100
Mixed Range Anxiety (MixedAMEVS)	vs.	current situation	80	100

From Tables 2.9 and 2.10, we can extract useful overall results for specific pricing scenarios and specific populations, gaining more insights about the aggregate behavior of customers. It is interesting to mention that the shape of the utility function for individual customers has such dramatic welfare effects on the whole grid. Therefore, a fully distributed approach seems suitable for the EV charging problem, as it manages EV charging differently for each individual, depending on his/her preferences. Uniform ap-

Table 2.10: EV and AMEVS adoption rates in various populations under TOU pricing

Population Type			Most Beneficial Adoption Rate (%)	Maximum Adoption Rate (%)
Low Range Anxiety (QuadAMEVS)	vs.	naive charging	100	100
High Range Anxiety (LinAMEVS)	vs.	naive charging	20	35
Mixed Range Anxiety (MixedAMEVS)	vs.	naive charging	60	100
Low Range Anxiety (QuadAMEVS)	vs.	current situation	100	100
High Range Anxiety (LinAMEVS)	vs.	current situation	40	87
Mixed Range Anxiety (MixedAMEVS)	vs.	current situation	70	100

proaches such as cost minimization mechanisms are expected to create herding of charging (i.e., all customers charge at the same time creating congestion on the grid) since the cost functions and prices are the same for each customer. Such approaches can be beneficial if different pricing schemes are imposed on different customers, by identifying their preferences. EV fleet aggregators or energy providers can encourage EV adoption and smart charging, if they can identify the preferences of their customers. In this way, customers can benefit from owning EVs and from adopting smart charging, without overloading the grid.

Another interesting result of this distributed approach is that it does not encourage herding behavior or conflicts of usage. Generally, in a large population the probability of destructive collisions is small, assuming that not every agent responds immediately and in the same way to every change in price. Even if all the individuals had identical behavior (worst case scenario, upper bound Figure 2.14), the peaks would still be lower than the ones created by uncoordinated charging. Supporting evidence is that since our system uses real-time prices, simultaneous increase in demand from many users would lead to spot price increase. Since we generally observe price reductions, it indicates that the system is becoming more balanced rather than unbalanced.

2.4.4 Sensitivity Analysis

Since both variations of AMEVS (LinAMEVS and QuadAMEVS) use exogenous parameters to express the EV owner’s preferences, we examine how these parameters (Table 2.11) influence the outcome of the algorithm. We showcase the results of the sensitivity analysis with respect to the peak reduction metric assuming Real Time Pricing (RTP). The influence of the parameters in all the metrics used throughout the paper is analogous with the peak reduction. The sensitivity analysis gives a similar outcome for the other

scenarios examined in this paper. For this analysis, we assume all variables constant over time and have therefore omitted the subscript t .

Table 2.11: Utility parameters

Parameter	Short Description	Value
η	weight of the EV charging in the linear utility function	
	Scenario 1	0.1
	Scenario 2	0.3
	Scenario 3	0.7
	Scenario 4	0.9
ω	coefficient of the quadratic utility function	
	Scenario 5	0.1
	Scenario 6	1
	Scenario 7	10
	Scenario 8	100

LinAMEVS Parametrization (Scenarios 1-4)

We now present the variability of results while the parametrization of the LinAMEVS changes. For this variation of the algorithm, we focus only on the parameter η since it is linearly dependent on the parameter β with the relationship $\beta + \eta = 1$. Parameter η indicates the weight that the EV charging demand has in the total utility function of the individual, compared to the weight that the household demand gets in this utility function (β). We expect that as EV charging gets higher weights in the utility function, the total power demand is driven more by EV charging, leading to higher peaks and a less volatile power demand curve. We draw similar conclusions from Table 2.12 where the peak demand increase is analogous to the weight of EV charging in the customer's utility function. Therefore, EV owners that put more value to EV charging than their household demand, are expected to have a more volatile demand curve, and to destabilize the grid more easily. It is more difficult to extract clear relationships between variable increase and PAR increase or decrease, since the peak-to-average ratio is influenced by factors such as driving profile which can affect volatility.


Figure 2.19 shows the decreasing peak demand reduction as a result of more value of EV charging in the total utility function (higher η). We see that the general decrease is similar to the results of other scenarios such as TOU.

QuadAMEVS Parametrization (Scenarios 5-8)

We now present the variability of results while the parametrization of the QuadAMEVS changes. We demonstrate scenarios with regards to parameter ω . This parameter in-

Table 2.12: Energy peak reduction - AMEVS RTP

	PAR red.(%)	Peak red.(%)
LinAMEVS-Scenario 1 ($\eta=0.1$) vs. Naive	-32.2	-9.9
LinAMEVS-Scenario 2 ($\eta=0.3$) vs. Naive	-16.9	-9.9
LinAMEVS-Scenario 3 ($\eta=0.7$) vs. Naive	1.7	-14.2
LinAMEVS-Scenario 4 ($\eta=0.9$) vs. Naive	-1.8	-25.2
LinAMEVS-Scenario 1 ($\eta=0.1$) vs. Observed	-33.7	7.4
LinAMEVS-Scenario 2 ($\eta=0.3$) vs. Observed	-18.3	7.4
LinAMEVS-Scenario 3 ($\eta=0.7$) vs. Observed	0.6	3.8
LinAMEVS-Scenario 4 ($\eta=0.9$) vs. Observed	-2.9	-5.4
LinAMEVS-Scenario 1 ($\eta=0.1$) vs. no EVs	-27.3	-17.7
LinAMEVS-Scenario 2 ($\eta=0.3$) vs. no EVs	-12.6	-17.7
LinAMEVS-Scenario 3 ($\eta=0.7$) vs. no EVs	5.3	-22.3
LinAMEVS-Scenario 4 ($\eta=0.9$) vs. no EVs	2.0	-34.0

**Figure 2.19:** Peak demand reduction as a function of the weight of EV charging in the linear utility function (η)

indicates how high the utility of the customer becomes by consuming certain amount of power, for constant α . As mentioned before, we keep α constant in a sufficiently low value so that we do not lose possible solutions of EV charging. Scenarios 5-8 refer to the parametrization of ω and we use omega values in multiplicative increments of 10 so that we get a sufficiently large spectrum. We examine the effects of the different parameters on peak reduction, assuming that the algorithm performs in an RTP environment. From the utility's form we expect that as ω increases, the more difficult it is for the EV owner to reach the utility saturation point. Therefore, for higher ω , the peak reduction will be lower as the demand for power will show higher peaks in an attempt to satisfy the customer's preferences (utility function). This expectation is confirmed by Table 2.13 where the peak demand reduction decreases when ω increases.

In Figure 2.20, we show the decreasing peak demand reduction as a result of higher ω . We see that by increasing ω , the peak demand reduction decreases in each situation, since the peaks become higher in order to satisfy the customer's utility functions.

Table 2.13: Energy peak reduction - AMEVS RTP

	PAR red.(%)	Peak red.(%)
QuadAMEVS-Scenario 9 ($\omega=0.1$) vs. Naive	6.3	13.4
QuadAMEVS-Scenario 10 ($\omega=1$) vs. Naive	6.3	13.4
QuadAMEVS-Scenario 11 ($\omega=10$) vs. Naive	7.8	9.7
QuadAMEVS-Scenario 12 ($\omega=100$) vs. Naive	9.7	3.9
QuadAMEVS-Scenario 9 ($\omega=0.1$) vs. Observed	7.6	20.4
QuadAMEVS-Scenario 10 ($\omega=1$) vs. Observed	7.6	20.4
QuadAMEVS-Scenario 11 ($\omega=10$) vs. Observed	9.1	17.0
QuadAMEVS-Scenario 12 ($\omega=100$) vs. Observed	10.9	11.7
QuadAMEVS-Scenario 9 ($\omega=0.1$) vs. no EVs	9.8	-2.0
QuadAMEVS-Scenario 10 ($\omega=1$) vs. no EVs	9.8	-2.0
QuadAMEVS-Scenario 11 ($\omega=10$) vs. no EVs	11.2	-6.3
QuadAMEVS-Scenario 12 ($\omega=100$) vs. no EVs	13.4	-13.2




Figure 2.20: Peak demand reduction as a function of the ω coefficient of the quadratic utility function

2.5 Policy Recommendations

From our IS artifact’s evaluation, we observed that individual preferences have a crucial influence both on individual demand profiles and on the distribution grid at an aggregate level. Therefore, it is important to assess the adoption rate of AMEVS in comparison with the current charging situation as described by the observed charging data (Netherlands, 2013). For example, in the previous section we saw that LinAMEVS shifts the peaks compared to naive charging. Therefore, if the policymakers encourage a moderate adoption rate of LinAMEVS, the benefits for the grid will be significant, since the customers that adopt LinAMEVS will be able to offset peaks resulting from EV customers that charge without any advanced charging mechanism. We demonstrate scenarios (Table 2.14) of heterogeneous populations with low, medium, and high AMEVS adoption rates, assuming customers with high range anxiety (LinAMEVS) and customers with low range anxiety (QuadAMEVS).

Table 2.14: Simulation scenarios

Scenario #	Short Description
Scenario 1	High Adoption Rate Populations (90% AMEVS, 10% Observed)
Scenario 2	Medium Adoption Rate Populations (50% AMEVS, 50% Observed)
Scenario 3	Low Adoption Rate Populations (10% AMEVS, 90% Observed)

In Scenario 1, we assume an AMEVS adoption rate of 90%, in Scenario 2, an adoption rate of 50%, and in Scenario 3, an adoption rate of 10% (the scenario of 100% adoption reflects homogeneous populations and is presented in Section 2.4). All these three scenarios are exposed to Real Time Prices (RTP). The results would be analogous if Time-of-Use (TOU) prices were chosen. All these scenarios are compared with the observed charging benchmark and the results are shown in Table 2.15. The observed charging situation reflects the current charging scenario where no advanced charging mechanism is in place (Section 2.4.1).

Table 2.15: Policy recommendations

Scenario Type	Peak reduction (%)
Scenario 1 - Low Range Anxiety (QuadAMEVS)	27.9
Scenario 1 - High Range Anxiety (LinAMEVS)	-0.3
Scenario 2 - Low Range Anxiety (QuadAMEVS)	15.5
Scenario 2 - High Range Anxiety (LinAMEVS)	14.8
Scenario 3 - Low Range Anxiety (QuadAMEVS)	3.1
Scenario 3 - High Range Anxiety (LinAMEVS)	5.1

From Table 2.15 we see that in populations with high range anxiety, high adoption rates are not beneficial for the distribution grid, as the peak demand is driven by the majority of the population which has linear preferences. However, in populations with low range anxiety, high adoption rates should be encouraged because they have positive effects on the peak reduction. In populations with medium adoption rates, high range anxiety customers have a higher peak reduction compared to the observed charging scenario. This results from their ability to offset peaks, thereby significantly flattening the aggregate demand curve. The result is similar in low adoption scenarios where customers with linear preferences (higher range anxiety) yield a higher peak reduction. A schematic overview of the previous scenarios is provided in Figure 2.21. The darker boxes indicate higher benefits (on each dimension) for the distribution grid compared to the lighter boxes. Energy policymakers have to identify the preferences of their EV customer portfolio (fleet) and, based on their low or high values, to incentivize high or low EVs and AMEVS adoption in order to maximize the benefits for the distribution grid (in terms of peak demand reduction in this case).




Figure 2.21: Policy recommendations: schematic overview across range anxiety and adoption rate dimensions. Dark colored areas indicate higher benefits for the distribution grid.

2.6 Conclusions & Future Work

Electric Vehicles will likely become a significant part of the transportation system. They are more efficient than conventional internal combustion engine cars and could be essential to a sustainable transportation system. If they are properly integrated in the market, they may yield significant benefits for individual commuters as well as for the environment.

However, the uncontrolled introduction of EVs may put the electricity grid under critical strain since the energy that was previously produced by conventional fuels to run internal combustion engine cars, will now have to be produced by electricity. When commuters arrive at home in the evening, they typically cook dinner, take baths, and engage in other activities that already cause daily demand peaks. If they also plug in their vehicles at that time, the extra demand could easily destabilize the grid. However, EVs carry storage capacities that can bring great benefits to the electricity grid if used properly. The storage available in EVs can be used as a buffer so that electricity demand is reallocated to periods of lower demand, alleviating the strain on the grid. This feature, combined with a correct charging schedule, can be advantageous for the transportation system and the electricity grid. In other words, by reallocating demand due to charging, no extra infrastructure will be needed to cover the new demand created by EV charging, increasing sustainability levels.

To accomplish this, we propose an Adaptive Management of EV Storage (AMEVS) algorithm to mitigate the negative influence of large scale EV integration and to enhance

the robustness and reliability of the grid. Our algorithm is a fully distributed approach that schedules the charging of each EV individually, using reinforcement learning and optimization techniques. It personalizes charging based on each individual's preferences and range anxiety, achieving redistribution of electricity demand without violating a customer's need for mobility service. In heterogeneous populations, we observe significant peak demand reduction as a result of AMEVS adoption. In homogeneous populations, we observe peak demand reduction, if the customers have low range anxiety. In homogeneous populations with high range anxiety the maximum AMEVS adoption depends on the charging profiles of the rest of the population. In our data set it is calculated as 40% compared to customers that use naive charging.

As a result of peak demand reduction, the use of AMEVS can reduce average energy prices for all customers, compared to naive charging. Consequently, using AMEVS EVs support grid sustainability as peaks are significantly mitigated and there is no need to construct extra electricity grid infrastructure to accommodate increased EV charging demand. Thus, AMEVS can be used to promote the adoption of EVs, as EV owners no longer have to decide when to charge their EV and minimize their costs.

At a fleet owner's level, we see the importance of individual preferences to a balanced charging regime. We saw that an increasing heterogeneity of preferences towards energy consumption leads to a higher peak demand reduction. Therefore, it is important for energy policymakers to identify the population they are targeting and design their EV adoption incentives under the individual preferences prism. Our second finding is that there seems to be an alignment between individual and overall benefit maximization in populations with lower range anxiety (QuadAMEVS). Simulations with customer populations with low range anxiety show that maximizing their own benefit leads to peak demand reduction that benefits the overall grid. In populations with high range anxiety (LinAMEVS), both customers and the grid benefitted under certain EV penetration rates (35% in our data set). Higher penetration rates of EVs in high range anxiety populations revealed the need for additional physical grid infrastructure to accommodate extra demand. In mixed populations (MixedAMEVS), we saw that the benefit from low range anxiety populations is able to offset any negative effects of high range anxiety customers, leading to an overall positive effect on the grid. In future, we plan to examine the effect of AMEVS on micro-grids and smart neighborhoods and to explore how a consumer social network (smart neighborhood) influences individual consumption behavior.

Appendix—Summary of Notation and Abbreviations

Table 2.16: Summary of Notation

Symbol	Definition
$\mathbf{A} = \{A_i\}$	discrete set of all possible actions in RL
C_t	battery's state of charge in time slot t
CA_t	individual charging availability per time slot t
$\mathbf{D} = \{D_t\}$	total (household and EV) demand vector per individual
$\mathbb{E}[Dist_t]$	expected distance needed for t
$F(I_1, ..., I_N)$	mapping function in input module to produce expected distance for t
$\mathbf{I} = \{I_1, ..., I_N\}$	input set for input module
$\mathbf{L} = \{l\}$	discrete set of all consumption levels $l \in \{1, ..., L\}$
n	number of inputs in AMEVS artifact
P	energy price per consumption unit
$R(t, l, A_i)$	reward function for each consumption level l at time t after an action A_i
$\mathbf{S} = \mathbf{T} \times \mathbf{L} = \{(t, l) t \in \mathbf{T}, l \in \mathbf{L}\}$	discrete set of all possible states in RL
SoC_{min}	minimum state of charge
$\mathbf{T} = \{t\}$	discrete set of time intervals, $t \in \{1, ..., T\}$
$U(x_{h,t}, x_{c,t})$	utility function
$X_{max,t}$	upper bound of $x_{c,t}$
$W(x_{h,t}, x_{c,t})$	individual benefit obtained from consumption $x_{c,t}$ and $x_{c,t}$
\mathbf{x}_c	vector of EV charging consumption for each t
\mathbf{x}_h	vector of household consumption for each t
α	individual customer's utility parameter
β	weight of individual household demand in utility function
γ	discount factor of the <i>learning module</i>
η	weight of EV charging demand in utility function
Θ	absolute value of negative reward ($-\Theta$) in RL
ρ	capacity/distance rate
ω	level of satisfaction per energy consumption unit

Table 2.17: Summary of Abbreviations

Abbreviation	Definition
EPEX	European Power Exchange
HC	Heuristic Charging
LF	Load Factor
PAR	Peak-to-Average Ratio
RL	Reinforcement Learning
RTP	Real Time Pricing
TOU	Time-of-Use

Chapter 3

Sustainable Electric Vehicle Charging: A Data-driven Approach¹

3.1 Introduction

Sustainability is a major driver for modern societies that aspire to minimize their negative impact on the environment using technological advancements (Watson et al., 2010). It is defined as the “development that meets the needs of the present without compromising the ability of future generations to meet their needs” (Brundtland, 1987). Therefore, in its core includes a sustainable electricity grid. The modern electricity grid, or smart grid, is characterized by a large scale integration of renewable sources. These renewables, apart from clean energy sources, are also volatile and highly dependent on the weather conditions. Therefore, they may easily destabilize the grid, threatening its reliability. Consequently, there is a need for reducing this volatility, while benefiting from the sustainable renewable sources.

The main challenge in balancing the smart grid is to mitigate peak demand. Peak demand is one component of the total electricity demand, additional to the base load. Base load power is provided around the clock and typically comes from large nuclear or coal-fired plants. It has a constant profile and cannot be easily ramped up or switched off. Peak demand is the extra consumption that results for example in hot summer afternoons

¹Parts of this chapter have appeared in the following peer reviewed conference proceedings: Valogianni, K., Ketter, W., Collins, J. & Zhdanov, D.(2014). Enabling Sustainable Smart Homes: An Intelligent Agent Approach. In Proceedings of the International Conference on Information Systems, (ICIS 2014), Auckland, New Zealand, Association for Information Systems (AIS). Valogianni, K., Ketter, W., Collins, J. & Zhdanov, D (2015). Sustainable Electric Vehicle Charging: A Data-driven Approach. In 25th ICIS Workshop of Information Technologies and Systems (WITS 2015), Fort Worth, Texas.

and is typically covered by power plants that can be switched on for shorter periods, such as gas turbines (Kempton and Tomić, 2005). This demand is difficult to be matched with the supply, because of its stochastic and volatile nature. More importantly, peak demand determines the total capacity that must be available in the grid, since the grid must be able to accommodate the peak demand at all times, even though this peak might be instantaneous. One solution to cover peak demand is to install additional generation and transmission infrastructure. This solution is costly and unsustainable. Scaling up supply would require building costly new power plants and infrastructure that decreases efficiency and consequently reduces environmental sustainability. Modern societies would be more sustainable if they could totally offset this peak demand or adapt energy use to availability of sustainable sources. Evidence that the usage of information systems (IS) or investments in smart solutions are financially more attractive than the investment in generation capacity or grid capacity, has been provided by Schmidt and Busse (2013).

A key role in the smart grid management is played by the energy customers, who are now able to adapt consumption to availability and also produce energy using for example photovoltaic panels (PVs) or wind turbines. Both of these production sources can be installed and operated in a household and are renewable sources, which means they have zero marginal cost for the customers. These features make them quite popular among customers. These new energy customers are consequently involved in both consuming and producing energy, and are commonly called prosumers (Lampropoulos et al., 2010). Prosumers mostly reside in smart homes, where with the help of information and communication technology (ICT), they can control their appliances' energy use. Smart homes appeal to the end customers because of the ease and comfort they offer and their sustainability. For example by adjusting heating based on households' occupancy, the residents can save significant amounts of energy and money. Furthermore, many of the smart home residents, in their effort to be sustainable, own an electric vehicle (EV).

The combination of EVs and smart homes has the potential to support a sustainable way of living, now and in the future, since smart home owners can use energy coming from renewable sources (PVs or wind turbines) to charge their EVs. This way, both their electricity expenses are reduced, since renewable sources have zero marginal cost, and the sustainability levels increase, since part of conventional energy is substituted by renewable. However, in order for smart home owners to make optimal decisions for their energy consumption and EV charging, they need to be aware of the information available such as future energy prices, renewable source availability, weather conditions, driving schedule etc. All this information may exceed the limits of human cognitive ability (Simon, 1979) and energy consumers cannot pay attention to energy issues all the

time. For example, a prosumer in order to charge her car at a low cost and to make use of renewable energy, has to know when the sun will be shining (so that the PV is generating energy), when the prices are going to be low (consuming energy) or high (producing energy) and combine all these factors with her driving preferences and desired state of charge of the EV battery. Consequently, it is difficult for humans to keep track of all this information and make optimal decisions. Therefore, we propose an intelligent software agent that can process the large amount of available information and facilitate customer decisions within the smart grid.

The software agent has the role of an intelligent decision support system (Hevner et al., 2004) that is able to explore and understand the complex environment of the smart grid and suggest an optimized consumption for energy customers. Our proposed intelligent agent is implemented as a Mobility Integrated Energy Management artifact (MIEM) that represents a household (smart home) combined with an EV and a PV. At the core of this agent there is an energy information system (Energy IS) that processes prices, individual preferences and weather information obtained by the environment and offers personalized household consumption and EV charging suggestions. Its main objective is to minimize the total energy consumption (household, EV and PV) cost and maximize the use of renewable energy. The agent must satisfy household consumers' electricity needs and at the same time benefit from energy price variations. In this chapter, we put our focus on combining mainly the EV with the PV, trying to make use of the storage in the EV's battery and benefit from renewable energy.

3.2 Mobility Integrated Energy Management Model

MIEM artifact offers a new solution to the existing problem of managing household consumption, EV charging and PV production in an effective and efficient manner. Therefore, from the Design Science point of view, it is categorized under the solution group "Improvement: New Solutions for Known Problems" (Gregor and Hevner, 2013). This new solution brings the maximum possible benefits to the individual household customers using available information (arrival and departure times, weather conditions, prices, etc.). With respect to relevance it is tied to environmental sustainability as well as smart grid and energy information systems (Energy IS) (Watson et al., 2010), where the use of information brings reduction to the amount of energy consumed by the smart home. With regard to rigor, our IS artifact is connected to the general Green IS research stream and to the general power systems area. An illustration of the proposed design is given in Figure

3.1. In the Appendix we include tables with the notation used to describe MIEM, as well as the abbreviations used in this chapter.

First, the artifact receives as inputs the individual arrival and departure preferences of the household owner (1). Then, based on these preferences and on the household demand data, it creates the aggregate demand profile, comprised by household and EV electricity demand (2). Thirdly, it receives the current weather conditions and historical PV generation data as inputs (3) and using a supervised learning algorithm, creates a PV generation forecast (4). Combining this forecast with the electricity prices (5), MIEM artifact creates a total electricity demand profile for the household that minimizes the overall costs and ensures maximum use of renewable energy (6). The previous steps are repeated so that MIEM updates its inputs in real-time, having always the latest information about user's behavior and preferences. The detailed actions taken by MIEM within the smart grid are described by the activity diagram depicted in Figure 3.2.

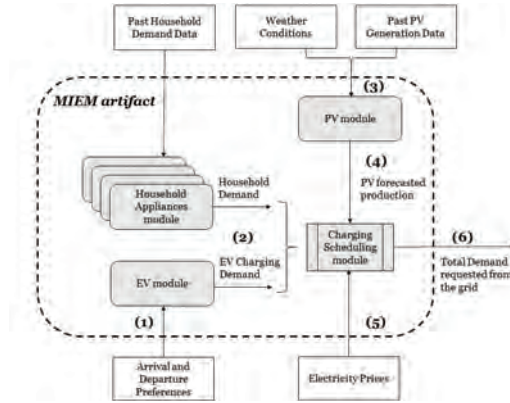


Figure 3.1: Mobility Integrated Energy Management Artifact Overview

In the evaluation section we assess the presented IS artifact with respect to validity, utility, quality and efficacy (Hevner et al., 2004). To this end, we demonstrate through simulation experiments the performance of the artifact under realistic conditions. The presented IS artifact engages in two different processes: a) forecasting the PV generation and b) scheduling the EV charging. For these two processes receives as inputs the weather conditions (historical observations and weather forecast), in the form of an instance space \mathbf{W} (steps ④ and ⑤ in Figure 3.2), the home owner's arrival and departure preferences $\theta = \{t_d^n, t_a^n\}$ where $n \in \{1, \dots, N\}$ is the index of driving activities, t_d^n, t_a^n are the times of departure and arrival back home for each respective activity n for each EV owner (step



Figure 3.2: MIEM activity diagram.

① in Figure 3.2) and the energy prices $\mathbf{P} = \{P_t\}$ over a predefined planning horizon T (step ⑦ in Figure 3.2).

3.2.1 Photovoltaic generation forecasting

As shown in Figure 3.1, to schedule EV charging in a cost beneficial way, MIEM artifact needs an accurate forecast of the PV generation. Having an accurate forecast of the generation will incentivize EV charging during periods that the PV is producing energy, so that it reduces the electricity costs and increases renewable usage. Therefore, we introduce a supervised learning algorithm in the PV module, which is trained on weather data and is

able to forecast the PV generation. Specifically, we use an ensemble decision-tree-learning algorithm (*random forest learning*) (Breiman, 2001) which has as its task to learn from the weather instance space \mathbf{W} . This space includes a set of attributes $W_1, \dots, W_{|\mathbf{W}|}$, where $|\mathbf{W}|$ is the cardinality of the weather data set. Each weather attribute has a different value for each time slot t , $W_{1,t}, \dots, W_{|\mathbf{W}|,t}$. The weather instance space is accompanied with a set of PV generation values $\mathbf{G} = \{g_t\}$. The combination of $(W_{1,t}, \dots, W_{|\mathbf{W}|,t}, g_t)$ is one example that is used by the learning algorithm. Having multiple examples over time t expedites the learning process, leading to a higher accuracy.

The end goal of the learning algorithm is to learn to assign weather conditions to corresponding PV generation values, $\mathbf{G} = \{g_t\}$ with a certain probability. In other words, the algorithm learns that when for example there is high overcast and low temperature, there is lower probability of high PV generation and it will assign this data entry to a low PV generation label, e.g. $g_t = 0.1$ kWh. The learning algorithm learns a function $f : \mathbf{W} \rightarrow \mathbf{G}$ such as the predictive accuracy is high without creating overfitting. Due to the nature of our dependent variable (PV generation in kWh), which is a continuous variable, we implement a regression-tree-learning ensemble algorithm, which means that the algorithm assigns labels to the input entries based on linear function derived from the training set. We denote the outcome of the algorithm as \hat{g}_t and it stands for the forecasted PV generation in each time slot t . Formally we have ((6) in Figure 3.2):

$$\hat{g}_t = f(W_{1,t}, \dots, W_{|\mathbf{W}|,t}) \quad (3.1)$$

where $W_{1,t}, \dots, W_{|\mathbf{W}|,t}$ are the weather attributes during time slot t and $f : \mathbf{W} \rightarrow \mathbf{G}$ is the ensemble of regression trees that needs to be learned by the algorithm. The forecasting error of the algorithm is the difference between the observed PV generation during the time slot t and the forecasted PV generation for the same time period: $\epsilon_t = g_t - \hat{g}_t$.

We selected the ensemble regression-tree learning algorithm against the traditional decision tree-learning methods because the latter are prone to overfitting depending on the training set size. Therefore, to prevent overfitting they require *pruning*, which means shortening the branches of the tree to reduce dependencies on the training set. Ensemble learning with decision trees does not entail the danger of overfitting, since many different decision trees are created from the same data set, using each time a random selection of features. This way, the algorithm combines all created decision trees and produces the learned outcome without the influence of overfitting. Furthermore, experimenting with both methods we found lower mean-squared and mean-average-percentage errors for the ensemble learning method.

3.2.2 Electric Vehicle Charging Scheduling

Having the PV generation forecast \hat{g} as input from the PV module, the Charging Scheduling module is able to determine the EV charging profile $\mathbf{x}_c = \{x_{c,t}\}$ over time so that the total electricity cost is minimized. The individual household electricity cost is denoted as $c(\hat{g}_t, x_{c,t}) = (x_{c,t} - \hat{g}_t) \cdot P_t$ (step ⑧ in Figure 3.2). The charging scheduling module has to solve the following problem per individual EV customer over a horizon T (step ⑨ in Figure 3.2):

$$\mathbf{x}_c^* = \underset{\mathbf{x}_c}{\operatorname{argmin}} \mathbb{E} \left\{ \sum_{t=1}^T c(\hat{g}_t, x_{c,t}) \cdot CA_t \right\} \quad (3.2)$$

subject to the constraints (3.3)-(3.5). This minimization problem expresses the artifact's goal to minimize the expected cost over a horizon T , for each individual EV owner. The variable CA_t is the charging availability of the household customer (if the customer is driving he/she cannot charge) and is calculated from the set of arrival and departure preferences $\theta = \{t_d^n, t_a^n\}$: since n is the index of each driving activity, t_d^n indicates the time that the EV owner in question departs for this particular driving activity (e.g. work) and t_a^n indicates the time that this EV owner arrives back home. Therefore, the charging availability per time slot t is calculated for each individual as $CA_t = 0 \quad \forall t \in [t_d^n, t_a^n]$ otherwise $CA_t = 1$ (step ③ in Figure 3.2).

$$0 \leq x_{c,t} \leq X_{max,t} \quad \forall t \in \{1, \dots, T\} \quad (3.3)$$

The upper bound $X_{max,t}$ represents the maximum energy that the customer agent can charge from the network per timeslot t . This represents the main network constraint and depends on the characteristics of the residential connection.

Constraint (3.4) expresses the necessity to charge at least as much as it is expected to be needed for the next time interval. This is imposed by the need that the state of charge (SoC_t) should be equal to sum of the previous state of charge (SoC_{t-1}) and the amount charged ($x_{c,t}$), reduced by the amount needed for driving $\mathbb{E}\{Dist_t\} \cdot \rho$. The variable $\mathbb{E}\{Dist_t\}$ is the expected driving demand (in km) for the time slot t (step ② in Figure 3.2) and ρ is the capacity over distance factor (kWh/km) provided by the battery specifications.

$$SoC_t = SoC_{t-1} + x_{c,t} - \mathbb{E}\{Dist_t\} \cdot \rho \quad \forall t \in \{1, \dots, T\} \quad (3.4)$$

Solving this equation with respect to $x_{c,t}$, we get a constraint for the minimization problem:

$$\begin{aligned}
x_{c,t} &= SoC_t - SoC_{t-1} + \mathbb{E}\{Dist_t\} \cdot \rho \quad \forall t \in \{1, \dots, T\} \\
SoC_0 &= 0
\end{aligned} \tag{3.5}$$

Constraint (3.6) ensures that charging only occurs when the customer is not driving and is available for charging.

$$x_{c,t} \leq CA_t \cdot M \quad \forall t \in \{1, \dots, T\} \tag{3.6}$$

where M is a sufficiently large number, SoC_t is the state of charge on timeslot t , and SoC_0 is the starting point for the state of charge. These variables are calculated from the driving data described later in the Data Description section. All the aforementioned constraints ensure that the agents do not violate the customer's comfort and have the EV always charged to cover the driving needs of the coming day. Defining cost as $c(\hat{g}_t, x_{c,t}) = (x_{c,t} - \hat{g}_t) \cdot P_t$, (3.2) becomes:

$$\mathbf{x}_c^* = \underset{\mathbf{x}_c}{\operatorname{argmin}} \mathbb{E}\left\{\sum_{t=1}^T ((x_{c,t} - \hat{g}_t) \cdot P_t \cdot CA_t)\right\} \tag{3.7}$$

subject to (3.3)-(3.6). The prices P_t at each hour are announced one time horizon T in advance, so that the electricity customers can plan their consumption for the coming time horizon T .

3.3 Experimental Evaluation

In order to evaluate MIEM's performance, we create simulations built with real-world data. The goal of our experiments is to show how MIEM increases the sustainability levels and assists smart home owners in their complex decisions. Our evaluation environment simulates the smart grid conditions and specifically the retail electricity market (based on assumptions made in Power TAC (Ketter et al., 2016)). The electricity customers are exposed to prices which reflect the energy availability (detailed information is given in the Data Description section). Each household has an individual household consumption profile, and a PV production profile. The PV module forecasts this production profile based on weather conditions using the ensemble learning algorithm presented before. Furthermore, each household customer has a particular driving profile, based on which the artifact schedules the charging to provide an optimal recommendation (with respect to reduced electricity costs).

Below we present the simulation parameters' values used in our numerical simulation experiments. By Level of Charging, X_{max} , we denote the maximum charging speed (in kW) that the EV can charge from the grid. This is now set to the single phase charging speed, which is 3.3 kW (230 VAC, 16 A) and is used in real world for the lowest level charging. It allows charging from a household socket and does not even require a three phase installation. We chose this particular charging speed, since we investigate the household consumption and charging behavior and we assume that there are no other (fast) charging options available, as it happens currently in most countries. Regarding the planning horizon we assume $T = 168h$ (1 week). We base this choice on the weekly repetition of activities which shows up both in the driving preferences and in the household consumption profiles.

3.3.1 Data Description

As presented in Figure 3.1, the input data needed for an effective simulation are: a) the household demand profiles b) the driving preferences, c) the weather conditions combined with d) PV training data and e) the electricity prices that the customers are facing in the market. For each input we present the calibration data set.

Household Demand Profiles

Regarding the household demand profiles, we use household consumption data from the Netherlands obtained in collaboration with a European Utility Company. This data set includes detailed consumption per 15 minutes for 24 different households². The data was collected in January 18-29, 2010. We aggregate this consumption data over hourly intervals to maintain the granularity of our analysis to be 1h. The measurements are gathered in 2010. Based on these households we create large customer populations drawing randomly one of the consumption profiles. In Figure 3.3 we see the daily average for these 24 households. We observe that the maximum demand we have is around 3kWh which is not quite frequent in the data. One can understand that the household profile is just a portion of the overall electricity profile in a smart home with EVs and PVs. The household consumption data set does not include the information of EV ownership. However, combining data from the Netherlands both related to household consumption and EV charging we can train the models to reflect the Dutch electric mobility scene.

²The number of households was the maximum number that could be made available by the Utility Company.

Driving Preferences

To extract arrival and departure preferences, we calibrate our simulation with data obtained from the Dutch Central Bureau of Statistics (Central Bureau of Statistiek, CBS). This data set includes driving demand of individuals based on their profession, the age group they belong, the different driving motives they have, etc. Consequently, we can estimate the arrival and departure time of each individual depending on the driving motives she has per day. Furthermore, we can estimate the driving demand in km. This demand distribution is displayed in Figure 3.3 for a 24 hour horizon. We see that most of the driving happens during the day, and peaks in the morning hours when most of the commuters are driving to work. Regarding charging availability we assume that customers can charge the EV's battery when they are not only at home but also at work ("standard" charging with direct billing to the customer), which is nowadays implemented by companies in order to encourage their employees to drive "green."




Figure 3.3: Left pane: Household demand distribution. Right pane: Distribution of driving demand of the population over a 24h horizon.

Weather Conditions

The weather conditions data is necessary in combination with PV generation data so that we train the ensemble learning algorithm. To create the training tuples $(W_{1,t}, \dots, W_{|W|,t}, g_t)$ we use weather data from the same region and time period that we derive the PV generation data. The weather data set includes the attributes displayed in Table 3.1. Table 3.1 shows an example of weather data from Saturday, April 25th, 2015 (detailed information obtained from weather substation K21D). We should note that the weather and PV generation data have a 20 minute granularity, which increases learning accuracy. In the end we aggregate the learned generation to hourly intervals (by averaging observations over an hour) since this is the overall granularity of our simulation.

In Figure 3.4 we display the temperature and the dew point for the whole day of April 25th, 2015. We see an increase of temperature during afternoon. We could expect that

Table 3.1: Weather attributes sample data, Saturday, April 25th, 2015, weather substation K21D.

Time	Temp (°C)	Wind- chill (°C)	Dew Point (°C)	Humid. (%)	Pressure (hPa)	Visib. (km)	Wind Dir.	Wind Speed (km/h)	Gust Speed	Precip.	Events	Cond.
12:10	7.0	3.4	3.0	76	1009.7	16.1	ESE	22.2	33.3	N/A	-	Overcast
12:30	7.0	3.6	3.0	76	1009.7	16.1	ESE	20.4	31.5	N/A	-	Overcast
12:50	7.0	3.6	3.0	76	1010.0	16.1	ESE	20.4	27.8	N/A	-	Overcast
01:10	7.0	3.8	3.0	76	1010.0	16.1	ESE	18.5	25.9	N/A	-	Overcast
01:30	7.0	3.6	3.0	76	1010.0	16.1	ESE	20.4	27.8	N/A	-	Overcast
01:50	7.0	4.0	3.0	76	1009.7	16.1	ESE	16.7	31.5	N/A	-	Overcast
...

this increase of temperature is related with an increase in the PV production, because high temperature could indicate a lot of solar radiation.

PV generation data

The PV generation data used for training the learning algorithm are obtained from a household PV installation in Hudson, WI, USA. We chose the data from this PV installation because the data set was highly granular allowing for accurate training of the ensemble learning mechanism. This PV installation is used by individual electricity consumers to produce energy both for supplying their own household but also for selling the surplus back to the grid and making profits. Since the PV installation is located in the USA, it is meant to serve a higher household electricity demand (than typical European households). Therefore, to avoid inconsistencies, we normalize this data to serve an average Dutch household and be in accordance with the rest of our data (the normalization is done so that the ratio PV generation/electricity consumption is the same in Europe and US). For demonstration purposes we train our algorithm on data collected during the period April 25th, 2015 to June 23rd, 2015. The original data set for this 2-month period includes 5 minute observations and in total 17,281 data points. Due to the granularity restriction of the weather data (20 minutes), we aggregate the PV generation output to 20 minute intervals to train the learning algorithm. After having the learned PV generation, this is summed up to 1 hour intervals so that we get more illustrative results. In Figure 3.4 we present exemplary PV generation curves for a week with 20 minute granularity (April 25th-May 1st, 2015). We see that the PV installation starts generating energy early in the morning when the sun rises. This coincides with the increase in the temperature shown in the left pane of Figure 3.4.

Since we train our algorithm on data from April 25th, 2015 to June 23rd we expect results biased towards high PV generation, since during these two months the PV generation is at its highest levels. Therefore, besides this summer scenario we show a winter




Figure 3.4: Left pane: Temperature and Dew Point graph for April 25th, 2015. Right pane: PV generation curves for a week (April 25th-May 1st, 2015).

scenario as well, reflecting the PV generation during the months December to end of January, when the PV generation was at its lowest level during the year 2015 for these particular installations.

Electricity Prices

We evaluate our artifact under variable pricing regimes, since we believe that electricity customers should be exposed to the actual energy availability. This way the smart grid managers are able to provide incentives for consumption during high supply periods and counter incentives (high prices) during shortage periods. Therefore, we create an example of variable pricing scheme based on the European Power Exchange (EPEX) intraday price-curve over weekly horizons. The Dutch retail energy prices account on 42% for the wholesale energy price³ (EPEX prices) while the rest 58% represents the distribution network fees, the energy taxes and VAT. Therefore, we create the retail price scheme that accounts both for wholesale prices, distribution network fees and taxes: $P_t = \frac{EPEXprice_t}{0.42} (\text{€/kWh})$ assuming $EPEXprice_t$ is the price given by EPEX. The average price curve is shown in Figure 3.5.

3.4 Numerical Results

To evaluate the performance of the proposed artifact, we run the simulation for 5 weeks (sufficient time for the output to converge) following the steps shown in Figure 3.2 and we calculate the steady state electricity consumption curve for the individuals. Since the final outcome has a high dependency on the PV generation forecast accuracy, we show, first, the performance of the ensemble learning algorithm in the PV module.

³<http://www.nuon.nl/energieprijzen/> [Date Accessed: July, 20th, 2015]




Figure 3.5: Example variable pricing scheme based on EPEX prices.

3.4.1 PV forecasting

In order to forecast the PV generation from the weather data we are using *random forest learning*. Based on random forest learning, we create ensembles of regression trees connecting weather conditions to PV generation data. A single regression tree (Breiman, 2001) is creating a regression function relating weather data to PV generation output: $f : \mathbf{W} \rightarrow \mathbf{G}$, as explained in Section 3.2.1. This regression tree is creating a partition of the data set depending on certain cut-off points. In other words, it can identify that if the temperature is above a certain threshold and the wind is blowing with a certain speed, then the PV generation will have a certain output. The ensemble method has the advantage that it creates multiple trees and combines them in order to increase the forecasting accuracy. Single regression trees require *pruning*, so that they are not biased by overfitting on the training set. Ensemble methods, like random forest, do not require pruning because the ensemble is deciding by itself about the depth of the tree preventing overfitting.

In order to train the random forest learning on the data, first, we have to decide on the optimal *leaf size* of each tree included in the ensemble. By *leaf size* we mean each end node of the tree. To do so, we display the mean-squared-error (MSE) for different leaf sizes as a function of the number of trees in the ensemble (Figure 3.6, left pane). Secondly, we must select the appropriate number of trees included in the ensemble. For this process we use the metric *out-of-bag mean squared error* (out-of-bag MSE) which reflects the mean-squared-error of the part of the training data points that have not been used in the ensemble training.

The out-of-bag MSE is used in the machine learning literature (Breiman, 2001) to indicate the MSE calculated from observations (weather data and the matched PV gen-

eration in our case) that are not used for training. In ensemble learning, due to random permutations, a part of the training data set is not actually used in training, and it can be used for improving the accuracy of the method. In the right pane of Figure 3.6 we show the out-of-bag MSE as a function of the number of trees in the ensemble. We see that after 50 trees the out-of-bag MSE does not improve significantly. The out-of-bag MSE does not mean that a next value is simulated by previous values, it means that there is a part of the initial data set that is not used for training and can be used as a testing sample (out of bag).




Figure 3.6: Left pane: Leaf size as a function of the number of trees in the ensemble. Right pane: Out-of-bag mean squared error as a function of the number of trees in the ensemble.

Since our weather training data includes a lot of features, we select the features that are most the important for forecasting the PV generation values. Therefore, we see in Figure 3.7 that the humidity is the most important feature for forecasting the PV generation. This is logical since the humidity indicates presence of cloud coverage, hence reduced solar radiation. To evaluate the accuracy of the ensemble learning algorithm

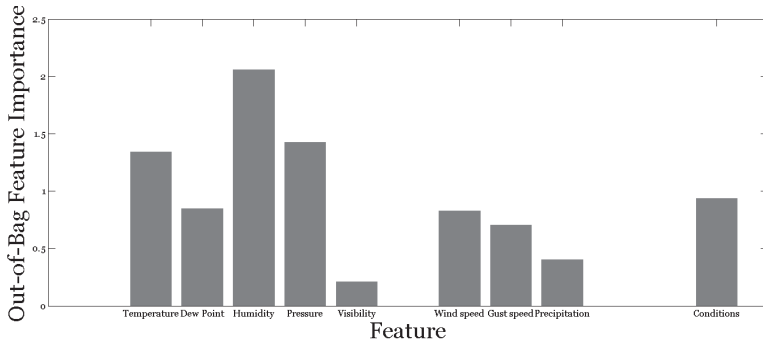


Figure 3.7: Feature importance.

we use the out-of-bag MSE which is 0.5 for 50 trees in the ensemble. Furthermore, we

calculate the mean average percentage error (MAPE) in a test set of 5 weeks of data. This error has a value of -0.006% , which means that in general our learning method under-predicts the PV generation. Calculating the root-mean-squared error (RMSE) for the same test set (5 weeks period), we get a value of 1.15 kWh. Considering the magnitude of the PV generation of this installation, the RMSE is relatively low, allowing for accurate forecasting.

3.4.2 Peak Demand Reduction

Having forecasted the PV generation, the Charging Scheduling module receives the electricity prices and decides on the EV charging scheduling so that costs are minimized. First, we present the steady state electricity consumption curve calculated from 5 weeks of simulation during the months May and June, 2015, when the PV generation was at its peak (Figure 3.8). We compare the produced curve with real-world charging data obtained from the Netherlands during the period January 2013 - December 2013. This data set includes charging observations of 10,462 EV owners in the whole country (obtained in collaboration with a grid infrastructure company). We observe, firstly, that the demand is shifted during low price periods (IS artifact Summer and Winter scenarios in Figure 3.8). Furthermore, we find that PV generation is enough to cover the household and EV charging needs during the day time and also to create surplus that the household can sell back to the grid. The summer scenario is the most optimistic one, since during summer the PV generation is at its maximum level. We compare this scenario with the winter scenario of December-January when the PV generation was at its lowest level. We observe that in the winter the household owners do not have surplus that can sell back to the grid. However, they use PV generation to reduce their expenses from buying electricity from the grid and increase their sustainability levels. In Table 3.2 we compare the electricity




Figure 3.8: Electricity demand curve after adopting the IS artifact and surplus resulting from the PV generation.

costs and the renewable usage of a household with an EV without an IS artifact (current

situation- data from the Netherlands), a household with our IS artifact in summer, when PV generation is high, a household with our IS artifact in winter, when PV generation is the lowest and finally a household in summer and winter but in a perfect information regime. The perfect information scenarios show the upper bound of the comparison, since our IS artifact relies on forecasting.

Table 3.2: Electricity Costs and Renewable Usage.

	Average Daily Electricity Cost (€)	Average Daily Renewable Usage (kWh)	Percentage of Daily Renewable Usage (%)
Current Situation - Real-world charging	5.05	0	0
IS artifact - Summer Scenario	-0.63	54.06	93.30
IS artifact - Winter Scenario	2.76	13.51	23.32
Summer Scenario Perfect Information	-1.15	57.34	98.97
Winter Scenario Perfect Information	2.63	14.33	24.74

We see that the current situation is the most expensive for the households, but also the least sustainable, since no renewable energy is used for EV charging. The adoption of the proposed IS artifact brings reduction to the electricity costs that are higher in summer when the solar generation is high. This time of the year the renewable usage reaches the 93% of the total household consumption, increasing the sustainability levels. Comparing the summer and winter scenarios using the IS artifact with the respective perfect information scenarios we observe that, of course, the costs are even lower and the renewable usage higher. However, it would not be realistic to assume that the smart homes operate in perfect information conditions, where all smart homes have full information about weather and future PV generation and no forecasting mechanisms are required. The difference between IS artifact and perfect information scenarios is not very high, making the adoption of the IS artifact appealing to smart home owners.

3.5 Conclusion & Future Work

We propose an IS artifact that supports smart home owners with their complex decisions in the smart grid. The goal of our IS artifact is to ensure low costs and high renewable usage for the household owners. At the core of the artifact lies an intelligent agent that forecasts the PV generation based on the weather conditions and schedules EV charging so that costs are reduced and renewable usage is maximized. Evaluating our artifact on real-world data, we observe that costs can be reduced significantly both in the optimistic summer scenario and in the least optimistic winter scenario. At the same time we observe renewable usage increase with a maximum of 93%, for the most optimistic scenario.

The proposed results are generalizable to households that own an EV and a PV and fulfil the assumptions listed in this chapter. If these assumptions are violated, then potentially the outcome will not be the same. For example, if the EV is not parked at a place where it can absorb the electricity generated by the PV panel, the result on renewable usage and peak demand reduction might not be so apparent. This issue, can be resolved by assuming a battery placed at home (or a Tesla Wall). Regarding the numerical results of this chapter, they will be different if the artifact is applied on different data sets. However, it is expected that on any data set of this type, the artifact will have a positive impact on sustainability and peak demand reduction.

Another limitation of this chapter is that it does not assume any correlation between driving behavior and weather conditions. In reality, these two variables are related since the driving behavior is influenced by the weather conditions. Therefore, in the future we will integrate the correlation between the weather and driving behavior in the model. In addition, we will add a market layer to our analysis, in which the smart homes will be trading electricity in the market through an aggregator (energy provider). Finally, we will benchmark our artifact with the state of the art artifacts in this domain.

Appendix—Summary of Notation and Abbreviations

Table 3.3: Summary of Notation

Symbol	Definition
$c(\cdot)$	electricity cost function
CA_t	individual charging availability per time slot t
$\mathbf{G} = \{g_t\}$	PV generation values for each time slot t
M	sufficiently large number
n	index of EV owner's driving activities
$\mathbf{P} = \{P_t\}$	set with energy price per consumption unit for all time slots t
SoC_t	battery's state of charge in timeslot t
$\mathbf{T} = \{t\}$	discrete set of time intervals, $t \in \{1, \dots, T\}$
t_a^n	individual arrival time for driving activity with index $n \in \{1, \dots, N\}$
t_d^n	individual departure time for driving activity with index $n \in N$
$X_{max,t}$	upper bound of $x_{c,t}$
$\mathbf{W} = \{W_{1,t}, \dots, W_{ \mathbf{W} }\}$	set with weather conditions for all time slots t
\mathbf{x}_c	vector of EV charging consumption for each t
$\epsilon_t = g_t - \hat{g}_t$	forecasting error of the ensemble learning
$\theta = \{t_d^n, t_a^n\}$	individual departure and arrival preferences, accompanying each driving activity with index $n \in N$
ρ	capacity/distance rate

Table 3.4: Summary of Abbreviations

Abbreviation	Definition
EPEX	European Power Exchange
EV	Electric Vehicle
ICT	Information and Communication Technology
MAPE	Mean Average Percentage Error
MIEM	Mobility Integrated Energy Management
MSE	Mean Squared Error
PV	Photovoltaic
RMSE	Root Mean Squared Error

Chapter 4

Maximizing Social Welfare in Grid Resource Allocation for Electric Vehicle Charging¹

4.1 Introduction

Electricity markets are undergoing fundamental changes moving toward a new digitized era where consumers own *smart* appliances, reside in *smart* homes and can interact with the market operator via an ICT infrastructure (Ketter et al., 2016). This new formation of the electricity grid is known as *smart grid* (Amin and Wollenberg, 2005). The term smart grid is used to “describe a next-generation electrical power system that is typified by the increased use of communications and information technology in the generation, delivery and consumption of electrical energy”². What makes the smart grid different from its predecessor - the traditional grid - is the large scale integration of renewable sources and the active role electricity consumers have in it, not only by consuming, but also by producing electricity (photovoltaic panels, wind turbines, electric vehicle batteries, etc.).

¹Parts of this chapter have appeared in the following peer reviewed conference proceedings: Valogianni, K., Gupta, A., Ketter, W., Sen, S. & van Heck, E.(2016). Using Optimal Grid Resources for Coordinating Electrical Vehicle Charging. In Winter Conference on Business Intelligence (WCBI 2016), Snowbird, Utah.
Valogianni, K., Gupta, A., Ketter, W., Sen, S. & van Heck, E.(2015). Maximizing Social Welfare in Grid Resource Allocation for Electric Vehicle Charging. In Conference on Information Systems and Technology, (CIST 2015), Philadelphia, PA.

²<http://smartgrid.ieee.org/ieee-smart-grid> [Date Accessed: June 21st, 2015]

In this chapter, we focus on the electric vehicle (EV) integration in the smart grid. EVs are important components of an electricity market for two main reasons. Firstly, they are significant electricity consumers (for example a US household might consume on average 30 kWh³ a day and an EV battery needs from 24kWh to 80kWh to charge fully). This means that a large scale EV integration is going to put the grid's infrastructure under critical strain. An illustrative example is California's grid, which must now accommodate demand from 100,000 EVs⁴. For this purpose, the California grid operator has decided to expand the capacity of the grid, so that it can service these extra electricity customers. Undoubtedly, the EV integration needs to be scheduled properly so that the grid is not overloaded and the EV owners are serviced without problems via the existing infrastructure. The second reason that EVs are of particular importance is that they own batteries which can store electricity. So far, electricity cannot be stored (in large amounts), therefore the storage features of EVs are expected to revolutionize the logic behind electricity markets. Practically, now EV owners can charge their cars when prices are low, and feed electricity back to the grid when prices are high or when there is a shortage in supply.

We examine the scheduling of EV charging from the market operator's (grid operator) point of view. A grid operator must schedule EV charging so that the grid is not overloaded and the consumers are serviced with the lowest delay possible. A successful scheduling of EV charging ensures the grid's stability without installing new capacity infrastructure (which would be an unsustainable solution, since more raw materials need to be consumed). Reinforcing the grid infrastructure to accommodate electricity demand is the traditional way of coping with electricity peak demand on the grid. However, this solution requires more raw materials (such as copper, etc.), which makes it unsustainable and costly. We propose an auction-based mechanism that schedules EV charging and determines the prices to the customers in real-time. Auctions, unlike posted-price and capacity allocation mechanisms, are preferred when the demand is not known or easy to estimate (Bapna et al., 2003). Therefore, in this particular problem, in which the grid operator does not know at each point in time how many EVs will require charging, auctions will contribute to allocating the grid resources efficiently. Our decision variables determine how many EV customers to accept at each point in time and at what price, given the fixed capacity of the grid (since the grid has fixed capacity, and given the high EV charging demand, some customers might not be serviced).

³<http://www.eia.gov/tools/faqs/faq.cfm?id=97&t=3>

⁴<http://www.greentechmedia.com/articles/read/california-steps-up-again-on-electric-vehicles> [Date Accessed: June 21st, 2015]

4.2 Related Work

Auctions have been used in many different application domains as a means of distributing goods. These domains vary from eBay (Bajari and Hortacsu, 2003) and web capacity auctions (Bapna et al., 2003, 2008) to flowers (Kambil and van Heck, 1998; Lu et al., 2013) and wireless spectrum auctions (Cramton, 1997). In all these different auction mechanisms, the notion of smart markets (Bichler et al., 2010) is the key for allocating goods and payments efficiently. Auction mechanisms used in EV charging assume different objectives and behavioral characteristics on the customer's side. Acha et al. (2011) propose a centralized capacity coordination mechanism applicable to EV charging, under the profit maximization objective. Rigas et al. (2013) introduce a centralized mechanism for matching EV charging to various charging stations, accounting for spatial and temporal dimensions. Bhattacharya et al. (2014) extend the concept of second price auctions to be applicable to EV charging. The authors assume a revenue maximization objective and adopt the Vickrey-Clarke-Groves (VCG) mechanism (Vickrey, 1961; Clarke, 1971; Groves, 1973) to clear the auction, which increases the computation time. De Craemer et al. (2014) present a dual implementation for shifting EV charging over time based on a central auctioneer, whereas Robu et al. (2013) introduce an online auction mechanism for EV charging coordination. Stein et al. (2012) extend the previous mechanism by adding pre-commitment attributes in the auction. Kahlen and Ketter (2015) propose a centralized fleet management system which is responsible for coordinating the EV charging, while Vandael et al. (2013) describe a three-step top-down charging coordination mechanism.

All these manuscripts assume a profit maximizer auctioneer, but this assumption does not hold most of the times in the smart grid. The smart grid operator is not interested in making profits, but is mainly interested in servicing all the EV customers in the market without creating bottlenecks and without putting the infrastructure under strains (Wissner, 2011; Kanchev et al., 2011). In other words, the grid operator's main objective is to service the EV owners while decreasing volatility in the electricity demand, since increased volatility threatens grid's infrastructure due to high peaks. We take the stand point of the grid operator and propose a social welfare maximization mechanism that assists her in EV scheduling decisions, ensuring at the same time a less volatile demand curve. We define social welfare as *the delay cost the system (smart grid) has to suffer in order for all the EV owners to be serviced*. As delay cost, we define the cost that each individual incurs during the time that she is not be able to use the EV, because it is plugged-in. For some individuals that they need the EV for driving, every hour that they are unable to use their car is very costly. For other individuals, this charging time might not be of high

cost, since they might not need the car for driving. As system's delay cost, we define the overall delay cost of all the EV owners who are using the grid infrastructure for charging. Therefore, we model the EV owner's utility by incorporating a delay component, which indicates the urgency of the EV charging. EV owners who need their EV charged urgently, have a higher delay cost, compared to EV owners who are willing to wait a bit till their EV is charged. With this approach we account for individual delays suffered by the consumers and propose a mechanism to schedule the EV charging so that delays are minimized and the total delay cost in the system is mitigated. We demonstrate how these metrics change after applying the proposed mechanism.

4.3 Model Formulation and Structural Analysis

We approach the whole EV charging capacity allocation as a knapsack problem with the grid capacity being the "knapsack" in our case. We should note that currently there are different EV charging speeds available in the smart grid and these charging speeds have a different effect on the electricity peak demand (Valogianni et al., 2014b). They are also found in the literature as levels of charging and the main ones are: 3.3 kW, 7 kW, 24 kW and 43 kW (these numbers reflect the average of charging speed given by a charging pole)⁵. These charging levels are offered through charging poles, and it can be the case that each charging pole offers a different charging speed (e.g., there might be adjacent charging poles that offer different charging speeds). We will refer to these levels as charging speeds, since they practically represent different classes of service. Once an EV is plugged in and allocated to a charging speed, this charging speed cannot vary over time and it is constant throughout the whole charging session. A direct analogy can be found in the grid computing and internet literature (Bapna et al., 2008), where different classes represent different internet speeds.

Our goal is to allocate the EV owners' requests for charging to different charging speeds so that the grid benefits and the EV owners suffer as low delays as possible. By delay, we indicate the time that the EV is plugged in for charging and cannot be used by the EV owner for driving. The benefits to the grid are measured by the electricity peak demand reduction, since the peak demand is the main determinant for installing new infrastructure (Strbac, 2008). Thus, reducing peak demand means reducing the need for extra infrastructure and therefore, higher sustainability on the grid (Watson et al., 2010). Another metric to quantify the grid's benefits is the peak-to-average ratio (PAR) which indicates how volatile a curve is. Consequently, for a stable grid it is desirable to

⁵The last charging level is not currently available in the US and is available in some regions of Europe and in Japan.

have a low PAR which means low volatility and higher stability. The benefits for the consumers are measured by the total delay cost the system (all EV owners serviced by the same infrastructure) suffers. In other words, the consumers receive better quality of service if they suffer lower delay until their EV gets charged.

4.3.1 Assumptions

Our set-up assumes $i \in \{1, \dots, N\}$ EV customers (bidders) and a service operator (smart grid operator) who receives requests for charging. These requests include a total charging need in kWh, ω_i and an arrival and a departure time, t_a^i and t_d^i , respectively. Each request for an amount ω_i is accompanied with a bid for this amount b_i and a cost δ_i over the delay she might suffer. The requested energy quantity ω_i can be charged at one of the charging speeds (in kW) r_j , $j \in \{1, \dots, z\}$. When a request for charging ω_i is allocated to a charging speed r_j then, the binary variable $x_{i,j}$ becomes equal to 1 indicating this allocation. Respectively, after such an allocation, the completion time of this charging request is $\tau_j(\omega_i)$ and holds only if $x_{i,j} = 1$ (otherwise $\tau_j(\omega_i)$ does not exist):

$$\tau_j(\omega_i) = t_{0,i} + \frac{\omega_i}{x_{i,j} \cdot r_j}, \quad x_{i,j} \neq 0 \quad (4.1)$$

We call $\tau_j(\omega_i)$ the delay each consumer i has to suffer until she gets the car charged (so that the amount ω_i is loaded) and is measured as the time the EV gets plugged in till the time it is charged and ready to be used by the owner. This delay includes the time $t_{0,i}$ that the car i is plugged in but not getting charged (since the grid capacity is used by other EVs) and the time $\frac{\omega_i}{x_{i,j} \cdot r_j}$ that it takes to get charged at a certain charging speed r_j . Each delay has a different cost for each EV owner, therefore we denote this cost by δ_i and increasing delay cost indicates increasing urgency for charging completion. This implies that consumers with high delay costs are willing to pay a higher price for EV charging. Therefore, we assume that there is a direct analogy between the bids and the delay costs. In other words, consumers that submit high delay costs to the auctioneer have a higher valuation of the requested service. If a request ω_i is allocated to a high speed r_j , the duration for its completion decreases. Therefore, urgent requests must be allocated to high charging speeds (the urgency of a request is indicated by the delay cost, δ_i).

Each EV customer (bidder) i has a utility function over a request ω_i such that:

$$U(\omega_i, \tau_j(\omega_i), \delta_i, b_i) = \gamma_i \cdot \omega_i - \delta_i \cdot \tau_j(\omega_i) - b_i \quad (4.2)$$

By γ_i we denote the weight each EV owner i puts on receiving an electricity amount ω_i and this weight is not dependent on the speed this request gets allocated. The variable b_i is the bid each EV owner submits for a service ω_i (before the consumer knowing the charging speed allocation). In this utility function we assumed a linear relationship between utility and delay cost, since linear relationships between congestion and utility have been used in the literature (Sen et al., 2010). We plan to experiment with other structural forms in the future.

Normally, utility functions for consumers involved in scheduling processes include an extra component, which represents the congestion factor (Sen et al., 2010). This congestion factor, in our case is indicated by the delay $\tau_j(\omega_i)$ each consumer is suffering as a result of the presence of congestion (or not). Therefore, by using this metric, we account for congestion at an aggregate level, but we also account for the effect of this congestion to each individual separately. We introduce the coefficient δ_i to denote the emphasis (cost) each consumer puts on this delay (in utility units per time units). For some consumers a potential delay might not be important since they have no urgency in using their EV, whereas for others this delay might be important since they need their EV for driving.

4.3.2 Multiple Vickrey Auction

The nature of our particular application requires real-time decision-making capabilities, since the smart grid is a fast changing environment with lots of information flows (electricity prices, capacity available, consumer preferences, availability of electricity). Therefore, it is important for the grid operator to be able to allocate capacity and payments to the EV customers in real time. Prior research in other application domains has used the Vickrey-Clarke-Groves (VCG) mechanism for payment and capacity allocation (Dash et al., 2007; Dimakis et al., 2006; Krishna, 2002). This mechanism ensures incentive compatibility and is therefore, preferred in set-ups where the consumers' true valuations are not known. However, it has a significant downside which is a high computational complexity (Nisan et al., 2007). Therefore, it might create significant burdens in applying this mechanism in real-world problems, where computation of the solution in real time is required.

Typically, in the smart grid large numbers of customers bid in the auction at each point in time. Therefore, despite the incentive compatibility, the VCG mechanism is not suitable for this particular case. It is interesting to mention that 10 participants in the auction might increase the payment calculation time to minutes. Considering that in the state of California there are currently 100,000 EVs, one can understand that the VCG

mechanism will create burdens in reducing congestion in the grid. Therefore, we use as starting point the Multiple Vickrey Auction (MVA) approach proposed by Bapna et al. (2005), modified to be applicable in our particular social welfare maximization set-up and prove new theoretical properties. According to this mechanism if there are m units of a good for sale, then the m highest bids win and the $(m + 1)^{st}$ bid becomes the price paid for each of the units sold. In our case we allocate the bidders to classes based on their delay costs, since we assumed before that a higher delay cost is linked to a higher bid (consumers who value a service more are the consumers who have urgent deadlines and need their car charged quickly).

This mechanism while computationally tractable, does not ensure incentive compatibility. However, for large number of bidders, which is the case in our market, incentive compatibility is not an issue due the law of large numbers. Practically, because of the large numbers of people involved, the impact of each person's misaligned incentives will be negligible. Usually, in complex market mechanisms, such as the Federal Communications Commission (FCC) wireless spectrum auctions, the incentive compatibility is not of real concern. Furthermore, Bapna et al. (2005) prove that MVA is *ex-post* incentive compatible and therefore, there is no actual gain for the bidders if they are untruthful. A grid operator gains more significant advantages by being able to compute the solution in real-time than by guaranteeing incentive compatibility of the bidders.

Adapting the MVA logic to our set-up, we formulate the grid operator's problem as a welfare maximization problem. Since we defined the social welfare as reduced delays, the grid operator's problem is practically the aggregate delay cost minimization problem. The delay cost minimizing knapsack formulation with uniform pricing is presented below:

$$\min_{x_{i,j}} \sum_i \sum_j \delta_i \cdot \tau_j(\omega_i) \Leftrightarrow \min_{x_{i,j}} \sum_i \sum_j \delta_i \cdot (t_{0i} + \frac{\omega_i}{x_{i,j} \cdot r_j}) \quad (4.3)$$

where $\tau_j(\omega_i)$ is the delay each EV owner i suffers after being allocated to a charging speed r_j , with a respective delay cost δ_i . The constraints of our problem are:

$$x_{i,j} \cdot \delta_i > \delta_{i+1} \quad \forall i \in \{1, \dots, N\} \quad \text{and} \quad \forall j \in \{1, \dots, z\} \quad (4.4)$$

$$\sum_j x_{i,j} \leq 1 \quad \forall i \in \{1, \dots, N\} \quad (4.5)$$

$$\sum_i \sum_j x_{i,j} \cdot r_j \leq C \quad (4.6)$$

$$x_{i,j} = \{0, 1\} \quad \forall i \in \{1, \dots, N\} \quad \text{and} \quad \forall j \in \{1, \dots, z\} \quad (4.7)$$

$$\delta_i \geq 0 \quad \forall i \in \{1, \dots, N\} \quad (4.8)$$

Constraint (4.4) ensures that the EV owners are allocated to charging speeds based on their delay cost declarations (and consequently their bid submissions) in ascending order. The consumer who values the service the most gets the highest quality of service (charging speed). Constraints (4.5) and (4.6) indicate that the variables $x_{i,j}$ are binary and denote whether a service is allocated to a consumer or not. Constraint (4.7) ensures that the grid is stabilized by not exceeding the capacity available (in our case we set C lower than the maximum capacity available to allow for a buffer on the grid). Equation (4.8) indicates the assumption that delay costs, processed by the mechanism, have always positive values (negative delay costs are not accepted by the auction).

4.4 Results

In this section we present the theoretical properties of the proposed mechanism and empirical results after applying our mechanism on real-world data.

4.4.1 Theoretical Properties

Prior use of the MVA mechanism for capacity allocation (Bapna et al., 2005, 2008) showed that multiple classes for different service levels were optimal from a revenue maximization point of view. Consequently, we would expect that having multiple classes (charging speed levels) would be more beneficial for the grid operator. However, we see that the lowest delay cost in the system is caused by having only one charging speed level, which corresponds to the highest charging speed. This is an interesting result, given the nature of our problem. It practically means that it is optimal for the grid operator to have only the fastest charging speed in place, and schedule the EV owners in this single class. We will prove the optimality of the single charging class both in the case that the EV owners have the same delay costs $\delta_i = \delta \quad \forall i \in \{1, \dots, N\}$ and in the case that EV owners have different delay costs δ_i . The delay costs are the main differentiating factor across EV customers, therefore, we need to account for both cases that EV customers have the same (Proposition 1) and different delay costs (Proposition 2).

Proposition 1. Having one charging speed level (one class), corresponding to the fastest charging speed, is optimal from a social welfare maximization point of view, when EV owners have the same delay costs δ .

Proof. First, we will show that having one charging class is optimal for all the consumers. Assume the fastest charging speed is r_1 and takes time $\tau_1(\omega_i) = \frac{\omega_i}{r_1}$ to charge an EV customer's request ω_i at this charging level ($t_{0,i} = 0$ here because we assume that this is the first customer in the market). Let us choose any two charging classes $j, k > 1$ and their respective charging speeds r_j and r_k . For the customer 1, who is scheduled in the charging class j , we need charging time $\tau_j(\omega_1) = j \cdot \tau_1(\omega_1)$ and for the customer 2, scheduled in charging class k , we need charging time $\tau_k(\omega_2) = k \cdot \tau_1(\omega_2)$ (in both cases we assume $t_{0,1} = t_{0,2} = 0$). We can see their scheduling in Figure 4.1. Now, if we increase both their charging speeds to the levels $r_{j'} > r_j, r_{k'} > r_k$ with $j', k' > 1$, then the times they need for charging will be respectively $\tau_{j'}(\omega_1) = j' \cdot \tau_1(\omega_1)$ and $\tau_{k'}(\omega_2) = k' \cdot \tau_1(\omega_2)$. But since $r_{j'} > r_j$ then $j' < j$ (r_1 is the highest charging speed) we have lower charging times $\tau_{j'}(\omega_1) < \tau_j(\omega_1)$. Similarly, $\tau_{k'}(\omega_2) < \tau_k(\omega_2)$. Then, the total delay the system suffers (all EV owners together) in the first case is $\tau_k(\omega_2) + \tau_j(\omega_1) > \tau_{k'}(\omega_2) + \tau_{j'}(\omega_1)$. Now, if we make the assumption that $r_{j'} = r_{k'}$, then we have $\tau_k(\omega_2) + \tau_j(\omega_1) > \tau_{k'}(\omega_2) + \tau_{j'}(\omega_1) > \tau_{j'}(\omega_2) + \tau_{j'}(\omega_1)$, which holds for every $j' > 1$. This shows that having one charging class is optimal both for each EV customer and for the system, overall.




Figure 4.1: Illustration of charging scheduling in classes j, k .

As a second step, we will show that this single class needs to be the fastest charging level. Assume we have N customers in the system and we schedule them all at the lowest charging level. Then, their charging will consume a capacity of $N \cdot r_z = C$, where r_z is the lowest charging level. This charging will last time $\tau_z(\omega_i) \quad \forall i \in \{1, \dots, N\}$ and therefore

the total delay cost the system has to suffer (let us call it λ_1) is:

$$\lambda_1 = \sum_{i=1}^N \delta \cdot \tau_z(\omega_i) = \delta \cdot \sum_{i=1}^N \tau_z(\omega_i) \quad (4.9)$$

If r_1 is the fastest charging speed, then scheduling all customers in this level will require time $\tau_1(\omega_i) = \frac{\tau_z(\omega_i)}{z} \implies \tau_z(\omega_i) = z \cdot \tau_1(\omega_i)$. Therefore, Equation (4.9) becomes:

$$\lambda_1 = \sum_{i=1}^N \delta \cdot \tau_z(\omega_i) = \delta \cdot z \cdot \sum_{i=1}^N \tau_1(\omega_i) \quad (4.10)$$

If at charging level r_z we could charge all N customers, now at level r_1 we can charge $\frac{N}{z}$ customers each time. This charging lasts time $\sum_{i=1}^{\frac{N}{z}} \tau_1(\omega_i)$ for the first $\frac{N}{z}$ customers ($t_{0,1} = 0$ for this first customer group), $\sum_{i=1}^{\frac{N}{z}} \tau_1(\omega_i) + \sum_{i=\frac{N}{z}+1}^{2 \cdot \frac{N}{z}} \tau_1(\omega_i)$ for the second $\frac{N}{z}$ customers ($t_{0,i} = \sum_{i=1}^{\frac{N}{z}} \tau_1(\omega_i), \forall i \in \{\frac{N}{z}+1, \dots, 2 \cdot \frac{N}{z}\}$ because they have to wait till the first $\frac{N}{z}$ are charged and the grid capacity is free and the rest is actual charging time $\tau_1(\omega_i)$), etc. So the total delay cost (λ_2) is:

$$\begin{aligned} \lambda_2 = & \delta \cdot \sum_{i=1}^{\frac{N}{z}} \tau_1(\omega_i) + \delta \cdot \sum_{i=1}^{\frac{N}{z}} \tau_1(\omega_i) + \delta \cdot \sum_{i=\frac{N}{z}+1}^{2 \cdot \frac{N}{z}} \tau_1(\omega_i) \\ & + \delta \cdot \sum_{i=1}^{\frac{N}{z}} \tau_1(\omega_i) + \delta \cdot \sum_{i=\frac{N}{z}+1}^{2 \cdot \frac{N}{z}} \tau_1(\omega_i) + \delta \cdot \sum_{i=2 \cdot \frac{N}{z}+1}^{3 \cdot \frac{N}{z}} \tau_1(\omega_i) + \dots \end{aligned} \quad (4.11)$$

Equation (4.11) becomes:

$$\lambda_2 = \delta \cdot \left(\sum_{i=1}^{\frac{N}{z}} \tau_1(\omega_i) + \sum_{i=1}^{2 \cdot \frac{N}{z}} \tau_1(\omega_i) + \dots + \sum_{i=1}^N \tau_1(\omega_i) \right) \quad (4.12)$$

From (4.10) and (4.12) we have $\lambda_1 > \lambda_2 \Leftrightarrow \delta \cdot z \cdot \sum_{i=1}^N \tau_1(\omega_i) > \delta \cdot \left(\sum_{i=1}^{\frac{N}{z}} \tau_1(\omega_i) + \sum_{i=1}^{2 \cdot \frac{N}{z}} \tau_1(\omega_i) + \dots + \sum_{i=1}^N \tau_1(\omega_i) \right)$ which holds for every $z > 1$. This means that the total delay cost the system suffers is the lowest if we have a single charging class and this charging class corresponds to the highest charging speed. \square

Proposition 2. Having one charging speed level (one class), corresponding to the fastest charging speed, is optimal from a social welfare maximization point of view, when EV owners have different delay costs δ_i .

Proof. The single class optimality proof is the same as in Proposition 1. We will show that the fastest class is optimal as single class. Assume we have N customers charging requests ω_i before they get allocated to a charging speed j . We schedule all N customers at the lowest charging level. Then, their charging will consume a capacity of $N \cdot r_z = C$, where r_z is the lowest charging level. This charging will last time $\tau_z(\omega)$ and therefore the total delay cost the system has to suffer is (similarly to Equation (4.9)):

$$\lambda_1 = \sum_{i=1}^N \delta_i \cdot \tau_z(\omega_i) = z \cdot \sum_{i=1}^N \delta_i \cdot \tau_1(\omega_i) \quad (4.13)$$

If at charging level r_z we could charge all N customers, now at level r_1 we can charge $\frac{N}{z}$ customers each time. This charging lasts time $\sum_{i=1}^{\frac{N}{z}} \tau_1(\omega_i)$ for the first $\frac{N}{z}$ customers ($t_{0,1} = 0$ for this first customer group), $\sum_{i=1}^{\frac{N}{z}} \tau_1(\omega_i) + \sum_{i=\frac{N}{z}+1}^{2 \cdot \frac{N}{z}} \tau_1(\omega_i)$ for the second $\frac{N}{z}$ customers ($t_{0,i} = \sum_{i=1}^{\frac{N}{z}} \tau_1(\omega_i), \forall i \in \{\frac{N}{z}+1, \dots, 2 \cdot \frac{N}{z}\}$ because they have to wait till the first $\frac{N}{z}$ are charged and the grid capacity is free and the rest is actual charging time $\tau_1(\omega_i)$), etc. So the total delay cost (λ_2) is:

$$\begin{aligned} \lambda_2 = & \sum_{i=1}^{\frac{N}{z}} \delta_i \cdot \tau_1(\omega_i) + \sum_{i=1}^{\frac{N}{z}} \delta_i \cdot \tau_1(\omega_i) + \sum_{i=\frac{N}{z}+1}^{2 \cdot \frac{N}{z}} \delta_i \cdot \tau_1(\omega_i) \\ & + \sum_{i=1}^{\frac{N}{z}} \delta_i \cdot \tau_1(\omega_i) + \sum_{i=\frac{N}{z}+1}^{2 \cdot \frac{N}{z}} \delta_i \cdot \tau_1(\omega_i) + \sum_{i=2 \cdot \frac{N}{z}+1}^{3 \cdot \frac{N}{z}} \delta_i \cdot \tau_1(\omega_i) + \dots \end{aligned} \quad (4.14)$$

Equation (4.14) becomes:

$$\lambda_2 = \sum_{i=1}^{\frac{N}{z}} \delta_i \cdot \tau_1(\omega_i) + \sum_{i=1}^{2 \cdot \frac{N}{z}} \delta_i \cdot \tau_1(\omega_i) + \dots + \sum_{i=1}^N \delta_i \cdot \tau_1(\omega_i) \quad (4.15)$$

From (4.13) and (4.15) we have:

$$\lambda_1 > \lambda_2 \Leftrightarrow z \cdot \sum_{i=1}^N \delta_i \cdot \tau_1(\omega_i) > \sum_{i=1}^{\frac{N}{z}} \delta_i \cdot \tau_1(\omega_i) + \sum_{i=1}^{2 \cdot \frac{N}{z}} \delta_i \cdot \tau_1(\omega_i) + \dots + \sum_{i=1}^N \delta_i \cdot \tau_1(\omega_i) \quad (4.16)$$

which holds for every $z > 1$. This means that the total delay cost the system suffers is the lowest if we have a single charging class and this charging class corresponds to the highest charging speed. \square

In Proposition 3, we will show that the customer with the highest delay cost needs to be scheduled first, so that the social welfare to be maximized.

Proposition 3. Charging needs do not influence the scheduling and the customer with the highest delay cost δ_i needs to be scheduled first.

Proof. Let us assume that we have $j \in \{1, \dots, z\}$ charging levels r_j . Assume we have two customers with utility functions $U_1 = U(\omega_1, \tau_j(\omega_1), \delta_1, b_1)$ and $U_2 = U(\omega_2, \tau_j(\omega_2), \delta_2, b_2)$ respectively and delay costs δ_1 and δ_2 . Let us assume that $\delta_1 > \delta_2$.

We will show that the system (both of the EV customers) has the highest utility, if we schedule the customer with the highest delay cost (customer 1) first. We prove this by assuming the opposite: that the total utility is the lowest, if we schedule customer 1 first. In this case we have that the delay for the first customer will be $\tau_j(\omega_1) = \frac{\omega_1}{r_j}$, $t_{0,1} = 0$ because this customer is scheduled first. The delay for the second customer will be $\tau_{j'}(\omega_2) = t_{0,2} + \frac{\omega_2}{r_{j'}}$. Here, we assumed that customer 2 is scheduled at a charging speed $r_{j'}$ that might be equal or different from r_j . Also, $t_{0,1} = t_{0,2} = 0$ in the case that $j \neq 1$ and $j' \neq 1$ (both customers start charging at the same time) or $t_{0,2} = |\tau_1(\omega_1) - \tau_1(\omega_2)|$, if both are charged at the highest speed (one has to wait for the other). If $t_{0,1} = t_{0,2} = 0$, both customers can charge at the same time, so there is no need to prioritize them. If both are charged at the highest speed ($j = 1$), then we have $t_{0,1} = 0$ and $t_{0,2} = |\tau_1(\omega_2) - \tau_1(\omega_1)|$. We will examine this case ($j = 1$), because, this is the situation where prioritization is needed. Therefore, for the rest of the proof, we assume $j = 1$ which is also shown in Propositions 1 and 2 that it is optimal for the EV owners. This means that customer 2 needs to wait for customer 1 to be fully charged and then start charging (Figure 4.2). In the opposite case that we schedule customer 2 first, we have $t'_{0,1} = |\tau_1(\omega_1) - \tau_1(\omega_2)| = t_{0,2}$ and $t'_{0,2} = 0$.

Since $\tau_1(\omega_1) < \tau_1(\omega_2)$ and also we assumed that the total utility of the system is lower if we schedule customer 1 first, we have the total utility of the system:

$$\begin{aligned}
 \sum_{i=1}^2 U_i &< (\sum_{i=1}^2 U_i)' \iff \\
 (\gamma_1 \cdot \omega_1 - \delta_1 \cdot \tau_1(\omega_1) - b_1) + (\gamma_2 \cdot \omega_2 - \delta_2 \cdot \tau_1(\omega_2) - b_2) &< \\
 (\gamma_1 \cdot \omega_1 - \delta_1 \cdot \tau_1(\omega_1)' - b_1) + (\gamma_2 \cdot \omega_2 - \delta_2 \cdot \tau_1(\omega_2)' - b_2) &\iff \\
 -\delta_1 \cdot \frac{\omega_1}{r_1} - \delta_2 \cdot (t_{0,2} + \frac{\omega_2}{r_1}) &< -\delta_1 \cdot (t'_{0,1} + \frac{\omega_1}{r_1}) - \delta_2 \cdot \frac{\omega_2}{r_1} \iff \delta_1 < \delta_2
 \end{aligned} \tag{4.17}$$

which is not true because of our initial assumption that $\delta_1 > \delta_2$.

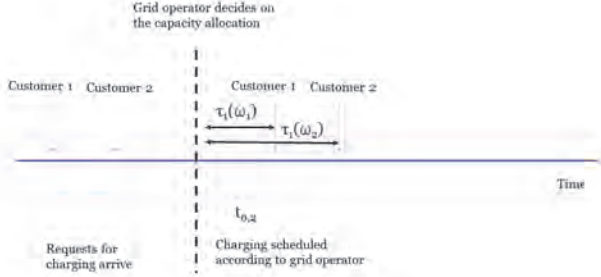


Figure 4.2: Illustration of scheduling for 2 customers (customer 1 scheduled first).

We will show now that this is generalizable to 3 EV customers with utility functions $U_1 = U(\omega_1, \tau_1(\omega_1), \delta_1, b_1)$ and $U_2 = U(\omega_2, \tau_1(\omega_2), \delta_2, b_2)$ and $U_3 = U(\omega_3, \tau_1(\omega_3), \delta_3, b_3)$ and delay costs $\delta_1 > \delta_2 > \delta_3$. In this case we have 6 combinations and we have to compare the scheduling case U_1, U_2, U_3 with all the others U_1, U_3, U_2 ; U_2, U_1, U_3 ; U_2, U_3, U_1 ; U_3, U_1, U_2 ; U_3, U_2, U_1 . We must show that the utility for the system (3 EV customers) is higher for the case U_1, U_2, U_3 compared to all the others.

To increase clarity in the calculations we set as $\Phi_i = \gamma_i \cdot \omega_i - b_i$, so that each utility function U_i becomes:

$$U_i = U(\omega_i, \tau_1(\omega_{i,1}), \delta_i, b_i) = \gamma_i \cdot \omega_i - \delta_i \cdot \tau_1(\omega_i) - b_i = \Phi_i - \delta_i \cdot \tau_1(\omega_i) \quad (4.18)$$

Generalizing Figure 4.2 to the 3-customer case we have Figure 4.3. Let us assume that if we schedule Customer 1 first, Customer 2 second and Customer 3 third, the delay for the first customer will be 1 time unit $\tau_1(\omega_1) = 1$ and for the second customer will be 2 time units $\tau_1(\omega_2) = 2$ and the delay for the third customer will be 3 time units $\tau_1(\omega_3) = 3$ (this holds for any numbers for the delay but we use these for demonstration purposes). We must show that all the following equations hold at the same time:

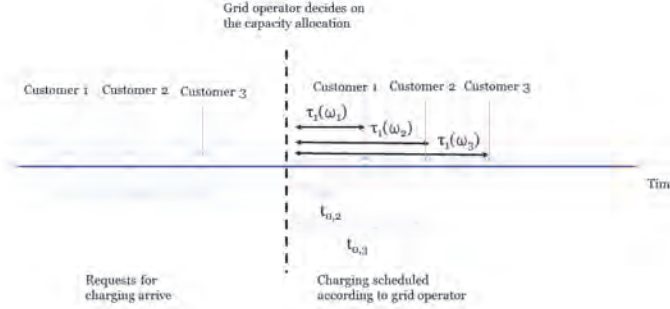


Figure 4.3: Illustration of scheduling for 3 customers.

$$\left\{ \begin{array}{l} \Phi_1 - \delta_1 + \Phi_2 - 2\delta_2 + \Phi_3 - 3\delta_3 > \Phi_1 - \delta_1 + \Phi_3 - 2\delta_3 + \Phi_2 - 3\delta_2 \\ \Phi_1 - \delta_1 + \Phi_2 - 2\delta_2 + \Phi_3 - 3\delta_3 > \Phi_2 - \delta_2 + \Phi_1 - 2\delta_1 + \Phi_3 - 3\delta_3 \\ \Phi_1 - \delta_1 + \Phi_2 - 2\delta_2 + \Phi_3 - 3\delta_3 > \Phi_2 - \delta_2 + \Phi_3 - 2\delta_3 + \Phi_1 - 3\delta_1 \\ \Phi_1 - \delta_1 + \Phi_2 - 2\delta_2 + \Phi_3 - 3\delta_3 > \Phi_3 - \delta_3 + \Phi_1 - 2\delta_1 + \Phi_2 - 3\delta_2 \\ \Phi_1 - \delta_1 + \Phi_2 - 2\delta_2 + \Phi_3 - 3\delta_3 > \Phi_3 - \delta_3 + \Phi_2 - 2\delta_2 + \Phi_1 - 3\delta_1 \end{array} \right. \quad (4.19)$$

To prove (4.19) we assume that the opposite holds in all cases. In other words we assume:

$$\left\{ \begin{array}{l} \Phi_1 - \delta_1 + \Phi_2 - 2\delta_2 + \Phi_3 - 3\delta_3 < \Phi_1 - \delta_1 + \Phi_3 - 2\delta_3 + \Phi_2 - 3\delta_2 \\ \Phi_1 - \delta_1 + \Phi_2 - 2\delta_2 + \Phi_3 - 3\delta_3 < \Phi_2 - \delta_2 + \Phi_1 - 2\delta_1 + \Phi_3 - 3\delta_3 \\ \Phi_1 - \delta_1 + \Phi_2 - 2\delta_2 + \Phi_3 - 3\delta_3 < \Phi_2 - \delta_2 + \Phi_3 - 2\delta_3 + \Phi_1 - 3\delta_1 \\ \Phi_1 - \delta_1 + \Phi_2 - 2\delta_2 + \Phi_3 - 3\delta_3 < \Phi_3 - \delta_3 + \Phi_1 - 2\delta_1 + \Phi_2 - 3\delta_2 \\ \Phi_1 - \delta_1 + \Phi_2 - 2\delta_2 + \Phi_3 - 3\delta_3 < \Phi_3 - \delta_3 + \Phi_2 - 2\delta_2 + \Phi_1 - 3\delta_1 \end{array} \right. \iff \left\{ \begin{array}{l} \delta_2 < \delta_3 \\ \delta_1 < \delta_3 \\ 2\delta_1 < \delta_2 + \delta_3 \\ \delta_1 + \delta_2 < 2\delta_3 \\ \delta_1 < \delta_3 \end{array} \right. \quad (4.20)$$

The equations $\delta_2 < \delta_3$ and $\delta_1 < \delta_3$ do not hold because of our initial assumption $\delta_1 > \delta_2 > \delta_3$. Also we show below that $2\delta_1 < \delta_2 + \delta_3$ and $\delta_1 + \delta_2 < 2\delta_3$ do not hold either because of the same initial assumption:

From the assumption $\delta_1 > \delta_2 > \delta_3$ we have that:

$$\delta_1 > \delta_3 \quad (4.21)$$

and

$$\delta_1 > \delta_2 \quad (4.22)$$

Adding (4.21) and (4.22) yields:

$$2\delta_1 > \delta_2 + \delta_3 \quad (4.23)$$

which means that $2\delta_1 < \delta_2 + \delta_3$ in Equation (4.20) is false.

Following the same logic we have:

$$\delta_1 > \delta_3 \quad (4.24)$$

and

$$\delta_2 > \delta_3 \quad (4.25)$$

Adding (4.24) and (4.25) yields:

$$\delta_1 + \delta_2 > 2\delta_3 \quad (4.26)$$

which means that $\delta_1 + \delta_2 < 2\delta_3$ in Equation (4.20) is false, because of the initial assumption that $\delta_1 > \delta_2 > \delta_3$. Along the same lines the proof can be generalized to N customers participating in the auction.

In conclusion, we showed that:

- The charging requirements do not play a role in scheduling EV charging and only the *delay costs* are the ones that influence the result. Therefore, there is no added value to the grid and the consumers from creating multiple charging classes since only delay costs matter. This confirms the proof of Proposition 1.
- The grid operator needs to schedule the customer with the highest delay cost first, since this yields the lowest overall delay to the system, hence the highest social welfare.

□

In summary, Propositions 1 and 2 show that both in the case that the customers have the same or different delay cost, it is optimal to have only one service class (charging speed) and this service class needs to be the highest charging speed, so that the overall delay in the electricity market is minimized. In Proposition 3 we show that given Proposition 1 and 2, the only factor that plays a role in the optimal scheduling is the delay cost that each EV customer is submitting to the market.

4.4.2 Data Description

We apply the proposed auction mechanism on real-world charging data obtained from the Netherlands during the period January 2013 - December 2013. This data set⁶ includes charging observations from 1500 charging poles in the whole country. It has recordings of 10,462 EV owners and includes in total 231,976 transactions with the grid operator. In Figure 4.4, we display the box plot of the steady state EV charging demand over a 24h horizon. We observe that a lot of EV charging happens during the day time. One would expect that most EV owners would prefer to charge their cars during the night so that they can drive the next morning. However, in the Netherlands there are currently strong incentives for EV owners to charge during the day at their employer’s premises. EV owners might even be able to charge for free if they “charge at work”. This explains the fact that most of the demand shows up during the day. We use this data set to test our mechanism. However, the auction mechanism can be tested on any other data set, since Propositions 1-3 hold irrespective of the data they are applied to.




Figure 4.4: EV Charging over a 24 hour horizon.

In Figure 4.5 we display some (anonymized) exemplary average daily charging profiles drawn randomly from our data set. We observe that the peak hours coincide for this EV owners, so it is necessary to impose a mechanism for preventing congestion.

Our data set includes detailed requests from EV owners to the grid for charging. Each EV owner has a unique ID, and each transaction with the grid is time-stamped and accompanied with the requested amount of energy. Since currently there is no optimal resource allocation mechanism implemented on the grid operator’s side, our data does not include bids for price and delay cost. Therefore, for the EV owners price and delay cost bids we assume that they come from a beta distribution parametrized in various ways, $\delta_i \sim \text{Beta}_{\alpha,\beta}$,

⁶Anonymized for confidentiality reasons.




Figure 4.5: Typical average daily charging profiles.

$b_i \sim \text{Beta}_{\alpha, \beta}$. The beta distribution is chosen because for various parametrizations of α and β it yields different commonly used distributions such as uniform, Gaussian, etc. The probabilities of each bid, $b_i \in (0, b_{\max}]$ or delay cost, $\delta_i \in (0, \delta_{\max}]$ drawn from this distribution are:

$$f(b_i; \alpha, \beta) = \frac{1}{B(\alpha, \beta)} \cdot b_i^{\alpha-1} \cdot (1 - b_i)^{\beta-1} \quad (4.27)$$

or

$$f(\delta_i; \alpha, \beta) = \frac{1}{B(\alpha, \beta)} \cdot \delta_i^{\alpha-1} \cdot (1 - \delta_i)^{\beta-1} \quad (4.28)$$

where $B(\alpha, \beta)$ is the beta function $B(\alpha, \beta) = \int_0^1 u^{\alpha-1} \cdot (1 - u)^{\beta-1} du$. By b_{\max} and δ_{\max} we denote the maximum values of price and delay cost bids that show up in an auction. We show in the Robustness Check section that our results are robust independent of the type of distribution the bids are drawn from. In Figure 4.6 we display some probability density functions for various parametrizations used for the price and delay cost bids in the interval $b_i \in (0, 1]$, $\delta_i \in (0, 1]$. For all these distribution types we conduct simulation experiments and display the results in the Robustness Check section.

4.4.3 Empirical Evaluation

We build a simulation environment, as shown in Figure 4.7 in which the grid operator has a limited capacity C and the EV owners bid for amounts of energy so that they charge their cars. The grid operator uses the model presented by (4.3)-(4.8) to schedule EV charging. We evaluate the MVA mechanism with regard to its ability to stabilize the grid and to reduce delays the customers suffer while charging their EVs. For the first criterion (grid stabilization) we use the peak-to-average ratio (PAR) metric and for the second




Figure 4.6: Beta probability density functions for price and delay costs bids in the interval $b_i \in (0, 1]$, $\delta_i \in (0, 1]$.

criterion we examine how many customers are delayed. To demonstrate the mechanism's performance we use the following benchmarks.

Benchmarks

To compare our method with regard to its ability to stabilize the grid, we use the real-world data described in section 4.4.2. This data set depicts the current (naive) situation where the EV owners can charge their cars based on their behavioral characteristics. We refer to this benchmark as *naive* charging. On the other hand, to evaluate our mechanism's ability to reduce delays in the system, we implement the same MVA mechanism (described by equations (4.3)-(4.8)) but with two charging levels (service classes) available. In this *2-class MVA* mechanism, for the second service level holds $r_2 = \frac{r_1}{2}$. We show that with the *2-class MVA* the delays are increased and as a result more customers remain unserved⁷.

Peak Demand Reduction

To compare the MVA mechanism with the naive charging we run the auction for 24 hours. We randomly draw M number of bidders per auction. We show how the aggregate charging demand is distributed over time using the naive charging benchmark, and we compare this with applying the MVA mechanism. We run 100 experiments and in each experiment both the single-class MVA mechanism and naive charging are used. In Figure 4.8 we show the demand redistribution in the 100 experiments.

⁷It would not be possible to use the *naive* charging for this comparison, since naive charging is totally uncontrolled without any capacity constraints and therefore, no delays occur. Instead the grid suffers high peak demand.

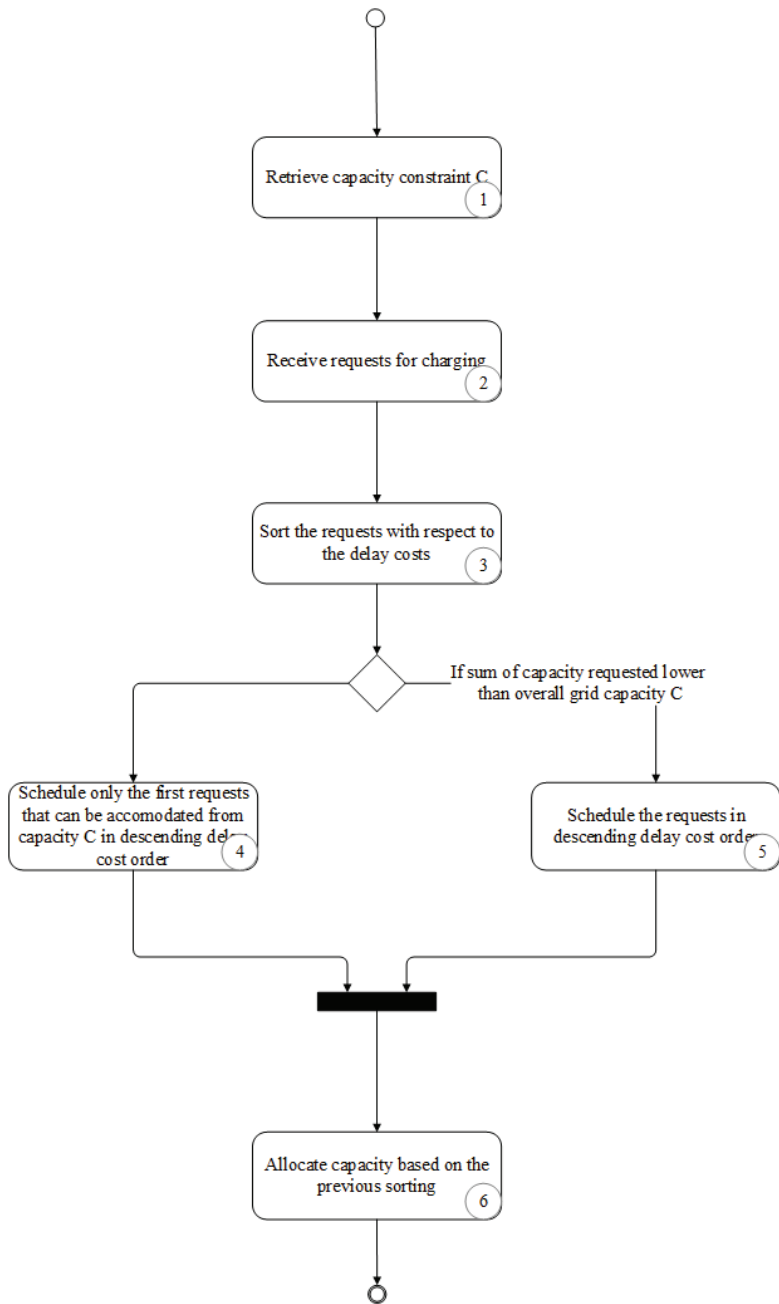


Figure 4.7: Activity diagram of the simulation environment.




Figure 4.8: Electricity demand comparison after adopting MVA and naive charging.

To quantify the comparison we show the peak-to-average ratio (PAR) and the peak demand metrics for each both charging mechanisms (Table 4.1). If we denote by $\mathbf{y} = (y_1, \dots, y_T)$ the temporal vector of the electricity demand curve, the peak-to-average ratio is calculated as: $(PAR = \frac{y_{peak}}{y_{rms}} = \frac{y_{peak}}{\sqrt{\frac{1}{T} \sum_{t=1}^T y_t^2}})$. PAR is also known as “crest factor” and indicates how extreme the peaks in a waveform are. PAR reduction is important because much of the cost of energy supply is driven by peak demand. This metric, besides academic literature, is used by the US Energy Information Administration (EIA)⁸ to measure the effect of peaks on power demand. In order to have higher sustainability in the electricity grid we need lower PAR. We see that the PAR is reduced in the MVA case compared

Table 4.1: PAR and Peak Metrics for MVA and Naive Charging

	Average PAR	Average Peak	Max PAR	Max Peak
Naive	2.04	91.31	6.20	565.65
MVA	1.43	63.82	2.84	521.89
Reduction (%)	29.91	30.11	54.19	7.74

to the naive charging. However, the MVA case does not yield a totally flat curve. This happens because of the granularity of the customers’ demand. In other words, once a customer is accepted for charging, she cannot receive half service and therefore the total demand is not entirely flat.

A reduced PAR by 29.91% means that the grid operator is able to achieve a less volatile electricity demand and mitigate the need for reinforcing the grid infrastructure. A peak reduction of 30.11% means that the grid operator is able to reduce the instantaneous peaks in the demand by 30.11%. The peak demand is an important metric for the grid

⁸<http://www.eia.gov/todayinenergy/detail.cfm?id=15051> [Date Accessed: June 21st, 2015]

operator since this demand is the determinant of the capacity installed. Specifically, if there is a high peak that lasts for a very short period of time, the grid needs to be able to service this peak and therefore needs to have the necessary capacity available. Thus, reducing peak demand supports the physical infrastructure of the smart grid.

Delay Reduction

Comparing the single-class MVA with the 2-class MVA, we show how our mechanism creates reduced delay in the system and as a result more EV owners are serviced. We run 100 experiments and in each experiment both single-class and 2-class MVA mechanism are used. Each auction lasts 24 hours and in each auction we draw randomly M participants from the data set. We observe from Table 4.2 that the average number of requests delayed with the single-class MVA is 17.47 ($\mu = 17.47$, $\sigma = 24.79$) and the average number of requests delayed with the 2-class MVA is 64.46 ($\mu = 64.46$, $\sigma = 70.88$). This result validates the Propositions 1 and 2, which state that one charging class (the fastest available) is optimal with respect to social welfare maximization. The 2-class MVA assumes two charging speeds available, unlike the single-class MVA, and the requests that remain unserved are higher (lower social welfare). In Figure 4.9, we present the distribution of the delayed requests using both single-class MVA and 2-class MVA.




Figure 4.9: Number of Delayed Requests for single-class MVA and 2-class MVA.

We found that by implementing a 2-class MVA the number of requests that remains unserved is higher by 72.89% compared to the single-class MVA. This means that from a social welfare point of view, it is beneficial for the grid operator and the EV owners to have only one class of service (charging speed) implemented. Ideally this charging speed should be the highest possible allowed by the infrastructure, so that the delays are

Table 4.2: Average and Maximum Delayed Requests for single-class and 2-class MVA

	Average Number of Delayed Requests	Maximum Number of Delayed Requests
2-class MVA	64.46	240
single-class MVA	17.47	93
Reduction (%)	72.89	61.25

minimized. The minimum charging speed is also determined by the infrastructure and it corresponds to the Level 1 charging⁹ (3.3-3.6 kW per hour).

4.4.4 Robustness Check

Since we assumed that our bids and delay costs are drawn from a beta distribution with various parametrizations, we examine the robustness of our result for different values of α and β . In Figure 4.10 we display the power demand, resulting from MVA adoption for the parametrizations of beta distribution: $(\alpha = 2, \beta = 4)$, $(\alpha = 1, \beta = 1)$, $(\alpha = 2, \beta = 2)$, $(\alpha = 4, \beta = 4)$, $(\alpha = 4, \beta = 2)$, $(\alpha = 0.75, \beta = 0.75)$ as presented in Figure 4.6. For clarity of the picture we display only the means of the produced demand (the standard deviations follow the same trend as well, hence are omitted). We observe that our results are quite robust to the different parametrizations of the bid and delay cost functions.




Figure 4.10: Electricity Demand for single-class MVA with various parametrizations.

In Figure 4.11 we display the distribution of the number of delayed requests for the parametrizations of beta distribution: $(\alpha = 2, \beta = 4)$, $(\alpha = 1, \beta = 1)$, $(\alpha = 2, \beta = 2)$, $(\alpha = 4, \beta = 4)$, $(\alpha = 4, \beta = 2)$, $(\alpha = 0.75, \beta = 0.75)$. From the graphs of Figure 4.11,

⁹<http://evobsession.com/electric-car-charging-101-types-of-charging-apps-more/> [Date Accessed: March 22nd, 2016]

we can conclude that our results are robust with respect to the bids distribution and therefore, the number of delayed requests is not dependent on the shape of the bid and delay cost distribution. This indicates that our mechanism can produce stable results independent from the way EV owners bid.




Figure 4.11: Number of Delayed Requests for single-class MVA with various parametrizations.

4.5 Conclusions & Future Work

We presented a social welfare maximization mechanism to optimally allocate smart grid resources so that EV charging is coordinated. The proposed mechanism assists the smart grid operator in scheduling the EV charging, so that as many EV owners as possible are serviced at the lowest delay. We proved that it is optimal for the EV owners to implement only one service class (charging speed), since the presence of multiple service classes increases the overall delay cost in the system. Furthermore, we showed that scheduling the EV owners with the largest delay costs first results in the lowest overall delay. To validate our results, we presented both theoretical and empirical evaluation. Applying the proposed mechanism on our data set from the Netherlands, we found an average peak demand and a PAR reduction of around 30%. Furthermore, we observed a reduction of the delayed requests by 70% compared to a scenario where two service classes (charging speeds) are implemented. Finally, we showed that our results are robust to different bid and delay cost distributions.

Reducing the peak demand and the PAR are two important achievements for the grid operator, who strives to stabilize the grid. An increased EV adoption in the grid is expected to threaten its stability and reliability. Reduced peak demand means that

the infrastructure needed to accommodate the current demand can be reduced, since the peak demand is the determinant of extra capacity installation. Additionally, reduced PAR means a less volatile electricity demand, which makes it more predictable and stable. On the other hand, reduced delay cost for the EV owners indicates high quality of service and high social welfare. EV owners rely on the EV charging for satisfying their commuting preferences and therefore any potential delay is a violation of their preferences and reduces their welfare.

The presented results are generalizable to all capacity allocation settings which make the same assumptions about the customers' utility function and the service classes available. The theoretical results presented in Propositions 1-3 hold for any auction setting with these assumptions, since they are based on analytical solutions and are independent from any particular data sets. The numerical results can generalize to other data sets, however the exact numbers of delayed requests or PAR reductions will depend on the particular data set.

This chapter assumes that EV customers have a certain structure for their utility function. A potential change of this structural form of the utility function might influence the results. Therefore, we are planning to tackle this limitation by proving the above propositions assuming different structural forms for the utility function. Another limitation of the proposed mechanism is that the mechanism does not guarantee ex-ante incentive compatibility. Therefore, we will show how this is influencing the results in a large EV customer population.

Furthermore, we will expand the mechanism to profit maximizer auctioneers, such as energy providers who aim at making profits through selling power (Peters et al., 2013). Finally, we will validate the presented mechanism with a real-world experiment conducted via a mobile application (Koroleva et al., 2014), where the EV owners will be able to bid for charging power via the app.

4.6 Acknowledgements

This research is supported by the *Vereniging Trustfonds Erasmus University Rotterdam* (Ref. No 97090.16/14.0574/evt).

Appendix—Summary of Notation and Abbreviations

Table 4.3: Summary of Notation

Symbol	Definition
b_i	bid for customer i
b_{max}	maximum bid for an auction
$B(\alpha, \beta)$	beta function
C	grid's overall capacity
$f(\cdot; \alpha, \beta)$	probability density function of beta distribution
i	EV owner
j	charging speed index
k	randomly selected charging speed index
m	number of units for sale in MVA
M	number of customers participating in the auction
N	maximum number of customers
r_j	charging speed j
t_a^i	arrival time for customer i
t_d^i	departure time for customer i
$t_{0,i}$	time that the car i is plugged in but not charging
T	time horizon
$U(\omega_i, \tau_j(\omega_i), \delta_i, b_i)$	utility function
$x_{i,j}$	binary variable indicating that a request ω_i has been allocated to a charging speed r_j
$\mathbf{y} = (y_1, \dots, y_T)$	temporal vector of the electricity demand
y_{peak}	peak demand of an electricity demand curve
y_{rms}	root mean square of an electricity demand curve
z	maximum number of charging speeds (index of the lowest charging level)
α	beta distribution parameter
β	beta distribution parameter
γ_i	individual customer's utility parameter
δ_i	delay cost each customer i suffers
δ_{max}	maximum delay cost for an auction
λ_1	total delay cost if all customers are scheduled in r_z charging speed
λ_2	total delay cost if all customers are scheduled in r_1 charging speed
μ	mean of delayed requests
σ	standard deviation of delayed requests
$\tau_j(\omega_i)$	delay of a charging request of amount ω_i in the charging speed r_j
Φ_i	utility function component equal to $\omega_i \cdot x_i - b_i$
ω_i	requested charging amount from a customer i

Table 4.4: Summary of Abbreviations

Abbreviation	Definition
EV	Electric Vehicle
FCC	Federal Communications Commission
MVA	Multiple Vickrey Auction
PAR	Peak-to-average ratio
rms	Root mean square
VCG	Vickrey-Clarke-Groves

Chapter 5

A Dynamic Pricing Mechanism to Coordinate Electric Vehicle Charging¹

5.1 Introduction

Electric vehicles (EVs) have the potential to significantly improve the energy efficiency and reduce the carbon intensity of our transportation system (International Energy Agency, 2013). The hose at a filling station delivers energy from a local storage tank to a vehicle fuel tank at a rate of over 10 MW, while EV chargers draw energy from the shared electricity grid, typically at a maximum rate of 25 kW. But 25 kW is about half the total capacity of the electric service in most U.S. homes (European homes have in general smaller or equal capacity), far higher than the power draw of any common household device. Current electricity grids are not designed to support the load of large numbers of EVs charging their batteries during the early evening hours (Verzijlbergh et al., 2012), at the same time when electricity demand peaks due to energy-intensive household activities like cooking and cleaning (Ipakchi and Albuyeh, 2009). What is needed is a way to

¹Parts of this chapter have appeared in the following peer reviewed conference proceedings: Valogianni, K., Ketter, W. & Collins, J. (2015). A Multiagent Approach to Variable-Rate Electric Vehicle Charging Coordination. In Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), pp. 1131-1139, May, 4-8, 2015, Istanbul, Turkey.
Valogianni, K., Ketter, W., Collins, J. & Adomavicius, G. (2015). A Dynamic Pricing Mechanism to Coordinate Electric Vehicle Charging. In Winter Conference on Business Intelligence (WCBI 2015), Snowbird, Utah.
Valogianni, K., Ketter, W., & Collins, J. (2015). A Hybrid Mechanism to Coordinate Electric Vehicle Charging. In 2015 AAMAS Workshop - Trading Agents Analysis and Design (TADA 2015), Istanbul, Turkey.

coordinate the charging of large numbers of EVs in a way that minimizes stress on the grid, and perhaps makes the best use of available renewable energy.

EV charging coordination can be either centralized (top-down) or decentralized (bottom-up). The benefits of top-down coordination mechanisms are that they easily satisfy the constraints imposed by the coordinator (e.g. smart grid manager), leading to a balanced system. However, there are significant shortcomings in these type of approaches. The most important challenge is that often the coordinator must intervene and exogenously control the EV battery, violating the EV driver's comfort. Bottom-up approaches on the other hand, have as major benefit that customers' individual comfort is not violated and the agents have the freedom to schedule their EV charging based on their individual preferences. However, the main disadvantage is that since the same price signals are provided to all customer agents, the EV charging schedules coincide. Specifically, since all agents are cost minimizers, they tend to shift power demand to the cheaper time instants, creating new peaks.

We propose a multiagent method that aligns the objectives of smart grid managers or energy retailers with the objectives of EV owners. We are especially interested in using market-based mechanisms for coordination, because they support distributed decision making among self-interested agents. Therefore, we have designed a hybrid pricing mechanism to achieve charging coordination through the use of price functions. This hybrid pricing scheme combines the features of a decentralized approach with the top-down features that a smart grid operator needs in order to manage grid stability and achieve a desirable match between energy production and consumption. Our approach combines two types of agents:

1. Energy providers or smart grid managers who want to minimize capacity investment by redistributing peak demand or who want to shape demand over time to follow the generation profile of renewable sources.
2. EV owners who receive price signals and modify their EV charging activities to satisfy their individual preferences, including cost minimization and risk (of running out battery (Franke et al., 2011)) reduction.

The key to our approach is the use of price *functions*, specifically prices that vary with charging rate, rather than simple price values. We show that simple time-varying prices, in the absence of other top-down control mechanisms, lead to herding behavior among self-interested EV agents. On the other hand, if prices vary not only by time but also by rate (measured in kW), then self-interested EV agents will adjust their charging rates

over time to minimize their costs, allowing the price-setting agent to shape the overall demand profile of the EV agent population.

The remainder of this chapter is organized as follows. First, we present relevant literature, addressing EV charging coordination. Second, we describe our algorithm and show how the decision processes of the two types of agents interplay in a multiagent simulation. Later on, we outline the assumptions of our multiagent simulation together with the data used to build it. Furthermore, we present the effect of our algorithm on the energy peak demand and we compare its performance with other commonly used benchmarks. Finally, we conclude by providing a summary of our results and describing future steps.

5.2 Related Work

EV charging, if not controlled properly, is anticipated to bring extreme peak load on the electricity grid (Masoum et al., 2010) and put the infrastructure under critical stress. Therefore, significant research has been dedicated to the EV charging coordination challenge. The related work can be divided into centralized (*top-down*) and decentralized (*bottom-up*) coordination mechanisms.

Kahlen et al. (2014) present a centralized mechanism managed by a fleet operator that aims to coordinate EV charging and make profits. Vandael et al. (2013) describe a three-step approach to coordinate the EV charging in a top-down fashion. Gerding et al. (2011) tackle the coordination problem with an online mechanism that accounts for individual preferences in the form of time availability and bids for power, and schedules EV charging accordingly. Building on this work, Stein et al. (2012) introduce a pre-commitment mechanism for EV charging coordination. De Craemer et al. (2014) present a dual implementation for shifting EV charging based on a central auctioneer. Kwak et al. (2014) present a top-down coordination framework in the context of a household, where different appliances' functionality can be shifted. All these approaches have the potential to achieve balance on the grid but most of the times they do not satisfy individual comfort and require direct control, which might not be easy to implement in practice.

Valogianni et al. (2014a) describe a decentralized EV charging mechanism that aims to reduce peak load. Similar bottom-up mechanisms are also implemented in (Gottwalt et al., 2011) but in the smart home context. In these situations individual comfort is not violated but since all agents are cost minimizers, they tend to shift power demand to the cheaper time instants, creating new peaks in the power demand and thus *herding*.

Both top-down and bottom-up approaches do not address the herding issue in EV charging because they assume the same signals offered to the agents for changing their behavior. Also, all these approaches price all charging speeds (slow or fast charging) in the same way or just ask for a premium in the fast charging case. Energy policy makers need more information in order to decide how to price the different charging speeds. And certainly the prices should not be the same for all charging speeds because fast charging creates higher instantaneous peaks in the demand, stressing the grid infrastructure. Therefore, we propose a hybrid mechanism in which prices are a function of charging rate (kW), can mitigate herding and achieve a desired demand profile. This mechanism additionally to benefiting from its hybrid nature, provides an answer to pricing charging rates so that grid overload is reduced.

5.3 Hybrid Coordination

The proposed hybrid coordination mechanism combines distributed, independent decision making with a top-down control mechanism to shape aggregate power demand. We assume that each individual EV owner is represented by an intelligent agent responsible for EV charging, installed in the EV's charging controller. The agent interacts with the user by estimating arrival and departure preferences and expected driving distances. This approach broadens the decision spectrum and overcomes bounded rationality barriers (Simon, 1979).

The control agent might represent a grid operator or energy portfolio manager (Peters et al., 2013). It acts by broadcasting price signals to the EV agents and monitoring their aggregate consumption. This agent is given a desired aggregate demand profile over some time horizon, and uses a learning component to adjust price signals, adapting to the EV agent population it faces. The price function adjustment is made through a learning factor λ that varies among control agents. Our approach requires no vehicle-to-grid (V2G) (Kempton et al., 2009) capability to achieve the desired demand curve, making it compatible with current grid infrastructure that does not support large scale V2G. Figure 5.1 provides an overview of the hybrid charging coordination mechanism.

5.3.1 EV Driver's Agent Decision Problem

We assume that EV agents $i \in \mathbf{I}$ (where $\mathbf{I} = \{1, \dots, I\}$) are self-interested (they represent their owners preferences) and wish to minimize energy cost over a time horizon T . The time horizon T is discretized to time intervals $t \in \mathbf{T}$, where $\mathbf{T} = \{1 \dots T\}$. Energy cost

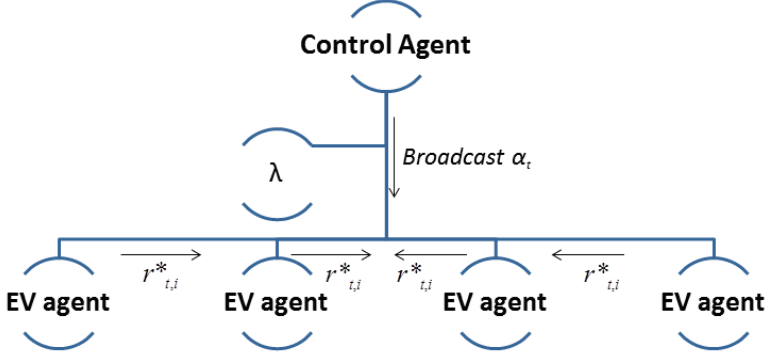


Figure 5.1: Hybrid coordination mechanism - multiagent implementation.

over T is the sum of costs for each interval

$$\sum_{t=1}^T c_t = \sum_{t=1}^T e_t \cdot P_t(\cdot) \quad (5.1)$$

where c_t is the cost of energy during time t , e_t is the energy consumed in kWh during the interval, and $P_t(\cdot)$ is the (possibly rate-dependent) price of energy during this time. If we assume time intervals of one hour and charging at a constant rate r_t in kWh, then $e_t = r_t \cdot 1$. The decision function over T is then

$$\mathbf{r}^* = \underset{\mathbf{r}}{\operatorname{argmin}} \sum_{t=1}^T r_t \cdot P_t(r_t) \quad (5.2)$$

subject to constraints (5.3)-(5.5):

$$0 \leq r_t \leq r_{\max} \quad \forall t \in \mathbf{T} \quad (5.3)$$

where r_{\max} is the highest allowable rate, commonly 25 kW.

The choice of charging rate r_t by the EV agents may be influenced by the user's *range anxiety* (Franke et al., 2011). Range anxiety is the fear that the battery's state of charge will be insufficient for unexpected driving needs. Typically, people with higher range anxiety prefer higher charging rates, allowing them to achieve a higher state of charge over a given time interval. This is an issue for EV owners due to long charging times and low density of charging facilities in most areas.

$$r_t = SoC_t - (SoC_{t-1} - E_t) \quad \forall t \in \mathbf{T} \quad (5.4)$$

E_t determines how much energy the agent should charge to cover the driving needs for the next time instant t , so $E_t, \forall t \in \mathbf{T}$ accounts for the charging needs over the whole planning horizon T . This constraint ensures that the state of charge in the battery will be at least the amount satisfying the driving needs, without violating individual comfort. SoC_t indicates the battery's state of charge at each time instant t . We assume that SoC is at its minimum value at the beginning of the time period:

$$SoC_0 = SoC_{min} \quad (5.5)$$

Each EV agent i has also a set of preferences θ_i , including n arrival $t_{a,i}^n$ and departure times $t_{d,i}^n$ over the horizon T : $\theta_i = \{t_{a,i}^n, t_{d,i}^n\} \forall n \in \mathbf{N}$, where \mathbf{N} is the set of intervals during which the vehicle is connected to a charger based on the user's driving profile. These preferences should always be satisfied by the decision function (5.2) of a self-interested agent so that individual comfort is not violated. Therefore, equation (5.2) becomes:

$$\mathbf{r}^* = \underset{\mathbf{r}}{\operatorname{argmin}} \sum_{n=1}^N \sum_{t=t_{a,i}^n}^{t_{d,i}^n} r_t \cdot P_t(r_t) \quad (5.6)$$

subject to constraints (5.3)-(5.5).

5.3.2 Smart Grid Manager's Decision Problem

The grid manager's agent (control agent) advertises prices for each time period over some time horizon to all EV agents. These prices can vary across time, and may also depend on the charging rate. Without this rate-dependent approach we observe herding, in which self-interested agents always charge at their maximum rates when price is the lowest. One possible formulation is the linear function

$$P_t(r_t) = P_{0,t} + \alpha_t \cdot r_t \quad (5.7)$$

where r_t is the charging rate (power consumption) during timestep t and $P_{0,t}$ is the price for zero demand and can either be constant, or be set as e.g. a) the wholesale price of electricity at time t or b) another variable price that is known ahead of time. The control agent's goal is to determine α_t at each timestep t that will produce the desired aggregate demand profile. The coefficient α_t determines the slope of the price curve with respect to charging rate (power).

To achieve the desired aggregate power demand vector \mathbf{D} , the control agent sets prices so that the summation of power demand over the EV agents comes as close as possible to the desired demand ($\mathbf{D} \approx \sum_{i=1}^I \mathbf{D}_i$). Since the EV drivers' preferences are unknown to the grid manager, it is unlikely to achieve an exact match of the desired aggregate demand and the summation of individual demands (i.e $\mathbf{D} = \sum_{i=1}^I \mathbf{D}_i$). Therefore, in Section 5.3.3 we present a learning component whereby the control agent observes the outcome of its actions on the EV driver population and adjusts its future actions accordingly.

In order to estimate the initial values of α_t , the control agent takes the view of an EV agent. Substituting the price function (5.7) in (5.2) we have:

$$\mathbf{r}^* = \underset{\mathbf{r}}{\operatorname{argmin}} \sum_{t=1}^T r_t \cdot (P_{0,t} + \alpha_t \cdot r_t) \quad (5.8)$$

which has optimal solution: $\mathbf{r}^* = (r_1^*, \dots, r_T^*)$ for a time horizon T . Since the solution is bounded by constraint (5.3) we have

$$r_t^* \leq r_{t,max} \Rightarrow \sum_{t=1}^T r_t^* \leq \sum_{t=1}^T r_{t,max} \quad (5.9)$$

We now show that fixed (not rate-dependent) prices lead to herding, while rate-dependent price functions can spread demand over time.

Theorem 1. Assume a self-interested agent population who wishes to charge EV batteries by adding an amount of energy E over a time interval T , which is divided into a sequence of discrete intervals $t \in T$. We assume that such a self-interested agent will act to first minimize its cost c , and second to acquire its desired energy E sooner rather than later. Let $c_t = r_t \cdot P_t$ be a continuous cost function over a range of charging rates $[0, r_{max}]$. If P_t is constant, $P_t = \xi_t$, where ξ_t is a constant price (in monetary units/kWh) during a given time interval t , and P'_t is an increasing function of charging rate r_t , $P'_t(r_t) = P_{0,t} + \alpha_t \cdot r_t$, then P'_t reduces the “herding” of self-interested charging agents over multiple time intervals compared to P_t .

Proof. For price function $P_t = \xi_t$, the cost function is $c_t = r_t \cdot \xi_t$. The optimal charging rate r_t^* for a *self-interested* agent is always either zero or equal to the maximum charging rate $r_{t,max}$, since there is no price incentive for the agents to change their charging rate. If ξ_t is constant over time, then the agent's overall cost is $c = \xi \cdot E$ regardless of when the charging takes place. Therefore all such agents will immediately charge at r_{max} for E/r_{max} time periods. If large numbers of such vehicles are connected during the same

time period, this will lead to herding. If ξ_t varies over time intervals, then such an agent will acquire as much energy as possible during the lowest-cost interval, followed by the next lowest-cost interval, and so on, until it has acquired E . This is illustrated in the top portion of Figure 5.2.

For price $P'_t(r_t) = P_{0,t} + \alpha_t \cdot r_t$, the cost function becomes: $c_t = r_t \cdot (P_{0,t} + \alpha_t \cdot r_t)$ and the optimal charging rate $r_t^{*'}$ for a *self-interested* agent is E/T . This is illustrated in the middle portion of Figure 5.2. The agent can arrive at this value incrementally as follows: divide E into an arbitrary number of small increments, and add each to the time period with the lowest price. If $P_{0,t} = \hat{P}$ and $\alpha_t = \hat{\alpha}$ are fixed, then this will always be the time period with the lowest allocated charge rate. The result will be constant-rate charging at a rate of E/T over the entire interval.

Furthermore, with price function $P'_t(r_t) = P_{0,t} + \alpha_t \cdot r_t$ the optimal charging can be exogenously determined by a central operating party (grid manager or energy retailer), through adjusting α_t across time, as shown in the lower portion of Figure 5.2. \square

In Figure 5.2 we show an illustrative example of Theorem 1. Assume we have 5 hours




Figure 5.2: Illustration of Theorem 1.

to charge, max charge rate is 10 kW, and we need 25 kWh during this charging period. With the flat-price scheme (top panel, $\alpha_t = 0 \forall t \in \mathbf{T}$), we get 10kWh at 0.07/kWh, 10kWh at 0.09/kWh, and 5kWh at 0.11/kWh, for a total cost of 2.15. No charging is done during the first and last time periods.

With the linear price functions (bottom two panels), we set $P_{0,t} = 0.05 \forall t \in \mathbf{T}$. The horizontal axis in each time period is in kW, from 0 to 10kW. So if we charge at 5 kW during a period, we pay half the maximum price for the period. We can solve for the minimum cost by finding the price/kWh that gives us the total energy we need, in this case about 0.07/kWh, for a total cost of 1.75. The charge rates in this example are (2.5, 5, 10, 3.5, 4). The total cost would be lower if the price we find is above the maximum price of one or more of the time periods.

5.3.3 Learning Component

The smart grid manager agent (control agent) needs to adapt to changes observed in the EV agent population, since we assume no prior knowledge related to the EV driver portfolio. Therefore, it needs to learn from observations related to EV agents' behavior and adapt the price signals accordingly, so that it achieves the desired aggregate demand profile \mathbf{D} . We introduce a learning component in its decision algorithm that helps the hybrid coordination mechanism converge to the desired profile \mathbf{D} without having knowledge about the EV agent population. This makes our coordination mechanism highly flexible since any potential additions of agents with different preferences or drop-outs of existing agents, can be observed online and the mechanism can adapt its behavior.

Specifically, a control agent observes and stores the deviations of the actual consumers profile and the intended profile that it wanted to achieve. Based on these observations it updates the error function over horizons T , $\sum_{t=1}^T \epsilon_t = \sum_{t=1}^T D_t - \sum_{t=1}^T \sum_{i=1}^I r_{i,t}^*$ and adjusts the value of α for the next period T based on the agent's *learning factor*, $\lambda > 1$, so that the aggregated demand profile created by the individuals approximates the intended demand profile. The learning factor λ varies across control agents and we experiment with different values in our simulation. Additionally, if $\sum_{t=1}^T \epsilon_t < 0$ it means that the produced aggregate result is higher because of higher charging rate of the individuals and using (5.7) we have to reduce charging rate $r_{i,t}^*$, and thus increase α_{t+T} : $\alpha_{t+T} = \lambda \cdot \alpha_t$. In the opposite case ($\sum_{t=1}^T \epsilon_t > 0$), the value of α_{t+T} needs to be decreased, so $\alpha_{t+T} = \frac{1}{\lambda} \cdot \alpha_t$. In summary, the learning component updates the next value of α_{t+T} based on the following rule:

$$\alpha_{t+T} = \begin{cases} \lambda \cdot \alpha_t & : \sum_{t=1}^T \epsilon_t < 0 \\ \frac{1}{\lambda} \cdot \alpha_t & : \sum_{t=1}^T \epsilon_t > 0 \end{cases} \quad (5.10)$$

This decision rule is repeated by the control agent until the error term $\sum_{t=1}^T \epsilon_t$ reduces to the desired error level $\sum_{t=1}^T \epsilon_{t,min}$. In Figure 5.3 we present the activity diagram of the coordination mechanism and the respective steps as they are implemented in our simulation. In Table 5.1 we summarize the proposed hybrid coordination mechanism in pseudo-code form.

5.4 Multiagent Simulation

To evaluate our coordination mechanism we create a multiagent simulation which consists of *self-interested* EV agents and a smart grid manager agent (control agent) who is responsible for keeping the aggregate demand closer to a stable level (desired aggre-

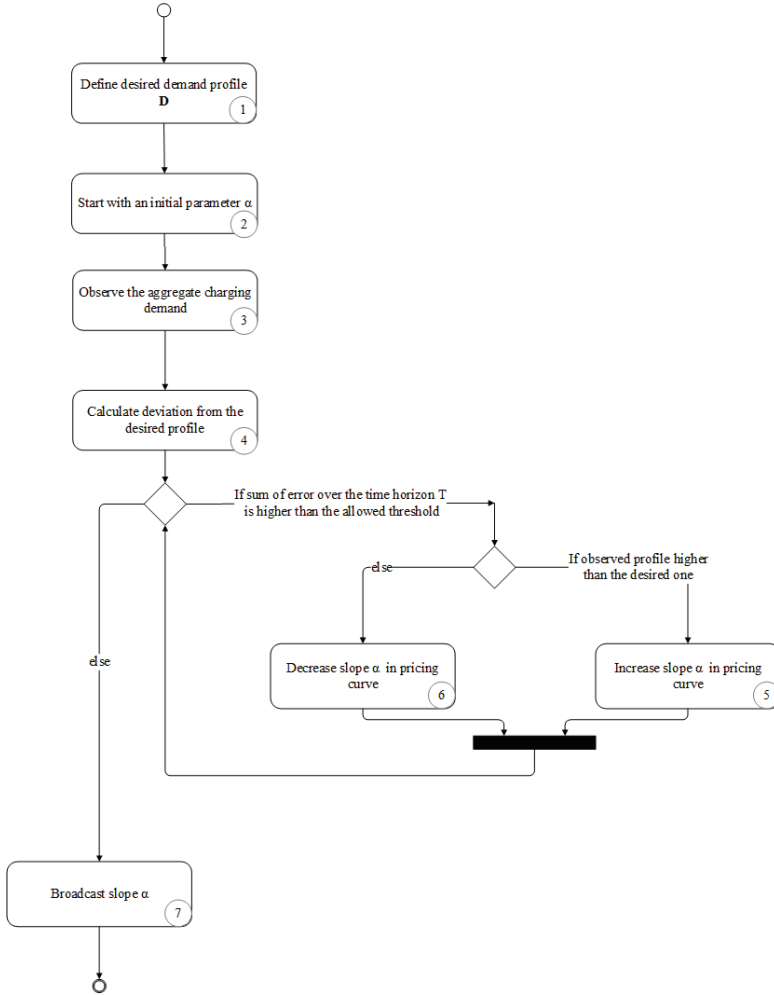


Figure 5.3: Activity diagram of the hybrid coordination mechanism.

gate profile, \mathbf{D}). Our simulation environment is built according to the smart markets paradigm (Bichler et al., 2010) and Power TAC’s specifications (Ketter et al., 2013b,a) since we aim to evaluate the mechanism within Power TAC’s simulation platform. These assumptions approximate electricity markets quite realistically, and this will be helpful in extending the results to the real world. The main assumptions are listed below. Appendix presents a summary of the notation and abbreviations used.

Table 5.1: Hybrid coordination pseudo-code

Hybrid coordination	
1	Initialization
2	Define desired aggregate demand profile $\mathbf{D} = \{D_t\}, \forall t \in \mathbf{T}$
3	Start with an initial value of parameter $\alpha_t, \forall t \in \mathbf{T}$
4	Broadcast α_t to the self-interested agents $i \in I$
5	Observe aggregate charging rate $\sum_{i=1}^I r_{t,i}^*$ as calculated by each agent: $\mathbf{r}_i^* = \text{argmin}_{\mathbf{r}_i} \sum_{t=1}^T r_{t,i} \cdot (P_{0,t} + \alpha_t \cdot r_{t,i})$
6	Calculate error: $\sum_{t=1}^T \epsilon_t = \sum_{t=1}^T D_t - \sum_{t=1}^T \sum_{i=1}^I r_{t,i}^*$
7	while $\sum_{t=1}^T \epsilon_t \geq \sum_{t=1}^T \epsilon_{t,min}$ do :
8	if $\sum_{t=1}^T \epsilon_t < 0$
9	$\alpha_{t+T} = \lambda \cdot \alpha_t$
10	else
11	$\alpha_{t+T} = \frac{1}{\lambda} \cdot \alpha_t$
12	endif
13	Observe aggregate charging rate $\sum_{i=1}^I r_{t,i}^*$ as calculated by each agent: $\mathbf{r}_i^* = \text{argmin}_{\mathbf{r}_i} \sum_{t=1}^T r_{t,i} \cdot (P_{0,t} + \alpha_t \cdot r_{t,i})$
14	Calculate error: $\sum_{t=1}^T \epsilon_t = \sum_{t=1}^T D_t - \sum_{t=1}^T \sum_{i=1}^I r_{t,i}^*$
15	end
16	return α_t

5.4.1 Scenarios & Assumptions

In order to demonstrate the performance of the algorithm, we will examine scenarios where the EV agents face prices given by $P_t(r_t) = P_{0,t} + \alpha_t \cdot r_t$. We create scenarios with both rate-independent ($\alpha_t = 0$) and linearly rate-dependent ($\alpha_t \neq 0$) prices. The constant factor of price function (5.7), $P_{0,t}$, may get either the average of the wholesale price over a day or the corresponding retail price of each hour (R_t), representing the generation cost of this particular amount of charging energy and taxes and network fees. We will use the latter option, since with this assumption price function (5.7) accounts for both the power generation cost, taxes and network fees and for an extra price factor $\alpha_t \cdot r_t$ which is analogous to charging rate r_t . This factor can be interpreted as the premium the EV agents have to pay on top of the retail price to obtain a particular charging rate r_t . The scenarios examined are presented in Table 5.2. The following assumptions draw the

Table 5.2: Simulation Scenarios.

	Attributes	
Rate-independent scenario	$\alpha_t = 0$	$P_{0,t} = R_t$
Linear scenario	$\alpha_t \neq 0$	$P_{0,t} = R_t$

boundaries of our simulation environment and determine our mechanism's goal:

- The simulation includes *self-interested* EV agents that do not exchange information among each other. They only interact with the grid (via the control agent).
- The interaction of the EV agents with the grid is limited to receiving price signals (retail prices) and decide on EV charging rate and duration, based on these prices.
- The *self-interested* EV agents have preferences regarding departure and arrival times, which are derived by the data set described in Section 5.4.2.
- All EV agents are located and driving under the same distribution network, to avoid procurement of charging power from other distribution networks.
- The granularity of the designed simulation is 1 hour, since typically the EV charging rate is calculated in hourly intervals.
- The planning horizon is 1 week ($T=168h$) because there seems to be repetition of driving habits and overall consumer behavior within weekly intervals. However, the algorithm can be adjusted to produce results for different planning horizons.
- The grid manager is in control of steering the aggregate EV charging consumption towards a desired profile, which might either be a less volatile demand curve or a demand curve that follows the production pattern of a renewable production unit (e.g. wind turbine).
- The goal of the hybrid coordination mechanism is to reach a targeted aggregate charging profile.

5.4.2 Data Description

In this section we present the data sets used to calibrate our simulation and evaluate the performance of the hybrid coordination mechanism. All data refers to same region in the Netherlands and is collected during 2012-2013. Our simulation environment can be calibrated with data from other areas without affecting the mechanism's functionality.

Individual Preferences

We bootstrap our simulation with arrival and departure preferences obtained by the Central Bureau of Statistics (CBS)² in the Netherlands. This data includes different population clusters (full time employees, part-time employees, students, retired persons, etc.)

²www.cbs.nl

with a variation of habits and driving behaviors (business commuting, leisure time driving, vacation, visits to relatives, shopping etc.). For each individual we get a driving profile with certain activities and driving demand for each activity, combined with arrival and departure times. The aggregate driving demand in (km) of our population is displayed in Figure 5.4.




Figure 5.4: Aggregate driving demand of the EV agents population.

Energy Prices

The results for the Rate-independent Scenario, where the prices are fixed over a time period, are produced using as an example of wholesale prices offered by the European Power Exchange (EPEX) adjusted to account for network fees, taxes and VAT for the Netherlands (44% of the retail price³). These prices are the values of $P_{0,t} = R_t$ in both Linear and Rate-independent Scenarios. In Figure 5.5 we show 3 weeks of retail price data, as it is used in the simulation. A time period of a week reveals both the daily repetition of activities and the differentiation across weekdays and weekends. In addition, 3 weeks of data is enough to capture repetition of activities across multiple weekdays and weekends. This data serves as a basis for constructing prices for longer time horizons (than 3 weeks).

Learning Factor

For the learning factor of the control agents we use values in the spectrum of $\lambda \in [1, M]$, where M is a sufficiently large number. Agents with $\lambda = 1$ are considered zero-intelligence agents that show no learning ability, whereas agents with $\lambda = M$ show the highest learning

³<http://www.nuon.nl/energie/energieprijzen-vergelijken/opbouw-energieprijs.jsp> [Date Accessed: 02/05/2014]




Figure 5.5: Retail prices (€/kWh) over 3 weeks.

ability and are the most desirable control agents. High learning factor λ , indicates higher adaptability of the control agent to the EV agent portfolio changes and thus potentially quicker convergence to the desired profile **D**. However, there is no proven direct analogy between learning factor and convergence, since the behavior of the EV agents includes stochasticity that cannot always be accounted for.

5.4.3 Benchmarks

The output of the hybrid coordination mechanism is the Linear Scenario, calibrated with the population's preferences. To evaluate its performance we compare it with the following benchmarks.

Benchmark 1 - Rate-independent Scenario

This is the baseline scenario of our analysis since it assumes *self-interested agents* that minimize their costs based on a given variable retail price signal, which does not depend on the charging rate (Section 5.4.2). The household demand combined with the power demand of this scenario is depicted in Figure 5.6. We present the EV charging demand combined with the household demand, because this is the demand that the grid faces from each household. Additionally, on this graph it is more clear that EV charging fills the valleys created by the household demand during early morning hours. Firstly, here we observe that the EV agents are price sensitive and are solely driven by the high variations in prices. This makes them consume significant portion of the daily power demand during low price periods (early morning and late night) whereas, they mitigate EV charging when prices are high (noon and evening hours). Secondly, we observe that herding of charging




Figure 5.6: Rate-independent Scenario - Combination of EV charging and household demand.

is present because every agent gets the same price signals and besides small differences in preferences, the power demand of all agents coincides, creating new peaks during low price periods. This herding is exactly what our algorithm aims to prevent by adjusting the price signals and partially redistribute the peaks across the whole time horizon. This redistribution mitigates volatility of aggregate demand, which is highly beneficial for the smart grid's infrastructure.

Benchmark 2 - Real-world Charging

As a second benchmark we use real-world EV charging data obtained in collaboration with EV charging infrastructure company in the Netherlands. The data set accounts for EV charging during 2012-2013 across the whole country. The steady state curve of this data, combined with the household demand is presented in Figure 5.7. From this graph we verify our initial assumption, that most of the people, without any control in EV charging, just plug their EV once they return home from work (around 6 pm) increasing peak demand.

Benchmark 3 - Aggregate Demand without EVs

This benchmark represents the total power demand of our population, assuming that there is no EV charging involved. It is crucial to compare the performance of the algorithm with this benchmark because the goal of a successful EV integration policy is to prevent extra peaks on the already volatile aggregate household demand. Therefore, we want to see how close the algorithm's results are compared to this benchmark. We do not expect the algorithm to reduce peak demand since extra power demand is added, coming from




Figure 5.7: Combination of real-world EV charging and household demand.

EV charging. It is desirable though, to show that during peak hours the EV charging does not create higher volatility. The data for this benchmark comes from households in the Netherlands obtained in collaboration with a European energy utility company. In Figure 5.8 we show some individual power demand curves (anonymized) and in Figure 5.9 the aggregate household demand of the population.




Figure 5.8: Typical household power demand of the EV driver agents population.

5.5 Numerical Results

In this section, we implement the hybrid coordination mechanism in the multiagent simulation described in Section 5.4 and present indicative performance results. We are mostly interested in the impact of the algorithm on the aggregate demand curve. Specifically, the




Figure 5.9: Aggregate household demand of the population over 4 month period.

peak demand of this curve is the determinant of installing extra capacity on the network. Therefore, the grid managers using coordination mechanisms strive to mitigate this peak demand. A second important factor is the demand's volatility. Reduced volatility of this curve protects the grid from critical strains.

5.5.1 Impact on Power Demand

Applying the algorithm to our EV agent population for a typical price function ($\alpha_t = 1, \forall t \in \mathbf{T}$): $P_t(r_t) = P_{0,t} + r_t$ we get the steady state power demand displayed in Figure 5.10. We observe that the EV charging demand is more evenly distributed on top of the household demand without having significant herding during low price periods.




Figure 5.10: Power demand after applying the hybrid coordination mechanism - Linear Scenario ($\alpha_t = 1 \forall t \in \mathbf{T}$).

Comparing the Linear Scenario with the Rate-independent Scenario and the real-world charging demand we get Figure 5.11. In this graph we notice first, that the Rate-independent scenario shifts most of the charging during low price time intervals (early in the morning or late at night), whereas during the day and specifically during high price periods, it does not charge at all. Consequently, significant herding is present because all of the *self-interested* EV agents congest to charge during the low price periods. That explains the high peak of 4kW around 2am-3am. This outcome is aligned with the results of Gottwalt et al. (2011) where they use bottom-up cost minimization in the smart home context. They also observe significant herding during these time periods.

Secondly, in Figure 5.11 we observe that the real-world charging mostly shows up during business hours despite the high prices. This happens because the current situation in Europe allows EV drivers to plug their car in their employers' premises and charge it there while working. Other EV drivers leave their EV charging the whole night to cover their range anxiety. This situation is undesirable because the daily peaks around 6pm-8pm increase even more with EV charging.

Finally, we observe that in the Linear Scenario where we put a price function on the charging rate (charging speed) the *self-interested* EV agents schedule their charging in a way that prevents extreme peaks but also covers the driving needs. This happens because increasing charging speed leads to increasing costs. To measure the impact of our




Figure 5.11: Comparison of Rate-independent, Linear and Real-world scenarios.

mechanism on the smart grid we use the peak-to-average ratio (PAR) metric: ($PAR = \frac{r_{peak}}{r_{rms}} = \frac{r_{peak}}{\sqrt{\frac{1}{T} \sum_{t=1}^T r_t^2}}$), which is also known as *crest factor* and measures the intensity of peaks or valleys in a curve. Secondly, we will measure the peak reduction created by our algorithm in comparison to the other benchmarks. Table 5.3 summarizes these metrics (negative reduction indicates increase). We observe that the Linear Scenario, which uses




Figure 5.12: Comparison of Linear Scenario with constant $\alpha_t, \forall t \in \mathbf{T}$ and Variable Scenario with $\alpha_t = R_t, \forall t \in \mathbf{T}$.

our mechanism, reduces the peak demand compared to all the other benchmarks. Of course compared to the household demand it is not possible to reduce peak demand, because we add extra demand that is attributed to the EV charging. Similar are the results for the PAR reduction (volatility reduction). It is interesting to note that the Linear Scenario reduces PAR compared to the plain household demand, resulting in a less volatile curve. Therefore, there are strong incentives for the energy policy makers to introduce such kind of coordination mechanisms.

Table 5.3: Energy Peak and PAR Reduction

	PAR reduction (%)	Peak reduction (%)
Linear vs. Rate-independent	16.00	15.02
Linear vs. Real-world	9.61	16.73
Linear vs. Household	11.40	-36.40

5.5.2 Shaping Aggregate Power Demand

Besides the scenarios shown before where α_t has a fixed price $\forall t \in \mathbf{T}$, we present here a scenario where we give variable values to α_t . In this Variable Scenario we set $\alpha_t = R_t$. In Figure 5.12 we display the first iteration of the algorithm. This iteration is practically the first observation the control agent gets from the EV agent population. Depending on this observation it will adjust the α_t values to reach the desired profile \mathbf{D} . From Figure 5.12 we can also observe that by changing α_t over time the aggregate demand curve becomes smoother compared to the Linear Scenario where α_t had a constant value $\forall t \in \mathbf{T}$. After




Figure 5.13: Demand after learning EV agents behavior - Variable Scenario with $\alpha_t = R_t, \forall t \in \mathbf{T}$.

the first iteration, we see how the control agent adjusts the α_t values to reach the goal. We assume a learning factor $\lambda = 10$ since we want the algorithm to converge quickly. We set the error threshold $\epsilon_{t,min} = 0.2, \forall t \in \mathbf{T}$ since lower than this level cannot be achieved by the agents. This happens because they have as hard constraints to satisfy the EV drivers's needs and therefore, they have to deviate from the desired profile to have the battery charged for their owners. The algorithm converges after 21 iterations and in Figure 5.13 we show how the weekly charging demand changes after 21 iterations.

5.5.3 Sensitivity Analysis

Since price coefficient α_t drives the outcome, we provide some indicative results for this parameter's sensitivity. In Figure 5.14, we present results in the spectrum of $\alpha_t \in [1, 4]$. As expected, increasing α_t decreases average charging rate. The interesting result of this graph is that increasing the value of α_t by increments of 1, yields small changes in the average charging rate. Therefore, we can confirm the assumption that higher learning factors λ on the control agent's side are crucial for achieving the desired convergence.

5.6 Conclusions & Future Work

We presented a hybrid mechanism that coordinates EV charging. It combines the decentralized decision making on the EV agents' side with a central coordination party that ensures convergence of the aggregate EV charging to the desired (coordinated) outcome. Our mechanism is based on *price functions* for EV charging rates that create incentives for charging at low rates (low speed charging) when the prices are high and at high rates (high




Figure 5.14: Sensitivity Analysis.

speed charging) when the prices are lower. The control agent does not require any prior knowledge of the EV agents portfolio to set the right prices since it learns their behavior online. Therefore, the mechanism is highly dynamic and can adjust quickly to exogenous shocks or portfolio changes. We show that the proposed mechanism prevents *herding* in EV charging, which is present in many coordination mechanisms and also distributes the EV charging demand in a way that peaks and volatility are reduced.

In future, we plan to investigate the integration of vehicle-to-grid (V2G) (Kempton et al., 2009) in our mechanism. Furthermore, we plan to extend the price functions to other forms and evaluate their performance. Finally, we aim to test this mechanism in a real-world experiment using a mobile application.

5.7 Acknowledgements

This research is supported by the *Vereniging Trustfonds Erasmus University Rotterdam* (Ref. No 97090.16/14.0574/evt).

Appendix—Summary of Notation and Abbreviations

Table 5.4: Summary of notation

Symbol	Definition
c_t	electricity cost in time instant t
$\mathbf{D} = \{D_t\}$	desired aggregate power demand vector
E_t	estimated driving demand in time instant t
$\mathbf{I} = \{i\}$	discrete set of EV driver agents
M	upper bound of the learning factor λ
$\mathbf{N} = \{n\}$	set of intervals during which the EV is connected to a charger
$P_t(r_t)$	charging rate price in time instant t
$P_{0,t}$	constant factor of the linear price function per time instant t
r_{peak}	peak charging rate in a demand curve
r_{rms}	root mean square charging rate
r_t	charging rate per time instant t
$r_{t,max}$	maximum charging rate per time instant t
R_t	retail price of power per time instant t
SoC_t	EV battery's state of charge during time t
$t_{a,i}^n$	arrival time of agent i for activity n
$t_{d,i}^n$	arrival time of agent i for activity n
$\mathbf{T} = \{t\}$	discrete set with time instants
α_t	charging rate coefficient in function $P_t(r_t)$
ϵ_t	error factor in time instant t
$\epsilon_{t,min}$	error factor threshold in time instant t
θ_i	set of preferences for agent i
λ	learning factor in the hybrid coordination
ξ	constant value for price

Table 5.5: Summary of Abbreviations

Abbreviation	Definition
EV	Electric Vehicle
EPEX	European Power Exchange
PAR	Peak-to-average ratio
rms	Root mean square

Chapter 6

Conclusion

To answer our main research question, we approached the problem from three different angles: the EV owner's point of view, the smart grid manager's point of view and a combination of both (hybrid approach). These three different perspectives allowed us to investigate and design incentives for EV customers so that they satisfy their individual preferences, reduce their electricity costs and contribute to reducing peak electricity demand. The common denominator in all three studies presented is the environmental sustainability. Following the Green IS main principles, we argue that using available information to reduce electricity peaks is beneficial for the people (reduced electricity costs) and for the environment (reduced peaks and carbon emissions). We saw that the different mechanisms are suitable for different types of markets and yield benefits to the stakeholders involved. In all situations the smart grid's stability (measured in peak demand and volatility reduction) is supported, and environmental sustainability is improved. Depending on which stakeholder's point of view is taken, the benefits change. Therefore, it is important for energy policymakers to identify the characteristics of each particular market and their objectives and then select the appropriate coordination mechanism.

Each chapter presents an answer to this thesis' research question, depending on the focal point and the market conditions. Below, we present this answer as it resulted from this dissertation. In Chapter 2, where the focus is the individual EV customer, we see that an IS artifact capable of achieving a charging coordination, should be able to identify the individual preferences of the EV owners and schedule the charging satisfying those preferences. It should be designed so that it represents each EV owner in the electricity market, and learn from the observed driving and consumption behavior so that it provides personalized charging recommendations. These recommendations if tailored to individuals, will lead to a bottom-up coordination, conditionally that the individual preferences are significantly different. In Chapter 3, the focus is still put on the individual EV customer,

but now the sustainability objective and the renewable usage are examined. In this chapter, an IS artifact capable of reducing peak demand and benefiting from renewable usage must be equipped with an accurate forecasting algorithm that can estimate the renewable generation (of a photovoltaic panel) from the weather conditions. In order for this artifact to be successful in this objective, the individual preferences need to be accounted for and the EV charging must be charged so that the renewable usage is maximized. In Chapter 4, the focus is shifted to the grid operator assisted by an IS artifact. In this chapter the EV scheduling needs to be performed properly so that the grid is not overloaded. In order for this IS artifact to be successful in coordinating the charging, the delays in the market need to be accounted for, as well as the different charging speeds available on the grid. In Chapter 5, we combine the points of view of the grid operator and the EV customer and we show that an IS artifact that can account for the benefits of both sides and coordinate EV charging, must adjust the electricity prices dependent on the charging speed. The grid operators and the EV customers have conflicting objectives (grid stabilization and cost minimization respectively) therefore, defining the price functions of charging speed properly is the main feature of the proposed IS artifact.

6.1 Synopsis of Main Findings

First, we summarize the main findings of each separate chapter and their associated contributions.

6.1.1 Chapter 2

The first step in answering the main research question was to take the standpoint of the individual EV owner. In this chapter, we investigated how a decentralized EV charging mechanism must be designed so that it satisfies the EV owner's preferences and at the same time mitigates peak electricity demand. These two goals are conflicting because typically, individual EV owners tend to fully charge their EV batteries, which is usually more than what they actually need. On the other hand, the smart grid manager desires to reduce this tendency of charging during peak hours, so that the grid is not overloaded.

We proposed an IS artifact that learns the individual characteristics of each EV owner and schedules EV charging so that the benefits for each owner are maximized. We observed that depending on the preferences, the EV charging schedule is different and has a different effect on the electricity grid. An important determinant of the EV charging schedule is *range anxiety* (Franke et al., 2011). We integrated this anxiety in the individual preference

model and scheduled EV charging so that individual preferences are always satisfied. We observed that low range-anxiety populations yield more benefits to the electricity grid, whereas high range-anxiety populations need different incentives so that they support environmental sustainability.

6.1.2 Chapter 3

An extension of Chapter 2 was the integration of the IS artifact in a smart home combined with an EV. In Chapter 3, we implemented a different decentralized mechanism which accounts for household consumption shifting, on top of EV charging scheduling. This mechanism focuses on the combination of household appliances and an EV in a smart home environment. We implemented a supervised learning algorithm that estimated the solar panel output of the smart home from the weather conditions. This estimation was used by the EV so that the renewable charging of the car was increased as much as possible.

We showed that a smart home is able to benefit from coupling a solar panel with an EV in terms of electricity costs. A solar panel runs at zero marginal cost, reducing significantly the costs for the household. In addition, we showed that the overall sustainability levels increase, since the EVs can use their batteries to store part of the renewable energy which cannot be used at the time when it is produced. With this approach, both smart home owners and the smart grid benefit from the integration of renewable sources.

6.1.3 Chapter 4

The second step to answer this dissertation's research question was to examine how a centralized mechanism would coordinate EV charging and what implications this has for the EV owners and the electricity grid. Therefore, in this study we took the standpoint of the smart grid manager who wants to schedule EV charging so that as many EV customers as possible are serviced given a limited grid capacity.

A grid operator must schedule EV charging so that the grid is not overloaded and the consumers are serviced with the lowest delay possible. A successful scheduling of EV charging ensures the grid's stability without installing new capacity infrastructure (which would be an unsustainable solution). Reinforcing the grid infrastructure to accommodate electricity demand is the traditional way of coping with electricity peak demand on the grid. However, this solution requires more raw materials (such as copper, etc.), which makes it unsustainable and costly. We proposed an auction-based mechanism that schedules EV charging and determines the prices to the customers in real time. Auctions, unlike

posted-price and capacity allocation mechanisms, are preferred when the demand is not known or easy to estimate (Bapna et al., 2003). Therefore, in this particular problem, in which the grid operator does not know at each point in time how many EVs will require charging, auctions contribute to allocating the grid resources efficiently.

We assumed the same EV customer preference model as in Chapter 2, enriched with *delay costs*. These are the costs, suffered by the EV owners for not charging during the time periods they prefer (individual preference violation). The *delay costs* are important in this study, since the EV charging is scheduled by an external coordinator (grid manager) and therefore the EV owners might have to suffer delays in their charging, depending on the charging congestion. In Chapter 2, these *delay costs* were set to zero, since each EV owner could schedule their charging without encountering any delays (decentralized charging).

6.1.4 Chapter 5

Finally, having examined the centralized and decentralized EV charging coordination mechanisms, we proposed a hybrid approach. This mechanism combines benefits from the decentralized and centralized methods and has in its core a dynamic pricing scheme. It relies on price functions, specifically prices that vary with the charging rate, rather than simple price values. It assumes that the energy price each consumer faces is a function of the charging rate (speed) when charging a car. Consequently, a more impatient consumer that wants an EV charged quickly will select a higher charging rate (speed) and will pay a higher price per energy unit (e.g. price/kWh). On the other hand, a more patient consumer can select a lower charging rate and reduce his/her cost for EV charging. With this approach, the grid manager can create counter-incentives for high demand periods in EV charging and also incentives for EV charging during low demand periods.

In addition, this mechanism is adaptive and can quickly adjust the price functions to the needs of the particular consumer portfolio. This is done purely by observing and learning consumer behaviors, without assuming any prior knowledge. Incorporating prior knowledge about the customer portfolio would naturally expedite the convergence process, and whenever this prior knowledge is available, it can be integrated in the IS artifact. However, in cases where this prior knowledge cannot be obtained easily, the mechanism can perform well, without this prerequisite. We showed that using the proposed mechanism the electricity demand curve becomes smoother and has reduced peaks. Furthermore, individual EV owners have their preferences always satisfied and the grid manager achieves convergence of the demand curve to a desired profile.

6.2 Combination of IS artifacts

In each chapter of this dissertation we presented a different IS artifact which was meant to assist a particular stakeholder in the electricity market. Therefore, combining an IS artifact which represents the EV customer, ensuring also renewable usage maximization (Chapter 2 and 3) and an IS artifact on the grid operator's side scheduling the charging in a balanced (for the grid) way (Chapter 4) will create an environment where both sides are using learning mechanisms to satisfy their objectives. Therefore, we need to add an IS artifact that will be overseeing these two independent processes - cost minimization on the EV customer's side and grid stabilization on the grid's side - and ensure that they act in a coordinated manner for the electricity market (Chapter 5). Therefore, it is not enough for each side of the electricity market to use IS artifacts for satisfying their objectives. There needs to be a coordination party which will ensure that the market converges to positive outcomes for both sides (i.e. not herding).

In order for the above combination to be feasible, some assumptions across chapters must be aligned. In all chapters we assumed that the EV customers own just one car which is purely electric, so that we avoid substitution effects from using conventional cars. Furthermore, we assumed that the EV customers are facing variable prices depending on the availability of energy resources. In addition, we assumed that all customers drive within the same local distribution network. Therefore, most of the major assumptions are the same across all chapters. Some assumptions though, need to be modified in order to achieve a feasible combination of artifacts. For example, in Chapter 4, a specific assumption about the EV customers' utility function is made (including delay costs) which is assumed to be zero in Chapters 2 and 3. Another assumption which needs to be relaxed is that in Chapter 2 the EV customers are utility maximizers, whereas in Chapter 5 they are assumed to be cost minimizers. Aligning these assumptions will make a combination of IS artifacts feasible. However, in each chapter we focus on a particular stand-point of the electricity market and provide results using this point of view. This allows for a more in-depth analysis and more clear insights to the involved stakeholders.

6.3 Generalizability and Methodology Discussion

The results presented in each particular chapter can be generalized to domains that fulfil the same assumptions. In Chapter 2 and 3, the IS artifacts representing the EV customers can be generalized to other market places where the customers must make decisions based on complex and fast changing information flows. However, electricity markets are special

in the sense that electricity is a perishable good which can only be stored in small parts (e.g., with EV batteries). This fact, in combination with the high volatility, makes the electricity market domain quite different from other markets and requires investigation of the particular assumptions. In Chapter 4, the IS artifact representing the grid operator can be used in other capacity allocation problems, such as grid computing, inventory problems, etc. The main two differentiating factors are the perishable nature of electricity, which goes beyond the traditional perishable nature of other goods such as fish or flowers (electricity needs to be stored immediately), and the fact that EVs are mainly designed for commuting and not for storing electricity. These two factors, make the balancing of an electricity market more challenging than any other perishable good market. Finally, in Chapter 5, the hybrid coordination mechanism can be applied to other capacity allocation problems where the speed of an offered service is considered. Specifically, when learning mechanisms are used by both the market operator and the customers, the proposed mechanism can be used without major adjustments.

In each chapter the presented IS artifacts are built using different methodology in order to achieve an EV charging coordination. In Chapter 2 we used mathematical modeling and reinforcement learning to model the IS artifact representing the EV customers. This methodology can be quite powerful in exploring the data space and learning from data in an effective way. The main weakness of this approach is that it relies on some assumptions with regard to the customer behavior in order to be modeled mathematically, and these assumptions might not always reflect the actual customer behavior. Another weakness of this method, is that if the reward function is not set properly, it might lead to over-fitting on the training set. To prevent our method from over-fitting we used random permutation of the reward matrix. In Chapter 3, random forest learning is used to forecast the PV generation from weather data. This approach has no over-fitting, but the forecasting accuracy might not be as high as in other methods. In our particular context, the forecasting accuracy is suitable, however the accuracy can be increased by adding more variables in the model, such as geo-location information. In Chapter 4, the multiple Vickrey auction mechanism is used to schedule the EV charging. This mechanism has the major advantage of computational tractability which makes it very appealing to fast changing environments where the auction has a lot of participants and the services must be allocated to the customers quickly. It is not ex-ante incentive compatible though, unlike other auction mechanisms which are not, however, computationally tractable. In our case, the incentive compatibility is not a real issue for the auctioneer since it is more important to achieve a capacity allocation quickly. In Chapter 5, we are using a combination of mathematical modeling and learning to coordinate the EV charging through price

functions. This mechanism is quite flexible and dynamic and can achieve convergence in the electricity market. However, if this mechanism does not adapt properly to the customers' reactions, it might lead to increase of prices without inducing the expected consumption curve. Therefore, the learning component needs to be calibrated properly to achieve the desired convergence.

6.4 Theoretical Contribution and Practical Impact

From an academic point of view, this thesis contributes to the discussion about designing IS artifacts for EV charging coordination. Chapter 2 demonstrates how personalized IS artifacts should be built so that they satisfy individual objectives but also lead to an overall positive outcome for the electricity market (charging coordination). This win-win situations cannot be achieved unless some specific conditions are met: a) the customers are in variable pricing regimes b) their preferences are sufficiently different to lead to a redistribution of charging demand over time. Chapter 3 shows how an IS artifact can be designed so that it benefits from renewable generation, maximizing the usage of renewable sources. The key component of achieving the renewable usage maximization is an accurate forecasting method which can estimate the PV generation from weather data, allowing for a beneficial scheduling of EV charging. Chapter 4 presents an auction mechanism which assumes a social planner (grid operator) in place. This social planner strives for servicing as many customers as possible given a certain capacity. We prove some properties which hold in the electricity market and assist the social planner in scheduling the EV charging. Except for the theoretical properties, testing the mechanism on real-world data provides in-depth insights about the auction's applicability. Chapter 5, proposes a novel pricing mechanism which can incentivize EV customers to produce a desired electricity demand profile. We show that this mechanism achieves convergence to the desired outcome, independent from the customer portfolio it is applied to.

In terms of contribution to practice, this thesis provides useful insights to the electricity market stakeholders. First, it supports EV customers in their complex decision in the electricity markets. Therefore, the presented artifacts can be used by the automotive industry in order to make EV charging easier for the user. For example using insights from Chapter 2, automotive companies can design charging decision support systems embedded in the EV and make the EV more attractive to the end consumer. Second, this thesis provides an in-depth analysis of EV customer behavior in the electricity markets under various pricing regimes. This analysis can be used by the grid operators in order to design their pricing schemes in a way that the EV customers will accept them and also react to

them in the desired way: for example setting the right prices to flatten electricity demand or induce a demand profile following the generation of a renewable source. Except for grid operators, also companies in charge of renewable generation can benefit, since they can team-up with charging infrastructure companies and promote sustainable EV charging (Chapter 3). In addition, grid operators are supported by presenting a new mechanism to clear the EV charging market, not by maximizing profit, but maximizing welfare for all participants in the market (Chapter 4). Finally, grid operators are presented with a novel pricing scheme structure which can be effective in dynamically changing consumer portfolios. They can use the proposed pricing scheme (Chapter 5) to stabilize the market by inducing a desired behavior and preventing herding effects in charging. They can also use the proposed mechanism as a simulation test-bed to examine scenarios that prices might increase unreasonably and design regulation policies based on the proposed simulations and pricing schemes.

6.5 Limitations and Directions for Future Research

This work contributed to the Design Science and Green IS research streams by proposing new algorithms to facilitate a smoother EV integration in the electricity grid. There are some limitations that lead to directions for future research.

First, in Chapter 2, we assumed some functional forms for the customer's utility functions with respect to electricity valuation. These functional forms are purely based on theoretical assumptions of prior literature. Therefore, it would be necessary in the future to conduct real-world experiments to estimate the actual form these utility functions. We have conducted a first experiment, from which we collected data related to electricity consumption valuation. Some first analyses validate our theoretical assumptions, but a more in-depth analysis would be necessary.

Second, in Chapter 3, we examined the impact of our designed IS artifact on the individual level. It will be beneficial in the future, to conduct a study where the impact of the proposed IS artifact is examined in an energy portfolio. That new assumption is expected to create new insights for policy makers, since the saturation points of renewable sources within a portfolio will be discovered. These saturation points are very useful for portfolio managers who strive to maintain a balanced customer portfolio and create incentives for their customers. Furthermore, in this chapter, the correlation of weather conditions and driving behavior is not considered. Accounting for this correlation will increase the precision of our behavioral model and will lead to more accurate results with regard to the EV's ability to consume and store renewable energy from the solar panel.

Third, in Chapter 4, we investigated the welfare maximization scenario, in which the grid operator is interested in reducing the overall delay in the customer population. It would be insightful in the future, to examine the profit maximization scenario, in which the grid operator - in the role of an auctioneer - is interested in making profits when auctioning amounts of electricity for EV charging. The comparison of the social welfare and profit maximization scenarios will create useful implications for smart grid system operators and electricity providers that act as auctioneers. Another limitation of this chapter is that the auction mechanism is validated in simulation-based experiments. The true effectiveness and applicability of this mechanism would be revealed in a real experiment where electricity customers would have to bid for electricity.

Finally, in Chapter 5, it is shown in simulations that the proposed hybrid-coordination mechanism is effective. A real-world experiment would be helpful in demonstrating the applicability of the algorithm in practice. Furthermore, currently we assume linear price functions of charging speed. A future research path could be to experiment with different types of price functions and compare their effectiveness as well as examine the impact of vehicle-to-grid functionalities (Kempton and Tomić, 2005) on the proposed scheme.

Overall, relaxing the assumptions made in the particular chapters with respect to customer behavior would be a future direction that could be promising for future researchers. Understanding and supporting human behavior with intelligent algorithms is a dynamic and fast evolving field that requires continuous improvement. Therefore, future researchers could benefit from this dissertation in building new intelligent algorithms to assist human decision making processes.

References

- S. Acha, T. Green, and N. Shah. Optimal charging strategies of electric vehicles in the UK power market. In *2011 IEEE PES, Innovative Smart Grid Technologies (ISGT)*, pages 1–8, January 2011. doi: 10.1109/ISGT.2011.5759128.
- G. Adomavicius, A. Gupta, and D. Zhdanov. Designing intelligent software agents for auctions with limited information feedback. *Information Systems Research*, 20(4):507–526, 2009.
- A. Almuhtady, S. Lee, E. Romeijn, M. Wynblatt, and J. Ni. A degradation-informed battery-swapping policy for fleets of electric or hybrid-electric vehicles. *Transportation Science*, pages 609 – 618, 2014.
- M. S. Amin and B. F. Wollenberg. Toward a smart grid: Power delivery for the 21st century. *IEEE Power and Energy Magazine*, 3(5):34 – 41, September/October 2005.
- B. Avci, K. Girotra, and S. Netessine. Electric vehicles with a battery switching station: Adoption and environmental impact. *Management Science*, pages 1–23, 2014.
- P. Bajari and A. Hortacsu. The winner’s curse, reserve prices, and endogenous entry: empirical insights from ebay auctions. *RAND Journal of Economics*, pages 329–355, 2003.
- R. Bapna, P. Goes, and A. Gupta. Analysis and design of business-to-consumer online auctions. *Management Science*, 49(1):85–101, 2003.
- R. Bapna, P. Goes, and A. Gupta. Pricing and allocation for quality-differentiated online services. *Management Science*, 51(7):1141–1150, 2005.
- R. Bapna, P. Goes, and A. Gupta. Auctioning vertically integrated online services: computational approaches for real-time allocation. *Journal of Management Information Systems*, 25(3):65–98, 2008.

- E. A. Bender. *An introduction to mathematical modeling*. Courier Corporation, 2012.
- A. Bharadwaj, O. A. El Sawy, P. A. Pavlou, and N. Venkatraman. Digital business strategy: toward a next generation of insights. *Management Information Systems Quarterly*, 37(2):471–482, 2013.
- S. Bhattacharya, K. Kar, J. H. Chow, and A. Gupta. Extended second price auctions for plug-in electric vehicle (pev) charging in smart distribution grids. In *American Control Conference (ACC), 2014*, pages 908–913. IEEE, 2014.
- M. Bichler, A. Gupta, and W. Ketter. Designing smart markets. *Information Systems Research*, 21(4):688–699, 2010.
- S. Blumsack and A. Fernandez. Ready or not, here comes the smart grid! *Energy*, 37(1): 61–68, 2012.
- T. Brandt, S. Wagner, and D. Neumann. Road to 2020: Is-supported business models for electric mobility and electrical energy markets. *Proceedings of the 33rd International Conference on Information Systems (ICIS)*, 2012.
- T. Brandt, S. Feuerriegel, and D. Neumann. A household-oriented approach to the benefits of vehicle-to-grid-capable electric vehicles. In *Wirtschaftsinformatik*, page 105, 2013.
- L. Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- G. H. Brundtland. *Our common future: Report of the 1987 World Commission on Environment and Development*, volume 383. Oxford University Press Oxford, 1987.
- A. Chrysopoulos, C. Diou, A. Symeonidis, and P. Mitkas. Bottom-up modeling of small-scale energy consumers for effective demand response applications. *Engineering Applications of Artificial Intelligence*, 35(0):299 – 315, 2014. ISSN 0952-1976. doi: <http://dx.doi.org/10.1016/j.engappai.2014.06.015>. URL <http://www.sciencedirect.com/science/article/pii/S0952197614001377>.
- E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971.
- J. Corbett. Designing and using carbon management systems to promote ecologically responsible behaviors. *Journal of the Association for Information Systems*, 14(7):339–378, 2013.
- P. Cramton. The fcc spectrum auctions: An early assessment. *Journal of Economics & Management Strategy*, 6(3):431–495, 1997.

- V. Dao, I. Langella, and J. Carbo. From green to sustainability: Information technology and an integrated sustainability framework. *The Journal of Strategic Information Systems*, 20(1):63–79, 2011.
- R. K. Dash, P. Vytelingum, A. Rogers, E. David, and N. R. Jennings. Market-based task allocation mechanisms for limited-capacity suppliers. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, 37(3):391–405, 2007.
- R. A. Daziano and M. Achtnicht. Forecasting adoption of ultra-low-emission vehicles using bayes estimates of a multinomial probit model and the ghk simulator. *Transportation Science*, pages 1–13, 2013.
- K. De Craemer, S. Vandael, B. Claessens, and G. Deconinck. An event-driven dual coordination mechanism for demand side management of phev. *IEEE Transactions on Smart Grid*, 5(2):751–760, March 2014. ISSN 1949-3053.
- J. Dedrick. Green IS: concepts and issues for information systems research. *Communications of the Association for Information Systems*, 27(1):11–18, 2010.
- M. B. Dias, R. Zlot, N. Kalra, and A. Stentz. Market-based multirobot coordination: A survey and analysis. *Proceedings of the IEEE*, 94(7):1257–1270, 2006.
- A. Dimakis, R. Jain, and J. Walrand. Mechanisms for efficient allocation in divisible capacity networks. In *45th IEEE Conference on Decision and Control, 2006*, pages 1264–1269. IEEE, 2006.
- European Commission et al. Roadmap to a single european transport area—towards a competitive and resource efficient transport system (white paper). *Brussels (COM (2011) 144 final)*. Available online at <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do>, 2011.
- M. Fahrioglu and F. L. Alvarado. Designing incentive compatible contracts for effective demand management. *IEEE Transactions on Power Systems*, 15(4):1255–1260, 2000.
- T. Franke, I. Neumann, F. Bühler, P. Cocron, and J. Krems. Experiencing range in an electric vehicle: Understanding psychological barriers. *Applied Psychology*, 61(3): 368–391, 2011.
- E. Gerding, V. Robu, S. Stein, D. Parkes, A. Rogers, and N. Jennings. Online mechanism design for electric vehicle charging. In *The Tenth International Joint Conference on*

- Autonomous Agents and Multi-Agent Systems (AAMAS 2011)*, pages 811–818, May 2011. URL <http://eprints.ecs.soton.ac.uk/21907/>.
- E. Gerding, S. Stein, V. Robu, D. Zhao, and N. R. Jennings. Two-sided online markets for electric vehicle charging. *The Twelfth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2013)*, pages 989–996, May 2013.
- S. Gottwalt, W. Ketter, C. Block, J. Collins, and C. Weinhardt. Demand side management - a simulation of household behavior under variable prices. *Energy Policy*, 39:8163–8174, 2011.
- S. Gregor and A. R. Hevner. Positioning and presenting design science research for maximum impact. *Management Information Systems Quarterly*, 37(2):337–355, 2013.
- S. Gregor and D. Jones. The anatomy of a design theory. *Journal of the Association for Information Systems*, 8(5):312–335, 2007.
- T. Groves. Incentives in teams. *Econometrica*, 41(4):617–631, 1973.
- R. E. Hall and F. S. Mishkin. The sensitivity of consumption to transitory income: estimates from panel data on households, 1982.
- R. Halvgaard, N. K. Poulsen, H. Madsen, J. Jorgensen, F. Marra, and D. E. M. Bondy. Electric vehicle charge planning using economic model predictive control. In *IEEE International Electric Vehicle Conference (IEVC)*, 2012, pages 1–6. IEEE, 2012.
- Y. He, B. Venkatesh, and L. Guan. Optimal scheduling for charging and discharging of electric vehicles. *IEEE Transactions on Smart Grid*, 3(3):1095–1105, 2012.
- A. R. Hevner, S. T. March, J. Park, and S. Ram. Design Science in Information Systems research. *Management Information Systems Quarterly*, 28(1):75–105, 2004.
- H. S. Houthakker. Some calculations on electricity consumption in great britain. *Journal of the Royal Statistical Society. Series A (General)*, 114(3):359–371, 1951.
- International Energy Agency. *Global EV Outlook*. Organisation for Economic Co-operation and Development, Paris, 2013.
- A. Ipakchi and F. Albuyeh. Grid of the future. *IEEE Power and Energy Magazine*, 7(2):52–62, März 2009. ISSN 1540-7977. doi: 10.1109/MPE.2008.931384. URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4787536>.

- M. Kahlen and W. Ketter. Aggregating electric cars to sustainable virtual power plants: The value of flexibility in future electricity markets. In *Twenty-Ninth Conference on Artificial Intelligence (AAAI-15)*, pages 665–671, Austin, Texas, January 2015.
- M. Kahlen, W. Ketter, and J. van Dalen. Balancing with electric vehicles: A profitable business model. In *Proceedings of the 22nd European Conference on Information Systems*, pages 1–16, Tel Aviv, Israel, June 2014.
- A. Kambil and E. van Heck. Reengineering the dutch flower auctions: A framework for analyzing exchange organizations. *Information Systems Research*, 9(1):1–19, 1998.
- A. Kambil and E. van Heck. *Making Markets: How Firms Can Design and Profit from Online Auctions and Exchanges*. Harvard Business School Press, 2002.
- H. Kanchev, D. Lu, F. Colas, V. Lazarov, and B. Francois. Energy management and operational planning of a microgrid with a pv-based active generator for smart grid applications. *IEEE Transactions on Industrial Electronics*, 58(10):4583–4592, 2011.
- G. C. Kane and M. Alavi. Information technology and organizational learning: An investigation of exploration and exploitation processes. *Organization Science*, 18(5):796–812, 2007.
- W. Kempton and J. Tomić. Vehicle-to-grid power implementation: From stabilizing the grid to supporting large-scale renewable energy. *Journal of Power Sources*, 144(1):280–294, 2005. ISSN 03787753. doi: 10.1016/j.jpowsour.2004.12.022.
- W. Kempton, V. Udo, K. Huber, K. Komara, S. Letendre, S. Baker, D. Brunner, and N. Pearre. A test of vehicle-to-grid (v2g) for energy storage and frequency regulation in the pjm system. *Mid-Atlantic Grid Interactive Cars Consortium Jan*, 2009.
- W. Kempton, F. Marra, P. B. Andersen, and R. Garcia-Valle. Business models and control and management architectures for ev electrical grid integration. In *Electric Vehicle Integration into Modern Power Networks*, pages 87–105. Springer, 2013.
- W. Ketter, J. Collins, M. Gini, A. Gupta, and P. Schrater. Real-time tactical and strategic sales management for intelligent agents guided by economic regimes. *Information Systems Research*, 23(4):1263–1283, Dec. 2012. ISSN 1047-7047, 1526-5536.
- W. Ketter, J. Collins, and P. Reddy. Power TAC: A competitive economic simulation of the smart grid. *Energy Economics*, 39:262–270, 2013a.

- W. Ketter, M. Peters, and J. Collins. Autonomous agents in future energy markets: The 2012 Power Trading Agent Competition. In *Association for the Advancement of Artificial Intelligence (AAAI) Conference Proceedings*, pages 1298–1304, Bellevue, WA, July 2013b.
- W. Ketter, J. Collins, P. Reddy, and M. de Weerd. The 2015 Power Trading Agent Competition. Technical Report ERS-2015-001-LIS, RSM Erasmus University, 2015.
- W. Ketter, M. Peters, J. Collins, and A. Gupta. A multiagent competitive gaming platform to address societal challenges. *Management Information Systems Quarterly*, 40(2):447–460, 2016.
- D. Kirschen. Demand-side view of electricity markets. *IEEE Transactions on Power Systems*, 18(2):520–527, Mai 2003. ISSN 0885-8950. doi: 10.1109/TPWRS.2003.810692.
- D. Kirschen and G. Strbac. *Fundamentals of Power System Economics*. John Wiley & Sons, 2005.
- K. Koroleva, M. Kahlen, W. Ketter, L. Rook, and F. Lanz. Tamagocar: Using a simulation app to explore price elasticity of demand for electricity of electric vehicle users. *International Conference on Information Systems (ICIS)*, 2014.
- J. Kossahl, S. Busse, and L. M. Kolbe. The evolvement of energy informatics in the information systems community-a literature analysis and research agenda. In *European Conference on Information Systems*, 2012.
- V. Krishna. *Auction Theory*. Academic Press, 2002. ISBN 9780124262973.
- J.-y. Kwak, P. Varakantham, R. Maheswaran, Y.-H. Chang, M. Tambe, B. Becerik-Gerber, and W. Wood. Tesla: an extended study of an energy-saving agent that leverages schedule flexibility. *Autonomous agents and multi-agent systems*, 28(4):605–636, 2014.
- I. Lampropoulos, G. M. Vanalme, and W. L. Kling. A methodology for modeling the behavior of electricity prosumers within the smart grid. In *Innovative Smart Grid Technologies Conference Europe (ISGT Europe), 2010 IEEE PES*, pages 1–8. IEEE, 2010.
- N. Li, L. Chen, and S. H. Low. Optimal demand response based on utility maximization in power networks. In *Power and Energy Society General Meeting, 2011 IEEE*, pages 1–8. IEEE, 2011.

- Y. Lu, A. Gupta, W. Ketter, and E. V. Heck. Designing intelligent software agents for b2b sequential dutch auctions: A structural econometric approach. *International Conference on Information Systems (ICIS)*, 2013.
- H.-Y. Mak, Y. Rong, and Z.-J. M. Shen. Infrastructure planning for electric vehicles with battery swapping. *Management Science*, 59(7):1557–1575, 2013.
- A. Malhotra, N. P. Melville, and R. T. Watson. Spurring impactful research on information systems for environmental sustainability. *Management Information Systems Quarterly*, 37(4):1265–1274, 2013.
- S. March, A. Hevner, and S. Ram. Research commentary: an agenda for information technology research in heterogeneous and distributed environments. *Information Systems Research*, 11(4):327–341, 2000.
- A. Mas-Colell, M. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press New York, 1995.
- A. S. Masoum, S. Deilami, P. S. Moses, and A. Abu-Siada. Impacts of battery charging rates of plug-in electric vehicle on smart grid distribution systems. In *Innovative Smart Grid Technologies Conference Europe (ISGT Europe), 2010 IEEE PES*, pages 1–6. IEEE, 2010.
- N. P. Melville. Information systems innovation for environmental sustainability. *Management Information Systems Quarterly*, 34(1):1–21, 2010.
- N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- P. Palensky and D. Dietrich. Demand side management: Demand response, intelligent energy systems, and smart loads. *IEEE Transactions on Industrial Informatics*, 7(3):381–388, 2011.
- J. Peças Lopes, S. A. Polenz, C. Moreira, and R. Cherkaoui. Identification of control and management strategies for lv unbalanced microgrids with plugged-in electric vehicles. *Electric Power Systems Research*, 80(8):898–906, 2010.
- M. Peters, W. Ketter, M. Saar-Tsechansky, and J. E. Collins. A reinforcement learning approach to autonomous decision-making in smart electricity markets. *Machine Learning*, 92:5–39, 2013.

- M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, USA, 1st edition, 1994. ISBN 0471619779.
- S. Reichert. Considerations for highly efficient bidirectional battery chargers for e-mobility. In *E-Mobility. Technologien, Infrastruktur, Märkte.*, page 5, Leipzig, 2010. VDE-Verlag. ISBN 978-3-8007-3304-0.
- E. S. Rigas, S. D. Ramchurn, N. Bassiliades, and G. Koutitas. Congestion management for urban ev charging systems. In *2013 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pages 121–126. IEEE, 2013.
- V. Robu, E. H. Gerding, S. Stein, D. C. Parkes, A. Rogers, and N. R. Jennings. An online mechanism for multi-unit demand and its application to plug-in hybrid electric vehicle charging. *Journal of Artificial Intelligence Research*, 48:175–230, 2013.
- P. Samadi, A. Mohsenian-Rad, R. Schober, V. W. Wong, and J. Jatskevich. Optimal real-time pricing algorithm based on utility maximization for smart grid. In *2010 First IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pages 415–420. IEEE, 2010.
- J. Schmidt and S. Busse. The value of is to ensure the security of energy supply—the case of electric vehicle charging. *Proceedings of the Nineteenth Americas Conference on Information Systems, Chicago, Illinois*, 2013.
- S. Sen, Y. Jin, R. Guérin, and K. Hosanagar. Modeling the dynamics of network technology adoption and the role of converters. *IEEE/ACM Transactions on Networking (TON)*, 18(6):1793–1805, 2010.
- S. Shao, M. Pipattanasomporn, and S. Rahman. Demand response as a load shaping tool in an intelligent grid with electric vehicles. *IEEE Transactions on Smart Grid*, 2(4): 624–631, 2011.
- H. A. Simon. Rational decision making in business organizations. *The American Economic Review*, 69(4):493–513, 1979.
- H. A. Simon. *The sciences of the artificial*. MIT press, 1996.
- B. Sovacool and R. Hirsh. Beyond batteries: An examination of the benefits and barriers to plug-in hybrid electric vehicles (phevs) and a vehicle-to-grid (v2g) transition. *Energy Policy*, 37(3):1095–1103, 2009.

- S. Stein, E. Gerding, V. Robu, and N. R. Jennings. A model-based online mechanism with pre-commitment and its application to electric vehicle charging. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 669–676. International Foundation for Autonomous Agents and Multiagent Systems, 2012.
- G. Strbac. Demand side management: Benefits and challenges. *Energy Policy*, 36(12): 4419–4426, Dezember 2008. ISSN 03014215.
- G. Strbac, E. D. Farmer, and B. Cory. Framework for the incorporation of demand side in a competitive market. *IEEE Proceedings on Generation, Transmission and Distribution*, 143(3):232–237, 1996.
- R. S. Sutton and A. G. Barto. *Reinforcement Learning: An introduction*. Cambridge University Press, 1998.
- D. Tilson, K. Lyytinen, and C. Sørensen. Research commentary-digital infrastructures: the missing is research agenda. *Information Systems Research*, 21(4):748–759, 2010.
- J. Tomic and W. Kempton. Using fleets of electric-drive vehicles for grid support. *Journal of Power Sources*, 168(2):459–468, 2007. ISSN 03787753.
- K. Valogianni, W. Ketter, J. Collins, and D. Zhdanov. Effective management of electric vehicle storage using smart charging. *Proceedings of 28th AAAI Conference on Artificial Intelligence*, pages 472–478, 2014a.
- K. Valogianni, W. Ketter, J. Collins, and D. Zhdanov. Enabling sustainable smart homes: An intelligent agent approach. *Proceedings of the 35th International Conference on Information Systems (ICIS)*, 2014b.
- S. Vandael, B. Claessens, M. Hommelberg, T. Holvoet, and G. Deconinck. A scalable three-step approach for demand side management of plug-in hybrid vehicles. *IEEE Transactions on Smart Grid*, 4(2):720 – 728, 2013. ISSN 1949-3053.
- R. A. Verzijlbergh, M. O. Grond, Z. Lukszo, J. G. Slootweg, and M. D. Ilic. Network impacts and cost savings of controlled ev charging. *IEEE Transactions on Smart Grid*, 3(3):1203–1212, 2012.
- W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961.

- J. vom Brocke, R. T. Watson, C. Dwyer, S. Elliot, and N. Melville. Green information systems: Directives for the IS. *Communications of the Association for Information Systems*, 33:30, 2013.
- S. Wagner, M. Götzinger, and D. Neumann. Optimal location of charging stations in smart cities: A points of interest based approach. *Proceedings of the 34th International Conference on Information Systems (ICIS)*, 2013.
- J. G. Walls, G. R. Widmeyer, and O. A. El Sawy. Building an information system design theory for vigilant EIS. *Information Systems Research*, 3(1):36–59, 1992.
- J. G. Walls, G. R. Widmeyer, and O. A. El Sawy. Assessing information system design theory in perspective: how useful was our 1992 initial rendition? *Journal of Information Technology Theory and Application (JITTA)*, 6(2):43–58, 2004.
- G. Watkins. A third factor in the variation of productivity: The load factor. *The American Economic Review*, 5(4):753–786, 1915.
- R. T. Watson, M.-C. Boudreau, and A. J. Chen. Information systems and environmentally sustainable development: energy informatics and new directions for the IS community. *Management Information Systems Quarterly*, 34(1):23–38, 2010.
- M. Wissner. The smart grid—a saucerful of secrets? *Applied Energy*, 88(7):2509–2518, 2011.
- M. Wooldridge and N. Jennings. Intelligent agents: Theory and practice. *Knowledge engineering review*, 10(2):115–152, 1995.

Summary

The purpose of this dissertation is to investigate how intelligent algorithms can support electricity customers in their complex decisions within the electricity grid. In particular, we focus on how electric vehicle (EV) owners can be supported in their charging and discharging decisions, benefiting from the information available. We examine the problem from different standpoints and show the benefits for each involved stakeholder dependent on the market conditions.

In Chapter 2, we take the perspective of an individual EV owner and design an intelligent algorithm which learning from her preferences and driving and consumption information, proposes optimized charging and discharging recommendations. These recommendations are tailored to each individual EV owner and strive to satisfy her own preferences, while at the same time ensure financial benefits on the electricity bill. We observed that besides the EV owners, the proposed algorithm creates benefits for the electricity grid in the form of peak demand reduction. Specifically, when the preference heterogeneity increases, the peak demand is reduced significantly. This leads to an emergent charging coordination which results from different preferences and driving schedules.

In Chapter 3, we extend Chapter 2 by incorporating the EV within a smart home with a photovoltaic panel. The main goal of this study is to examine how accurate solar generation forecasting can be useful for charging the EV and make the best out of renewable sources. We propose a supervised learning algorithm which estimates the solar generation output from the weather conditions. We observe that the algorithm is capable of reducing the electricity costs on the customer side, since significant amount of EV charging is done with renewable energy and capable of increasing the levels of sustainability on the grid.

In Chapter 4, we examine the problem from the grid operator's point of view, taking a top-down approach. We propose an auction mechanism that has as its main goal to service as many EV owners as possible, given a certain grid capacity. We show that an

important parameter in scheduling the EV charging is the cost that each customer incurs from any potential delay. We prove the properties of the optimal EV charging scheduling and show that using the proposed mechanism reduces both the peak electricity demand and the overall delay in the grid.

In Chapter 5, we propose a hybrid mechanism which combines benefits from top-down and bottom-up approaches. This mechanism is based on dynamic price functions that are able to incentivize EV customers to delay their charging duration when there is no urgency. We show that the overall peak demand is reduced and that the herding effects, that might appear in traditional pricing schemes, are mitigated. Furthermore, the proposed mechanism is dynamic and learns from the EV customer portfolio, resulting to fast adaptability when the market conditions change.

Overall, this dissertation contributes to the academic literature with new algorithms that can leverage the power of data available and personalize EV charging recommendations. It also contributes to practice by providing useful insights to the involved stakeholders such as grid operators, energy utility companies, individual customers and automotive companies with respect to creating the right incentives for EV adoption. Finally, it adds to the very important discussion about sustainability, since it proposes ways to reduce carbon footprint and benefit the most from the available renewable sources.

Nederlandse Samenvatting

(Summary in Dutch)

Het doel van dit proefschrift is om te onderzoeken hoe intelligente algoritmes elektriciteit gebruikers kunnen helpen bij het maken van complexe beslissingen op het elektriciteitsnet. Ik leg hierbij de focus op hoe gebruikers van elektrische voertuigen (EV), gebruikmakend van de beschikbare informatie, kunnen worden bijgestaan in het maken van de beslissing om te op- of ontladen. Ik onderzoek het probleem vanuit verschillende standpunten en toon de voordelen, afhankelijk van de marktomstandigheden, voor elke betrokken belanghebbende aan.

In hoofdstuk 2 kijk ik vanuit het perspectief van een individuele EV gebruiker en ontwerp een intelligent algoritme dat in staat is om te leren van zijn voorkeuren en rij- en verbruiksgegevens, en presenteert geoptimaliseerde op- en ontladaanbevelingen. Deze aanbevelingen zijn afgestemd op elke individuele EV gebruiker specifiek, streven ernaar om aan zijn voorkeuren te kunnen voldoen en terwijl zijn elektriciteitsrekening te verlagen. Naast de EV gebruikers zie ik dat de uitgewerkte algoritmes ook voordelig zijn voor het elektriciteitsnet daar de piekvraag wordt gereduceerd. Zo zal een hogere voorkeursheterogeniteit zorgen voor een significante vermindering van de piekvraag en een gedistribueerde laadcoördinatie, afkomstig uit verschillende voorkeuren en soorten rijgedrag.

In hoofdstuk 3 bouw ik verder op hoofdstuk 2 door middel van de integratie van de EV in een slimme woning met een zonnepaneel. Het hoofddoel van deze studie is om te onderzoeken hoe accurate voorspellingen van zonne-energie nuttig kunnen zijn voor het laden van de EV en tegelijkertijd hernieuwbare energiebronnen op de beste manier te kunnen gebruiken. Ik ga hierbij uit van een gecontroleerd leeralgoritme, die een voorspelling maakt van de opgewekte zonne-energie op basis van de weersvoorspellingen. Het algoritme is in staat om de gebruiker zijn elektriciteitskosten te verminderen, daar een aanzienlijk deel van het laden van de EV gebruik maakt van hernieuwbare energie, terwijl

de duurzaamheid van het net hierbij wordt verhoogd.

In hoofdstuk 4 analyseer ik het probleem vanuit het standpunt van de netbeheerder door middel van een top-down design. Dit doe ik aan de hand van een veilingmechanisme, welke als hoofddoel heeft om zoveel als mogelijk EV gebruikers tot dienst te zijn, rekening houden met de capaciteit van het net. Een belangrijke parameter in het plannen van het laden van de EV is de kost die elke gebruiker dreigt te lopen omwille van een mogelijke vertraging. Ik toon de eigenschappen om EVs optimaal op te laden aan en laat zien dat door middel van het voorgestelde mechanisme zowel de piekvraag naar elektriciteit als de algehele vertraging verminderd wordt.

In hoofdstuk 5 stel ik een hybride mechanisme voor, die de voordelen van top-down en bottom-up designs combineert. Dit mechanisme is gebaseerd op dynamische prijsfuncties die in staat zijn om EV gebruikers zonder urgentie te stimuleren de duur van het laden te verlengen. Ik toon aan dat de algehele piekvraag verminderd wordt en dat mogelijke congesties uit traditionele prijsmethodes worden beperkt. Bovendien is het gepresenteerde mechanisme dynamisch en leert het van de EV gebruiker zijn portfolio, wat resulteert in een snel aanpassingsvermogen wanneer marktomstandigheden veranderen.

In conclusie draagt dit proefschrift bij tot de academische wereld door middel van nieuwe algoritmes die de kracht van de beschikbare data benutten en EV laadaanbevelingen personaliseren. Het draagt verder bij tot de praktijk door nuttige inzichten te verstrekken aan de betrokken belanghebbende zoals netbeheerders, energie nutsbedrijven, individuele EV gebruikers en automobiel bedrijven met betrekking tot het creëren van de juiste drijfveren voor de adoptie van EVs. Tot slot draagt het bij tot de zeer belangrijke discussie rond duurzaamheid, aangezien het manieren aandraagt om de maatschappelijke ecologische voetafdruk te verminderen en het meeste uit de beschikbare hernieuwbare energiebronnen te halen.

Περίληψη στα Ελληνικά

(Summary in Greek)

Σκοπός της παρούσας διατριβής είναι να ερευνήσει πώς ευφυείς αλγόριθμοι μπορούν να προσφέρουν υποστήριξη στις πολυδιάστατες αποφάσεις των καταναλωτών ηλεκτρικής ενέργειας, σε σύγχρονα δίκτυα ηλεκτρικής ενέργειας. Συγκεκριμένα, επικεντρώνεται στο πώς οι ιδιοκτήτες ηλεκτρικών οχημάτων μπορούν να υποστηριχθούν στις αποφάσεις φόρτισης και αποφόρτισης του ηλεκτρικού οχήματος, βάσει των διαθέσιμων πληροφοριών της αγοράς. Το πρόβλημα εξετάζεται από τις οπτικές γωνίες όλων των εμπλεκόμενων μερών, δεδομένων διαφορετικών συνθηκών της αγοράς ηλεκτρικής ενέργειας.

Στο κεφάλαιο 2 υιοθετείται η οπτική του καταναλωτή ηλεκτρικής ενέργειας ο οποίος είναι ιδιοκτήτης ενός ηλεκτρικού οχήματος και παρουσιάζεται η σχεδίαση ενός ευφυούς αλγορίθμου, ο οποίος μαθαίνοντας από τις προτιμήσεις του καταναλωτή που αντιπροσωπεύει στην αγορά και τις πληροφορίες κατανάλωσης, προτείνει βελτιστοποιημένες ενέργειες φόρτισης και αποφόρτισης του οχήματος. Αυτές οι προτεινόμενες ενέργειες είναι σχεδιασμένες για κάθε καταναλωτή αποκλειστικά έτσι, ώστε να ικανοποιούν τις προτιμήσεις του, ενώ παράλληλα εξασφαλίζουν μειώσεις στα έξοδα ηλεκτρικής ενέργειάς του. Παρατηρείται ότι εκτός των καταναλωτών, επωφελείται και το ηλεκτρικό δίκτυο, καθώς η μέγιστη ζήτηση του δικτύου μειώνεται σημαντικά, ελαττώνοντας κατά συνέπεια πιθανή συμφόρηση εξαιτίας υψηλής ζήτησης. Συγκεκριμένα παρατηρείται, ότι αυξανόμενης της ετερογένειας των προτιμήσεων από πλευράς καταναλωτών, δημιουργείται μείωση της μέγιστης ζήτησης που οδηγεί σε συγχρονισμό της φόρτισης των διαφόρων οχημάτων στο δίκτυο.

Στο κεφάλαιο 3 εξετάζεται η οπτική του καταναλωτή στο πλαίσιο ενός έξυπνου σπιτιού. Βασικός σκοπός αυτού του κεφαλαίου είναι να διερευνήσει πόσο χρήσιμες μπορούν να φανούν προβλέψεις που αφορούν στην παραγωγή ηλεκτρικής ενέργειας από φωτοβολταϊκά πάνελ σε συνδυασμό με την φόρτιση ενός ηλεκτρικού οχήματος. Παρουσιάζεται η σχεδίαση ενός ευφυούς συστήματος μηχανικής μάθησης που εκτιμά την ποσότητα ηλεκτρικής ενέργειας που

παράγεται από ένα φωτοβολταϊκό πάνελ, χρησιμοποιώντας δεδομένα καιρικών συνθηκών. Παρατηρείται ότι το προτεινόμενο σύστημα είναι ικανό να μειώσει τα έξοδα των καταναλωτών που το υιοθετούν, καθώς σημαντικό ποσοστό της ενέργειας που απαιτείται για τη φόρτιση του οχήματος προέρχεται από ανανεώσιμες πηγές (φωτοβολταϊκό πάνελ) και όχι από το δίκτυο. Συνεπώς, τα επίπεδα εκμετάλλευσης ανανεώσιμων πηγών στο δίκτυο αυξάνονται σημαντικά.

Στο κεφάλαιο 4 υιοθετείται η οπτική του διαχειριστή δικτύου ηλεκτρικής ενέργειας που επιθυμεί την ομαλή ένταξη ηλεκτρικών οχημάτων στο δίκτυο. Παρουσιάζεται ένα μηχανισμός δημοπρασίας ο οποίος στοχεύει στην ένταξη όσο το δυνατόν περισσότερων οχημάτων στο δίκτυο χωρίς να διαταράσσεται η σταθερότητα και η αξιοπιστία του, δεδομένης μια συγκεκριμένης χωρητικότητας του δικτύου. Παρατηρείται ότι μια σημαντική παράμετρος στον προγραμματισμό της φόρτισης ηλεκτρικών οχημάτων με σκοπό την ομαλή λειτουργία του δικτύου είναι το κόστος που δημιουργεί σε κάθε καταναλωτή μια πιθανή καθυστέρηση. Αποδεικνύονται οι μαθηματικές συνθήκες που πρέπει να πληρούνται, ώστε να γίνεται βελτιστοποιημένος προγραμματισμός της φόρτισης των ηλεκτρικών οχημάτων. Παρατηρείται επίσης ότι ο προτεινόμενος μηχανισμός δημοπρασίας μειώνει τη μέγιστη ζήτηση στο δίκτυο και την συνολική καθυστέρηση που ενδέχεται να υποστούν οι καταναλωτές.

Στο κεφάλαιο 5 προτείνεται ένας υβριδικός μηχανισμός που ενσωματώνει την οπτική του καταναλωτή και του δικτύου, παρουσιάζοντας πλεονεκτήματα και των δυο. Ο προτεινόμενος μηχανισμός βασίζεται σε δυναμικές συναρτήσεις τιμολόγησης ηλεκτρικής ενέργειας, ικανές να δημιουργήσουν κίνητρα για τους καταναλωτές έτσι, ώστε να φορτίζουν τα οχήματά τους, όταν υπάρχει πλεόνασμα ηλεκτρικής ενέργειας και το αντίθετο. Παρατηρείται ότι με τη χρήση αυτού του μηχανισμού, η μέγιστη ζήτηση του δικτύου μειώνεται σημαντικά καθώς επίσης και τα φαινόμενα συμφόρησης τα οποία μπορεί να εμφανιστούν σε περιπτώσεις δυναμικής τιμολόγησης. Επιπλέον, ο παρών μηχανισμός είναι ευέλικτος και ικανός να 'μαθαίνει' από τα δεδομένα των καταναλωτών, έχοντας γρήγορη προσαρμοστικότητα σε ταχέως μεταβαλλόμενες συνθήκες αγοράς.

Συνοψίζοντας, η παρούσα διατριβή συνεισφέρει στην ακαδημαϊκή έρευνα παρουσιάζοντας καινοτόμους αλγόριθμους, ικανούς να επωφεληθούν από την αφθονία δεδομένων που είναι διαθέσιμα στις αγορές ηλεκτρικής ενέργειας και να προτείνουν βελτιστοποιημένες ενέργειες φόρτισης των ηλεκτρικών οχημάτων. Όσον αφορά στην πρακτική εφαρμογή, οι προτεινόμενοι αλγόριθμοι ρίχνουν φως στις πολυδιάστατες αποφάσεις των εμπλεκόμενων μερών, όπως οι διαχειριστές δικτύου ηλεκτρικής ενέργειας, οι πάροχοι ηλεκτρικής ενέργειας, οι καταναλωτές και οι βιομηχανίες παραγωγής ηλεκτρικών οχημάτων, σχετικά με την ομαλή

ενσωμάτωση ηλεκτρικών οχημάτων στο δίκτυο. Τέλος, συνεισφέρει στον πολύ σημαντικό τομέα αξιοποίησης των ανανεώσιμων πηγών ενέργειας και στη μείωση των εκπομπών διοξειδίου του άνθρακα.

Curriculum Vitae



Konstantina Valogianni was born in 1988 in Larisa, Greece. She studied Electrical and Computer Engineering in Aristotle University of Thessaloniki, Greece, where she specialized in computer software and electronics.

In September 2011, Konstantina started her Ph.D. with the department of Technology & Operations Management at Rotterdam School of Management, Erasmus University. In 2014, she spent three months as a visiting scholar at Carlson School of Management, University of Minnesota. Her re-

search interests include electricity markets, artificial intelligence, data-driven algorithmic design, and intelligent agents. Alongside with her Ph.D. Konstantina has worked as a researcher for the FP7 EU-funded project Cassandra. Her work has appeared in the proceedings of various conferences such as the Conference of Advancement on Artificial Intelligence (AAAI), International Conference on Autonomous Agents and Multiagent Systems (AAMAS), International Conference on Information Systems (ICIS), Conference on Information Systems and Technology (CIST), Winter Conference in Business Intelligence (WCBI).

Apart from her research activities, Konstantina has been involved in teaching of the master courses “Designing Business Applications”, “Next Generation Business Applications” and “Programming for Managers”, and has supervised several MSc theses.

As of September 2016, Konstantina will be employed as an Assistant Professor of Information Systems at IE Business School, Madrid, Spain.

Portfolio

Publications

Papers under Review

- Valogianni, K., Ketter, W., Collins, J. & Zhdanov, D. Sustainable Demand Side Management for Electric Vehicles using Personalized Learning Agents. *2nd* Review Round.
- Valogianni, K. & Ketter, W. Data-Driven Electricity Customer Decision Support: Evidence from a Real-world Pilot. *2nd* Review Round.

Working Papers

- Valogianni, K., Ketter, W., Collins, J. & Zhdanov, D. Sustainable Electric Vehicle Charging: A Data-driven Approach. Under final preparation for submission.
- Valogianni, K., Gupta, A., Ketter, W., Sen, S. & van Heck, E. Maximizing Social Welfare in Grid Resource Allocation for Electric Vehicle Charging. Under preparation for submission.
- Valogianni, K., Ketter, W., Collins, J. & Adomavicius, G. A Dynamic Pricing Mechanism to Coordinate Electric Vehicle Charging. Under preparation for submission.

Conference Proceedings

- Valogianni, K., Gupta, A., Ketter, W., Sen, S. & van Heck, E.(2016). Using Optimal Grid Resources for Coordinating Electrical Vehicle Charging. In Winter Conference on Business Intelligence (WCBI 2016), Snowbird, Utah.
- Valogianni, K., Gupta, A., Ketter, W., Sen, S. & van Heck, E.(2015). Maximizing Social Welfare in Grid Resource Allocation for Electric Vehicle Charging. In

Conference on Information Systems and Technology, (CIST 2015), Philadelphia, PA.

- Valogianni, K., Ketter, W. & Collins, J. (2015). A Multiagent Approach to Variable-Rate Electric Vehicle Charging Coordination. In Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), pp. 1131-1139, May, 4-8, 2015, Istanbul, Turkey.
- Valogianni, K., Ketter, W., Collins, J. & Adomavicius, G. (2015). A Dynamic Pricing Mechanism to Coordinate Electric Vehicle Charging. In Winter Conference on Business Intelligence (WCBI 2015), Snowbird, Utah.
- Valogianni, K., Ketter, W., Collins, J. & Zhdanov, D.(2014). Enabling Sustainable Smart Homes: An Intelligent Agent Approach. In Proceedings of the International Conference on Information Systems, (ICIS 2014), Auckland, New Zealand, Association for Information Systems (AIS).
- Valogianni, K., Ketter, W., Collins, J. & Zhdanov, D.(2014). Adaptive Learning Agents for Electric Vehicle Customer Decision Support. In Conference on Information Systems and Technology, (CIST 2014), San Francisco, California.
- Valogianni, K., Ketter, W., Collins, J. & Zhdanov, D.(2014). Effective Management of EV Storage using Smart Charging. In Proceedings of the Twenty-Eighth Conference of Advancement on Artificial Intelligent (AAAI 2014), pp. 472-478, July, 27-31, Quebec, Canada.
- Valogianni, K., Ketter, W., & Collins, J. (2014). Learning to Schedule Electric Vehicle Charging given Individual Customer Preferences. In Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014), Lomuscio, Scerri, Bazzan, Huhns (eds.) pp.1591-1592, May 5-9, 2014, Paris, France.
- Valogianni, K., Ketter, W., de Weerd, M. & Collins, J. (2013). Influence of the Customers Risk Attitude on the Introduction of Electric Vehicles. In Winter Conference on Business Intelligence (WCBI 2013), Snowbird, Utah.
- Valogianni, K., Ketter, W., de Weerd, M. & Collins, J. (2012). Analysis of Smart Grid Balancing using Realistic Customer Models. In Conference on Information Systems and Technology (CIST 2012), Phoenix, Arizona.

- Kahlen, M., Valogianni, K., Ketter, W. & van Dalen, J. (2012). A Profitable Business Model for Electric Vehicle Fleet Owners. In IEEE - International Conference on Smart Grid Technology (SG-TEP 2012), Economics and Policies, pp. 1-5, Nuremberg, Germany.
- Richter, A., van der Laan, E., Ketter, W. & Valogianni, K.. (2012). Transitioning from the traditional to the smart grid: Lessons learned from closed-loop supply chains. In IEEE - International Conference on Smart Grid Technology (SG-TEP 2012), Economics and Policies, pp. 1-7, Nuremberg, Germany.
- Collins, J., Ketter, W., Valogianni, K. & de Weerd, M. (2012). Constructing and operating a large - scale simulation of retail electric power markets. In 8th Symposium on Statistical Challenges in Electronic Commerce Research (SCECR 2012), Montreal, Canada.

Workshop Proceedings

- Valogianni, K., Ketter, W., Collins, J. & Zhdanov, D (2015). Sustainable Electric Vehicle Charging: A Data-driven Approach. In 25th ICIS Workshop of Information Technologies and Systems (WITS 2015), Fort Worth, Texas.
- Valogianni, K., Ketter, W., & Collins, J. (2015). A Hybrid Mechanism to Coordinate Electric Vehicle Charging. In 2015 AAMAS Workshop - Trading Agents Analysis and Design (TADA 2015), Istanbul, Turkey.
- Valogianni, K., Ketter, W., & Collins, J. (2013). Smart Charging of Electric Vehicles using Intelligent Software Agents. In 23rd ICIS Workshop of Information Technologies and Systems (WITS 2013), Milano, Italy.
- Valogianni, K., Ketter, W., & Collins, J. (2013). Smart Charging of Electric Vehicles. In 2013 AAAI Workshop - Trading Agents Analysis and Design (TADA 2013), Bellevue, Washington.
- Valogianni, K., Ketter, W., de Weerd, M. & Collins, J. (2013). Modeling Electric Vehicle Customers and Learning a Heuristic Charging Behavior. In Autonomous Agents and Multiagent Systems (AAMAS) Workshop - Agent Technologies and Energy Systems (ATES 2013), Minneapolis, Minnesota.
- Valogianni, K., Ketter, W., de Weerd, M. & Collins, J. (2012). Increasing Social Welfare and Individual Savings using Economic Incentives for Electric Vehicles.

In 22nd ICIS Workshop of Information Technologies and Systems (WITS 2012), Orlando, Florida.

Research Visit

Decision and Information Sciences Department, University of Minnesota (September 2014-December 2014)

Teaching

“Programming for Managers”, MSc Business Information Management – elective course

“Next Generation Business Applications”, MSc Business Information Management – elective course

“Designing Business Applications”, MSc Business Information Management – core course

MSc thesis supervision both within the MSc Business Information Management and the Masters in Management trajectory

PhD Courses

Research Proposal Writing

ERIM/CentER Workshop on Information Management Research

English

Machine Learning

Applied Econometrics

Publishing Strategy

Statistical Methods

Social Networks and Market Competition

Multi Agent Systems Research

Behavioral Decision Theory

Interaction Performance Training / Coaching

Language Skills and Certifications

Cambridge Certificate of Proficiency in English (CPE)

TOEFL

GMAT

Zentrale Mittelstufenprüfung (ZMP) - Goethe Institut

ERASMUS RESEARCH INSTITUTE OF MANAGEMENT (ERIM)

ERIM PH.D. SERIES RESEARCH IN MANAGEMENT

The ERIM PhD Series contains PhD dissertations in the field of Research in Management defended at Erasmus University Rotterdam and supervised by senior researchers affiliated to the Erasmus Research Institute of Management (ERIM). All dissertations in the ERIM PhD Series are available in full text through the ERIM Electronic Series Portal: <http://repub.eur.nl/pub>. ERIM is the joint research institute of the Rotterdam School of Management (RSM) and the Erasmus School of Economics at the Erasmus University Rotterdam (EUR).

DISSERTATIONS LAST FIVE YEARS

Abbink, E.J., *Crew Management in Passenger Rail Transport*, Promotor(s): Prof.dr. L.G. Kroon & Prof.dr. A.P.M. Wagelmans, EPS-2014-325-LIS, <http://repub.eur.nl/pub/76927>

Acar, O.A., *Crowdsourcing for Innovation: Unpacking Motivational, Knowledge and Relational Mechanisms of Innovative Behavior in Crowdsourcing Platforms*, Promotor(s): Prof.dr.ir. J.C.M. van den Ende, EPS-2014-321-LIS, <http://repub.eur.nl/pub/76076>

Akin Ates, M., *Purchasing and Supply Management at the Purchase Category Level: strategy, structure and performance*, Promotor(s): Prof.dr. J.Y.F. Wynstra & Dr. E.M. van Raaij, EPS-2014-300-LIS, <http://repub.eur.nl/pub/50283>

Akpınar, E., *Consumer Information Sharing*, Promotor(s): Prof.dr.ir. A. Smidts, EPS-2013-297-MKT, <http://repub.eur.nl/pub/50140>

Alexander, L., *People, Politics, and Innovation: A Process Perspective*, Promotor(s): Prof.dr. H.G. Barkema & Prof.dr. D.L. van Knippenberg, EPS-2014-331-S&E, <http://repub.eur.nl/pub/77209>

Almeida e Santos Nogueira, R.J. de, *Conditional Density Models Integrating Fuzzy and Probabilistic Representations of Uncertainty*, Promotor(s): Prof.dr.ir. U. Kaymak & Prof.dr. J.M.C. Sousa, EPS-2014-310-LIS, <http://repub.eur.nl/pub/51560>

Bannouh, K., *Measuring and Forecasting Financial Market Volatility using High-frequency Data*, Promotor(s): Prof.dr. D.J.C. van Dijk, EPS-2013-273-F&A, <http://repub.eur.nl/pub/38240>

Ben-Menahem, S.M., *Strategic Timing and Proactiveness of Organizations*, Promotor(s): Prof.dr. H.W. Volberda & Prof.dr.ing. F.A.J. van den Bosch, EPS-2013-278-S&E, <http://repub.eur.nl/pub/39128>

Benning, T.M., *A Consumer Perspective on Flexibility in Health Care: Priority Access Pricing and Customized Care*, Promotor(s): Prof.dr.ir. B.G.C. Dellaert, EPS-2011-241-MKT, <http://repub.eur.nl/pub/23670>

Berg, W.E. van den, *Understanding Salesforce Behavior using Genetic Association Studies*, Promotor(s): Prof.dr. W.J.M.I. Verbeke, EPS-2014-311-MKT, <http://repub.eur.nl/pub/51440>

Betancourt, N.E., *Typical Atypicality: Formal and Informal Institutional Conformity, Deviance, and Dynamics*, Promotor(s): Prof.dr. B. Krug, EPS-2012-262-ORG, <http://repub.eur.nl/pub/32345>

Bliek, R. de, *Empirical Studies on the Economic Impact of Trust*, Promotor(s): Prof.dr. J. Veenman & Prof.dr. Ph.H.B.F. Franses, EPS-2015-324-ORG, <http://repub.eur.nl/pub/78159>

Blitz, D.C., *Benchmarking Benchmarks*, Promotor(s): Prof.dr. A.G.Z. Kemna & Prof.dr. W.F.C. Verschoor, EPS-2011-225-F&A, <http://repub.eur.nl/pub/22624>

Boons, M., *Working Together Alone in the Online Crowd: The Effects of Social Motivations and Individual Knowledge Backgrounds on the Participation and Performance of Members of Online Crowdsourcing Platforms*, Promotor(s): Prof.dr. H.G. Barkema & Dr. D.A. Stam, EPS-2014-306-S&E, <http://repub.eur.nl/pub/50711>

Brazys, J., *Aggregated Macroeconomic News and Price Discovery*, Promotor(s): Prof.dr. W.F.C. Verschoor, EPS-2015-351-F&A, <http://repub.eur.nl/pub/78243>

Burger, M.J., *Structure and Cooption in Urban Networks*, Promotor(s): Prof.dr. G.A. van der Knaap & Prof.dr. H.R. Commandeur, EPS-2011-243-ORG, <http://repub.eur.nl/pub/26178>

Byington, E., *Exploring Coworker Relationships: Antecedents and Dimensions of Interpersonal Fit, Coworker Satisfaction, and Relational Models*, Promotor(s): Prof.dr. D.L. van Knippenberg, EPS-2013-292-ORG, <http://repub.eur.nl/pub/41508>

Camacho, N.M., *Health and Marketing: Essays on Physician and Patient Decision-Making*, Promotor(s): Prof.dr. S. Stremersch, EPS-2011-237-MKT, <http://repub.eur.nl/pub/23604>

Cancurtaran, P., *Essays on Accelerated Product Development*, Promotor(s): Prof.dr. F. Langerak & Prof.dr. G.H. van Bruggen, EPS-2014-317-MKT, <http://repub.eur.nl/pub/76074>

Caron, E.A.M., *Explanation of Exceptional Values in Multi-dimensional Business Databases*, Promotor(s): Prof.dr. H.A.M. Daniels & Prof.dr. G.W.J. Hendrikse, EPS-2013-296-LIS, <http://repub.eur.nl/pub/50005>

Carvalho, L. de, *Knowledge Locations in Cities: Emergence and Development Dynamics*, Promotor(s): Prof.dr. L. Berg, EPS-2013-274-S&E, <http://repub.eur.nl/pub/38449>

Cox, R.H.G.M., *To Own, To Finance, and To Insure - Residential Real Estate Revealed*, Promotor(s): Prof.dr. D. Brounen, EPS-2013-290-F&A, <http://repub.eur.nl/pub/40964>

Deichmann, D., *Idea Management: Perspectives from Leadership, Learning, and Network Theory*, Promotor(s): Prof.dr.ir. J.C.M. van den Ende, EPS-2012-255-ORG, <http://repub.eur.nl/pub/31174>

Deng, W., *Social Capital and Diversification of Cooperatives*, Promotor(s): Prof.dr. G.W.J. Hendrikse, EPS-2015-341-ORG, <http://repub.eur.nl/pub/77449>

Desmet, P.T.M., *In Money we Trust? Trust Repair and the Psychology of Financial Compensations*, Promotor(s): Prof.dr. D. de Cremer, EPS-2011-232-ORG, <http://repub.eur.nl/pub/23268>

Dollevoet, T.A.B., *Delay Management and Dispatching in Railways*, Promotor(s): Prof.dr. A.P.M. Wagelmans, EPS-2013-272-LIS, <http://repub.eur.nl/pub/38241>

Doorn, S. van, *Managing Entrepreneurial Orientation*, Promotor(s): Prof.dr. J.J.P. Jansen, Prof.dr.ing. F.A.J. van den Bosch, & Prof.dr. H.W. Volberda, EPS-2012-258-STR, <http://repub.eur.nl/pub/32166>

Douwens-Zonneveld, M.G., *Animal Spirits and Extreme Confidence: No Guts, No Glory?* Promotor(s): Prof.dr. W.F.C. Verschoor, EPS-2012-257-F&A, <http://repub.eur.nl/pub/31914>

Duca, E., *The Impact of Investor Demand on Security Offerings*, Promotor(s): Prof.dr. A. de Jong, EPS-2011-240-F&A, <http://repub.eur.nl/pub/26041>

Duyvesteyn, J.G. *Empirical Studies on Sovereign Fixed Income Markets*, Promotor(s): Prof.dr P.Verwijmeren & Prof.dr. M.P.E. Martens, EPS-2015-361-F&A, hdl.handle.net/1765/79033

Duursema, H., *Strategic Leadership: Moving Beyond the Leader-Follower Dyad*, Promotor(s): Prof.dr. R.J.M. van Tulder, EPS-2013-279-ORG, <http://repub.eur.nl/pub/39129>

Eck, N.J. van, *Methodological Advances in Bibliometric Mapping of Science*, Promotor(s): Prof.dr.ir. R. Dekker, EPS-2011-247-LIS, <http://repub.eur.nl/pub/26509>

Elmes, A., *Studies on Determinants and Consequences of Financial Reporting Quality*, Promotor: Prof.dr. E.Peeck, EPS-2015-354-F&A, <http://hdl.handle.net/1765/79037>

Ellen, S. ter, *Measurement, Dynamics, and Implications of Heterogeneous Beliefs in Financial Markets*, Promotor(s): Prof.dr. W.F.C. Verschoor, EPS-2015-343-F&A, <http://repub.eur.nl/pub/78191>

Eskenazi, P.I., *The Accountable Animal*, Promotor(s): Prof.dr. F.G.H. Hartmann, EPS-2015-355-F&A, <http://repub.eur.nl/pub/78300>

Essen, M. van, *An Institution-Based View of Ownership*, Promotor(s): Prof.dr. J. van Oosterhout & Prof.dr. G.M.H. Mertens, EPS-2011-226-ORG, <http://repub.eur.nl/pub/22643>

Evangelidis, I., *Preference Construction under Prominence*, Promotor(s): Prof.dr. S.M.J. van Osselaer, EPS-2015-340-MKT, <http://repub.eur.nl/pub/78202>

Faber, N., *Structuring Warehouse Management*, Promotor(s): Prof.dr. MB.M. de Koster, Prof.dr. Ale Smidts, EPS-2015-336-LIS, <http://repub.eur.nl/pub/78603>

Fernald, K., *The Waves of Biotechnological Innovation in Medicine: Interfirm Cooperation Effects and a Venture Capital Perspective*, Promotor(s): Prof.dr. E.Claassen, Prof.dr. H.P.G.Pennings & Prof.dr. H.R. Commandeur, EPS-2015-371-S&E, <http://hdl.handle.net/1765/79120>

Fourne, S.P., *Managing Organizational Tensions: A Multi-Level Perspective on Exploration, Exploitation and Ambidexterity*, Promotor(s): Prof.dr. J.J.P. Jansen & Prof.dr. S.J. Magala, EPS-2014-318-S&E, <http://repub.eur.nl/pub/76075>

Gharehgozli, A.H., *Developing New Methods for Efficient Container Stacking Operations*, Promotor(s): Prof.dr.ir. M.B.M. de Koster, EPS-2012-269-LIS, <http://repub.eur.nl/pub/37779>

Gils, S. van, *Morality in Interactions: On the Display of Moral Behavior by Leaders and Employees*, Promotor(s): Prof.dr. D.L. van Knippenberg, EPS-2012-270-ORG, <http://repub.eur.nl/pub/38027>

Ginkel-Bieshaar, M.N.G. van, *The Impact of Abstract versus Concrete Product Communications on Consumer Decision-making Processes*, Promotor(s): Prof.dr.ir. B.G.C. Dellaert, EPS-2012-256-MKT, <http://repub.eur.nl/pub/31913>

Gkougkousi, X., *Empirical Studies in Financial Accounting*, Promotor(s): Prof.dr. G.M.H. Mertens & Prof.dr. E. Peek, EPS-2012-264-F&A, <http://repub.eur.nl/pub/37170>

Glorie, K.M., *Clearing Barter Exchange Markets: Kidney Exchange and Beyond*, Promotor(s): Prof.dr. A.P.M. Wagelmans & Prof.dr. J.J. van de Klundert, EPS-2014-329-LIS, <http://repub.eur.nl/pub/77183>

Hekimoglu, M., *Spare Parts Management of Aging Capital Products*, Promotor: Prof.dr.ir. R. Dekker, EPS-2015-368-LIS, <http://hdl.handle.net/1765/79092>

Heij, C.V., *Innovating beyond Technology. Studies on how management innovation, co-creation and business model innovation contribute to firm's (innovation) performance*, Promotor(s): Prof.dr.ing. F.A.J. van den Bosch & Prof.dr. H.W. Volberda, EPS-2012-370-STR, <http://repub.eur.nl/pub/78651>

Heyde Fernandes, D. von der, *The Functions and Dysfunctions of Reminders*, Promotor(s): Prof.dr. S.M.J. van Osselaer, EPS-2013-295-MKT, <http://repub.eur.nl/pub/41514>

Heyden, M.L.M., *Essays on Upper Echelons & Strategic Renewal: A Multilevel Contingency Approach*, Promotor(s): Prof.dr.ing. F.A.J. van den Bosch & Prof.dr. H.W. Volberda, EPS-2012-259-STR, <http://repub.eur.nl/pub/32167>

Hoever, I.J., *Diversity and Creativity*, Promotor(s): Prof.dr. D.L. van Knippenberg, EPS-2012-267-ORG, <http://repub.eur.nl/pub/37392>

Hogenboom, A.C., *Sentiment Analysis of Text Guided by Semantics and Structure*, Promotor(s): Prof.dr.ir. U. Kaymak & Prof.dr. F.M.G. de Jong, EPS-2015-369-LIS, <http://hdl.handle.net/1765/79034>

Hogenboom, F.P., *Automated Detection of Financial Events in News Text*, Promotor(s): Prof.dr.ir. U. Kaymak & Prof.dr. F.M.G. de Jong, EPS-2014-326-LIS, <http://repub.eur.nl/pub/77237>

Hollen, R.M.A., *Exploratory Studies into Strategies to Enhance Innovation-Driven International Competitiveness in a Port Context: Toward Ambidextrous Ports*, Promotor(s): Prof.dr.ing. F.A.J. Van Den Bosch & Prof.dr. H.W. Volberda, EPS-2015-372-S&E, hdl.handle.net/1765/78881

Hoogendoorn, B., *Social Entrepreneurship in the Modern Economy: Warm Glow, Cold Feet*, Promotor(s): Prof.dr. H.P.G. Pennings & Prof.dr. A.R. Thurik, EPS-2011-246-STR, <http://repub.eur.nl/pub/26447>

Hoogervorst, N., *On The Psychology of Displaying Ethical Leadership: A Behavioral Ethics Approach*, Promotor(s): Prof.dr. D. de Cremer & Dr. M. van Dijke, EPS-2011-244-ORG, <http://repub.eur.nl/pub/26228>

Hout, D.H. van, *Measuring Meaningful Differences: Sensory Testing Based Decision Making in an Industrial Context; Applications of Signal Detection Theory and Thurstonian Modelling*, Promotor(s): Prof.dr. P.J.F. Groenen & Prof.dr. G.B. Dijksterhuis, EPS-2014-304-MKT, <http://repub.eur.nl/pub/50387>

Houwelingen, G.G. van, *Something To Rely On*, Promotor(s): Prof.dr. D. de Cremer & Prof.dr. M.H. van Dijke, EPS-2014-335-ORG, <http://repub.eur.nl/pub/77320>

Hurk, E. van der, *Passengers, Information, and Disruptions*, Promotor(s): Prof.dr. L.G. Kroon & Prof.dr. P.H.M. Vervest, EPS-2015-345-LIS, <http://repub.eur.nl/pub/78275>

Hytonen, K.A., *Context Effects in Valuation, Judgment and Choice: A Neuroscientific Approach*, Promotor(s): Prof.dr.ir. A. Smidts, EPS-2011-252-MKT, <http://repub.eur.nl/pub/30668>

Iseger, P. den, *Fourier and Laplace Transform Inversion with Applications in Finance*, Promotor(s): Prof.dr.ir. R. Dekker, EPS-2014-322-LIS, <http://repub.eur.nl/pub/76954>

Jaarsveld, W.L. van, *Maintenance Centered Service Parts Inventory Control*, Promotor(s): Prof.dr.ir. R. Dekker, EPS-2013-288-LIS, <http://repub.eur.nl/pub/39933>

Jalil, M.N., *Customer Information Driven After Sales Service Management: Lessons from Spare Parts Logistics*, Promotor(s): Prof.dr. L.G. Kroon, EPS-2011-222-LIS, <http://repub.eur.nl/pub/22156>

Kappe, E.R., *The Effectiveness of Pharmaceutical Marketing*, Promotor(s): Prof.dr. S. Stremsch, EPS-2011-239-MKT, <http://repub.eur.nl/pub/23610>

Karreman, B., *Financial Services and Emerging Markets*, Promotor(s): Prof.dr. G.A. van der Knaap & Prof.dr. H.P.G. Pennings, EPS-2011-223-ORG, <http://repub.eur.nl/pub/22280>

Khanagha, S., *Dynamic Capabilities for Managing Emerging Technologies*, Promotor(s): Prof.dr. H.W. Volberda, EPS-2014-339-S&E, <http://repub.eur.nl/pub/77319>

Kil, J., *Acquisitions Through a Behavioral and Real Options Lens*, Promotor(s): Prof.dr. H.T.J. Smit, EPS-2013-298-F&A, <http://repub.eur.nl/pub/50142>

Klooster, E. van 't, *Travel to Learn: the Influence of Cultural Distance on Competence Development in Educational Travel*, Promotor(s): Prof.dr. F.M. Go & Prof.dr. P.J. van Baalen, EPS-2014-312-MKT, <http://repub.eur.nl/pub/51462>

Koendjibiharie, S.R., *The Information-Based View on Business Network Performance: Revealing the Performance of Interorganizational Networks*, Promotor(s): Prof.dr.ir. H.W.G.M. van Heck & Prof.dr. P.H.M. Vervest, EPS-2014-315-LIS, <http://repub.eur.nl/pub/51751>

Koning, M., *The Financial Reporting Environment: The Role of the Media, Regulators and Auditors*, Promotor(s): Prof.dr. G.M.H. Mertens & Prof.dr. P.G.J. Roosenboom, EPS-2014-330-F&A, <http://repub.eur.nl/pub/77154>

Konter, D.J., *Crossing Borders with HRM: An Inquiry of the Influence of Contextual Differences in the Adoption and Effectiveness of HRM*, Promotor(s): Prof.dr. J. Paauwe & Dr. L.H. Hoeksema, EPS-2014-305-ORG, <http://repub.eur.nl/pub/50388>

Korkmaz, E., *Bridging Models and Business: Understanding Heterogeneity in Hidden Drivers of Customer Purchase Behavior*, Promotor(s): Prof.dr. S.L. van de Velde & Prof.dr. D. Fok, EPS-2014-316-LIS, <http://repub.eur.nl/pub/76008>

Kroezen, J.J., *The Renewal of Mature Industries: An Examination of the Revival of the Dutch Beer Brewing Industry*, Promotor(s): Prof.dr. P.P.M.A.R. Heugens, EPS-2014-333-S&E, <http://repub.eur.nl/pub/77042>

Kysucky, V., *Access to Finance in a Cross-Country Context*, Promotor(s): Prof.dr. L. Norden, EPS-2015-350-F&A, <http://repub.eur.nl/pub/78225>

Lam, K.Y., *Reliability and Rankings*, Promotor(s): Prof.dr. Ph.H.B.F. Franses, EPS-2011-230-MKT, <http://repub.eur.nl/pub/22977>

Lander, M.W., *Profits or Professionalism? On Designing Professional Service Firms*, Promotor(s): Prof.dr. J. van Oosterhout & Prof.dr. P.P.M.A.R. Heugens, EPS-2012-253-ORG, <http://repub.eur.nl/pub/30682>

Langhe, B. de, *Contingencies: Learning Numerical and Emotional Associations in an Uncertain World*, Promotor(s): Prof.dr.ir. B. Wierenga & Prof.dr. S.M.J. van Osselaer, EPS-2011-236-MKT, <http://repub.eur.nl/pub/23504>

Legault-Tremblay, P.O., *Corporate Governance During Market Transition: Heterogeneous responses to Institution Tensions in China*, Promotor: Prof.dr. B. Krug, EPS-2015-362-ORG, <http://repub.eur.nl/pub/78649>

Lenoir, A.S. *Are You Talking to Me? Addressing Consumers in a Globalised World*, Promotor(s) Prof.dr. S. Puntoni & Prof.dr. S.M.J. van Osselaer, EPS-2015-363-MKT, , <http://hdl.handle.net/1765/79036>

Leunissen, J.M., *All Apologies: On the Willingness of Perpetrators to Apologize*, Promotor(s): Prof.dr. D. de Cremer & Dr. M. van Dijke, EPS-2014-301-ORG, <http://repub.eur.nl/pub/50318>

Li, D., *Supply Chain Contracting for After-sales Service and Product Support*, Promotor(s): Prof.dr.ir. M.B.M. de Koster, EPS-2015-347-LIS, <http://repub.eur.nl/pub/78526>

Li, Z., *Irrationality: What, Why and How*, Promotor(s): Prof.dr. H. Bleichrodt, Prof.dr. P.P. Wakker, & Prof.dr. K.I.M. Rohde, EPS-2014-338-MKT, <http://repub.eur.nl/pub/77205>

Liang, Q.X., *Governance, CEO Identity, and Quality Provision of Farmer Cooperatives*, Promotor(s): Prof.dr. G.W.J. Hendrikse, EPS-2013-281-ORG, <http://repub.eur.nl/pub/39253>

Liket, K., *Why 'Doing Good' is not Good Enough: Essays on Social Impact Measurement*, Promotor(s): Prof.dr. H.R. Commandeur & Dr. K.E.H. Maas, EPS-2014-307-STR, <http://repub.eur.nl/pub/51130>

Loos, M.J.H.M. van der, *Molecular Genetics and Hormones: New Frontiers in Entrepreneurship Research*, Promotor(s): Prof.dr. A.R. Thurik, Prof.dr. P.J.F. Groenen, & Prof.dr. A. Hofman, EPS-2013-287-S&E, <http://repub.eur.nl/pub/40081>

Lovric, M., *Behavioral Finance and Agent-Based Artificial Markets*, Promotor(s): Prof.dr. J. Spronk & Prof.dr.ir. U. Kaymak, EPS-2011-229-F&A, <http://repub.eur.nl/pub/22814>

Lu, Y., *Data-Driven Decision Making in Auction Markets*, Promotor(s): Prof.dr.ir. H.W.G.M. van Heck & Prof.dr. W. Ketter, EPS-2014-314-LIS, <http://repub.eur.nl/pub/51543>

Manders, B., *Implementation and Impact of ISO 9001*, Promotor(s): Prof.dr. K. Blind, EPS-2014-337-LIS, <http://repub.eur.nl/pub/77412>

Markwat, T.D., *Extreme Dependence in Asset Markets Around the Globe*, Promotor(s): Prof.dr. D.J.C. van Dijk, EPS-2011-227-F&A, <http://repub.eur.nl/pub/22744>

Mees, H., *Changing Fortunes: How China's Boom Caused the Financial Crisis*, Promotor(s): Prof.dr. Ph.H.B.F. Franses, EPS-2012-266-MKT, <http://repub.eur.nl/pub/34930>

Mell, J.N., *Connecting Minds: On The Role of Metaknowledge in Knowledge Coordination*, Promotor: Prof.dr.D.L. van Knippenberg, EPS-2015-359-ORG, <http://hdl.handle.net/1765/78951>

Meuer, J., *Configurations of Inter-firm Relations in Management Innovation: A Study in China's Biopharmaceutical Industry*, Promotor(s): Prof.dr. B. Krug, EPS-2011-228-ORG, <http://repub.eur.nl/pub/22745>

Micheli, M.R., *Business Model Innovation: A Journey across Managers' Attention and Inter-Organizational Networks*, Promotor(s): Prof.dr. J.J.P. Jansen, EPS-2015-344-S&E, <http://repub.eur.nl/pub/78241>

Mihalache, O.R., *Stimulating Firm Innovativeness: Probing the Interrelations between Managerial and Organizational Determinants*, Promotor(s): Prof.dr. J.J.P. Jansen, Prof.dr.ing. F.A.J. van den Bosch, & Prof.dr. H.W. Volberda, EPS-2012-260-S&E, <http://repub.eur.nl/pub/32343>

Milea, V., *News Analytics for Financial Decision Support*, Promotor(s): Prof.dr.ir. U. Kaymak, EPS-2013-275-LIS, <http://repub.eur.nl/pub/38673>

Naumovska, I., *Socially Situated Financial Markets: A Neo-Behavioral Perspective on Firms, Investors and Practices*, Promotor(s): Prof.dr. P.P.M.A.R. Heugens & Prof.dr. A. de Jong, EPS-2014-319-S&E, <http://repub.eur.nl/pub/76084>

Nielsen, L.K., *Rolling Stock Rescheduling in Passenger Railways: Applications in short term planning and in disruption management*, Promotor(s): Prof.dr. L.G. Kroon, EPS-2011-224-LIS, <http://repub.eur.nl/pub/22444>

Nuijten, A.L.P., *Deaf Effect for Risk Warnings: A Causal Examination applied to Information Systems Projects*, Promotor(s): Prof.dr. G.J. van der Pijl, Prof.dr. H.R. Commandeur & Prof.dr. M. Keil, EPS-2012-263-S&E, <http://repub.eur.nl/pub/34928>

Osadchiy, S.E., *The Dynamics of Formal Organization: Essays on bureaucracy and formal rules*, Promotor(s): Prof.dr. P.P.M.A.R. Heugens, EPS-2011-231-ORG, <http://repub.eur.nl/pub/23250>

Ozdemir, M.N., *Project-level Governance, Monetary Incentives, and Performance in Strategic R&D Alliances*, Promotor(s): Prof.dr.ir. J.C.M. van den Ende, EPS-2011-235-LIS, <http://repub.eur.nl/pub/23550>

Peers, Y., *Econometric Advances in Diffusion Models*, Promotor(s): Prof.dr. Ph.H.B.F. Franses, EPS-2011-251-MKT, <http://repub.eur.nl/pub/30586>

Peters, M., *Machine Learning Algorithms for Smart Electricity Markets*, Promotor(s): Prof.dr. W. Ketter, EPS-2014-332-LIS, <http://repub.eur.nl/pub/77413>

Porck, J., *No Team is an Island: An Integrative View of Strategic Consensus between Groups*, Promotor(s): Prof.dr. P.J.F. Groenen & Prof.dr. D.L. van Knippenberg, EPS-2013-299-ORG, <http://repub.eur.nl/pub/50141>

Porrás Prado, M., *The Long and Short Side of Real Estate, Real Estate Stocks, and Equity*, Promotor(s): Prof.dr. M.J.C.M. Verbeek, EPS-2012-254-F&A, <http://repub.eur.nl/pub/30848>

Poruthiyil, P.V., *Steering Through: How organizations negotiate permanent uncertainty and unresolvable choices*, Promotor(s): Prof.dr. P.P.M.A.R. Heugens & Prof.dr. S.J. Magala, EPS-2011-245-ORG, <http://repub.eur.nl/pub/26392>

Pourakbar, M., *End-of-Life Inventory Decisions of Service Parts*, Promotor(s): Prof.dr.ir. R. Dekker, EPS-2011-249-LIS, <http://repub.eur.nl/pub/30584>

Pronker, E.S., *Innovation Paradox in Vaccine Target Selection*, Promotor(s): Prof.dr. H.J.H.M. Claassen & Prof.dr. H.R. Commandeur, EPS-2013-282-S&E, <http://repub.eur.nl/pub/39654>

Protzner, S. *Mind the gap between demand and supply: A behavioral perspective on demand forecasting*, Promotor(s): Prof.dr. S.L. van de Velde & Dr. L. Rook, EPS-2015-364-LIS, <http://repub.eur.nl/pub/79355>

Pruijssers, J.K., *An Organizational Perspective on Auditor Conduct*, Promotor(s): Prof.dr. J. van Oosterhout & Prof.dr. P.P.M.A.R. Heugens, EPS-2015-342-S&E, <http://repub.eur.nl/pub/78192>

Retel Helmrich, M.J., *Green Lot-Sizing*, Promotor(s): Prof.dr. A.P.M. Wagelmans, EPS-2013-291-LIS, <http://repub.eur.nl/pub/41330>

Rietveld, N., *Essays on the Intersection of Economics and Biology*, Promotor(s): Prof.dr. A.R. Thurik, Prof.dr. Ph.D. Koellinger, Prof.dr. P.J.F. Groenen, & Prof.dr. A. Hofman, EPS-2014-320-S&E, <http://repub.eur.nl/pub/76907>

Rijsenbilt, J.A., *CEO Narcissism: Measurement and Impact*, Promotor(s): Prof.dr. A.G.Z. Kemna & Prof.dr. H.R. Commandeur, EPS-2011-238-STR, <http://repub.eur.nl/pub/23554>

Rösch, D. *Market Efficiency and Liquidity*, Promotor: Prof.dr. M.A. van Dijk, EPS-2015-353-F&A, <http://hdl.handle.net/1765/79121>

Roza-van Vuren, M.W., *The Relationship between Offshoring Strategies and Firm Performance: Impact of innovation, absorptive capacity and firm size*, Promotor(s): Prof.dr. H.W. Volberda & Prof.dr.ing. F.A.J. van den Bosch, EPS-2011-214-STR, <http://repub.eur.nl/pub/22155>

Rubbaniy, G., *Investment Behaviour of Institutional Investors*, Promotor(s): Prof.dr. W.F.C. Verschoor, EPS-2013-284-F&A, <http://repub.eur.nl/pub/40068>

Schoonees, P. *Methods for Modelling Response Styles*, Promotor: Prof.dr. P.J.F. Groenen, EPS-2015-348-MKT, <http://repub.eur.nl/pub/79327>

Shahzad, K., *Credit Rating Agencies, Financial Regulations and the Capital Markets*, Promotor(s): Prof.dr. G.M.H. Mertens, EPS-2013-283-F&A, <http://repub.eur.nl/pub/39655>

Sousa, M.J.C. de, *Servant Leadership to the Test: New Perspectives and Insights*, Promotor(s): Prof.dr. D.L. van Knippenberg & Dr. D. van Dierendonck, EPS-2014-313-ORG, <http://repub.eur.nl/pub/51537>

Spriet, R., *Vehicle Routing with Uncertain Demand*, Promotor(s): Prof.dr.ir. R. Dekker, EPS-2013-293-LIS, <http://repub.eur.nl/pub/41513>

Staat, J.L., *Leading Public Housing Organisation in a Problematic Situation: A Critical Soft Systems Methodology Approach*, Promotor(s): Prof.dr. S.J. Magala, EPS-2014-308-ORG, <http://repub.eur.nl/pub/50712>

Stallen, M., *Social Context Effects on Decision-Making: A Neurobiological Approach*, Promotor(s): Prof.dr.ir. A. Smidts, EPS-2013-285-MKT, <http://repub.eur.nl/pub/39931>

Tarakci, M., *Behavioral Strategy: Strategic Consensus, Power and Networks*, Promotor(s): Prof.dr. D.L. van Knippenberg & Prof.dr. P.J.F. Groenen, EPS-2013-280-ORG, <http://repub.eur.nl/pub/39130>

Teixeira de Vasconcelos, M., *Agency Costs, Firm Value, and Corporate Investment*, Promotor(s): Prof.dr. P.G.J. Roosenboom, EPS-2012-265-F&A, <http://repub.eur.nl/pub/37265>

Troster, C., *Nationality Heterogeneity and Interpersonal Relationships at Work*, Promotor(s): Prof.dr. D.L. van Knippenberg, EPS-2011-233-ORG, <http://repub.eur.nl/pub/23298>

Tsekouras, D., *No Pain No Gain: The Beneficial Role of Consumer Effort in Decision-Making*, Promotor(s): Prof.dr.ir. B.G.C. Dellaert, EPS-2012-268-MKT, <http://repub.eur.nl/pub/37542>

Tuijl, E. van, *Upgrading across Organisational and Geographical Configurations*, Promotor(s): Prof.dr. L. van den Berg, EPS-2015-349-S&E, <http://repub.eur.nl/pub/78224>

Tuncdogan, A., *Decision Making and Behavioral Strategy: The Role of Regulatory Focus in Corporate Innovation Processes*, Promotor(s): Prof.dr.ing. F.A.J. van den Bosch, Prof.dr. H.W. Volberda, & Prof.dr. T.J.M. Mom, EPS-2014-334-S&E, <http://repub.eur.nl/pub/76978>

Uijl, S. den, *The Emergence of De-facto Standards*, Promotor(s): Prof.dr. K. Blind, EPS-2014-328-LIS, <http://repub.eur.nl/pub/77382>

Vagias, D., *Liquidity, Investors and International Capital Markets*, Promotor(s): Prof.dr. M.A. van Dijk, EPS-2013-294-F&A, <http://repub.eur.nl/pub/41511>

Veelenturf, L.P., *Disruption Management in Passenger Railways: Models for Timetable, Rolling Stock and Crew Rescheduling*, Promotor(s): Prof.dr. L.G. Kroon, EPS-2014-327-LIS, <http://repub.eur.nl/pub/77155>

Venus, M., *Demystifying Visionary Leadership: In search of the essence of effective vision communication*, Promotor(s): Prof.dr. D.L. van Knippenberg, EPS-2013-289-ORG, <http://repub.eur.nl/pub/40079>

Vermeer, W., *Propagation in Networks: The impact of information processing at the actor level on system-wide propagation dynamics*, Promotor: Prof.mr.dr. P.H.M. Vervest, EPS-2015-373-LIS, <http://repub.eur.nl/pub/79325>

Visser, V.A., *Leader Affect and Leadership Effectiveness: How leader affective displays influence follower outcomes*, Promotor(s): Prof.dr. D.L. van Knippenberg, EPS-2013-286-ORG, <http://repub.eur.nl/pub/40076>

Vlam, A.J., *Customer First? The Relationship between Advisors and Consumers of Financial Products*, Promotor(s): Prof.dr. Ph.H.B.F. Franses, EPS-2011-250-MKT, <http://repub.eur.nl/pub/30585>

Waltman, L., *Computational and Game-Theoretic Approaches for Modeling Bounded Rationality*, Promotor(s): Prof.dr.ir. R. Dekker & Prof.dr.ir. U. Kaymak, EPS-2011-248-LIS, <http://repub.eur.nl/pub/26564>

Wang, T., *Essays in Banking and Corporate Finance*, Promotor(s): Prof.dr. L. Norden & Prof.dr. P.G.J. Roosenboom, EPS-2015-352-F&A, <http://repub.eur.nl/pub/78301>

Wang, Y., *Information Content of Mutual Fund Portfolio Disclosure*, Promotor(s): Prof.dr. M.J.C.M. Verbeek, EPS-2011-242-F&A, <http://repub.eur.nl/pub/26066>

Wang, Y., *Corporate Reputation Management: Reaching Out to Financial Stakeholders*, Promotor(s): Prof.dr. C.B.M. van Riel, EPS-2013-271-ORG, <http://repub.eur.nl/pub/38675>

Weenen, T.C., *On the Origin and Development of the Medical Nutrition Industry*, Promotor(s): Prof.dr. H.R. Commandeur & Prof.dr. H.J.H.M. Claassen, EPS-2014-309-S&E, <http://repub.eur.nl/pub/51134>

Wolfswinkel, M., *Corporate Governance, Firm Risk and Shareholder Value*, Promotor(s): Prof.dr. A. de Jong, EPS-2013-277-F&A, <http://repub.eur.nl/pub/39127>

Yang, S., *Information Aggregation Efficiency of Prediction Markets*, Promotor(s): Prof.dr.ir. H.W.G.M. van Heck, EPS-2014-323-LIS, <http://repub.eur.nl/pub/77184>

Zaerpour, N., *Efficient Management of Compact Storage Systems*, Promotor(s): Prof.dr.ir. M.B.M. de Koster, EPS-2013-276-LIS, <http://repub.eur.nl/pub/38766>

Zhang, D., *Essays in Executive Compensation*, Promotor(s): Prof.dr. I. Dittmann, EPS-2012-261-F&A, <http://repub.eur.nl/pub/32344>

Zwan, P.W. van der, *The Entrepreneurial Process: An International Analysis of Entry and Exit*, Promotor(s): Prof.dr. A.R. Thurik & Prof.dr. P.J.F. Groenen, EPS-2011-234-ORG, <http://repub.eur.nl/pub/23422>

SUSTAINABLE ELECTRIC VEHICLE MANAGEMENT USING COORDINATED MACHINE LEARNING

The purpose of this dissertation is to investigate how intelligent algorithms can support electricity customers in their complex decisions within the electricity grid. In particular, we focus on how electric vehicle (EV) owners can be supported in their charging and discharging decisions, benefiting from the information available. We examine the problem from different standpoints and show the benefits for each involved stakeholder, dependent on the market conditions. In the first essay, we take the perspective of an individual EV owner and design an intelligent algorithm, which learning from her preferences, driving and consumption information, proposes optimized charging and discharging recommendations. In the second essay, we extend the first one by incorporating the EV within a smart home with a photovoltaic panel. The main goal of this study is to examine how accurate solar generation forecasting can be useful for charging the EV and make the best out of renewable sources. We propose a supervised learning algorithm which estimates the solar generation output from the weather conditions. In the third essay, we examine the problem from the grid operator's point of view, taking a top-down approach. We propose an auction mechanism which has as its main goal to service as many EV owners as possible, given a certain grid capacity. In the fourth essay, we propose a hybrid mechanism which combines benefits from top-down and bottom-up approaches. This mechanism is based on dynamic price functions that are able to incentivize EV customers to delay their charging duration when there is no urgency. Overall, this dissertation contributes to the academic literature with new algorithms that can leverage the power of data available and personalize EV charging recommendations. It also contributes to practice by providing useful insights to the involved stakeholders such as grid operators, energy utility companies, individual customers and automotive companies with respect to creating the right incentives for EV adoption. Finally, it adds to the very important discussion about sustainability, since it proposes ways to reduce carbon footprint and benefit the most from the available renewable sources.

ERiM

The Erasmus Research Institute of Management (ERIM) is the Research School (Onderzoekschool) in the field of management of the Erasmus University Rotterdam. The founding participants of ERIM are the Rotterdam School of Management (RSM), and the Erasmus School of Economics (ESE). ERIM was founded in 1999 and is officially accredited by the Royal Netherlands Academy of Arts and Sciences (KNAW). The research undertaken by ERIM is focused on the management of the firm in its environment, its intra- and interfirm relations, and its business processes in their interdependent connections.

The objective of ERIM is to carry out first rate research in management, and to offer an advanced doctoral programme in Research in Management. Within ERIM, over three hundred senior researchers and PhD candidates are active in the different research programmes. From a variety of academic backgrounds and expertises, the ERIM community is united in striving for excellence and working at the forefront of creating new business knowledge.

ERIM PhD Series Research in Management

Erasmus Research Institute of Management - ERiM
Rotterdam School of Management (RSM)
Erasmus School of Economics (ESE)
Erasmus University Rotterdam (EUR)
P.O. Box 1738, 3000 DR Rotterdam,
The Netherlands

Tel. +31 10 408 11 82
Fax +31 10 408 96 40
E-mail info@erim.eur.nl
Internet www.erim.eur.nl

