

MARKUS PETERS

Machine Learning Algorithms for Smart Electricity Markets

Essays on Autonomous Electricity Broker Design,
Probabilistic Preference Modeling, and Competitive
Benchmarking



Machine Learning Algorithms for Smart Electricity Markets

**Essays on autonomous electricity broker design, probabilistic
preference modeling, and competitive benchmarking**

Machine Learning Algorithms for Smart Electricity Markets

Essays on autonomous electricity broker design, probabilistic preference modeling,
and competitive benchmarking

Machine learning algoritmen voor slimme elektriciteitsmarkten
Essays over designs voor autonome elektriciteits brokers, preferentiemodellering,
en competitive benchmarking

Thesis

to obtain the degree of Doctor from the
Erasmus University Rotterdam
by the 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
Friday 13 February 2015 at 13:30 hours

by

Markus Peters

born in Neuss, Germany



Doctoral Committee

Promotor: Prof.dr. W. Ketter

Other members: Prof.dr.ir. H.W.G.M. van Heck
Prof.dr. T.M. Heskes
Dr. M. Saar-Tsechansky

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://hdl.handle.net/1765/1>

ERIM PhD Series in Research in Management, 332

ERIM reference number: EPS-2014-332-LIS

ISBN 978-90-5892-384-4

© 2014, Markus Peters

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, ISO 14001.

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.



Foreword

When I started my PhD trajectory in 2011, it was for the joy of contemplating questions that few have thought about before. What I had not expected was the additional joy that comes from exploring the hidden world of those who do science for a living. I have been fortunate to be guided by true experts in my explorations, and I gladly acknowledge their contributions to this dissertation.

I am grateful to my promotor Wolf Ketter for providing the vision upon which this work is built, for his enthusiasm in supervising it, and for becoming a good friend in the process. Wolf has a rare ability to see opportunities where others see risks, and he was among those who dared to believe in a thirty-three year old entering graduate student right from the start. His positive attitude is contagious, and his matchless ability to connect with people has opened doors for this work on more than one occasion. Wolf generates more ideas than a generation of doctoral students could possibly explore, but he also gives his students the liberty and the means to develop his ideas into something that excites them. Thank you!

Maytal Saar-Tsechansky was among those who guided my first ventures into academic research and publishing, and I was delighted about the opportunity to collaborate with her again on the final article of this dissertation. Maytal patiently helped me circumnavigate the pitfalls of scientific writing, and she never ran out of challenging questions to hone our arguments. The reader will find that Chapters 2 and 3 are significantly better for her contributions! I am particularly grateful for Maytal's willingness, without a moment's hesitation, to travel the long way from Austin to Rotterdam to serve on my doctoral committee.

John Collins co-authored two of the articles in this dissertation (Chapters 2 and 4), and he is one of the leaders of the Power TAC project that provides their greater context. John is an inexhaustible source of smart grid and software engineering knowledge, which he contributed to many interesting discussions throughout my time at Erasmus University. Remarkably, his recent retirement from the University of Minnesota has not lessened his zest for action in the least, and he continues to translate his vast experience into Power TAC innovations whenever he is not converting his home into the most self-sufficient, sustainable

patch of Wisconsin. It has been a pleasure to work with a scholar of John's expertise and dedication.

Perry Groot and Tom Heskes guided me through the technically most difficult parts of this work, and they are co-authors of the article in Chapter 3. I am thankful for Perry's efforts to keep me from despairing over the intricacies of efficient Bayesian inference. On more than one occasion, his uncanny ability to recall the appendix of just the right article has proven indispensable. I appreciate Tom's feedback on this work, and his willingness to serve on my doctoral committee.

Eric van Heck, Patrick Groenen, and Oliver Kramer have been generous in their willingness to serve as members of my doctoral committee. In his capacities as head of the Business Information Management (BIM) group at the Rotterdam School of Management, and as secretary of the inner doctoral committee, Eric has additionally provided me with valuable guidance on positioning this dissertation in the Information Systems landscape.

Alok Gupta kindly provided strategic guidance on the Competitive Benchmarking article in Chapter 4. In many ways, this technically inconspicuous chapter proved the most difficult to write, and I am thankful for his perspectives on it.

Many other colleagues have influenced my thinking about the subjects presented in this dissertation or have indirectly contributed to the presented results. In particular, I would like to thank the members of the Learning Agents Research Group (LARGE), including Micha Kahlen, Yixin Lu, Laurens Rook, Tommi Tervonen, Konstantina Valogianni, Jan van Dalen, and Gertjan van den Burg for interesting presentations and discussions on topics in Machine Learning, Learning Agents, and Decision Analysis.

The empirical work in Chapter 4 would not have been possible without the contributions of numerous Power TAC participants worldwide. In particular, I would like to thank the Power TAC development team at Erasmus University, Govert Buijs, Jeroen Ruigrok van der Werven, and Tijn Witsenburg. In addition to their work on Power TAC, these colleagues have operated the computing infrastructure on which all of the experiments in this dissertation were conducted.

The Department of Technology & Operations Management at the Rotterdam School of Management is a great place to work, and more colleagues are to be commended for this than can possibly be named in a foreword. I would like to extend my thanks to my fellow doctoral students and to the department's faculty for creating this uniquely positive atmosphere. In particular, I would like to thank Bas Giesbers, Morteza Pourakbar, Konstantina Valogianni, and Christina Wessels for being great roommates and neighbors, and Cheryl Blok-Eiting, Carmen Meesters-Mirasol, and Ingrid Waaijer for their tireless efforts in administering the department's affairs.

I am thankful for the comprehensive support I have received from the Erasmus Research Institute of Management (ERIM). Throughout my studies, I benefitted from a unique line of services ranging from course guidance, to conference funding, to grant writing, and to finally publishing this book. I especially appreciate the support of the colleagues who provide these services on a daily basis: Natalija Gersak, Kim Harte, Eliane Haseth, Miho Izuka, Tineke van der Vhee, Marisa van Iperen, and Veerle van Laere.

Finally, I am indebted to my family for their support. To my mom and dad for raising me curious, and for encouraging me to remain so; to Christian for volunteering to be my paranimf; and to Ariane for being my dearest scientist. Thank you!

Markus Peters
Rotterdam, August 2014

Contents

| | |
|---|-----------|
| Foreword | i |
| Contents | v |
| 1 Introduction | 1 |
| 1.1 Research Question and Main Contributions | 3 |
| 1.2 Smart Electricity Markets | 5 |
| 1.2.1 Transmission Systems | 6 |
| 1.2.2 Distribution Systems | 8 |
| 1.2.3 Challenges and Opportunities | 9 |
| 1.3 Machine Learning Algorithms | 11 |
| 1.3.1 Reinforcement Learning | 12 |
| 1.3.2 Supervised and Unsupervised Learning | 12 |
| 1.4 Outline | 14 |
| 2 Autonomous Decision-Making in Smart Electricity Markets | 17 |
| 2.1 Introduction | 17 |
| 2.2 Smart Electricity Market Simulation | 19 |
| 2.3 RL and Strategies for High-Dimensional State Spaces | 26 |
| 2.3.1 Reinforcement Learning | 28 |
| 2.3.2 Function Approximation | 29 |
| 2.3.3 Feature Selection and Regularization | 30 |
| 2.4 Learning Strategies | 32 |
| 2.4.1 SELF | 33 |
| 2.4.2 Reference Strategies | 35 |
| 2.5 Experimental Evaluation | 36 |
| 2.5.1 Manual broker construction | 36 |

| | | |
|----------|---|-----------|
| 2.5.2 | Feature Selection | 38 |
| 2.5.3 | Market Stability / Guarding against Overfitting | 41 |
| 2.5.4 | Regularization | 46 |
| 2.5.5 | Impact of customer characteristics | 46 |
| 2.5.6 | Summary of Results | 50 |
| 2.6 | Related Work | 51 |
| 2.7 | Conclusions and Future Work | 52 |
| 3 | A Preference Model for Autonomous Decision-Making | 55 |
| 3.1 | Introduction | 55 |
| 3.2 | Background and Related Work | 58 |
| 3.2.1 | Information Systems | 58 |
| 3.2.2 | Computer Science | 59 |
| 3.2.3 | Psychology | 60 |
| 3.2.4 | Marketing and Econometrics | 60 |
| 3.3 | Gaussian Process Trade-off Model (GTM) | 60 |
| 3.4 | Fast Bayesian Inference in GTM | 69 |
| 3.4.1 | Structured Gaussian Processes | 70 |
| 3.4.2 | Learning Trade-off Evaluations | 71 |
| 3.4.3 | Learning User Characteristics | 75 |
| 3.4.4 | Learning Hyperparameters | 77 |
| 3.5 | Empirical Evaluation | 78 |
| 3.5.1 | Datasets and Benchmark Methods | 78 |
| 3.5.2 | Model Scalability and Predictive Accuracy | 81 |
| 3.5.3 | Dimensionality Characteristics | 84 |
| 3.5.4 | Sparse Approximation Quality | 85 |
| 3.6 | Discussion and Conclusions | 88 |
| 4 | Competitive Benchmarking of Electricity Brokers | 91 |
| 4.1 | Introduction | 91 |
| 4.2 | Information Systems Research for Societal Challenges | 94 |
| 4.3 | Competitive Benchmarking | 99 |
| 4.3.1 | Competitive Benchmarking ALIGNMENT | 100 |
| 4.3.2 | Competitive Benchmarking PLATFORM | 103 |
| 4.3.3 | Competitive Benchmarking PROCESS | 104 |
| 4.3.4 | Interaction effects of ALIGNMENT, PLATFORM, and PROCESS | 109 |

| | | |
|----------|---|------------|
| 4.4 | The Power TAC Case: Competitive Benchmarking for Sustainable Energy Systems | 112 |
| 4.4.1 | Power TAC ALIGNMENT | 112 |
| 4.4.2 | The Power TAC PLATFORM | 114 |
| 4.4.3 | The Power TAC PROCESS | 115 |
| 4.4.4 | Case Study Evaluation | 117 |
| 4.4.5 | Interim Balance | 124 |
| 4.5 | Discussion | 124 |
| 4.5.1 | Related Work | 124 |
| 4.5.2 | Best Practices | 127 |
| 4.6 | Conclusions | 129 |
| 5 | Conclusions | 133 |
| 5.1 | Summary of Main Findings and Implications | 133 |
| 5.2 | Limitations and Future Research | 138 |
| | Summary | 141 |
| | Samenvatting | 143 |
| | Bibliography | 145 |
| | Appendices | 164 |

Chapter 1

Introduction

“Civilization advances by extending the number of important operations which we can perform without thinking about them.”

Alfred N. Whitehead, as quoted in Friedrich A. von Hayek, *The Use of Knowledge in Society*

Recent advances in information and communication technologies (ICT) have caused unprecedented changes in how firms and their customers transact business. Surely the most familiar face of these changes is the growing volume of business-to-consumer (B2C) transactions conducted electronically. Estimated at \$4.2 trillion for the G20 states in 2016, these transactions stand for a growing share of economic activities worldwide. For example, in the United Kingdom more than 12% of GDP will be earned through electronic B2C transactions by 2016 (Dean et al., 2012). Less visible, but arguably more incisive, are similar developments in the business-to-business (B2B) sphere, for which market studies in the United States estimate that it now has more than twice the size of its B2C counterpart (Hoar et al., 2012).

It is hardly surprising to find businesses harnessing straightforward advantages of electronic commerce like lower transaction costs. But at the same time that businesses are using ICT to convert and transform *individual* business models (Anderson, 2009), a fundamental shift is taking place in how economic activity is coordinated *overall*. In many domains, **smart markets** – an emerging class of markets based on computational intelligence – are now replacing existing hierarchical forms of coordination (Bichler et al., 2010, Kambil and van Heck, 2002, McCabe et al., 1991).

One important instance of this shift can be seen in the form of modern wholesale elec-

tricity markets, which have evolved out of a system of regional monopolistic electricity providers. In North America¹, vertically integrated monopolies dominated the generation, transport, and distribution of electricity until the late 1990s when the Federal Energy Regulatory Commission (FERC) mandated that every generator was to be given equal grid access. Previously, the network-based electricity business had been perceived and regulated as a natural monopoly in its entirety. But suddenly ICT advances made the separation of transmission networks from users feasible, and policymakers were keen on unleashing market forces on the electricity business to the benefit of consumers (Kirschen and Strbac, 2005). Most present-day wholesale electricity markets fall into the smart markets category. For example, their clearing process usually involves a computational procedure known as security-constrained unit commitment (Stoft, 2002), which ensures that power is produced by the lowest-cost generators while respecting complex technical side-constraints such as generators' ramping times and system-wide security margins.

But while wholesale electricity markets have revolutionized the coordination of economic activity between large generators, transmission system operators, and electricity distributors, other parts of the electricity business have remained largely unaffected. For instance, most American electricity is still generated in central power plants from fossil (68%) or nuclear (20%) fuels in response to inelastic consumer demand (US Energy Information Administration, 2013a). The average US power plant converts only one third of its primary fuel into usable electricity, while 6% of generated electricity is lost to the aging power lines that connect generators with consumers (US Department of Energy, 2003). Generation from renewable or decentralized sources could alleviate these problems, but they are difficult to integrate when consumers have little information about their electricity usage and few incentives to invest in smart appliances that adapt to the changing availability of wind and sun. Technologies for mitigating these issues – real-time metering, bidirectional communication, home automation, etc. – are now commercially available. But intelligently *coordinating* the intricate interplay between individual behaviors and tight physical constraints remains difficult in complex modern power systems with millions of self-interested participants.

Researchers have consequently advocated the implementation of **smart retail electricity markets** as a way of providing this coordination (Ketter et al., 2014) based on an information layer (sometimes called “Internet of Energy”) residing on top of the physical infrastructure. Smart markets make use of this information layer, and they leverage ICT innovations to provide consumers with additional information, and to economically incentivize behaviors that

¹Below, North American examples will be used to illustrate key issues in power systems. While there are a number of differences compared to, e.g., European power systems in terms of their physical structure, the structure of their supporting markets, and the strategic agenda that policy makers follow in developing them, these differences are inconsequential to the main argument, except where noted. See Coll-Mayor et al. (2007), for example, for a comparison between the European Union and the United States in this regard.

are aligned with overarching goals such as the integration of renewable sources. A critical part of this vision are retail electricity **brokers**², Information Systems (IS) artifacts that intermediate in the smart markets connecting retail customers with large-scale generators, either autonomously or in support of a human decision-maker. Brokers serve in many of the same capacities as current electricity retailers, but their IS-based nature allows them to provide participants with real-time information and fine-grained economic incentives that are currently unfeasible. For example, brokers could offer electric vehicle owners in certain areas temporarily reduced charging rates in exchange for the option to use their batteries as local buffer against solar production drops when cloud covers are erratic. Through such targeted use of information, brokers encourage more efficient use of existing infrastructures, and they enable behaviorally driven change (Watson et al., 2010).

What is currently still unclear, is how brokers are best designed and evaluated. Here, design refers to both, the design process (design as a verb) and the concrete design that follows from this process (design as a noun). This dissertation aims to contribute to answering this question. In the remainder of the introduction we proceed as follows. Section 1.1 contains the research question and a short summary of the main contributions. Section 1.2 provides an introduction to the electricity domain as background for the following chapters. Section 1.3 similarly introduces Machine Learning, the science of building computer programs that improve through experience (Mitchell, 1997), and the key reference discipline in our studies. And Section 1.4 provides an outline of the remainder of the dissertation.

1.1 Research Question and Main Contributions

The goal of this work is to advance the development of IS design theories for brokers, and its main research question is:

How should IS-based brokers for retail electricity markets be designed?

We aim to answer this question through a theory for design and action (IS theory type V; Gregor, 2006), and we follow a design science approach (Simon, 1996, Hevner et al., 2004, Gregor and Jones, 2007) to establish it. A theory for design and action is a prescriptive theory that informs the choices of artifact designers. As such, its principal components are constructs (e.g., brokers, markets, customers) and relationships between those constructs

²We use the shorter term brokers instead of retail electricity brokers henceforth. In the power systems literature, brokers are sometimes also referred to as aggregators, or load serving entities (LSEs).

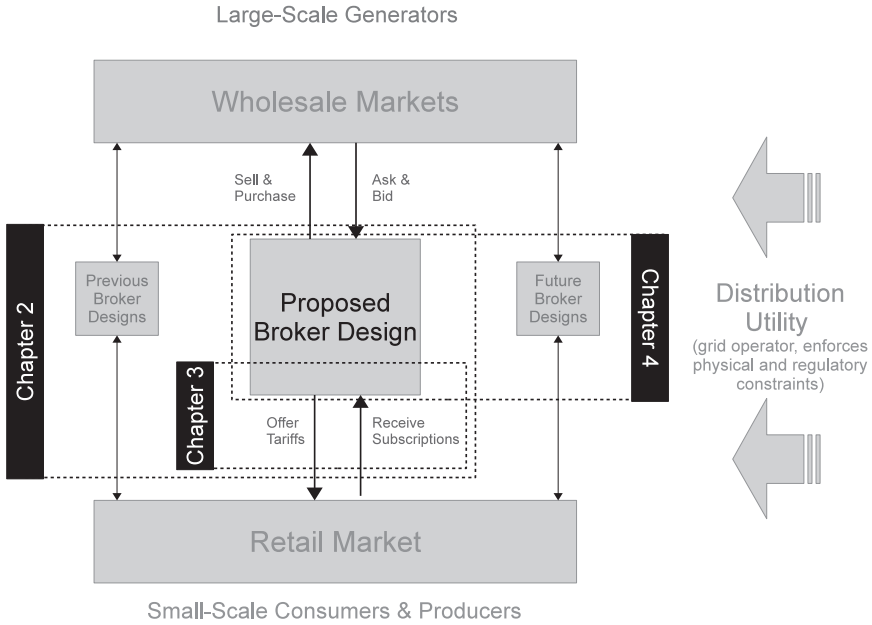


Figure 1.1. Areas of contribution for the three studies that comprise the dissertation

that are prescriptive in nature, backed by causal explanations, and empirically testable. Our contribution is divided into three separate studies (see also Figure 1.1):

Chapter 2 – Broker Design: We propose and evaluate a new design theory for brokers based on ideas from Machine Learning. The unit of analysis in this study is the broker itself, and we derive one particular design theory. Our work improves over existing designs, e.g., (Reddy and Veloso, 2011b) through a more realistic problem setup, and through improved performance in spite of the additional challenges we consider.

Chapter 3 – Preference Learning: We study in greater detail the broker’s core problem of learning from past customer choices as the basis for future decisions. The unit of analysis in this study is the broker itself, and we derive one particular design theory. Specifically, we propose and evaluate a probabilistic model that addresses important peculiarities of preference learning in electricity markets, and we demonstrate the performance of this model on electricity tariff choice tasks.

Chapter 4 – Competitive Benchmarking: We propose and study a novel research method for accelerating progress on future broker designs and similar complex IS design tasks where the underlying real-world phenomena evolve rapidly, and where the social cost

of failure is high. The unit of analysis in this study is the overall design task, and we use a case study approach to demonstrate the efficacy of Competitive Benchmarking for broker design. Our data furthermore allow us to provide preliminary insights into the societal consequences of deploying competing brokers in retail electricity markets.

Scientifically, this work contributes to the implementation of smart markets in energy retailing, and to the nascent IS research stream on Energy Informatics (Watson et al., 2010). Its managerial and societal relevance lie in its contributions to the design and evaluation of brokers. Brokers incentivize consumers to act within the grid's operational bounds, and thereby provide “an opportunity to create shared value – that is, a meaningful benefit for society that is also valuable to the business” (Porter and Kramer, 2006). This intelligent form of intermediation is one possible business model for utility companies in future Smart Grids.

1.2 Smart Electricity Markets

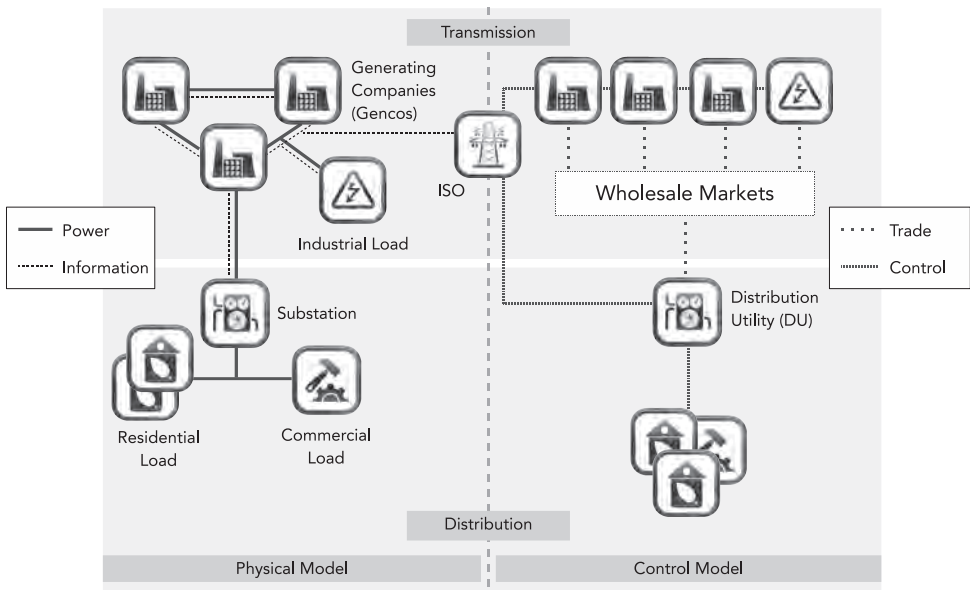


Figure 1.2. Physical and control structure of the electric grid (today). The top half of the model shows the high-voltage transmission level, the bottom half the low-to-intermediate-voltage distribution level.

To better understand the business environment that brokers operate in, let us consider the structure of current power systems and some associated challenges. The model in Figure 1.2 shows the physical and control structures resulting from the introduction of whole-

sale electricity markets.³ As is customary in power systems, the model is separated into a **transmission system** where large-scale generators like coal and nuclear power plants feed high-voltage electricity into a transmission network, and a **distribution system** responsible for regional electricity provisioning to commercial and residential end-customers. Some large industrial loads capable of operating at high voltage levels (e.g., blast furnaces) may be connected directly to transmission networks, but the bulk of loads in terms of volume and number are served through distribution systems. The two types of system vary widely in the number and sophistication of connected stakeholders, the sensing abilities built into the physical infrastructure, the current level of ICT deployment, and the dominant control paradigm. These differences are further characterized in Table 1.1.

1.2.1 Transmission Systems

Transmission systems are highly evolved in terms of sensory capabilities and ICT deployment. The Independent System Operators (ISO) running these systems and other participants such as large industrial loads and generating companies have significant incentives for investing in state-of-the-art sensing, optimization, and decision-support capabilities due to the tremendous costs and benefits behind their capital-intensive businesses (including the cost of outages). A key difference between electricity provisioning and other businesses is the presence of two separate layers of control that participants engage in: economic and physical control.

Economic control at the transmission level is provided through the aforementioned wholesale markets where generating companies sell energy to aggregated distribution-level loads or large industrial users of electricity. Transactions in these markets are tied to a specific delivery time and location, because electricity is currently not stored at scale. Ideally, forward markets would provide allocation ahead of time and spot markets would provide efficient real-time allocation of residual imbalances. But due to ramping and other physical constraints of power plants, generators can only sell their capacity forward, i.e., there is no conventional, competitive spot market for electricity. The upshot of all wholesale market transactions is a tentative schedule of production and consumption that does not yet account for real-time deviations from market participants' forecasts (e.g., plant and line failures),

³The model is an idealized representation and significant regional differences exist. For example, not all parts of the United States have introduced competitive wholesale markets, yet. The Independent System Operator (ISO) referenced in the model is a Regional Transmission Operator (RTO) in some parts of the US, and a Transmission System Operator (TSO) in most of Europe. In some markets, independent third parties take over parts of the ISO's responsibilities, as in the case of Scheduling Coordinators (SC) in California (Cameron and Cramton, 1999). And, at the time of this writing, approximately half of the American states have tentatively introduced competitive elements at the retail level. While these differences are important for understanding concrete power systems and their respective performances, they are inconsequential to our argument.

| Attribute | Transmission System | Distribution System |
|----------------------------------|---|---|
| Number of nodes and stakeholders | Low - Approximately <ul style="list-style-type: none"> • 10^2 large-scale generators • 10^3 transmission lines • 10^4 substations | High - Approximately <ul style="list-style-type: none"> • 5×10^4 distribution feeders • 10^7 customer meters • 5×10^8 appliances |
| Stakeholder sophistication | per system High - Commercial generators use sophisticated forecasting and optimization routines to compute optimal bids in wholesale markets; Independent System Operators (ISO) use advanced power flow analysis tools for operations and contingency planning | per system Low-to-Medium - Distribution Utilities (DU) use relatively simple forecasting schemes; Residential customers make ad-hoc consumption decisions and potentially sporadic, manual tariff decisions; Commercial customers may exhibit some sophistication in tariff negotiations and use of controllable capacities |
| Sensing ability | Medium-to-High - Mature commercial sensing devices are available and widely deployed; increasing deployment of Phasor Measurement Units (PMU) for wide-area monitoring | Low-to-Medium - Commercial customer on-site sensing devices available but questions remain with respect to standards, optimal level of functionality, etc.; DUs are gradually installing feeder sensors |
| ICT deployment | High - ISOs use sophisticated Energy Management Systems (EMS), Supervisory Control and Data Acquisition (SCADA) systems, and Phasor Data Concentrators (PDCs) | Low-to-Medium - Customer-site deployment efforts are still tentative; uncertainty with respect to standards, regulatory framework, required functionality, and business cases; DUs use Distribution Management Systems (DMS) but proprietary technology makes them difficult to extend to new applications |
| Control paradigm | Market-based / Direct - Long-term matching of supply and demand, as well as ancillary services such as regulation, load-following, and various types of reserves; Real-time operation starts from previous market transaction, applies direct control to handle real-time deviations and contingencies | Direct - Predominantly passive management to match supply to demand; some use of direct control for load shedding |

Table 1.1. Summary of current transmission and distribution system characteristics. Numeric examples for a typical U.S. system are taken from (Widgren et al., 2004)

physical network constraints, and some system-level considerations, such as reliability reserves or contingency planning.

Based on the preliminary wholesale market allocation, the ISO manages the real-time operation of the grid (for which markets are too slow) including all remaining system-level considerations (for which markets do not account) through **physical control**. Physical control is a highly involved, distributed process executing on multiple timescales that aims to optimize economic performance and power quality under physical constraints.⁴ Parts of this process are, in turn, based on market mechanisms. For example, the ISO's acquisition of reserves from generators is mediated through wholesale markets.

1.2.2 Distribution Systems

Distribution systems are comparatively less evolved due to their smaller volumes of energy supplied over wider areas, and due to the large number of unsophisticated participants connected to them. While equipment for real-time customer metering, and for sensing and automatically controlling distribution feeders is commercially available, incentives to invest in it remain low in many parts of the world. Residential and commercial customers have historically enjoyed an anytime supply of electricity and are wary of paying for new metering equipment that will take years to amortize, and that has the potential to complicate their lives. Distribution Utilities (DU), the monopolistic retailers and operators of distribution networks, are subject to strict economic regulation of their offerings and have little incentives for inducing behavioral changes among their customers, or to improve their service quality beyond mandated levels.

Consequently, most distribution networks follow a hierarchical control approach where customers subscribe to fixed or simple time-of-use tariffs, make consumption decisions independently of the availability of electricity, and receive monthly bills based on sporadic meter readings. Little use is made of economic control through retail markets, and physical control is at most exercised over larger commercial customers who allow the DU to remotely control parts of their loads (curtailment). The DU forecasts the aggregate consumption of its customers based on historical data and procures offsetting generation commitments in the wholesale markets. These markets then invoke the right number of generators through security-constrained unit commitment.

While this simple approach to distribution system operation has worked well for decades, it has lately been criticized for several reasons. First, shielding customers from true electric-

⁴Sub-activities include security-constrained economic dispatch, unit commitment, allocation and invocation of several types of reserves, governor control, voltage control, automatic generation control (AGC), and controlled load reduction (load shedding). See Kirschen and Strbac (2005) for further details.

ity prices leads to **inefficient allocations** (Watson et al., 2012). For example, customers have no incentives to shift their non-time-critical capacities to periods of high production from renewable sources (demand response). Second, capacity planning within the grid is based on peak demand, with the consequence that 10% to 18% of American power systems' capacity is currently utilized less than 1% of the time. Such **under-utilized capacity** is both environmentally harmful (e.g., excess line corridors) and costly (Kassakian and Schmalensee, 2011). Responsive customer demand has the potential to reduce peak demand by up to 20% until the end of the decade (Department of Energy, 2012). And third, unresponsive demand in combination with competitive wholesale markets leads to opportunities for exercising **market power** on the side of highly reactive marginal generators. It is therefore important that at least some customers can react to electricity scarcity in real time to avoid market failures like the California energy crisis (Borenstein, 2002, Cramton, 2003).

1.2.3 Challenges and Opportunities

From an IS vantage point, the differences between transmission and distribution systems are unsurprising. Transmission networks with their comparatively small number of sophisticated participants and extensive real-time sensory and control capabilities are naturally amenable to central optimization and dispatch. The methods required to perform such optimizations are highly involved, to be sure. But the problem of transmission systems operations is essentially one of optimizing power flows among a small number of strategizing participants under physical side-constraints and it remains well within the realm of central optimization.

In distribution systems, on the other hand, central optimization is impeded by a number of practical and conceptual issues:

- The DU as system operator has virtually no information about the dynamic electricity consumption of individual grid participants. Currently available data include sporadic meter readings and aggregate loads on distribution feeders, but the granularity of those data is insufficient for real-time control schemes. Much less than real-time *visibility* do distribution systems provide real-time *controllability* of customer capacities. Some commercial customers may have load curtailment arrangements in place, but this level of control remains elusive for the majority of electricity customers without dedicated data connections and smart appliances.
- Customers have few incentives to adapt their behavior. While it seems unlikely that customers would be willing to monitor, for example, electricity prices in real time, there is significant untapped potential for smart appliances that optimize their electric-

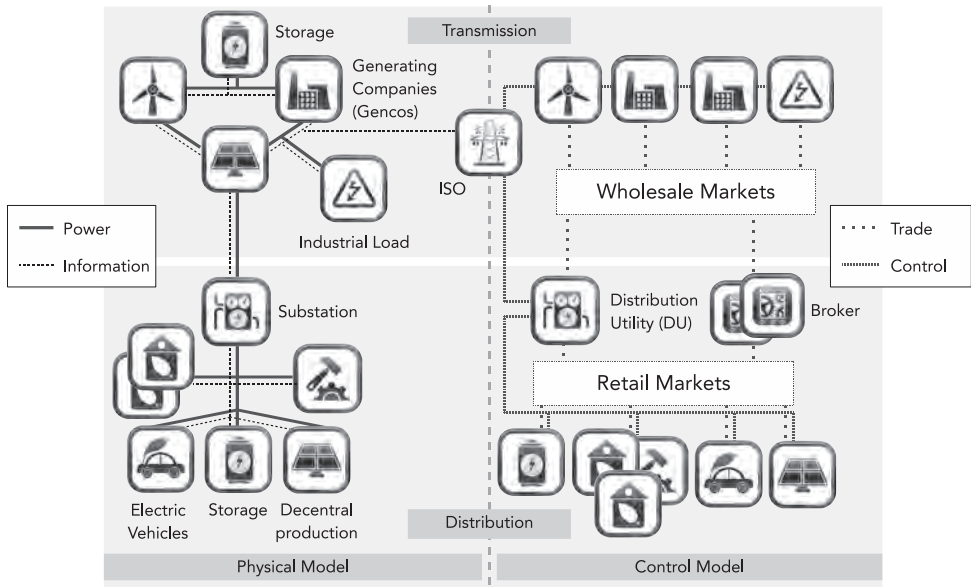


Figure 1.3. Physical and Control Structure of the Electric Grid (Smart Grid)

ity usage profile while respecting their users' preferences.⁵

- Even if more granular data *was* available, a centralized control scheme would likely not scale to the point of making optimal decisions involving tens of millions of customers and hundreds of millions of appliances in real time.

Alleviation for these issues could come in the form of **Smart Grids** (see Figure 1.3), with among other features (Massoud Amin and Wollenberg, 2005):

Sensing and Control Capabilities at the network level and among customers through the large-scale deployment of advanced metering capabilities.

Renewables and Decentralized Generation which could reduce the environmental impact of electricity production.

Electricity Storage which can increase asset utilization, match the non-dispatchable power flows from renewable sources with demand, and enable new grid topologies like Smart Energy Neighborhoods (Ibrahim et al., 2008, Sioshansi et al., 2009).

⁵A 2010 market research report by a commercial consulting firm concludes that “the U.S. smart appliance market will expand from \$1.42 billion in 2011 to \$5.46 billion in 2015, representing a nearly 40 percent growth rate. Clothes washers and dryers are expected to make up 36 percent of the market while refrigerators and freezers are forecast to comprise 24 percent of the market. Further, Whirlpool expects to make all appliances smart grid capable by 2015.” (EPRI - Electric Power Research Institute, 2011)

Retail Competition which provides the institutional foundation for new approaches to distribution system control. Retail competition decouples economic control (through tariffs or forward contracts between independent retailers and customers) from the real-time operation of the distribution grid, and it provides retailers with incentives for offering customers competitive, innovative electricity services.

These Smart Grid features can facilitate a market-based alternative to current distribution system control schemes by increasing visibility and controllability. For example, real-time sensing capabilities in feeders and households could be used for allocating infrastructure costs based on *peak* usage (the true cost driver), not *overall* usage (the information that is currently available). Decentralized generation could satisfy electricity demand with small ramping times and without the need for costly long-range transmission. And electricity storage in batteries or electric vehicles could provide highly responsive ancillary services that otherwise require the idle operation of large generators.

If there was no mention of transmission system innovations in the previous paragraph, it is because the greatest remaining opportunities for smart markets lie at the distribution level. Transmission systems hold a great number of important challenges, to be sure,⁶ but the notion of making systems smarter to make them better is inherently limited at the transmission level (Joskow, 2012). Transmission systems are a balancing mechanism of last resort, and coordination should be established locally before delegating it to the transmission level. Significant smartness and flexibility must therefore reside at the edges of the grid. This insight is mirrored by industry estimates that 70% of required Smart Grid investments will accrue at the distribution level (EPRI - Electric Power Research Institute, 2011). High-impact opportunities for smart retail electricity markets also arise from the closeness between distribution systems and customers. Given the relative ease of influencing a customer appliance compared to building new power plants, more effective use should be made of the human factor instead of treating it as a disturbance term, as is current practice in power systems engineering (Palensky and Dietrich, 2011).

⁶The 2003 Northeast Blackout, for example, was at least partially caused by a lack of human and technical communication between neighboring systems operators (Fox-Penner, 2005, Joskow, 2012). How precisely transmission-level Energy Management Systems (EMS) should be designed to optimally support users in complicated, real-time decision-making tasks therefore is an important question worthy of scientific study. Furthermore, “the impact of ... ICT systems on controllability and observability [of the Smart Transmission Grid] is poorly understood ... [and] the exact relationship between critical ICT components and the power system needs improved methods for analysis.” (Vanfretti et al., 2011)

1.3 Machine Learning Algorithms

A question that immediately follows from the preceding discussion is how the increased visibility and controllability afforded by Smart Grids can be turned into practical, new control mechanisms. The scale and complexity of electric grids cast doubts on the prospects for manual control, but IS-based mechanisms are equally unlikely to be effective if built from a limited set of fixed, predefined rules cast into computer algorithms. While it might be possible to formulate *goals* for the behavior of such algorithms, it is entirely unclear which of the myriad possible actions (sensing, changing prices, acquiring small-scale generation capacities, etc.) should be effected at what time to *achieve* those goals.

A promising alternative to predefining behaviors is the use of **Machine Learning**, a collection of techniques for the “[construction of] computer programs that automatically improve with experience.” (Mitchell, 1997) By observing activities in their environment, often labeled as more or less desirable, these techniques learn to recognize and react to regularities without an explicit set of rules fixed at design time. Machine learning techniques have been applied successfully in complex, dynamic settings ranging from autonomous game-playing to self-driving cars (Kober et al., 2013).

A comprehensive introduction to Machine Learning is beyond the scope of this introduction. Instead, we give an outline of the major task categories, and refer the reader to Bishop (2006) and Murphy (2012) for further details.

1.3.1 Reinforcement Learning

Reinforcement Learning (RL) deals with control tasks where the goal is to maximize long-term rewards by repeatedly taking actions in a stochastic environment (Sutton and Barto, 1998). A Reinforcement Learning algorithm observes the state of the environment, chooses an action, and is subsequently rewarded or penalized. Then the environment progresses to the next state and the process repeats itself. In general, both the rewards and the state transitions made by the system can be stochastic, which means the algorithm has to learn about them. In particular, reward maximization could entail foregoing immediate rewards to instead explore the environment and to obtain higher rewards in the future.

Our main research question asks for a design theory for brokers in an environment with stochastic reactions from customers and other wholesale market participants. In Chapter 2, we will show how the RL framework can be applied to this challenge.

1.3.2 Supervised and Unsupervised Learning

In **Supervised Learning**, pairs of data $D = \{(\mathbf{x}_i, y_i)\}, i = 1, \dots, N$ are given, where the \mathbf{x}_i are the values of one or more input variables, and the y_i are the values of an associated output variable. Often, the \mathbf{x}_i are real-valued feature vectors and the y_i are categorical (classification) or continuous (regression). But in general, both \mathbf{x}_i and y_i can be vectors or more richly structured objects. The goal in supervised learning is to discover the mapping between the \mathbf{x}_i and the y_i so that the outputs for previously unobserved inputs $\{\mathbf{x}_{new}\}$ can be predicted accurately. **Unsupervised Learning** differs from supervised learning in that only unlabeled data $D = \{\mathbf{x}_i\}$ are given, and the goal is to characterize these data, or to find patterns of interest in them. It is therefore also referred to as **pattern recognition** or **knowledge discovery** (Webb, 2003, Bishop, 2006).

While little use is made of unsupervised learning in this dissertation, Chapter 3 presents a novel method for learning customer preferences in a supervised fashion. In the study, a learner observes customer choices and generalizes from these observations to improve future decisions involving customer preferences.

1.4 Outline

We now describe the setup we use to answer the main research question “How should IS-based brokers for retail electricity markets be designed?” (see also Figure 1.1)

Chapter 2 – Broker Design⁷

Our first step is to develop and evaluate a high-level design for brokers in retail electricity environments. An important insight from autonomous systems design is that endowing brokers with all possible successful behaviors at design time is futile, because these behaviors are impossible to enumerate and will quickly become obsolete. Instead, we aim to design brokers that *learn* effective behaviors by observing their environment, taking actions, monitoring the long-run consequences that these actions entail, and updating their behavior accordingly. This is essentially a problem in optimal control, and Reinforcement Learning (RL, see Section 1.3.1) offers a suitable framework for addressing it.

In retail electricity markets, brokers publish electricity tariffs to customers, who subscribe to them based on latent preferences. Different tariff terms incentivize different kinds of customer behaviors, and the resulting consumption and production add up to a portfolio of obligations that the broker seeks to cover by trading in the wholesale market. Both markets are competitive, and any net imbalances remaining after these trades are handled by the operator of the distribution grid, the Distribution Utility, at typically higher costs to the broker.

While RL has previously been used to learn electricity wholesale trading strategies (Nanduri and Das, 2007, Ramavajjala and Elkan, 2012), retail electricity trading has received considerably less attention. Reddy and Veloso (2011b) were the first to note that RL offers an appropriate broker design framework. The design we propose and evaluate in Chapter 2 improves over their work in that it requires significantly less manual tuning, better accommodates the rich data available in smart market environments, and better incorporates new types of information as market conditions change. We identify the most effective design conditional on prevailing market conditions. Because of the high dimensionality of a broker’s observations, we show how feature selection and regularization techniques can be leveraged for better performance under these conditions.

Our simulation-based evaluations of these designs use SEMS, a data-driven Smart Electricity Market Simulation, developed specifically for this study. SEMS is built on data from

⁷This work was published as: Markus Peters, Wolfgang Ketter, Maytal Saar-Tsechansky, and John E. Collins, A reinforcement learning approach to autonomous decision-making in smart electricity markets. *Machine Learning*, 92:539, 2013. Preliminary versions of this work appeared in, or were presented at: Peters (2012), Peters et al. (2012a,b,c,d).

the Ontario wholesale market, a complete micro-level model of appliance usage in private households (Gottwalt et al., 2011), and several benchmark designs for brokers proposed in the literature (Reddy and Veloso, 2011b).

Chapter 3 – Preference Learning⁸

A particular challenge for brokers are the idiosyncratic preferences (Lichtenstein and Slovic, 2006) of their (prospective) customers. To incentivize favorable behaviors, brokers must be capable of reasoning about customer responses to changes in tariff terms. But the prior experiences they can draw from to this end are limited, because each current customer makes only few choices, even though the sum of choices observed across all customers is large. Furthermore, the observed choices are usually subject to behavioral inconsistencies, such as inertia.

A broker must respect these inconsistencies, because the uncertainty arising from them is a crucial ingredient for autonomous decision-making: brokers should only make high-value decisions autonomously if past evidence suggests that they will be correct with high probability, and prompt their users for additional information otherwise. However, brokers cannot demand their human operators' attention too frequently and a preference model must therefore begin to make accurate predictions from limited training data.

To address these requirements, we propose a Bayesian preference model based on Gaussian processes (Rasmussen and Williams, 2006) in Chapter 3 that learns from limited data by pooling the data of similar users. Our model quantifies the certainty of its predictions as input to the broker's autonomous decision-making task, and it infers probabilistic user segments based on observed choices in the process. Probabilistic inference in nonparametric Bayesian models is often computationally expensive, but by combining properties of the broker's task with advances in sparse (Quinero-Candela et al., 2007) and structured (Saatchi, 2011) Gaussian processes we are able to reduce the costs of inference substantially. We evaluate our model on several real-world choice datasets used in an earlier study (Houlsby et al., 2012), and on electricity tariff choice data that we collected specifically for the purposes of this study on a commercial crowdsourcing platform.

⁸This work has been submitted for publication as: Markus Peters, Perry Groot, Wolfgang Ketter, Maytal Saatchi, Tom Heskes: A Scalable Preference Model for Autonomous Decision-Making Involving Consumer Choices. Preliminary versions of this work appeared in, or were presented at: Peters and Ketter (2012, 2013a,b,c, 2014).

Chapter 4 – Competitive Benchmarking⁹

In the last study, we shift our focus from *one particular* broker design to a method for accelerating the overall progress on broker design theories made by a community of competing researchers. With this study we aim to address the growing concern that challenges like sustainable electricity provisioning progress quicker than researchers' ability to counteract them (Hey et al., 2009). Design science studies like the first two studies of this dissertation are well-suited for initially identifying and studying promising designs, but their homogeneous setup limits their ability to quickly detect promising alternatives and possible social negatives. Agent-Based Modeling (ABM; Axelrod, 2006) and Agent-Based Computational Economics (ACE; Tesfatsion 2006) have previously been used to shed light on social negatives, but real-world brokers will inevitably evolve under competitive pressure, and neither of these methods account for this competitive co-evolution.

We propose to address this limitation through Competitive Benchmarking (CB), a novel method for IS artifact design in high-complexity environments. CB is rooted in the competitive research approach pioneered by the Trading Agents community (Ketter and Symeonidis, 2012) which challenges researchers to devise software agents for complex, uncertain environments like supply chains (Arunachalam and Sadeh, 2005) and advertisement auctions (Jordan and Wellman, 2010).

The first instantiation of CB as we define it here is the Power Trading Agent Competition (Power TAC, Ketter et al. 2014), in which more than a dozen research groups from four different continents now jointly devise, benchmark, and improve broker designs. The Power TAC platform models a competitive retail power market in a medium-sized city, in which consumers and small-scale producers may choose from among a set of alternative electricity providers, represented by competing brokers. These brokers are built by individual research groups with expertise in Artificial Intelligence, Electrical Engineering, Information Systems, Machine Learning, and other areas, and their heterogeneous design approaches have contributed to a rich repository of design ideas and executable broker artifacts.

Cornerstones of Power TAC's CB process are annual championships, and pilots that provide additional informal benchmarking opportunities. To date, pilots have been held at IJCAI 2011 in Barcelona, at AAMAS 2012 in Valencia, and at IEEE SG-TEP 2012 in Nuremberg. The first two official Power TAC championship were held at AAAI 2013 in Bellevue, WA, and at AAMAS 2014 in Paris.¹⁰ Using fine-grained records of the Power

⁹This work has been submitted for publication as: Wolfgang Ketter, Markus Peters, John E. Collins, Alok Gupta, Competitive Benchmarking: An Information Systems Research Method for Societal Challenges. A related research note is under review as: Wolfgang Ketter, Markus Peters, John E. Collins, Alok Gupta, Power TAC: Competitive IS Research on Sustainable Electricity Systems. Preliminary versions of this work appeared in, or were presented at: Ketter et al. (2013c), Peters et al. (2013a,b).

¹⁰AAAI = Conference of the Association for the Advancement of Artificial Intelligence; AAMAS = International

TAC community's results, we quantify performance differences between alternative broker designs, and between subsequent iterations of the same designs to give preliminary empirical evidence of CB's efficacy as a research method.

Chapter 5 – Conclusions

The final chapter of this dissertation concludes with a discussion of the impact and limitations of the three studies, and with directions for future work.

Chapter 2

A Reinforcement Learning Approach to Autonomous Decision-Making in Smart Electricity Markets

2.1 Introduction

Liberalization efforts in electricity markets and the advent of decentralized power generation technologies are challenging the traditional ways of producing, distributing, and consuming electricity. The Smart Grid “aims to address these challenges by intelligently integrating the actions of all users connected to it . . . to efficiently deliver sustainable, economic and secure electricity supplies” (ETPSG, 2010). This ambitious vision requires substantial advances in intelligent decentralized control mechanisms that increase economic efficiency, while keeping the physical properties of the network within tight permissible bounds (Werbos, 2009).

A fundamental objective of the Smart Grid is to maintain a tight balance of supply and demand in real-time. Presently, the task of balancing the output of large-scale power plants with customer demand is handled via centralized control mechanisms. The increasing penetration of small-scale production from renewable sources like solar and wind, however, introduces inherently intermittent, variable, and geographically dispersed supply, and renders real-time balancing significantly more challenging. In addition, proposals for Demand-side Management (DSM) and for tariffs with time-of-use or dynamic pricing complicate the prediction of

consumption patterns. Existing centralized control mechanisms are unable to accommodate this combination of intermittent and variable supply, a grid of staggering scale including vast numbers of small-scale producers, and dynamic changes in demand in response to price variations.

A promising approach to effective balancing in the Smart Grid is the introduction of **electricity brokers** (Ketter et al., 2012b), intermediaries between retail customers and large-scale producers of electricity. Brokers offer a distributed alternative to the centralized system of today's grid, facilitate localized markets that reduce inefficiencies from wide-area transmission, and attain socially desirable market outcomes in response to appropriate economic incentives. Because brokers serve as intermediaries, they must also trade in multiple interrelated (e.g., retail and wholesale) markets simultaneously – a structure that Bichler et al. (2010) refer to as *Smart Markets*. Smart Markets constitute a novel class of complex, fast-paced, data-intensive markets, in which participants ought to employ (semi-)autonomous trading agents in order to attain good trading results.

It is imperative that the design of an electricity broker agent can adapt to a wide variety of market structures and conditions. This is because there is considerable variability in the structure that a future Smart Electricity Market might have, and also because such flexibility is generally beneficial for high performance in dynamic environments. We present several important innovations beyond the class of autonomous electricity brokers for retail electricity trading that we presented in (Peters et al., 2012a). Our brokers can accommodate arbitrary economic signals from their environments, and they learn efficiently over the large state spaces resulting from these signals. Existing approaches (Reddy and Veloso, 2011a,b) are limited in the state space size they can accommodate, and are thus constrained in terms of the economic environments they can be deployed into. These works have also not considered customers' variable daily load profiles (instead, assuming fixed consumption), or the broker's wholesale trading – both core challenges for real-world electricity brokers. Our design alleviates these limitations.

The research we report here extends our previous work (Peters et al., 2012a) by exploring alternatives for the data-driven identification of particularly informative signals from the broker's data-rich Smart Electricity Market environment. We explore the role that feature selection and regularization techniques play in the broker's adaptation process. Specifically, we explore the benefits of two different feature selection procedures based on Genetic Algorithms and greedy forward selection, and compare them to L1-regularized online learning techniques over the full state space. We find that the inexpensive regularization approach yields satisfactory results under some market conditions; however, the more extensive feature selection techniques can be highly effective across different market regimes (Ketter

et al., 2012a) if adequate precautions are taken against environmental overfitting. Based on our empirical results, in this paper we also provide guidance on how such overfitting can be alleviated, as well as discuss approaches which are specialized to the Smart Grid challenge we study here. Our empirical evaluations are based on real-world electricity market data from the Ontario Wholesale Market and a revised model of individual customers' consumption decisions. The customer model we employ captures intra-day variability in demand and has been shown to yield realistic aggregate load curves, rendering our empirical results significantly more meaningful as compared to earlier studies, including our own work (Peters et al., 2012a). Our empirical results demonstrate that our broker design is highly effective and that it consistently outperforms prior approaches despite the additional challenges we consider.

More generally, research on autonomous electricity brokers for the Smart Grid constitutes a nascent, emerging field, in which most of the challenges are largely unexplored. Improving our understanding of methods that address these challenges has far-reaching implications to society at large. For example, our broker design contributes to current research on economic mechanism design for the Smart Grid by providing effective strategies against which such mechanisms can be validated, e.g. (de Weerd et al., 2011). Our extensive evaluation of feature selection techniques raises new and interesting questions about connections between overfitting and market stability in the presence of autonomous trading strategies. We also offer an example of how Machine Learning research can inform important developments in the future Smart Grid in Section 2.5.5. In addition to the development of a novel broker agent design, important objectives of this paper are to contribute to our understanding of key design decisions that enable broker agents to operate effectively in the Smart Grid, and to inform future work of challenges and promising research directions.

The paper is organized as follows. In Section 2.2 we give an overview of our Smart Electricity Market Simulation (SEMS). Section 2.3 describes foundations in Reinforcement Learning, feature selection, and regularization that our approach builds on. Section 2.4 introduces SELF, our class of Smart Electricity Market Learners with Function Approximation. A thorough empirical evaluation of our learners in comparison to strategies proposed in the literature follows in Section 2.5. In Section 2.6 we review relevant literature. Finally, we conclude with directions for future research.

2.2 Smart Electricity Market Simulation

Smart Electricity Markets aim to intelligently integrate the actions of customers, generating companies, and the Distribution Utility, cf. Figure 2.1. We developed SEMS, a data-driven

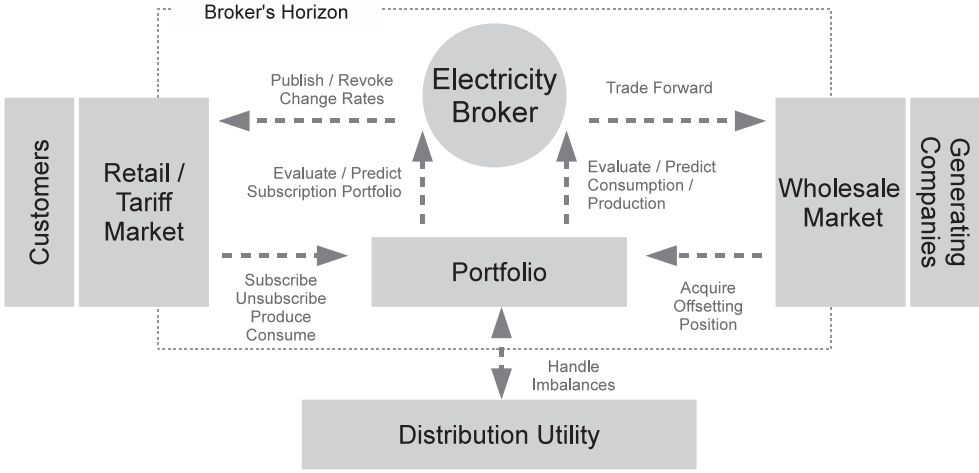


Figure 2.1. Smart Electricity Market structure

Smart Electricity Market Simulation, based on wholesale prices from a real-world electricity market¹ and a complete, micro-level model of appliance usage in private households. An important property of our simulation, with implications for the broker we design to operate in this environment, is to relax the assumption in previous work that consumers exhibit fixed demand (Reddy and Veloso, 2011a,b). Fixed demand simplifies the broker's task, however the resultant brokers may not offer an adequate response to the realities of electricity markets. In particular, a key challenge for real-world brokers is to effectively deal with *patterns* in consumer demand. This is important for effective grid balancing, as some patterns, such as high consumption during midday peak-hours, are significantly more costly for the broker to offset in the wholesale market (cf. Figure 2.2).

Below we outline the key elements of a Smart Grid, along with the models that represent them in our simulation.

- **Customers** $C = \{C_j\}$ are small-to-medium-size consumers or producers of electricity, such as private households and small firms. Each C_j denotes a group of one or more customers with similar characteristics and a joint, aggregate consumption profile. Customers buy and sell electricity through the *tariff market*, where they subscribe to standardized tariff offerings, including fixed-rate, time-of-use (ToU), and variable-rate tariffs. We describe our customer model in more detail below. Presently, only a small proportion of electricity is produced decentrally, and central production will

¹For this study we used data from Ontario's Independent System Operator, <http://www.ieso.ca>, which has also been used in a related study (Reddy and Veloso, 2011b).

continue to play a significant role in the near future. As a liberal upper bound consider that, of the 592 TWh of electricity produced in Germany in 2009, merely 75 TWh were produced decentrally under the country's Renewable Energy Act (12.6%) (Eurostat, 2011). Accordingly, the customers in our model act exclusively as consumers of electricity.

- **Generating Companies (GenCos)** are large-scale producers of energy, such as operators of fossil-fueled power plants and wind parks. GenCos are wholesalers of electricity production commitments. Because changes in power plant production levels have significant lead times, wholesale electricity is traded *forward* from several hours to several months in advance.
- The **Distribution Utility (DU)** is responsible for operating the electric grid in real-time. In particular, the DU manages imbalances between the total energy consumption and the total outstanding production commitments at any given time. To this end, the DU provides or absorbs energy on short notice and charges the responsible broker imbalance penalties. In SEMS, the DU charges balancing fees that are roughly twice as high as the long-term average cost of electricity in the wholesale market, and thus provides brokers a strong incentive to build easily predictable portfolios of subscribers.
- **Electricity Brokers $B = \{B_i\}$** are profit-seeking intermediaries, trading for their own account.² Brokers are retailers of electricity in the tariff market, and they offset the net consumption of their tariff subscribers by acquiring production commitments in either the tariff (small-scale producers) or wholesale market (GenCos). The **portfolio** of contractual arrangements that brokers obtain in this way is executed in real-time by the DU. Brokers aim to compile a portfolio of high-volume, high-margin tariff subscriptions with predictable consumption patterns, that can be offset with production commitments at a low cost. In SEMS, brokers publish one fixed-rate tariff at any given time. This design reflects the fact that fixed rates are currently still the dominant tariff model, mainly due to the absence of advanced metering capabilities among electricity customers. We are interested in the performance of methods for autonomous *retail electricity trading*. To this end, we endow both, our own strategies and our benchmark strategies, with a fixed wholesale trading algorithm based on Neural Network load forecasting, and brokers learn to develop a profitable retail trading strategy against this backdrop. Our choice of Neural Networks is mainly due to their good out-of-the-box performance in timeseries forecasting tasks. Alternatives, e.g. based on ARIMA

²Electricity Brokers are sometimes also referred to as *Load Serving Entities (LSEs)* or *Aggregators* in the electricity market literature.

models (Conejo et al., 2005), exist but we do not consider them further as they would impact the performance of all examined strategies in the same way.

The **SEMS Simulation Environment** is responsible for coordinating brokers, customers, and the DU. It manages the tariff market, and provides a wholesale market based on actual market data from Ontario's Independent System Operator. The wholesale market in SEMS determines prices by randomly selecting a window of sufficient size for the simulation run from almost ten years of real-world wholesale market pricing data. Figure 2.2 shows the long-term daily price distribution as well as daily price curves for 10 randomly selected days from that dataset. Once these prices have been determined, broker orders have no impact on them. Modeling brokers as price-takers is reflective of liberalized retail electricity markets, where an increasing number of small brokers compete against each other. For 2008, for example, the European Commission reported close to 940 non-main electricity retailers in Germany that shared 50% of the German market (Eurostat, 2011).




Figure 2.2. (a) Price distribution (mean \pm one standard deviation) for 10 years of price data from the Ontario wholesale market; (b) Price curves for 10 randomly selected sample days

A SEMS simulation runs over N timeslots $1, \dots, n, \dots, N$ which are structured as described in Figure 2.3.

1. Each broker B_i receives information about its current customers $C_n(B_i)$, the history of

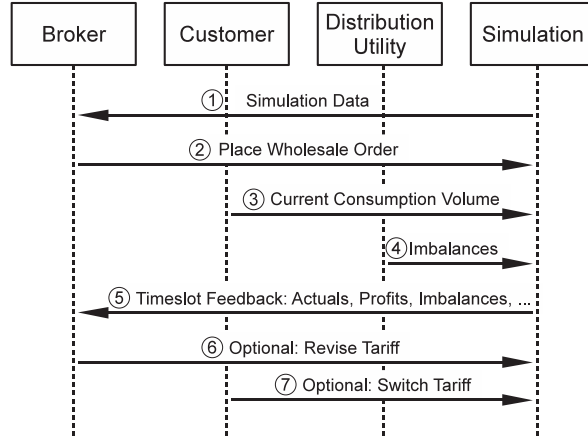


Figure 2.3. Sequence diagram for one simulation timeslot

wholesale prices W_1, \dots, W_{n-1} , the tariffs offered by all brokers at the end of the last timeslot $\mathbf{T}_{n-1} = \{\tau_{B_1}, \dots, \tau_{B_{|B|}}\}$, and its current cash account balance.³

2. Each broker indicates the volume of energy \hat{V}_n^c that it wishes to procure in the current timeslot. Note, that the broker has no previous knowledge of its customers' actual consumption nor of the wholesale prices for the current timeslot. There is no acquisition uncertainty; the indicated volume \hat{V}_n^c is always filled by the simulation.
3. Each customer C_j decides the volume of electricity $V_n^c(C_j)$ to consume given its current tariff, and announces this volume to the simulation. The volume consumed, $V_n^c(C_j)$, is derived from the corresponding customer's consumption model, which we describe below.
4. Based on the consumption decisions of its customers, its current tariff, and its acquisition in the wholesale market, each broker's cash account is credited (debited) with a trading profit (loss) $\tau^c(V_n^c) - \hat{V}_n^c \cdot W_n$, where $\tau^c(V_n^c)$ denotes the cost of consuming V_n^c under the current tariff τ^c to the customers (i.e., the revenue of the broker), and $\hat{V}_n^c \cdot W_n$ denotes the cost of procuring \hat{V}_n^c units of energy at the prevailing wholesale price W_n . Any imbalance between the broker's forecast, and the actual amount of energy consumed by its customers is made up for by the Distribution Utility. An imbalance penalty of I per unit of mismatch, or $|V_n^c - \hat{V}_n^c| \cdot I$ in total, is debited from the cash account of the broker for this service.
5. Each broker receives ex-post information on the actual aggregate consumption volume

³We summarize the mathematical notation used here and below in Table 2.1.

of its customers in the current timeslot V_n^c , its trading profit, its imbalance penalty, and its cash account balance at the end of the timeslot.

6. Each broker is queried if it wishes to change its offered tariff. A fixed amount reflecting administrative costs on the side of the broker is charged for each tariff update.
7. Each customer is queried if it wishes to subscribe to a different tariff.

Customers in SEMS are represented by a customer model, each instance of which represents the aggregate behavior of a group of customers. The customer model consists of a **consumption model**, which computes the amount of energy consumed in a given timeslot, and a **tariff evaluator**, which defines how customers select a tariff from a set of offered tariffs. Separating the consumption decision from the tariff selection decision in this way is economically well-motivated. In the short run, the electricity demand of private households is unresponsive to changes in price level. There is some empirical evidence for customers' willingness to *shift* electricity consumption over the day in response to changing electricity prices, e.g., (Herter et al., 2007). However, this phenomenon does not apply to our scenario of a fixed-rate tariff.

The **consumption model** in SEMS employs a micro-level simulation of electric appliance usage in private households based on the work of Gottwalt et al. (2011). This model incorporates statistical data on household types, household sizes, appliance saturation, seasonality on multiple timescales, vacations, etc. to replicate consumption decisions of a real-world customer population, and has been shown to yield realistic aggregate load curves. Our model makes use of all of these features, except for long-term seasonal effects which have no significant impact on the broker's primary challenges of short-term demand forecasting and balancing. Figure 2.4 shows weekly consumption profiles generated by our consumption model for populations of 10 and 1000 households, respectively. The profiles reflect the characteristic consumption peaks exhibited by private households around noon and during the early evening hours, seasonality effects between weekdays and weekends, as well as the typical consumption averaging behavior in large customer populations (Figure 2.4, Panel (b)) with significant reductions in the aggregate consumption's noise level.

Our **tariff evaluator** is based on current insights about customers' tariff selection behavior and works as follows:⁴ If the tariff that a customer is currently subscribed to is still

⁴Standard models of electricity tariff selection behavior are currently still an open question in Behavioral Economics and Smart Grid research. Our model captures two central characteristics of customer choice that are thought to lead to sub-optimal tariff choices by human decision-makers: *Inertia* or *switching probability* refers to customers' tendency to remain in their current tariff even if better alternatives surface, e.g. (Nicolaisen et al., 2001), and *customer irrationality* refers to sub-optimality resulting from a range of behavioral factors, e.g. (Wilson and Price, 2010).

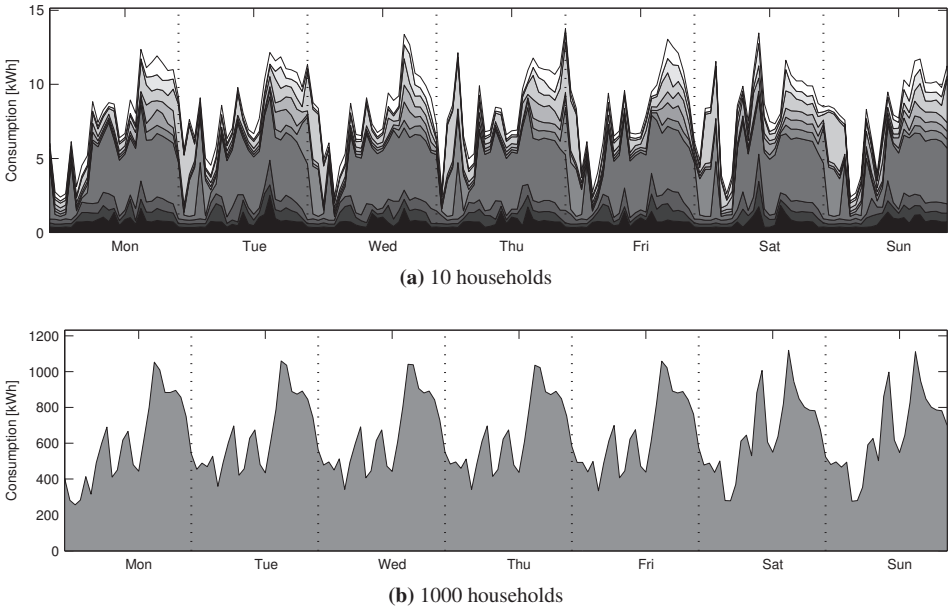


Figure 2.4. Simulated load curves from the SEMS customer model. Panel (a) shows the individual consumption profiles of 10 households, Panel (b) illustrates how the averaging effect leads to a smoother aggregate consumption pattern in a population of 1000 households.

available, the customer considers selecting a new tariff with a fixed probability q . With probability $1 - q$ it remains in its current tariff without considering any other offers. This behavior captures customers' *inertia* in selecting and switching to new tariffs. If the tariff that the customer is currently subscribed to is not available any longer, the customer selects a new tariff with probability 1. To select a new tariff, the customer ranks all tariffs according to their fixed rates; ties are broken randomly. A perfectly informed and rational customer would simply select the lowest-rate tariff from this ranking, because the lowest-rate tariff minimizes the expected future cost of electricity. In reality, however, customer decisions will tend to deviate from this theoretical optimum for reasons that include (1) customers do not possess perfect information about all tariffs, either because it is unavailable to them, or because they eschew the effort of comparing large numbers of tariffs; and (2) they make decisions based on non-price criteria such as trust and network effects that are absent from our model. We capture these deviations from a simple price rank-order using a Boltzmann distribution.

Assume a customer has to decide among a total of $|\mathbf{T}|$ tariffs. Then the probability of




Figure 2.5. CDF for the Boltzmann distribution. The parametrized discrete distribution is used to model departures from rationality in tariff selection.

selecting the r -th best tariffs is:

$$Pr(\text{Rank} = r) = \frac{e^{-r/\tau}}{\sum_{i=1}^{|\mathbf{T}|} e^{-i/\tau}}$$

Here, τ is the so-called *temperature* parameter with $\tau \in (0, \infty)$. The temperature can be interpreted as the customers' *degree of irrationality* relative to the theoretically optimal tariff decision. Consider the Cumulative Distribution Functions (CDF) depicted in Figure 2.5 for different values of τ . For $\tau \rightarrow 0$, only the best-ranked tariff has considerable mass, i.e., the tariff decision is perfectly rational. For $\tau \rightarrow \infty$, the distribution approaches a discrete uniform distribution, i.e., customers select their tariff at random.

2.3 RL and Strategies for High-Dimensional State Spaces

To operate effectively in the Smart Electricity Market outlined in Section 2.2, an electricity broker agent ought to learn from its environment in multiple ways. In particular, it must learn about potential customers and their behavior in terms of tariff selection and electricity consumption. A broker should also learn the behavior of its competitors, and derive tariff pricing policies that strike a balance between competitiveness and profitability. Furthermore, because a broker also acts in the wholesale market, it must learn ways to match its tariff

market actions with wholesale trading strategies in order to maximize its profit. Note, that the broker's only means of learning is its ability to act in the markets it trades in, and to observe the (long-term) consequences that its actions entail.

2.3.1 Reinforcement Learning

Reinforcement Learning (RL) offers a suitable framework to address the challenges faced by a broker acting in environments with unknown dynamics, and with the objective to collect the highest net present value over all current and future rewards. This can entail foregoing some immediate rewards for higher rewards in the future (Sutton and Barto, 1998). More formally, the Reinforcement Learning task we consider here is defined as a finite **Markov Decision Process** (MDP) with observable states and a known, fixed set of actions: $MDP = \langle \mathbf{S}, \mathbf{A}, P, R \rangle$ where \mathbf{S} denotes a finite set of states, \mathbf{A} denotes a finite set of actions, and P and R define the transition probability function and immediate reward function as follows:

$$P_a(s, s') = Pr(s_{n+1} = s' | s_n = s, a_n = a)$$

that is, $P_a(s, s')$ gives the probability of the environment choosing s' as the following state when s is the current state and the learner chooses action a . And

$$R_a(s, s') = E(r_{n+1} | s_n = s, a_n = a, s_{n+1} = s')$$

that is, $R_a(s, s')$ denotes the expected immediate reward received from the environment when choosing action a in state s and being sent to state s' by the environment thereafter. The solution to such an MDP is the **optimal policy** π^* that maximizes the net present value of all current and future expected immediate rewards, i.e.,

$$\pi^* = \arg \max_{\pi} \sum_{n=0}^{\infty} \gamma^n R_{a_n=\pi(s_n)}(s_n, s_{n+1})$$

where the learner follows the policy π that gives, for each state s , a corresponding action $a = \pi(s)$ to pursue. $0 \leq \gamma < 1$ denotes the discount parameter where smaller values of γ lead to greater emphasis on current rewards.

Many algorithms have been proposed for finding good policies (Szepesvári, 2010). For our agent, we use SARSA: a Temporal Difference (TD) algorithm, that is designed for online control problems, such as our retail electricity trading task. The algorithm starts out with some initial model of an **action-value function** $Q(s, a)$, which captures the learner's estimate of the net present value of being in state s , choosing action a next, and following the policy

| Symbol | Definition |
|--|---|
| $\mathbf{A} = \{a_i\}$ | Constant set of actions available to the SELF reinforcement learner |
| α_{max}, α' | Initial learning rate and decay of learning rate (1.0 = linear, 0.5 = square root, etc.) |
| $\mathbf{B} = \{B_i\}$ | Set of all competing brokers |
| $\mathbf{C} = \{C_j\}$ | Set of all retail customers, customers of broker B_i at time n |
| $\mathbf{C}_n(B_i)$ | |
| δ | Temporal difference in $Q(s, a)$ between subsequent observations |
| $\vec{e}(\lambda), \lambda$ | Eligibility trace vector, $\vec{e}(\lambda)$ has the same dimensionality as $\vec{\mathbf{F}}(s, a)$, $\vec{\theta}$ and measures the eligibility of $\vec{\theta}$'s elements for temporal difference updates based on recent observations; $0 \leq \lambda \leq 1$ determines the degree of recency, with greater values leading to a longer memory of observations |
| $\epsilon_{max}, \epsilon'$ | Initial exploration rate and decay of exploration rate (1.0 = linear, 0.5 = square root, etc.) |
| $\vec{\mathbf{F}}(s, a), \vec{\theta}$ | Vector of features of the state-action pair (s, a) and their weights in a linear action-value function, respectively |
| γ | MDP discount parameter |
| μ | Markup parameter of the benchmark strategies TableRL and Fixed |
| π, π^* | (Optimal) policy of a given MDP |
| $\Phi \in \{0, 1\}^n$ | Vector indicating the features actually employed out of the set of all available features |
| $\Psi \in R \subseteq \mathbb{R}^m$ | Vector of learning parameters such as α_{max}, λ , etc. |
| q | Customer switching probability; probability that a customer model considers a new tariff in any given timeslot |
| $Q(s, a)$ | Action-value of state s given action a |
| $r_n = R_a(s, s')$ | Immediate reward earned in timeslot n when the current state is s , the learner takes action a , and is sent to state s' by the environment |
| $\mathbf{S} = \{s_i\}$ | Discrete set of all possible states of the environment |
| \mathbf{T} | Set of all tariffs offered in the tariff market |
| τ | Customer irrationality $\tau \in [0; \infty)$, where greater values represent less rational or less informed tariff selection behavior |
| $P = P_a(s, s')$ | Transition probability of the environment moving the learner to state s' when in s and choosing action a |
| $V_n^c, V_n^c(C_j)$ | Actual net electricity consumption in time n and net consumption of customer C_j , respectively; from the perspective of one broker |
| \hat{V}_n^c | Estimate of one broker for its customers' electricity consumption at time n |
| W_n | Actual wholesale market price of electricity at time n |

Table 2.1. Summary of mathematical notation

implied by Q thereafter. The learner acts (approximately, except for occasional exploration) greedily with respect to the policy implied by Q , and updates Q with the true feedback it receives from the environment in each timeslot according to

$$Q(s, a) \leftarrow Q(s, a) + \alpha \underbrace{[r_{n+1} + \gamma Q(s_{n+1}, a_{n+1}) - Q(s_n, a_n)]}_{\text{temporal difference}} \quad (2.1)$$

where α denotes the learning rate. With probability ϵ , the learner chooses explorative actions instead of the greedy action implied by $Q(s, a)$ to investigate the value of other state-action pairs. In our experiments below we let α and ϵ decay over time to obtain stronger learning and more aggressive exploration towards the beginning of the simulation.⁵ In general, SARSA only converges to an exact estimate of Q when each state-action pair is visited an infinite number of times, and when the policy followed by the learner converges to a fixed policy. In our empirical evaluation we show that our learner performs well in spite of not fully meeting these theoretical requirements.

A key challenge of using RL for the problem we address here is the definition of an effective state space. Because it is not well understood which environmental features are useful for capturing changes in the action-value, it is beneficial to employ a wide array of features so as to avoid the exclusion of particularly relevant ones. However, even with a limited number of features, the state space quickly becomes too large to hold in memory. Furthermore, when the state space is large, the extent of exploration required for the learner to arrive at a reliable estimate of the action values $Q(s, a)$ for each $a \in \mathbf{A}$ becomes prohibitive. Previous work has dealt with this challenge by introducing *derived features* that combine multiple environmental features into a single feature for the learner (Reddy and Veloso, 2011a,b). However, these derived features are inherently less informative for learning, and there is no principled approach to constructing them. We address these challenges through a two-pronged strategy: (1) We employ **function approximation** to enable the learner to deal with potentially large state spaces; and (2) we explore the performance of **feature selection** and **regularization** techniques that reduce the (effective) state space size.

2.3.2 Function Approximation

Function approximation refers to a parametrized, functional representation of $Q(s, a)$ that allows the broker to explore the effectiveness of strategies over a wider array of potentially relevant states (see, e.g., (Rummery and Niranjan, 1994) for one of the earliest uses of func-

⁵We use α_{\max} and ϵ_{\max} to denote maximal rates, and α' and ϵ' to denote the degree of the parameter decay monomial, where $\alpha', \epsilon' > 0$ and a value of 1.0 stands for linear decay, 0.5 for square root decay, 2.0 for quadratic decay, and so forth.

tion approximation in online RL). The most common type of function approximation uses the representation

$$Q(s, a) = \vec{\theta} \vec{F}(s, a)^T$$

where $Q(s, a)$ is linear in $\vec{F}(s, a)$, a vector of selected *features* of the current state s given an action a . The reinforcement learner continually updates the weights in $\vec{\theta}$ to make Q more representative of the experiences gathered from the environment. With linear function approximation this gradient descent update of $\vec{\theta}$ takes the particularly simple form

$$\vec{\theta} \leftarrow \vec{\theta} + \alpha \delta \vec{e}(\lambda) \quad (2.2)$$

where α again denotes the learning rate, δ denotes the temporal difference (equivalent to the update term in Equation 2.1), and \vec{e} is the so-called *eligibility trace* which captures the weights eligible for a learning update based on the recently visited state-action pairs. The degree of recency is determined through the parameter $0 \leq \lambda \leq 1$ with $\lambda = 0$ representing updates only based on current observations, whereas greater values of λ introduce increasing degrees of memorization into the update process. Note, that the features in $\vec{F}(s, a)$ can themselves be nonlinear functions of features from the environment. Other types of function approximation have also been proposed instead of this linear scheme, e.g. (Busoniu et al., 2010, Pyeatt and Howe, 2001).

2.3.3 Feature Selection and Regularization

To improve our broker's learning performance in its information-rich Smart Market environment, we complement function approximation with a principled reduction in (effective) state space size. To this end, we explore different feature selection techniques as well as regularization, and examine their performance and resulting trade-offs in our setting. It is important to note that an MDP's state space size must be fixed and that the reduction referred to above is achieved in two fundamentally different ways: The *feature selection* techniques we study select a subset of relevant features offline, *before* the MDP is constructed. Once the relevant features are selected, the learning process proceeds on a fixed MDP with a state space that is reduced as compared to the space over all original candidate features. Regularization, on the other hand, aims to reduce the dimensions of the state space that are *effectively used*. That is, while the MDP is constructed using a large fixed state space, a regularized learner will likely assign zero weights to many of these features.

While both approaches have been studied extensively in supervised learning, e.g. (Guyon and Elisseeff, 2003), they have only recently been considered in RL (Loth et al., 2007, Painter-Wakefield and Parr, 2012, Parr et al., 2008, Petrik et al., 2010), and it is impor-

tant to understand their contributions to the effectiveness of the broker's actions and any implications for future work.

Feature selection here refers to methods that select informative projections $\vec{F}'(s, a)$ of the complete feature vector $\vec{F}(s, a)$ as basis for learning. Formally, let $\vec{F}(s, a)$ be a vector of n candidate features of the current state-action pair, and Ψ a vector of m learning parameters. Then

$$\mathbf{B}_{LinFA} = \{B_{LinFA}(\phi_1, \dots, \phi_n, \psi_1, \dots, \psi_m) \mid \Phi \in \{0, 1\}^n, \Psi \in R \subseteq \mathbb{R}^m\}$$

is a class of linear function approximation based RL brokers that use the feature $(\vec{F}(s, a))_i$ as part of their state space iff the indicator $\phi_i = 1$. Note, that we combine the feature selection task (i.e., finding good values of Φ) and the related parameter learning task (i.e., finding good values for Ψ), because they can be conveniently tackled simultaneously during the heuristic optimization described below.

We evaluate how well a particular broker $B \in \mathbf{B}_{LinFA}$ competes in a given environment by the **fitness function** $F : B \mapsto [0, 1]$ which measures the empirical average profit share that B captures in a given number of sample simulations. This procedure is also known as the **wrapper approach** to feature selection (Blum and Langley, 1997). The best broker B^* for the given environment is then $B(\arg \max_{\Phi, \Psi} F(B(\Phi, \Psi)))$.⁶

Optimizing F with respect to B is, in general, intractable due to the lack of structure in F and the size of its domain $\{0, 1\}^n \times R$. To alleviate this challenge, we employ one of two types of heuristic optimization:

- **Greedy Feature Selection / Hill Climbing:** Starting with a cohort $C^1 = \{B_1^1, \dots, B_n^1\}$ of all possible single-feature brokers, we determine their fitness values $F(B_1^1), \dots, F(B_n^1)$ and select the broker B^{1*} with the maximal F value. We then construct the next cohort C^2 by augmenting B^{1*} with each possible second feature and evaluate $F(B_1^2), \dots, F(B_{n-1}^2)$ for the $n - 1$ resulting brokers. We repeat the process until no further improvement in F values is achieved between cohorts, or until a predefined time limit is reached. We select the feature set Φ^* of the overall F -maximizing broker B^* . This process is commonly known as *forward selection* (Guyon and Elisseeff, 2003).

To select the broker's parameters we use a hill-climbing procedure where we first draw p parameter vectors Ψ_1, \dots, Ψ_p and corresponding gradients $\nabla \Psi_1, \dots, \nabla \Psi_p$ at random. We then construct the first cohort as $C^1 = \{B(\Phi^*, \Psi_1), \dots, B(\Phi^*, \Psi_p)\}$ and sub-

⁶Note, that a profit-maximizing fitness function does not preclude other social desiderata such as fairness, efficiency, or sustainability considerations from our Smart Market. By properly setting the market's *economic mechanisms*, market designers can create incentive structures that lead self-interested, profit-maximizing brokers to jointly aim towards socially desirable outcomes. Because we are primarily interested in the performance of autonomous retail electricity trading strategies themselves, we consider these economic mechanisms to be given and refer the reader to, e.g., (Dash et al., 2003, Parkes, 2007) for further details.

sequently develop each broker along the predetermined gradient until no further improvement is possible. For example $B(\Phi^*, \Psi_1)$ from C^1 is developed into $B(\Phi^*, \Psi_1 + \nabla \Psi_1)$ in C^2 , $B(\Phi^*, \Psi_1 + 2\nabla \Psi_1)$ in C^3 , and so forth. We finally select the learning parameters Ψ^* of the overall F -maximizing broker.

- **Genetic Algorithms:** Genetic Algorithms are a well-suited heuristic optimization procedure for our problem given their good performance over large binary domains, such as that of our feature indicators (De Jong, 1988). Starting with a randomly initialized cohort $C^1 = \{B_1^1, \dots, B_q^1\}$ we apply a Genetic Algorithm with Mutation and Crossover operators defined as usual, and a small number of elite individuals that is carried over from one cohort to the next (Liepins and Hilliard, 1989). As with the greedy procedure outlined above, the Genetic Algorithm runs until no further improvements in fitness value take place or a predefined time limit is reached. Note that the Genetic Algorithm modifies features and parameter values of the individuals concurrently, whereas our greedy approach selects features and parameters in two separate tasks.

Regularization, in contrast to feature selection, shrinks or penalizes the weights in $\vec{\theta}$ so as to obtain sparse inner products $\vec{\theta} \vec{F}(s, a)^T$. The resulting approximations are less prone to overfitting the peculiarities in action-values because strong, repeated evidence is required for a particular weight to become and remain non-zero. Regularized function approximations are also quicker to evaluate due to their inherent sparsity. One of the key advantages of regularization over feature selection is its natural integration into online learning processes which obviates the need for a separate offline learning phase.

Despite its seeming appeal, little theoretical groundwork has so far been done on the use of regularization in online RL. One of few exceptions is the work by Painter-Wakefield et al. (2012) who extend the regularized batch RL algorithm LARS-TD (Kolter and Ng, 2009) into L1TD, an L1-regularized online RL algorithm. L1TD adds the shrinkage operation

$$\vec{\theta} \leftarrow \text{sgn}(\vec{\theta}) \odot \max\{|\vec{\theta}| - v, 0\}$$

(with all operators defined component-wise) to the gradient descent update from Equation 2.2. The shrinkage operation effectively moves each component of $\vec{\theta}$ towards zero by v on each update, and the combined procedure can be shown to yield equivalent results to the L1-regularized regression formulation in the batch case (Painter-Wakefield and Parr, 2012).

2.4 Learning Strategies

In this section, we introduce SELF, our class of **S**mart **E**lectricity **M**arket **L**earners with **F**unction Approximation. A thorough empirical evaluation of our learners in comparison to strategies proposed in the literature follows in Section 2.5.

2.4.1 SELF

Our candidate strategy SELF is a class of SARSA reinforcement learners with linear function approximation. The state set of each SELF instance reflects selected aspects of its observable economic environment (e.g., its own tariff rate, competitors' tariff rates, market competitiveness indicators, etc.), and its action set contains possible actions in the tariff market. The learning objective is to find a policy π that approximately maximizes the learner's long-term reward in a Smart Electricity Market environment, while competing against other, both learning and non-learning, strategies.

As outlined in Section 2.3.1, one of the key challenges in our Smart Electricity Market setting is the definition of an effective state space for the learner to learn over. We address this challenging problem by defining a large set of candidate features that captures as much environmental detail as possible, and then applying feature selection and regularization techniques to identify a suitable subset of features that benefit learning. Table 2.2 shows a grid of features (vertical) and related encodings (horizontal), and shaded cells mark the feature/encoding pairs that are available to the SELF learner for learning.⁷ This example list of candidate features provides a good coverage of all economic information available from the environment. It is important to note that because a primary goal of our design is to substitute laborious, manual state space construction with principled optimization techniques, our methods can accommodate arbitrary additions to this feature set.

Another important element of an electricity broker design are the actions it can take in the retail market. Generally, a broker can either (a) set a new rate on its tariff, or (b) maintain its existing rate. The canonical model for this action set is a continuous or a discretized set of plausible target rates. However, our evaluations revealed that simultaneously learning the variability in the wholesale price-level *and* the variability among its competitors in this way overburdens the learner. To facilitate learning, we propose a set of economically meaningful actions for the broker to choose from. In particular, SELF brokers can choose among the discrete action set shown in Table 2.3, which is normalized relative to the prevailing wholesale market price level: A SELF broker can set its tariffs *relative* to other tariffs in the market. In

⁷In the table, *Plain* denotes the unencoded feature, *RBF* and *RBF(T)* denote Radial Basis Function encoding (optionally with thresholding) (Sutton and Barto, 1998), and *Bin* denotes a sign binary encoding which, given a real value x , transforms $x \mapsto (\mathbf{I}(\text{sgn}(x) = -1), \mathbf{I}(\text{sgn}(x) = 0), \mathbf{I}(\text{sgn}(x) = +1)) \in \{0, 1\}^3$.

| Feature / Encoding | Plain | RBF | RBF(T) | Bin | Description |
|------------------------|-------|-----|--------|-----|--|
| Bias | | | | | Constant 1.0 |
| ActionIndex | | | | | Index of the selected action |
| ActionOneInK | | | | | One-In-K representation of the selected action |
| BetterConsumptionRates | | | | | Number of better (lower) rates in the tariff market |
| CashGradient | | | | | Change in cash account balance over the last 48 hours |
| CustomerGradient | | | | | Change in number of customers over the last 48 hours |
| MarketBreadth | | | | | Range from lowest to highest rate in the tariff market |
| MarketShare | | | | | Percentage of all customers subscribed to SELF |
| MarkupLeader | | | | | Relative margin, as percentage of smoothed wholesale price, between SELF and the cheapest tariff in the market |
| NumberCustomers | | | | | Number of subscribed customers |
| RateChangeIndicator | | | | | 1 if selected action would result in a rate change, 0 otherwise |
| TargetMargin | | | | | Margin over smoothed wholesale price after performing a given action |
| WholesalePrice | | | | | Smoothed electricity wholesale price |
| WorseConsumptionRates | | | | | Number of worse (higher) rates in the tariff market |

Table 2.2. Candidate features for SELF state-action spaces

| Action | Margin over Wholesale Price |
|---------------|---|
| MarginLeader | Slightly lower than cheapest competitor |
| MarginAvg | Average of all competitors |
| MarginTrailer | Slightly higher than most expensive competitor |
| LowMargin | Constant 10% margin |
| HighMargin | Constant 20% margin |
| NoOp | Keep the current <i>tariff rate</i> . Could lead to changes in margin if wholesale prices change. |

Table 2.3. Available actions for SELF instances

doing so, the broker can choose among attacking its competitors (MarginLeader), positioning itself in the middle of the market (MarginAvg), or avoiding competition altogether by posting the most expensive tariff (MarginTrailer). Alternatively, rather than setting its tariffs relative to the market, the broker can set its tariffs in an *absolute* fashion, choosing between LowMargin and HighMargin, irrespective of the competing tariffs in the market. We chose the margins in Table 2.3 for their good observed performance in our experiments. The broker may also leave its current tariff unchanged (NoOp).

2.4.2 Reference Strategies

We evaluated SELF against the learning and non-learning strategies proposed in (Reddy and Veloso, 2011b). To address the need for a limited state space, the reference learning strategy uses derived features, referred to as *PriceRangeStatus* and *PortfolioStatus*. Importantly, the simulation model for which this strategy was evaluated did not include an explicit representation of a wholesale market, represented consumers demand as fixed throughout, and the brokers' only sources of electricity production commitments were small-scale producers. Brokers offer one *producer tariff* in addition to the consumer tariff used by the brokers in our study. These differences make some of the published results for this strategy difficult to interpret in the context of the market settings we consider here.⁸

The relevant benchmark strategies for evaluating our SELF Electricity Broker Agent are

- **Learning** a table-based reinforcement learner operating over the reduced, manually constructed state space outlined above. For clarity, we henceforth refer to the Learning strategy as **TableRL**.

⁸To incorporate these strategies in our simulation setting we used wholesale prices for producer prices, and suppressed actions pertaining to small-scale producer tariffs. We also excluded the *PortfolioStatus* feature, which is not meaningful for learning the TableRL strategy in our simulation model.

- **Fixed** a strategy which charges a constant markup μ over the smoothed wholesale price
- **Greedy** an adaptive strategy which charges either the highest rate in the market or an average rate, depending on the current PriceRangeStatus PRS_n . PRS_n is defined to be *Rational* if the difference between consumption and production rates in the market is at least μ (i.e., if the market charges a reasonable markup). In this case, the strategy opportunistically chooses the currently highest rate in the market. Otherwise, PRS_n is *Irrational* and the strategy chooses an average rate next.
- **Random** a strategy which chooses the next action at random

We refer the reader to (Reddy and Veloso, 2011b) for complete details on these strategies.

2.5 Experimental Evaluation

We evaluated our SELF broker against the benchmark strategies from Section 2.4.2 in a series of experiments. Each experiment ran over 10 simulated days (240 timeslots) since longer durations had very little impact on performance differences. The performance of each individual broker was computed as the share of the overall profits they captured. We repeated each experiment 70 times to obtain confidence intervals; all confidence intervals and significance claims reported below are at the 95% confidence level. The customer population was fixed to five customer groups based on our customer model, each representing the aggregate behavior of a *group* of ten households.⁹ Each customer model instance was parametrized with the same switching probability q and degree of irrationality τ as indicated below. Note, that the parameter settings only imply equal *levels* of switching probability and irrationality among customer groups, whereas the actual *decisions* vary among groups. The markup parameter μ of the reference strategies Fixed and TableRL was set to 0.10, at which we found that these strategies performed best.

2.5.1 Manual broker construction

We first constructed several instances of SELF manually by selecting learning parameters and features based on our best knowledge of the problem domain. One rather typical example configuration is summarized in Table 2.4 where gray cells mark candidate feature/encoding combinations, and black cells mark pairs that were actually used as part of the state space.

⁹We found that a larger numbers of customer groups had no significant impact on the results as they did not change the diversity of the population, while fewer customer groups produced an unrealistic “winner takes it all” competition (see also Section 2.5.5).

This particular instance uses features that reflect the broker’s own profitability (CashGradient) and customer base (NumberCustomers), the competitiveness of the market (MarketBreadth), as well as the broker’s own aggressiveness in the market (MarkupLeader) – arguably some of the fundamental variables in tariff pricing decisions.

| Feature | Plain | RBF | RBF(T) | Bin | Parameter | Value |
|------------------------|-------|-----|--------|-----|------------------|-------|
| Bias | | | | | α_{max} | 0.40 |
| ActionIndex | | | | | α' | 1.00 |
| ActionOneInK | | | | | ϵ_{max} | 0.20 |
| BetterConsumptionRates | | | | | ϵ' | 0.70 |
| CashGradient | | | | | γ | 0.90 |
| CustomerGradient | | | | | | |
| MarketBreadth | | | | | | |
| MarketShare | | | | | | |
| MarkupLeader | | | | | | |
| NumberCustomers | | | | | | |
| RateChangeIndicator | | | | | | |
| TargetMargin | | | | | | |
| WholesalePrice | | | | | | |
| WorseConsumptionRates | | | | | | |

Table 2.4. Configuration of a SELF instance constructed manually. Gray shading indicates all candidate features, black shading represents features that were manually selected using domain knowledge.

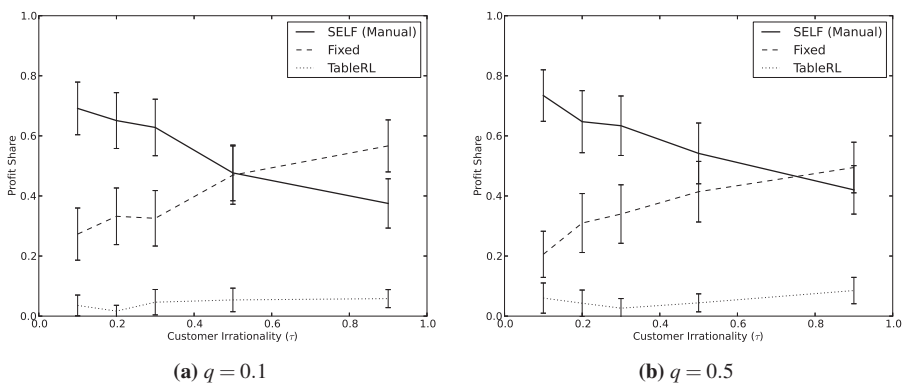


Figure 2.6. Performance of the manually constructed SELF instance from Table 2.4. While it is possible to manually select features that perform well over a limited parameter range, constructing a broker that performs universally well proves challenging.

The empirical performance of the manually constructed instance in competitions against a Fixed and a TableRL benchmark broker is shown in Figure 2.6. The tariff switching probability q was set to a low value ($q = 0.1$, left panel) and a moderate value ($q = 0.5$, right panel), and we recorded the broker’s performance while varying the customer irrationality parameter τ . SELF beats the reference strategies Fixed and TableRL by a statistically significant margin in many of these environments. It is interesting to note that TableRL’s performance lags not only behind SELF, but also behind the Fixed strategy. This does not contradict the good performance reported in (Reddy and Veloso, 2011b), as the settings we explore here differ from those for which TableRL was constructed (see Section 2.4.2). However, this result underscores the importance of a well-chosen state space, and the need for a broker design that is able to identify and accommodate any effective state space for a given environment.

Importantly, our results also demonstrate a common outcome for manually constructed SELF brokers: While it is possible to construct broker instances that perform very well under *some* market conditions, achieving robustness over a wide range of market conditions is exceedingly difficult. For high levels of customer irrationality, the performance of the manually-constructed broker approaches that of the Fixed strategy. This result may seem counter-intuitive, because even for the challenging case of customers choosing their tariffs at random, there is a winning strategy: by raising tariff rates, a broker can increase its profit margin without affecting its customer base. The diminishing performance of SELF for large values of τ stems from an implicit assumption behind its manually constructed state space. Recall that this broker’s state space is constructed from the number of subscribed customers, a profitability measure (CashGradient), a competitiveness measure (MarketBreadth), as well as a measure of the broker’s own aggressiveness in the market (MarkupLeader). This is a well-chosen feature set for capturing rational, consistent market conditions; however, these features are far less informative or even distracting in settings with significant randomness in customers’ choices.

In further experiments we analyzed the performance of manually constructed SELF instances for a wide array of settings by varying the simulation length, the number of customers, values of the markup parameter μ of the reference strategies, and the settings of the learning parameters. We omit details here for the sake of brevity, but we note that this SELF instance performs competitively in all cases except for pathological choices of learning parameters.

2.5.2 Feature Selection

To further improve the SELF broker's learning and to overcome the challenges that arise from manual feature and parameter selection, we explored the use of the feature selection approaches described in Section 2.3.3 to automatically adapt SELF instances to different market conditions. We fixed a customer population with relatively low switching probability ($q = 0.1$) and irrationality ($\tau = 0.1$) and employed greedy feature selection or a Genetic Algorithm to identify high performing SELF instances. Both methods were allotted a maximum of 24 hours for practical reasons; during the search, the fitness of each candidate instance was evaluated over 25 simulation runs.

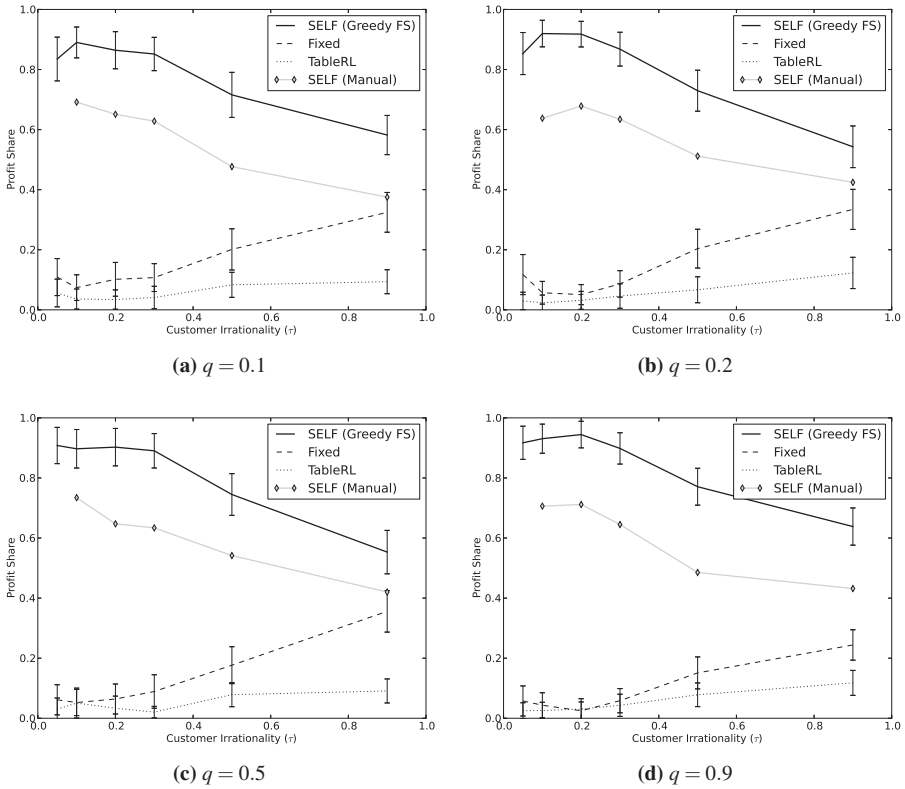


Figure 2.7. Performance of the SELF instance obtained through greedy feature selection, see also Table 2.5. Its performance compares favorably to the manually constructed instance over a wide range of environmental parameters.

Figure 2.7 shows the performance of the broker configuration obtained with the **greedy**

feature selection process. In contrast to the manually constructed instance, the data-driven feature-selection instance consistently outperforms the Fixed strategy over the full range of environmental settings, even while declining noticeably in high-noise environments (higher values of τ). The latter behavior may suggest an adaptation to the low- τ environment used for feature selection, and may also reflect inherently greater difficulty of deriving a profitable policy when customers exhibit random behaviors.

| Feature | Plain | RBF | RBF(T) | Bin | Parameter | Value |
|------------------------|-------|-----|--------|-----|------------------|-------|
| Bias | | | | | α_{max} | 0.73 |
| ActionIndex | | | | | α' | 0.22 |
| ActionOneInK | | | | | ϵ_{max} | 0.45 |
| BetterConsumptionRates | | | | | ϵ' | 0.04 |
| CashGradient | | | | | | |
| CustomerGradient | | | | | γ | 0.71 |
| MarketBreadth | | | | | | |
| MarketShare | | | | | | |
| MarkupLeader | | | | | | |
| NumberCustomers | | | | | | |
| RateChangeIndicator | | | | | | |
| TargetMargin | | | | | | |
| WholesalePrice | | | | | | |
| WorseConsumptionRates | | | | | | |

Table 2.5. SELF instance configuration obtained through greedy feature selection. The resulting state space is sparse, the parameter values hint at a strong overfitting effect, however.

Evidence for the former hypothesis can be found in the configuration summary in Table 2.5. The extremely high initial learning rate $\alpha_{max} = 0.73$, along with slow decay ($\alpha' = 0.22$), and a high exploration rate hint at strong overfitting. This result is a direct consequence of the relatively stable environment for which the SELF instance's features are optimized. In fact, in most simulation runs we observed under this configuration, SELF learned to price its tariff slightly below the Fixed strategy's rate very early on, and remained in that position for the rest of the simulation. While this policy does well in many settings, it will not likely perform well in environments with high levels of wholesale market volatility or customer irrationality (see the remarks on winning strategies in these environments above). We give a complete example of such a simulation run in Appendix A.1.

Support for the overfitting hypothesis also comes from the performance shown in Figure 2.7. Somewhat surprisingly, SELF's performance in high- τ environments first decreases with increasing values of q (closing gap on right side in panels (a) though (c)) but then improves significantly for $q = 0.9$ (widening gap in panel (d)). We interpret this as a low-bias /

high-variance phenomenon with high variance results away from the original feature selection environment. This insight gives rise to opportunities for increasing the robustness of the feature selection outcome, which we explore below.

| Feature | Plain | RBF | RBF(T) | Bin | Parameter | Value |
|------------------------|-------|-----|--------|-----|---------------------|-------|
| Bias | | | | | α_{max} | 0.66 |
| ActionIndex | | | | | α' | 0.37 |
| ActionOneInK | | | | | ε_{max} | 0.16 |
| BetterConsumptionRates | | | | | ε' | 0.62 |
| CashGradient | | | | | | |
| CustomerGradient | | | | | γ | 0.83 |
| MarketBreadth | | | | | | |
| MarketShare | | | | | | |
| MarkupLeader | | | | | | |
| NumberCustomers | | | | | | |
| RateChangeIndicator | | | | | | |
| TargetMargin | | | | | | |
| WholesalePrice | | | | | | |
| WorseConsumptionRates | | | | | | |

Table 2.6. SELF instance configuration obtained through GA feature selection, overfitting effects are less pronounced than with greedy feature selection but still noticeably present.

The performance of the SELF instance constructed using **Genetic Algorithm** feature selection is shown in Figure 2.8. This broker significantly outperforms its benchmarks, but it performs slightly worse than the SELF instance derived with greedy feature selection. Here, as in the greedy case, we find evidence of overfitting in the empirical results and in the learning parameters, with very high learning and exploration rates (Table 2.6) and high-variance behavior in environments that are different from those for which the broker was optimized. Moreover, the GA produces solutions to the feature selection problem that are significantly denser than the greedy solutions and it takes the algorithm longer to find them. In all our experiments, the GA feature selection exhausted the maximum allotted 24 hours, while greedy feature selection typically terminated within 5 - 10 hours, well below the limit. We therefore recommend the use of greedy feature selection over Genetic Algorithms and henceforth employ the greedy procedure exclusively.

2.5.3 Market Stability / Guarding against Overfitting

Our findings above have significant practical implications for the Smart Grid domain and beyond: Overfitting, either from automatically adapting autonomous trading strategies to certain environments or from comparable manual optimization, threatens the stability of

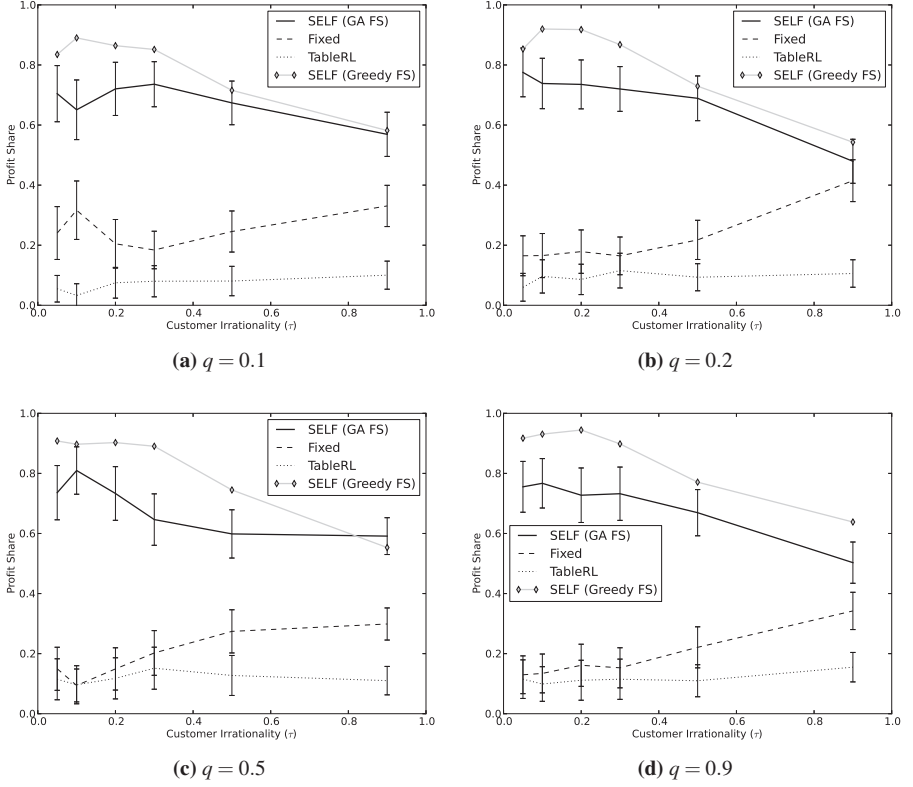


Figure 2.8. Performance of a SELF instance obtained through GA feature selection, see also Table 2.6. The more thorough search for the best state space does not result in better performance as compared to a simple greedy approach.

Smart Markets. As such, measures against overfitting are important to both designers of autonomous trading strategies and policy makers. In this section, we first consider two measures that can be built into the optimization process, **bootstrapping** and **noise injection**, before we turn our attention to regularization as an alternative to offline feature selection in Section 2.5.4.

In the experiments above, all SELF instances were initialized with zero weights for their value functions, corresponding to an initially random policy that evolves into a meaningful decision-making strategy over time. Initializing brokers with a random policy has two important implications in our setting. In the case of offline optimization, this initialization encourages the selection of features and parameters that allow for very fast learning early

on, at the cost of not generalizing well. In addition, this sets SELF instances at a significant disadvantage relative to non-learning strategies, such as Fixed, which can take reasonable actions from the very beginning.

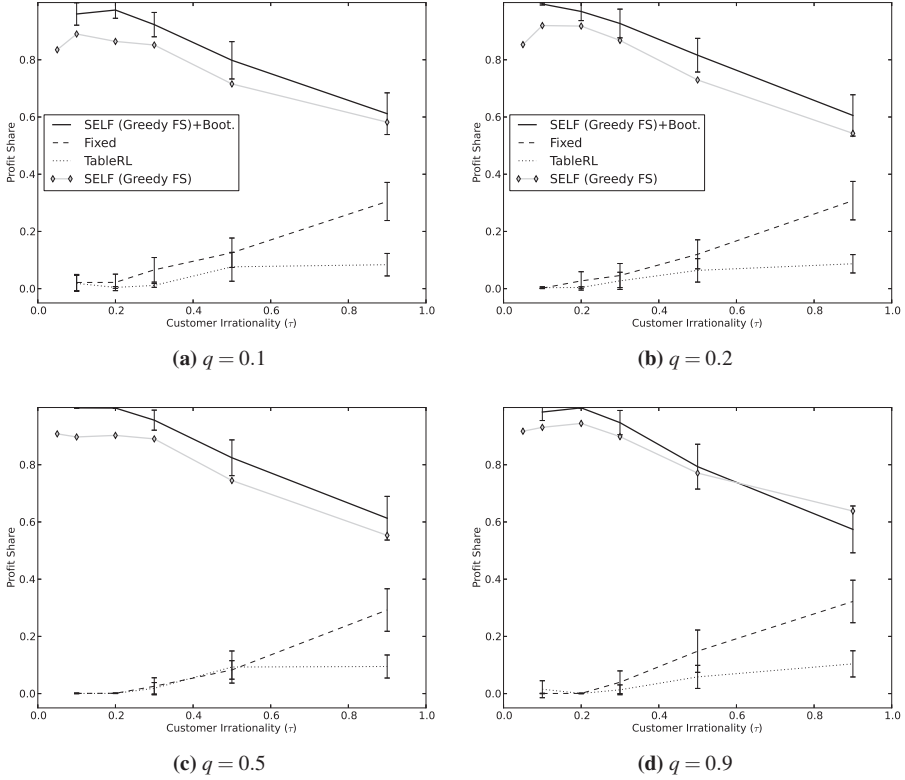


Figure 2.9. Performance of a SELF instance obtained through greedy feature selection with bootstrapping. Bootstrapping leads to a performance increase due to the lowered impact of costly initial explorations.

To counteract the overfitting tendency of greedy optimization, we explored the use of **bootstrapping** within the optimization process. Specifically, we first trained the candidate on one run of the simulation, as before. The fitness measure, however, was then evaluated on a second run of the simulation in which we bootstrapped the candidate using the policy learned during the first run. This procedure can be interpreted as a form of cross-validation of the broker's configuration.

An examination of the configurations obtained in this manner reveals that several automatically selected parameters now have more moderate values compared to when bootstrap-

ping is not used. Importantly, with bootstrapping the decay rates $\alpha' = 1.41$ and $\varepsilon' = 1.09$ are both super-linear, yielding a quick decline in learning and exploration, and correspond to a stable learned policy towards the end of each simulation. While the initial learning rate $\alpha_{max} = 0.37$ and the initial exploration rate $\varepsilon_{max} = 0.34$ are relatively high, they are both significantly lower than those produced without bootstrapping.¹⁰

In comparison with the non-bootstrapping case (Figure 2.7), the broker's performance shown in Figure 2.9 is promising in two important ways. First, as shown, the broker's performance towards the left of the graphs (low- τ , the environment for which the broker is optimized) is better than with the non-bootstrapping instance. This results from the broker taking informed actions from the beginning. In addition, as shown in Figure 2.9, while performance towards high- τ values declines, it does so in a consistent, predictable fashion across different values of q . Taken together, our findings suggest that bootstrapping is an effective countermeasure against the overfitting effects associated with plain greedy feature selection.

We now consider **noise injection**, which has long been recognized for its potential benefits in alleviating overfitting and improving generalization in supervised settings, e.g. (Bishop, 1995). A challenging choice with this approach is setting the level of noise to be injected into the learning process, such that generalizable patterns remain while spurious patterns are masked by noise. We propose that in competitive Smart Market settings this problem can be circumvented by introducing additional brokers into the market. In particular, in the experiments that we present here, we included in the environment additional brokers which follow a Random strategy. While purely random brokers cannot be expected to be present in actual Smart Electricity Markets, they are interesting to consider because they allow the separation of noise injection effects from the effects of additional competition, i.e., random traders inject noise while not capturing any significant market share.¹¹

The performances with noise injection are presented in Figure 2.10, and reveal several insightful differences to the performance of the broker learned with bootstrapping. Perhaps the most significant of those is that the broker's performance for high- τ regimes is not only drastically improved over brokers learned with bootstrapping and plain greedy feature selection, but it is also better over a wider range of environmental settings. At the same time, without the use of bootstrapping, the broker must initially learn a policy at the cost of uninformed exploration; hence, the profit share of the SELF instance is lower in the low- τ regimes for which it is optimized, as compared to when bootstrapping is used. Interestingly, the bene-

¹⁰The full configuration is given in Appendix A.3, Table A.1.

¹¹We also considered the case of additional smart trading strategies. Specifically, we let two instances of SELF compete against each other, as well as against the benchmarks used earlier to explore whether our broker's strategy remains stable under self-play. Our results indicate that both SELF strategies deliver stable performance over a wide range of environments. We refer the reader to Appendix A.2 for details.

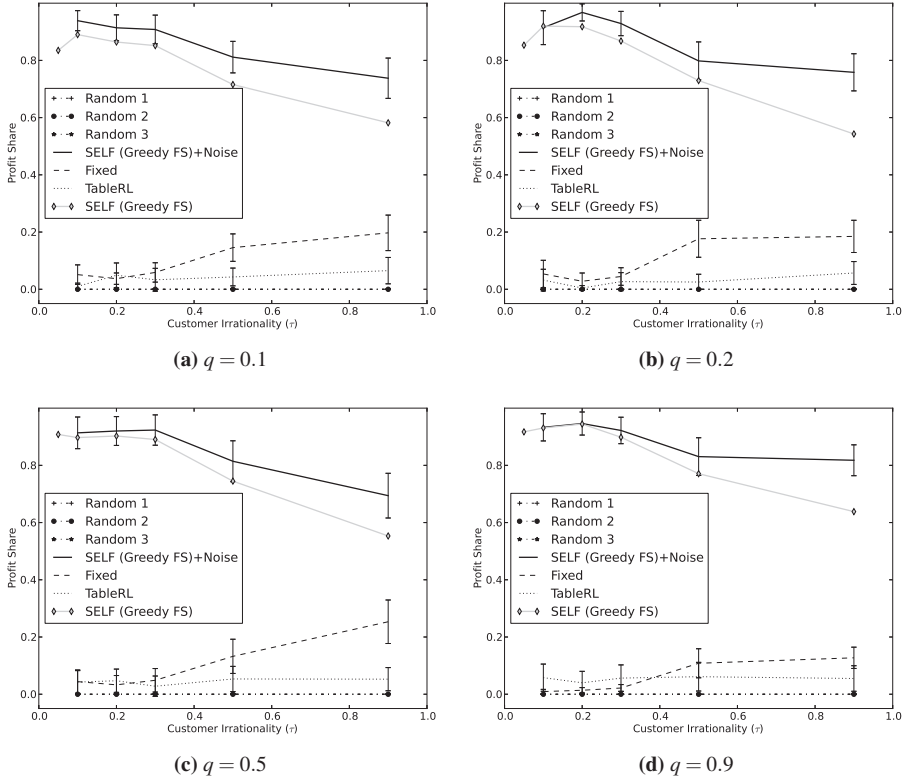


Figure 2.10. Performance of a SELF instance obtained through greedy feature selection with noise injection. Note, that all three random strategies are superimposed at the zero-line.

fits from bootstrapping and noise injection are complementary. In subsequent experiments, we show that the benefits of noise injection and bootstrapping can indeed be combined to obtain even more robust strategies that achieve both improved performance in their target environment, as well as lower-variance performance for environments that differ from the anticipated ones. We refer the reader to Appendix A.3 for further details on these studies.

Finally, our findings above raise several useful questions for Reinforcement Learning scholars as well as for designers of Smart Market mechanisms. We show that better and more robust learning performance can be achieved via the injection of additional noise from randomized brokers, who, importantly, do not capture market share at all. The mere presence of these brokers prompted the greedy optimization procedure to choose a more robust configuration and the resulting broker to act more profitably. While the introduction of purely

randomized traders is infeasible in Smart Markets, there may well be market designs that introduce *equivalent noise* to be perceived by brokers (e.g., trading delays, artificial noise on price and volume data, etc.). From a market design perspective, it is interesting to consider whether such distortions can in fact lead to *better* allocation outcomes in the presence of automated trading strategies. From a theoretical perspective, reinforcement learners typically face stochastic reward and transition functions, and they introduce additional randomness through their own action selection (exploration). However, to our knowledge, the potential benefits of noise injection into the learner's *observations* have so far only been explored for supervised learning tasks.

2.5.4 Regularization

As an alternative to the feature selection techniques discussed above we explored the use of weight regularization for our function approximator. Regularization automatically selects an effective subset from the set of all available features during online learning, and obviates the need for an explicit, offline feature selection phase. It is conceptually simple, and it can be easily integrated into the online learning process.

Figure 2.11 shows the performance of a SELF instance using regularization as roughly comparable to that of the manually constructed instance shown in Figure 2.6. While regularization is not uniformly effective, it is important to note that this level of performance has been achieved with far less domain knowledge and without a model of the environment against which the broker instance could be optimized. Our regularized broker performs well for environments with high customer switching probabilities ($q = 0.9$, right panel) and low levels of customer irrationality (τ small, left end of both panels). Both findings are intuitive: a regularized function approximation requires frequent, strong, and consistent feedback to form an effective policy; in our example such environments arise when customers exhibit high switching probabilities (high q , more frequent retail market feedback) and low levels of irrationality (low τ , more consistent retail market feedback). Thus, while regularization is not consistently effective, it appears to be a beneficial low-cost strategy for learning without explicit feature selection in stable market environments.¹²

2.5.5 Impact of customer characteristics

In a final experiment, we aim to highlight an interesting case of how Machine Learning research can further inform policy decisions in the Smart Grid domain. This does not aim to

¹²An anonymous reviewer offered the following potential explanation for the comparatively weak performance of regularization: Shrinkage tends to drive weights to zero in areas of the state space that the learner is currently not exploring (based on the sequential nature of exploration in online RL). In other words, the learner “forgets” about regions of the state space not recently visited. We find this explanation quite plausible.

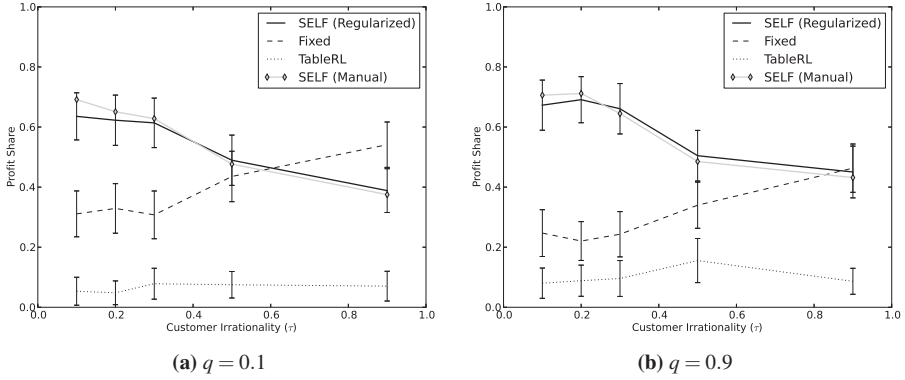


Figure 2.11. Performance of a SELF instance using regularization. The instance performs at par with a manually constructed broker instance but does not require domain knowledge.

offer a full-fledged policy analysis, but a demonstration of how Machine Learning research can contribute to shaping this important, challenging domain. Consider the concept of *microgrids*, self-organizing communities of small-scale electricity consumers and producers, who act as one customer in the retail market, and who only trade net imbalances (i.e., consumption they are unable to meet with own production, or production they cannot absorb with own consumption) in the upstream grid (cf. Figure 2.1).

Important open questions pertaining to microgrids are whether their introduction is beneficial to market stability; and, if so, what are the expected implications of the *number*, *size*, and *sophistication* of microgrids in the market. One can also interpret these factors as follows: the *number* of microgrids captures the *lumpiness* of decision making (where fewer microgrids corresponds to fewer independent decisions in the retail market); the *size* of microgrids reflects idiosyncratic risk, because the aggregated behavior of a larger number of households tends to appear smoother, cf. Figure 2.4; and, finally, the *sophistication* of microgrids pertains to the degree of rationality in their decision-making, and is inversely related to τ in our notation. Higher levels of sophistication can be a consequence of implementing advanced decision-support within each microgrid, but can also be related to the amount of information flowing from the upstream market into the microgrid. We expect the autonomous retail trading task to become measurably harder as the lumpiness of decision-making increases (up to the limiting case of “winner takes all” competition), as the size of each microgrid decreases (up to the limiting case of single households, as in traditional electricity retail markets), and as sophistication decreases.

We studied the implications of these factors on the performance of SELF, as it represents a state-of-the-art autonomous retail trading agent. The left column in Figure 2.12 shows increasing numbers of microgrids with low levels of sophistication (high τ values), whereas the right column presents microgrids with higher sophistication; each panel shows the performances of the broker across different sizes of microgrids. One insight from these results is that the number and sophistication of microgrids are key factors in the performance of our strategy, whereas the implications of each microgrid's size is less significant. Specifically, the performance curves for SELF and the alternative strategies are roughly flat, with no clear trend as microgrid sizes change. The only exceptions are the single-grid cases in panels (a) and (b), where SELF starts to perform consistently from about 10-20 households. This result is partly due to a degenerate “winner takes it all” competition as reflected in the wider confidence bounds, and partly due to the Fixed strategy generally ignoring consumption patterns and tariff selection decisions. A second useful insight is the decisive role that the sophistication of individual microgrid decisions plays in autonomous retail trading strategies' performance. For low sophistication (left column), SELF performs comparably to a simple Fixed strategy in most cases; only for relatively large numbers of microgrids does SELF perform significantly better than the alternative strategies. In contrast, the right column shows SELF outperforming other strategies almost across the board.

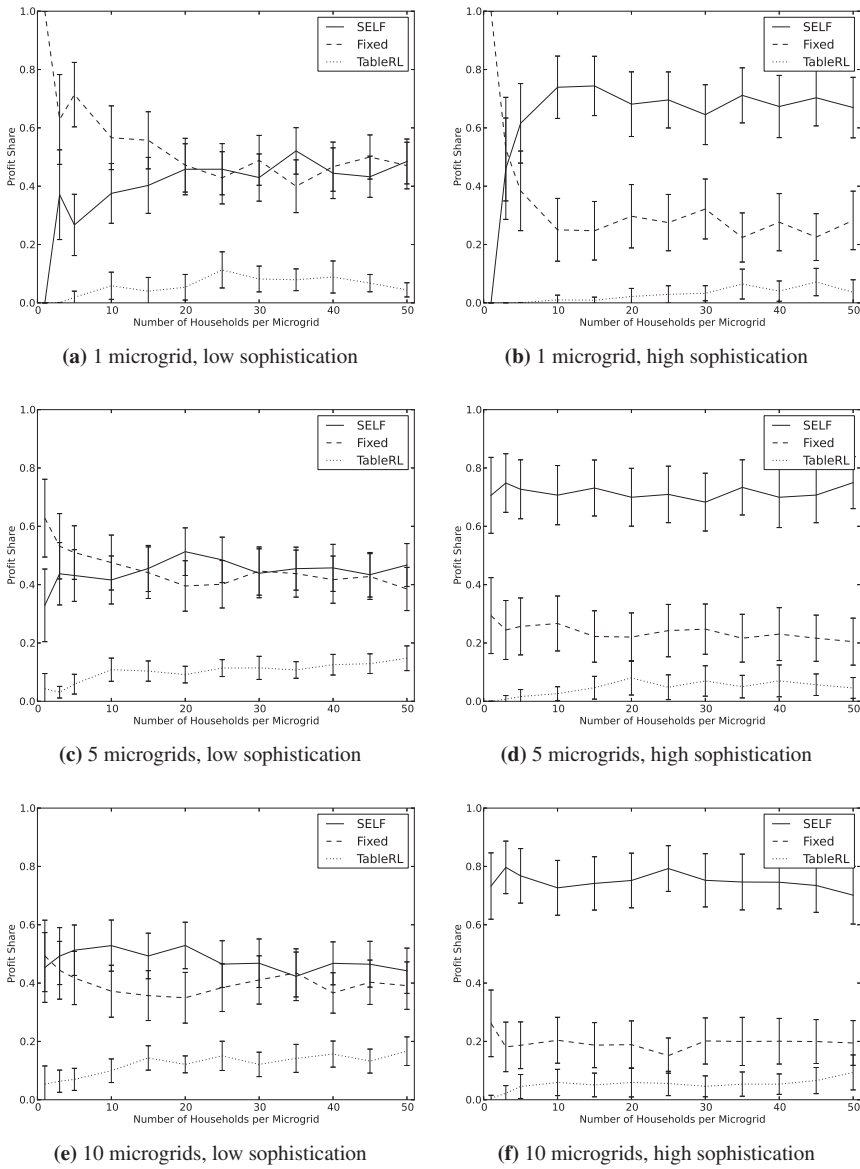


Figure 2.12. Impact of microgrid characteristics on broker performance, left column: low level of microgrid sophistication ($\tau = 1.0$) for increasing numbers of microgrids, right column: ditto for higher level of microgrid sophistication ($\tau = 0.1$)

It is important to note, that the superior performance of SELF is not merely beneficial to the self-interested broker. Rather, it also reflects higher overall economic welfare: consumers' own preference of SELF over competing strategies is indicative of its desirable, competitive tariff offerings, especially as compared to the customer-agnostic Fixed strategy. Furthermore, higher profit shares are a result of better balanced portfolios with lower imbalance charges, and ultimately contribute to a more efficient and stable grid. From a policy perspective, our results indicate that further research into decision-support systems for microgrids will likely benefit both consumers and the grid overall. Specifically, further inquiry is needed into the benefits that accrue to microgrids from using data-driven, decision-support technology; into whether microgrids would invest in such technology by their own choice; and, alternatively, if public-sector investments into novel decision-support systems could lead to increased economic welfare and a more stable Smart Grid in the future.

2.5.6 Summary of Results

Figure 2.13 summarizes the characteristics of the methods we proposed by juxtaposing their performance in a Smart Market environment, their capacity to generalize, and their computational requirements.

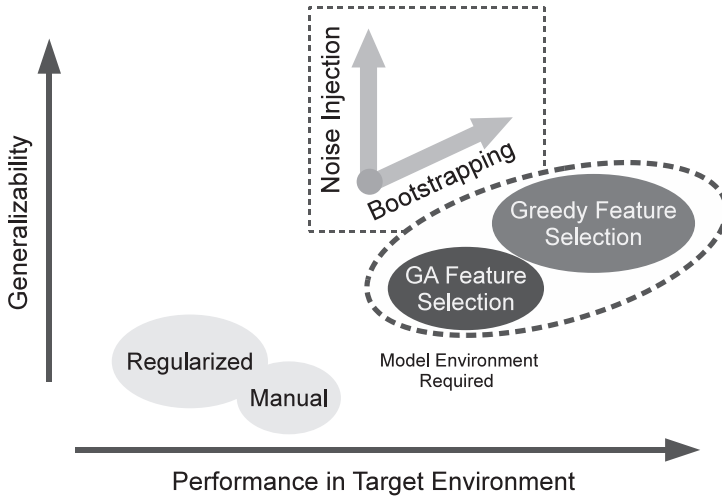


Figure 2.13. Summary of results. Darker colored ellipses indicate higher computational demand

The simple manual and regularized approaches performed reasonably well for the environments for which they were constructed, but their performance deteriorated quickly as environmental conditions changed. The regularized approach generalized slightly better and,

importantly, required significantly less domain knowledge in the design process. Both approaches are computationally efficient and can be applied without a model of the target environment.

When an environmental model is available, feature selection techniques can be leveraged to obtain significantly better performance in the target environment and beyond. Both, greedy and GA feature selection, led to strategies that generalized significantly better than, e.g., the simple regularized approach. Generalization is desirable because of potential shifts in a Smart Market environment, but also to effectively accommodate potential mismatches between the environmental model used for feature selection and the actual target environment. Both feature selection techniques require significant computation before executing the strategy, but have little impact on the derived strategies' runtime requirements. In our experiments we found greedy feature selection to deliver generally better, sparser results at lower computational costs and we therefore recommend it over GA feature selection for our application.

Finally, we demonstrated how bootstrapping and noise injection can be integrated into the feature selection process to improve performance (bootstrapping) in a given environment, and generalizability (noise injection). Importantly, we show that both techniques can be combined to benefit both objectives.

2.6 Related Work

To date, research on retail electricity trading has received relatively little attention. To our knowledge, Reddy and Veloso (2011b) were the first to suggest RL as an appropriate framework for constructing brokers for retail electricity markets. A key distinguishing feature of our approach is the automated, data-driven construction of the state space. In contrast, the strategies developed in (Reddy and Veloso, 2011b) are derived from manually constructed features and are limited in the number of economic signals they can accommodate as well as in their ability to incorporate new signals when the market environment changes. Another key distinction is that the brokers presented in (Reddy and Veloso, 2011b) are derived for an environment with fixed rates of electricity consumption and production for all market participants where brokers source electricity exclusively from small-scale producers. Consequently, the broker learns to steer towards an optimal *consumer/producer ratio* among its subscribers by changing tariff rates. These settings yield a broker which is unable to develop appropriate responses to any variability of consumption and production over time or between different customers.

Reinforcement Learning has been used on a wide range of problems in electronic com-

merce in which agents aim to learn optimal policies through interaction with the environment. For example, Pardoe et al. (2010) develop a data-driven approach for designing electronic auctions based on notions from RL. In the electricity domain, RL has primarily been used to derive wholesale trading strategies, or to build physical control systems. Examples of electricity wholesale applications include (Rahimiyan and Mashhadi, 2010), who derive bidding strategies for electricity wholesale auctions, and Ramavajjala et al. (2012) who study Next State Policy Iteration (NSPI) as an extension to Least Squares Policy Iteration (LSPI) (Lagoudakis and Parr, 2003) and demonstrate the benefits of their extension on the day-ahead commitment problem of a wind farm. Physical control applications of RL include load and frequency control within the electric grid and autonomous monitoring applications, e.g., (Venayagamoorthy, 2009).

Feature selection and regularization have been studied widely in supervised settings, e.g., (Guyon and Elisseeff, 2003), but have only recently gained momentum in the Reinforcement Learning community. In our experiments we implemented the L1 regularized version of LARS-TD by Painter-Wakefield (2012) due to its conceptual simplicity. An alternative approach is Sparse TD with Equi-Gradient Descent (EGD) by Loth et al. (2007) and we are planning on exploring its relative merits in future work. Wrapper approaches to feature selection are commonly used in RL as they are easily integrated as a pre-processing step to the actual RL task. For example, Whiteson et al. (2005) present FS-NEAT, an extension of the well-known NEAT algorithm, to incorporate feature selection capabilities. They demonstrate the benefit of this approach on two standard RL benchmarks. Another feature selection technique specifically targeted at RL applications is the LSTD-RP method by Ghavamzadeh et al. (2010). They extend the classic Least Squares TD (LSTD) algorithm (Bradtke and Barto, 1996) to work with random projections of high-dimensional state spaces, and show how their work translates to online RL settings by replacing LSTD with LSPI.

Whiteson et al. (2011) provide interesting insights into the role of *environment overfitting* in empirical evaluations of RL applications. They argue that *fitting*, i.e., the adaptation of a learner to environmental conditions known to be present in the target environment, is an appropriate strategy. *Overfitting*, i.e., the adaptation of the learner to conditions only present during evaluation, on the other hand, is inappropriate. Our experiments suggest techniques that strike a good balance between fit and performance levels for autonomous trading strategies.

2.7 Conclusions and Future Work

The Smart Grid vision relies critically on decentralized control methods that can balance electric grids in real-time. Developing an understanding of such methods is one of the cornerstones of an efficient, safe, and reliable Smart Grid, with far-reaching benefits for society at large.

We presented SELF, a novel design for autonomous Electricity Broker Agents built on insights from Reinforcement Learning, and from Machine Learning more generally. The key design objectives behind SELF are flexibility and robustness. We framed the broker challenge as optimal control problem and used RL with function approximation to derive robust long-term policies for our SELF brokers. Towards the flexibility objective, we explored the use of feature selection and regularization techniques to automatically adapt brokers to a broad range of market conditions. Using these techniques, SELF brokers can identify and accommodate arbitrary sets of informative signals from their environment, resulting in significantly better performances compared to previous designs. We also evaluated complementary bootstrapping and noise-injection methods to reduce overfitting, and we showed how their use leads to more robust, generalizable feature selection outcomes.

Our work formalizes a class of Smart Electricity Markets by means of our simulation model SEMS, which is a contribution in its own right. SEMS employs real-world wholesale market data and a complete, micro-level model of electric appliance usage in private households, making it a more realistic model of future Smart Electricity Markets than those used in previous studies. We demonstrated the efficacy of our broker design for a range of Smart Electricity Markets which varied substantially in terms of tariff choice behaviors among their customer populations. Our experimental results demonstrate that both, the broker's capacity to accommodate arbitrary state spaces, and its selection of informative features, are important for learning robust policies. Our SELF brokers are significantly more flexible in this regard than previously suggested strategies.

Research on autonomous electricity brokers for the Smart Grid is an emerging field. Hence, beyond the development of a novel broker agent design, we aimed to generate useful insights on key design decisions that enable broker agents to operate effectively in the Smart Grid. For instance, we studied the use of L1 regularization and found that it offers a viable alternative to manual broker construction under stable market conditions. We contrasted regularization with greedy and GA feature selection and found that a simple, greedy feature selection approach can yield significant performance improvements when a model of the environment is available, and when overfitting can be avoided. We presented effective strategies for counteracting overfitting, including an innovative approach for injecting effective noise via *random* broker strategies.

In future work it would be beneficial to further explore SELF's performance in increasingly sophisticated Smart Electricity Markets. Key features to explore include advanced tariff structures, renewable energy sources, and storage devices such as electric vehicles. A question of great practical import is whether the performance of the learner we present here, possibly extended with more advanced RL techniques, will translate to these more complex environments as well. Customers, for example, may eventually adopt electronic agents of their own. In fact, this is commonly thought to be a prerequisite for the success of more complicated tariff models. Our preliminary analysis in Section 2.5.5 gives reason to believe that the use of such agents might actually *benefit* broker agents if they act closer to perfect rationality. How broker agents cope with strategic departures from perfect rationality is, however, unclear.

Our noise injection experiments entail the question whether the extensive work on overfitting that the Machine Learning community has done can be connected to questions about market stability that are under study in the Finance field. Smart Market designs could, for example, be drastically improved if artificially introduced noise (e.g., trading delays) could indeed be proven to be generally beneficial in the presence of autonomous trading strategies. To our knowledge, this connection has not been studied previously.

Another related, and possibly complementary, objective is to derive policies that are *comprehensible* for human decision-makers. Comprehensible policies serve as a further safeguard against overfitting. But they also increase the trust of human decision-makers in autonomous trading strategies, an important precondition to the adoption of Machine Learning techniques in the safety-conscious Smart Grid domain.

We believe that our proposed strategies offer important benchmarks for future work and that this work offers a meaningful contribution to our understanding of key design decisions for broker agents to operate effectively in the Smart Grid.

Chapter 3

A Scalable Preference Model for Autonomous Decision-Making Involving Consumer Choices

3.1 Introduction

Enhancing individual and organizational performance through information technology is one of the fundamental promises of Information Systems (IS). Consumers can now choose among a wide variety of affordable products and services, and managers routinely make decisions based on rich, differentiated real-time information. However, while IS has become increasingly adept at provisioning decision-relevant information, autonomous decision-making remains elusive under all but the most highly structured circumstances. Two prominent application areas, dynamic flight pricing and automated credit approvals, now apply automated business rule engines to make most operative decisions quickly, cheaply, and reliably (Davenport and Harris, 2005, Baker, 2013). Yet, in less structured settings – from planning the next vacation to trading in complex multi-echelon markets – autonomous decision-making through software agents remains an open problem, and an active research area, e.g., (Adomavicius et al., 2009).

What hampers autonomous decision-making in unstructured settings is not the inability of information technology (IT) to select good decisions, but rather its inability to detect what *constitutes* a good decision in a user's eyes. In many settings, individuals are unaware of the mechanisms underlying their own preferences (Lichtenstein and Slovic, 2006), and data-

driven preference models that elicit unknown preferences by generalizing from observed choices are therefore particularly beneficial (Bichler et al., 2010).

As an example, consider **smart grids**, where IT is anticipated to play a key role in improving the efficiency of electricity distribution and use (Kassakian and Schmalensee, 2011). Particular challenges in this context are electric vehicles that are charged in varying locations, and the incorporation of intermittent and variable renewable electricity sources, such as solar and wind (Valogianni et al., 2014b). Data-driven modeling of electricity consumption preferences is essential for predicting consumption patterns, and for effectively incentivizing consumers towards sustainable behaviors (Watson et al., 2010, Peters et al., 2013c). For example, a preference model can learn and predict that a given user is unlikely to use her electric vehicle (EV) in the afternoon. If renewable energy from solar and wind is scarce, the system can offer the owner personalized incentives to make the battery's energy available to nearby consumers. The benefits from data-driven preference learning and autonomous decision-making in EV coordination can be significant, with estimated electricity cost reductions of 3 to 14%, and CO₂ emission reduction of up to 3.5% (Kahlen et al., 2014).

Prior work on preference learning has made significant advances in generating accurate predictions from noisy observations such as electricity meter readings that exhibit inconsistencies, heterogeneity, and bias (Kohavi et al., 2004, Evgeniou et al., 2005). In particular, recent non-parametric models automatically adapt to the complexity of real-world observations, and they embrace inconsistencies in human choices rather than imposing stringent rationality assumptions. The ability to accommodate inconsistencies in observed choices allows these models to distinguish between instances where estimates are certain enough to justify an autonomous action, and instances where the model should actively acquire additional evidence or transfer control to a human decision-maker (Saar-Tsechansky and Provost, 2004, Bichler et al., 2010).

However, additional progress is necessary for preference models to become widely adopted in practice. First, important domains such as energy commerce and healthcare require methods that are **computationally efficient**, and that scale gracefully to millions of observations. Contemporary electric distribution systems, for example, connect up to ten million consumer meters, each transmitting observations every few minutes (Widergren et al., 2004). Such meters produce large amounts of data that must be processed quickly and with high granularity (i.e., not aggregated). Observations must be processed quickly because consumers' needs are momentary, and the time window for incentivizing desirable behaviors is small. Similarly, observations must be processed at high granularity because fine-grained, local solutions to supply bottlenecks are superior to global solutions that require costly wide-area transmission. To achieve these goals, preference models must provide short and consistent

training times, and incorporate and act on new data in a timely manner. Second, preference models and corresponding inference methods should be **conceptually simple** and **easy to validate** to promote adoption, e.g., (DeLone and McLean, 2003). Significant progress has been made on automated inference in simpler predictive models, where automated tools now generate inference algorithms based on declarative descriptions, e.g., (Koller and Friedman, 2009). However, more complex settings, such as electricity preference learning, still entail manual implementation, which is challenging even for highly trained professionals. Both scalability and conceptual simplicity are important in a wide variety of applications, ranging from marketing (van Bruggen et al., 1998, Cui and Curry, 2005), to personalization (Murthi and Sarkar, 2003, Birlutiu et al., 2012), and public infrastructure, e.g., autonomous traffic congestion pricing (Yang, 1998).

As shown in our empirical results, existing approaches exhibit a clear trade-off between computational cost and scalability on the one hand, and predictive accuracy on the other. Simple methods have the benefits of exhibiting low computational cost and high scalability. However, they also offer substantially lower accuracy than more complex and computationally expensive ones. We present a novel approach, the **Gaussian process Trade-off Model (GTM)**, that strikes a new balance between scalability, predictive accuracy, and conceptual simplicity.

GTM offers an order-of-magnitude improvement in scalability, which makes it competitive with simple models in terms of computational cost. However, unlike these simple models, *GTM* exhibits predictive accuracies that compare favorably to the most accurate and computationally expensive approaches. *GTM* uses a conceptually simple inference procedure with important implications for ease of adoption in practice. It is a member of the Bayesian family of models, which have recently gained popularity in the preference modeling community, including in IS. Inference in Bayesian models thus far has been either fast but conceptually expensive (such as with variational methods, for example), or elegant but computationally slow (such as with sampling-based approaches; Montgomery and Smith 2009). Our proposed *GTM* makes use of a simple approximate inference scheme known as Laplace’s method, and we leverage common features of consumer choice settings, particularly the small number of relevant product attributes and limited attention to each attribute, towards vastly improved scalability.

We empirically demonstrate these improvements on three real-world consumer choice datasets. Specifically for this study, we collected an electricity tariff choice dataset on a commercial crowdsourcing platform based on data from a US retail electricity market. We use two choice datasets from other domains (political elections and automobile purchases) to confirm our findings. In our evaluations, we find that *GTM* is more than an order of mag-

nitude faster than some existing methods, and highly scalable in the number of observations. Moreover, the accuracy of *GTM*'s predictions compares favorably to other state-of-the-art preference models despite *GTM*'s greater simplicity and speed.

GTM introduces a new benchmark to the preference learning toolset. In consumer choice settings, where alternatives can be described by a small number of relevant attributes, it produces accurate estimates over an order of magnitude faster than previous methods, and it scales to the large numbers of users and choices common in IS applications. Its probabilistic design is especially suitable for applications that aim at autonomous decision-making. And its conceptual simplicity makes *GTM* significantly more accessible to modeling practitioners who wish to apply the model to new domains.

3.2 Background and Related Work

Human preferences have long been a subject of interest for scholars from a multitude of fields including Information Systems, Computer Science, Psychology, Marketing, and Econometrics.

3.2.1 Information Systems

Within IS, our work is most closely related to the literature on **recommender systems**, which are concerned with finding and recommending items of interest, e.g., in electronic commerce settings (Adomavicius and Tuzhilin, 2005). We adopt several of their usefulness criteria, in particular accuracy, computational tractability, and scalability for use in real-world IS. But the problem we study is different, and in some sense harder than the recommendation problem:

1. Our model targets autonomous decision-making and decision-making supported by software agents instead of giving recommendations, a difference that is reflected in its probabilistic design. Autonomous decisions are becoming increasingly important in fast-paced, data-intensive electronic marketplaces (Peters et al., 2013c), and we contend that these settings will require a principled quantification of uncertainty that allows software agents to return control to their users in cases of doubt. Some previous recommender systems studies make use of probabilistic models, e.g., (Liu et al., 2009), but these works are not easily generalized to the preference learning setting we consider, because they do not predict users' choices between arbitrary combinations of items.

2. Most recommender systems operate based on ratings which have several drawbacks in the context of preference learning. While reasoning based on ratings is significantly easier, there are four important motivations for studying direct, pairwise comparisons between alternatives. First, pairwise comparisons do not force users into self-consistency. Whereas rating systems suppress, e.g., intransitivity, *GTM* embraces inconsistencies and lets them unfold in the form of uncertainties. Second, pairwise choice situations are cognitively easier for users and evoke a qualitative reasoning mode in contrast to the quantitative reasoning evoked by ratings, making them an interesting subject of study in their own right (Lichtenstein and Slovic, 2006). Third, pairwise comparisons are often readily derived from users' actions, which reduces the users' burden compared to providing explicit training information. Finally, ratings are subjective in the sense that two users may share the same preferences but assign different ratings to alternatives (Liu et al., 2009).

Other related research includes work on personalization (Murthi and Sarkar, 2003) which studies methods for adapting products and services, often in digital environments where the associated costs are low. For example, Atahan and Sarkar (2011) propose a Bayesian method for the accelerated learning of user profiles through dynamic adaption of navigation options. Several other IS scholars have worked on discrete methods that span across multiple academic disciplines. For example, Huang et al. (2012) propose a discrete choice prediction method based on Machine Learning techniques, and Roy et al. (2008) combine ideas from Multi-Criteria Decision Analysis (MCDA) and Operations Research to learn utility functions. These works are related to ours in isolated aspects, but they generally follow different objectives. To our knowledge, our study is the first IS study to propose a preference learning model suitable for large user populations, that gives highly accurate predictions, principled probabilistic estimates, and in which learning is efficient and conceptually simple.

3.2.2 Computer Science

Researchers in **Artificial Intelligence** have a long-standing interest in preferences as a flexible alternative to goals, reviewed in (Brafman and Domshlak, 2009). It may not be clear from the outset whether a certain *goal* is attainable, but using a preference relation, a software agent can work towards the most favorable outcome that is attainable under given circumstances. Artificial Intelligence researchers realized the importance of preferences for autonomous decision-making agents early on (Maes, 1994), and they proposed a series of innovative representations for preference relations and associated reasoning schemes, e.g., CP-Nets (Boutilier et al., 2004). However, most of these representations make strong ratio-

nality assumptions, and reasoning consequently becomes difficult when they are confronted with actual human preferences.

Machine Learning takes a data-driven, predictive approach to preferences, and researchers in the field have proposed models based on a broad variety of learning techniques (Fürnkranz and Hüllermeier, 2011). Of particular interest to us is work on probabilistic preference models that yield uncertainty estimates based on noisy preference observations. Chu and Ghahramani (2005) were the first to model preference learning using Gaussian processes. However, their model lacks the ability to capture heterogeneity across users. More recent work (Guo et al., 2010, Birlutiu et al., 2012, Hounsby et al., 2012) has alleviated this shortcoming, but it does not easily scale to the large numbers of users and observations we are targeting. In Section 3.5, we benchmark our work against several of these models to demonstrate these differences.

3.2.3 Psychology

Scholars in Psychology have primarily taken an explanatory interest in human preferences. One of the key achievements of the discipline has been the establishment of positive theories of human preferences, e.g., (Tversky and Simonson, 1993, Lichtenstein and Slovic, 2006) next to the normative theories from Economics (von Neumann et al., 2007). These positive theories attest that human preferences are inconsistent, constructed as needed, and that the construction process depends on situational framing and environmental factors. The probabilistic nature of our model acknowledges these inconsistencies and lets them unfold in the form of uncertainties. These uncertainties are an important output of our model: a software agent can use them to return control to the user when uncertain about a high-value decision.

3.2.4 Marketing and Econometrics

Finally, our work is related to the Marketing and Econometrics literature where preference measurement methods such as conjoint analysis, Logit/Probit models, and several other discrete choice prediction techniques were pioneered (Greene, 2012). Early preference measurement was limited to population-level estimates, but more recent techniques accommodate heterogeneity across consumer segments (Allenby and Rossi, 1998, Evgeniou et al., 2007). The primary target audience of these models are human decision makers, and their outputs are therefore interpretable coefficients. By contrast, our work focuses on preference learning for IS, and has to consider scalability, incremental updates, and other practical issues that arise when moving from passive preference measurements to autonomous decision-making (Netzer et al., 2008). In Section 3.5, we illustrate these differences by benchmarking

| Symbol | Definition |
|----------------------------------|---|
| \circ | Hadamard (element-wise) matrix product |
| \otimes | Kronecker matrix product |
| $C = \{(u, x^1, x^2, y)\}$ | Choice situations: when presented with alternatives x^1 and x^2 , user u chose $y = +1$ (first alternative), or $y = -1$ (second alternative) |
| γ_u^c, Γ | User u 's possession of characteristic c ; $\Gamma \in \mathbb{R}^{n_U \times n_c}$ collects all γ_u^c |
| d_T, d_X | Dimensionality of elements in T, X |
| $f_u, f^c, f_t = f(t)$ | Users' latent evaluation of trade-offs and characteristic evaluation of trade-offs, evaluation of f at trade-off t , respectively |
| $\theta = \{l_d\}$ | Lengthscale hyperparameters |
| I | Identity matrix |
| K | Covariance matrix, $K \in \mathbb{R}^{(n_c n_T) \times (n_c n_T)}$ |
| $L : LL^T = W$ | Lower Cholesky factor of W |
| n_c | Fixed number of characteristics |
| n_e | Number of Eigenvalues used in low-rank approximations |
| n_T, n_U, n_X | Number of elements in T, U , and X |
| N, Φ | Probability density function (PDF), and cumulative distribution function (CDF) of the standard normal distribution |
| $p(C f), \nabla p(C f)$ | Likelihood and its Jacobian $\frac{\partial p(C f)}{\partial f_i}$ |
| $t, t^{(d)}, T$ | Trade-off t , its d -th element, and set of all trade-offs, respectively |
| U | Set of all users |
| $W = -\nabla \nabla \log p(C f)$ | Negative Hessian of the log likelihood |
| X | Set of all instances |
| $y \in \{-1, +1\}$ | Single choice: $y = +1$ (first alternative), or $y = -1$ (second alternative) |
| Z | Model evidence, also known as marginal likelihood |

Table 3.1. Summary of Mathematical Notation.

GTM against the Mixed Logit model, a well-established standard in Marketing and Econometrics.

3.3 Gaussian Process Trade-off Model (GTM)

Let $U = \{u_1, \dots, u_{n_U}\}$ denote a set of **users** and $X = \{x_1, \dots, x_{n_X} \mid x_i \in \mathbb{R}^{d_X}\}$ a set of **instances**, the objects or actions between which users choose. Each instance is described by d_X real-valued attributes. Thus, a set of **choices** is denoted as:

$$C = \{(u, x^1, x^2, y) \mid u \in U, x^i \in X, y \in \{+1, -1\}\}$$

Here, $(u, x^1, x^2, +1)$ means that user u prefers the first alternative (instance x^1) to the second alternative (instance x^2), whereas $(u, x^1, x^2, -1)$ means the opposite.¹ In contrast to models

¹Table 3.1 summarizes the mathematical notation used throughout the paper.

that additionally include users' demographic information, we aim to infer individuals with similar preferences directly from their choices. This alleviates the need for the common but problematic assumption that individual characteristics (such as demographics) correlate with preferences.

The goal of preference learning is to learn an order relation \succeq_u over instances for each user so as to predict previously unobserved choices, including those of previously unobserved users. Rather than operating directly on order relations \succeq_u , preference models often estimate latent functions from which the order relations can easily be inferred. For example, the standard discrete choice models by Thurstone (1927) and Bradley and Terry (1952) estimate functions $\tilde{f}_u : X \rightarrow \mathbb{R}$ that capture the utility $\tilde{f}_u(x)$ that user u derives from each instance x . When presented with a previously unobserved choice between instances x^1 and x^2 , these models will predict that $x^1 \succeq_u x^2$ if and only if $\tilde{f}_u(x^1) \geq \tilde{f}_u(x^2)$.

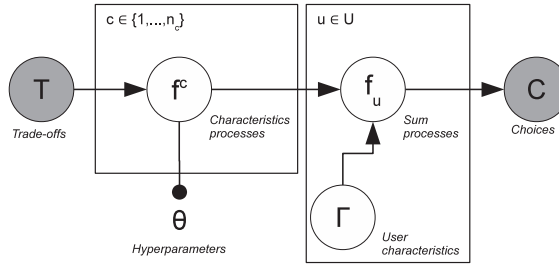
Two important disadvantages of models based on instance utilities are their absolute interpretation of utility independently of context, and the stringent rationality assumptions that follow from this treatment. When making decisions, individuals have been shown to focus on **trade-offs** resulting from their choices rather than on absolute outcomes, and thus perceive alternatives within the context in which they are presented (Tversky and Simonson, 1993). The assumption of utility models that individuals simply recall absolute, predetermined instance utilities $\tilde{f}_u(x)$, and the strict transitivity of \succeq_u implied by this assumption, are frequently violated in practice. In our model, we aim to capture trade-offs $t = \tau(x^1, x^2) \in T$ between alternatives and we estimate probability distributions over users' trade-off evaluations $f_u(t)$.

As an example, suppose that electricity tariffs (i.e., rates or plans) are characterized by their cost per kilowatt-hour and by whether the electricity is generated from renewable sources. Given user u is presented with a choice between two such tariffs:


$$x^1 = \left[32 \frac{\text{¢}}{\text{kWh}}, 1 \text{ (renewable)} \right] \text{ and } x^2 = \left[28 \frac{\text{¢}}{\text{kWh}}, 0 \text{ (non-renewable)} \right]$$

our goal is to predict whether the user will prefer the first ($y = +1$ and $x^1 \succeq_u x^2$) or the second tariff ($y = -1$ and $x^2 \succeq_u x^1$). The trade-off the user faces is $(x^1 - x^2) = (4 \frac{\text{¢}}{\text{kWh}}, 1)$, i.e., by choosing tariff x^1 , the user pays an additional $4 \frac{\text{¢}}{\text{kWh}}$ in exchange for a supply of renewable energy.² In this formulation of the trade-off, our goal is to classify whether the

²If suggested by domain knowledge, one can alternatively formulate the trade-off using percentage increases, $t = (\frac{32-28}{28} \frac{\text{¢}}{\text{kWh}}, 1)$, or any other relevant transformation. Such alternative transformations may increase the interpretability of the model's outputs, but they have a negligible effect on the predictive accuracy of our flexible, non-parametric model, as shown in our empirical evaluations. We conjecture that it is also possible to learn the τ mapping from the data using, e.g., warped Gaussian processes (Snelson et al., 2004).



(a)



(b)

Figure 3.1. Probabilistic Graphical Model of *GTM*. Panel (a): Users make choices based on their evaluations $f_u(t)$ of the associated trade-offs. Users' evaluations are linear combinations of n_c behavioral characteristics f^c which they possess to different degrees γ_u^c . Shaded circles represent observed data, white circles represent latent quantities of interest. The two **plates** in the figure denote replication of the enclosed elements for each characteristic c or user u , respectively. Panel (b): Graphical table of contents for this article.

user will perceive the trade-off as favorable ($f_u(t)$ positive) and choose the first tariff, or perceive the trade-off as unfavorable ($f_u(t)$ negative) and choose the second tariff instead. From a decision-theoretic perspective, our approach is inspired by the trade-off contrast principle (Tversky and Simonson, 1993), and by case-based decision theory (Gilboa and Schmeidler, 1995), which posits that a user's trade-off evaluation will resemble evaluations of similar trade-offs made by the user in the past. From a machine learning perspective, our approach is related to collaborative classification, because we aim to classify trade-offs as favorable or unfavorable based on users' latent evaluations $f_u(t)$.

Because human preferences are latent and inconsistent, and because observed choices may be biased and distorted by noise (Evgeniou et al., 2005), we cast the problem in probabilistic terms, which accommodates all of these properties. Panel (a) in Figure 3.1 outlines the generative process underlying our **Gaussian process Trade-off Model (GTM)**. Reading panel (a) from right to left: users make observable choices C between presented alternatives

based on their latent evaluation of trade-offs t , denoted as $f_u(t)$. Evaluations are modeled as linear combinations of n_c behavioral characteristics f^c which individuals possess to different degrees, denoted by $\gamma_u^c \in [0, 1]$:

$$f_u(t) = \sum_{c=1}^{n_c} \gamma_u^c \cdot f^c(t) \quad \text{with} \quad \sum_c \gamma_u^c = 1 \quad (3.1)$$

Γ denotes the $n_U \times n_c$ matrix of all γ_u^c . For now, we assume that Γ is known, and focus on the problem of efficiently obtaining probabilistic estimates of the f^c .

The evaluations f^c are latent, uncertain quantities, and we aim to infer their distributions. A common modeling choice is to assume that the f^c are members of a parametric class, such as the class of linear functions $f^c(t) = w^{cT} \phi(t)$, and then infer the distribution of the parameters w^c . However, this choice has two important limitations. First, in the absence of prior knowledge about the shape of the distribution of the f^c , restricting the distribution to a particular parametric class entails the risk of excluding promising candidates. In futuristic smart grid applications, for example, little is known about the eventual choice behavior of users. Data-driven models have an advantage over parametric models in this setting, because they impose fewer up-front restrictions. Second, a fixed number of parameters w makes the model prone to **underfitting** or **overfitting** if the complexity of relationships present in the data does not match the number of parameters. These model complexity challenges then have to be overcome using additional tweaks, such as regularization or cross-validation (Bishop, 2006).

Gaussian processes (GPs; Rasmussen and Williams, 2006) are a non-parametric, data-driven alternative to traditional parametric models, and they overcome these limitations by design.³ GPs are stochastic processes that model function distributions directly (i.e., without the need for explicit parameters w) in Bayesian statistics, and they provide automatic, data-driven control of model complexity while remaining computationally tractable. Our aim is to estimate function distributions of the form $p(f^c|C, T, \Gamma)$ that quantify the probability that a particular set of trade-off evaluations f^c generated the observed choices C , given trade-offs T and user characteristics Γ .⁴ Starting from some prior distribution $p(f^c)$, this distribution

³Note, that the term *non-parametric* is a misnomer as non-parametric models may in fact employ parameters, such as the length-scale hyperparameters we introduce below. The distinguishing feature of such models is that the number of parameters automatically increases with the training data.

⁴To simplify the notation, we henceforth omit conditioning on T and Γ .

can be updated to incorporate observed choices C by using Bayes' theorem as follows:

$$\underbrace{p(f^c|C)}_{\text{Posterior}} = \frac{\overbrace{p(C|f^c)}^{\text{Likelihood}} \times \overbrace{p(f^c)}^{\text{Prior}}}{\underbrace{\int p(C|f^c) \times p(f^c) df^c}_{\text{Marginal Likelihood}}} = \frac{p(C|f^c) \times p(f^c)}{Z} \quad (3.2)$$

An initial **prior** distribution over trade-off evaluations $p(f^c)$ may be either uninformative and consider all possible f^c somewhat probable, or it may be informative and reflect a strong *a priori* belief about the f^c . An effective approach for deriving informative priors is through incremental updates: a posterior $p(f^c|C)$ derived from observations up to a certain time can form a prior that is subsequently updated with newly arriving choices. Incremental updates are also beneficial because they greatly simplify the computation, and because they enhance the online, real-time performance of *GTM*. The **likelihood** $p(C|f^c)$ in Equation (3.2) relates the latent f^c to the choice observations C .⁵ The denominator Z is a scalar called marginal likelihood which ensures that $p(f^c|C)$ is a valid probability distribution.

While the Bayesian update in Equation (3.2) is commonly used, it is essential to understand how it proceeds on distributions over functions (i.e., on Gaussian processes f). We therefore give a two-dimensional demonstration of this update in Figure 3.2. Note, that dimensionality here refers to the *number* of unique trade-offs t at which f is evaluated. The updating process begins with Panel (a), which illustrates the prior belief that those functions f are most probable for which the joint evaluation (f_{t_1}, f_{t_2}) is close to $(0,0)$ and positively correlated.⁶ The correlated prior reflects a belief that f_{t_1} and f_{t_2} should be similar, e.g., because t_1 and t_2 are similar. We could make the prior less informative and thereby less influential on the estimates by increasing its variance. The likelihood in Panel (b) illustrates the probability that a function with joint evaluation (f_{t_1}, f_{t_2}) gave rise to the observations $C = \{(u, t_1, y_1 = -1), (u, t_2, y_2 = +1)\}$. The likelihood is highest where f_{t_1} is negative (the user evaluates trade-off t_1 as unfavorable) and where f_{t_2} is simultaneously positive. These evaluations agree with the observed rejection of trade-off t_1 and the acceptance of trade-off t_2 . Panel (c) illustrates the posterior belief $p(f_{t_1}, f_{t_2}|C)$ about the latent trade-off evaluations f , and it reflects the update of the prior from panel (a) based on the observed choices. Importantly, the posterior shown in panel (c) offers two valuable estimates: it allows us to predict future (unobserved) choices, and it allows us to quantify the uncertainty inherent in our current estimate.

⁵Technically, the likelihood model establishes the relationship between the f_u and C . However, because Γ is fixed for now, the f_u are deterministic combinations of the f^c .

⁶Henceforth, we use the shorthand f_t to denote $f(t)$.




Figure 3.2. Example of a Bayesian Gaussian Process Update. Update of a two-dimensional Gaussian process evaluated at t_1 and t_2 with observations $C = \{(u, t_1, y_1 = -1), (u, t_2, y_2 = +1)\}$. Darker colors indicate higher probabilities. The prior is a bivariate Gaussian centered on $(t_1, t_2) = (0, 0)$. The likelihood assigns high joint probabilities to functions f that evaluate to small values at t_1 and high values at t_2 , i.e., that are compatible with the observed choices. The posterior reflects an update of the prior based on the observed choices.




Figure 3.3. Illustration of the Probit Likelihood Model. Probit likelihood model applied to a range of one-dimensional trade-offs t . The evaluation function $f(t)$ is fixed in the example. The bold line shows the Probit likelihood that assigns higher probabilities $p(y = +1|f_t)$ from the interval $[0, 1]$ (dotted lines, right axis) to trade-offs with more positive evaluations.

Towards probabilistic inference in *GTM*, we must specify its three fundamental components: the likelihood model $p(C|f^c)$, the Gaussian process prior $p(f^c)$, and a scalable method to efficiently compute the posterior $p(f^c|C)$ via Equation (3.2). In the remainder of this section, we will address the first two components. We develop an approach for efficient inference in Section 3.4.

The **likelihood model** relates the latent trade-off evaluations f^c to the observed choices C . For a single observed choice y given a trade-off t , the likelihood specifies the probability $p(y|f_t^c)$ that f_t^c was the true evaluation function that gave rise to y . The f^c capture users' characteristic evaluations of trade-offs, and higher values of f^c imply a higher probability that the user chooses the first alternative (i.e., a higher probability of the event $y = +1$). Consequently, the likelihood model must be a sigmoidal function, mapping evaluations f_t^c from a real value to the interval $[0, 1]$ of valid probabilities $p(y|f_t^c)$. The two most prominent candidates used for such mappings are the Probit and Logit functions (Train, 2003). Given that both functions can be computed efficiently and that no significant differences exist between them in terms of predictive accuracy (Rasmussen and Williams, 2006), we follow earlier work (Chu and Ghahramani, 2005, Houlsby et al., 2012) and use the Probit likelihood:

$$p(y|f_t^c) = \begin{cases} \Phi\left(\sum_c \gamma_u^c f_t^c\right) & \text{if } y = +1 \\ 1 - \Phi(f_{u,t}) = \Phi(-f_{u,t}) & \text{if } y = -1 \end{cases} \quad (3.3)$$

where Φ denotes the cumulative distribution function of the standard normal distribution.⁷

⁷In some models, the Probit likelihood also includes a noise variance term, $p(y|\cdot) = \Phi\left(\frac{f_{u,t}}{\sigma_n}\right)$. However, because

The shape of our likelihood model is illustrated in Figure 3.3. Assuming that multiple choices i are independent of each other, the joint likelihood of all choices is given by:

$$p(C|f^c) = \prod_i \Phi(y_i \cdot f_{u_i, t_i})$$

We are now ready to formulate our **GP prior** $p(f^c)$. GPs are sets of random variables, any finite subset of which has a joint multivariate Gaussian distribution (Rasmussen and Williams, 2006, MacKay, 1998). In our model, the random variables correspond to the values of the f^c at finite sets of trade-offs T . When conditioning on finite sets of trade-offs, inferences about function distributions are reduced to updating multivariate Gaussians based on the observed choices. In Figure 3.2, we illustrated this update for $|T| = 2$. In general, a GP-distributed process f can be written as: $f(t) \sim \mathcal{GP}(m(t), k(t, t'))$ where $m(t) = E[f(t)]$ denotes the mean function, and $k(t, t') = \text{Cov}[f_t, f_{t'}]$ denotes the covariance function that specifies how strongly evaluations of f at t and t' are correlated. In what follows, we set $m(t) = 0$, reflecting indifference in the absence of other information. We employ squared exponential covariances of the form:

$$k(t, t') = \prod_{d=1}^{d_T} \exp \left(- \frac{(t^{(d)} - t'^{(d)})^2}{2 \cdot l_d^2} \right) \quad (3.4)$$

Under the squared exponential covariance function, evaluations $f_t, f_{t'}$ of trade-offs t, t' become less correlated as the distance between them increases. This gives preference to smooth functions f , which reflects the intuition that people make similar choices when confronted with similar trade-offs. The product structure of Equation (3.4) corresponds to the assumption that each dimension of a trade-off contributes independently to the covariance, a property that will be crucial for efficient posterior inference. $\{l_d \mid d = 1, \dots, d_T\} =: \theta$ denotes the **length-scale hyperparameters** of the squared exponential covariance function. These length-scales characterize, in each dimension, the magnitude at which a trade-off becomes material to the users. Because the length-scales depend on the measurement of trade-offs in each dimension (e.g., dollars vs. cents), we will learn them from the data in Section 3.4.4.

Evaluating k at all pairs of observed trade-offs (t_1, t_2) yields the covariance (kernel) matrix K necessary for posterior inference. That the cost of many important operations on K grows cubically in the number of unique trade-offs presents naïve inference methods with scalability challenges. However, in Section 3.4 we show that the structure of our preference learning task can be exploited to reduce this cost substantially.

our trade-off evaluation interpretation of the $f_{u,t}$ is invariant under scaling, we set $\sigma_n^2 = 1$ without loss of generality.

Having specified both the likelihood model and the GP prior in Equation (3.2), we can now obtain a posterior distribution $p(f^c \mid C)$, reflecting our updated belief about the latent f^c . This belief can be used to predict users' choices regarding unobserved trade-offs t_* . Specifically, the probability of a user u choosing the first alternative when presented with trade-off t_* is given by:

$$p(y_* = +1 \mid t_*, C) = \int \Phi(f_{u,t_*}) p(f_{u,t_*} \mid C) df_{u,t_*} \quad (3.5)$$

$$= \Phi\left(\frac{y \cdot E[f_{u,t_*}]}{\sqrt{1 + \text{Var}[f_{u,t_*}]}}\right) \quad (3.6)$$

where Equation (3.6) holds if $p(f_{u,t_*})$ is Gaussian (Rasmussen and Williams, 2006). An important advantage of our Bayesian model relative to traditional discrete choice models is that predictions are derived from *infinitely many* evaluation functions f_{u,t_*} , all of which impact the final prediction proportional to their posterior probability $p(f_{u,t_*} \mid C)$. In contrast, a traditional discrete choice model bases its predictions only on a single “true” utility function from a parametric class. This single utility function is obtained, e.g., through Maximum Likelihood, a procedure which tends to overfit the data (see also Panel (b) in Figure 3.2). As we will see, the Bayesian approach results in substantial accuracy improvements.

3.4 Fast Bayesian Inference in GTM

Computational efficiency and scalability are a critical prerequisites for the adoption of preference models in practical applications, where the number of trade-offs and users can be large. In *GTM*, efficiently computing the posterior $p(f^c \mid C)$ requires addressing two key challenges. First, the cost of many important operations on the covariance matrix K grows cubically in the number of unique trade-offs, thereby rendering them infeasible in many interesting applications. Second, under the Probit likelihood, the posterior is neither Gaussian nor analytically tractable, and it is thus necessary to approximate it. In this section, we develop a novel inference method for *GTM* that brings about conceptual simplicity and that meets the critical scalability requirements for preference learning in real-world applications.

Our development proceeds as follows: In Section 3.4.1, we show how particular assumptions regarding the trade-offs T lead to structured Gaussian process models, where K has Kronecker structure, and where many operations can be performed efficiently as long as the dimensionality d_T is low. In Section 3.4.2, we combine structured GPs and other computational shortcuts with an approximate inference scheme known as Laplace's method (MacKay, 1996) to efficiently obtain the approximate posterior distribution $p(f^c \mid C) \approx$

$q(f^c|C)$ of characteristic trade-off evaluations. In Section 3.4.3, we show how the user characteristics Γ can be learned from the data, and in Section 3.4.4, we provide heuristics for learning the hyperparameters θ .

3.4.1 Structured Gaussian Processes

To address the cubic computational cost of several important operations on K , we propose an approach which exploits common characteristics of consumer choice settings. In particular, consumer and econometric research has established that information overload causes consumers to focus on relatively small subsets of attributes and values when choosing amongst alternatives, e.g., (Caussade et al., 2005, Hensher, 2006). We show that when alternatives can be represented by a small number of attributes and values, it is possible to obtain matrices K which are large, but on which important operations can be performed efficiently. Subsequently, we consider settings in which (1) the number of users, instances, and observed choices is large and naïve methods are therefore computationally infeasible; (2) trade-offs can be represented by a small number of attributes; and (3) each attribute has a small number of values, or can be discretized.

In the setting we consider, trade-offs can be arranged on a d_T -dimensional grid. Together with the product structure of the covariance function in Equation (3.4) this yields **Kronecker** covariance matrices as follows. Suppose that $T \subset \mathbb{R}^d$, and let T_d denote the set of unique values that occur on the d -th attribute in T . In our electricity tariffs example, trade-offs can be characterized by (1) price differences per kWh, and (2) differences in renewable sources, so that we might have the following unique trade-off values: $T_1 = \{-0.10, -0.09, \dots, 0.09, 0.10\}$ and $T_2 = \{-1, 0, 1\}$. Not all possible combinations of trade-offs are usually observed ($|T| < |T_1| \cdot |T_2| = 63$), and the covariance matrix $\tilde{K} = [k(t, t')]_{t, t' \in T}$ is therefore much smaller than 63×63 . The key notion in structured GPs is to replace \tilde{K} with a larger matrix of the form:

$$K = K_1 \otimes \dots \otimes K_d$$

where \otimes denotes the Kronecker product.⁸ The entries K_d hold the covariance contributions of the d -th dimension and they are generally much smaller than \tilde{K} (in our example, $K_1 \in \mathbb{R}^{21 \times 21}$ and $K_2 \in \mathbb{R}^{3 \times 3}$). The Kronecker matrix K , on the other hand, holds the covariances

⁸For two arbitrarily sized matrices A, B , the Kronecker product is defined as:

$$A \otimes B := \begin{bmatrix} a_{11}B & \dots & a_{1n}B \\ \vdots & \ddots & \vdots \\ a_{m1}B & \dots & a_{mn}B \end{bmatrix}$$

between all trade-offs in the cartesian product $\times_d T_d$, and it is thus much larger (in our example, $K \in \mathbb{R}^{63 \times 63}$).

The significant computational savings that the Kronecker structure of K allows us to achieve follow from the fact that, instead of explicitly generating and manipulating \tilde{K} , we can now operate on the smaller K_d . In this setting, several important matrix operations involving K can be performed efficiently. Most importantly:

- Matrix-vector products of the form Kb can be computed at a cost that is linear in the size of b , instead of the quadratic cost entailed by standard matrix-vector products. Similar benefits are obtained for products of the form $(K \circ K)b$ where \circ denotes the Hadamard (element-wise) matrix product.
- Eigendecompositions of the form $K = Q^T \Lambda Q$ can be computed from the Eigendecompositions of the K_d :

$$Q = \bigotimes_{d=1}^D Q_d \quad \Lambda = \bigotimes_{d=1}^D \Lambda_d$$

at cubic cost in the size of the largest K_d . In particular, this allows us to determine the Eigenvectors to the n_e largest Eigenvalues of K efficiently.

Kronecker products have additional computational advantages, and we refer the reader to (van Loan, 2000, Gilboa et al., 2013) for further references. In addition, note that all operations can be implemented by considering only the set of unique observed or predicted trade-offs. This reduces the actively considered region from the large space covered by K to a manageable superset of T .⁹

3.4.2 Learning Trade-off Evaluations

Let us now address the problem that the posterior $p(f^c|C)$ is analytically intractable under the Probit likelihood. Discrete choice models often employ sampling-based methods to approximate the posterior (Allenby and Rossi, 1998, Train, 2003). However, sampling is slow, especially for high-dimensional models based on GPs. Alternatives include Laplace's method, Expectation Propagation, and Variational Bayesian methods, all of which seek to approximate $p(f^c|C)$ with a similar distribution $q(f^c|C)$ that can be computed and represented efficiently (Bishop, 2006).

⁹Unobserved trade-offs can be modeled through infinite noise variances in Equation 3.3. The corresponding likelihood terms then evaluate to indifference ($p = 0.5$), and their derivatives to zero. These zero values lead to even sparser matrices W and L in Algorithms 1 and 2 below, and they can easily be exploited using standard sparse matrix operations.




Figure 3.4. Laplace Approximation $q(f^c|C)$ of the True Posterior $p(f^c|C)$. The solid line shows the approximation $q(f^c|C)$ of the true posterior $p(f^c|C)$ for a one-dimensional marginal distribution. The approximation is centered on the mode \hat{f}^c of the true posterior and its variance is matched to a second-order Taylor expansion of the true posterior at that point.

In this paper, we use Laplace’s method because it is computationally fast and conceptually simple – two of our key design objectives. Laplace’s method aims to approximate the true posterior p with a single Gaussian q , centered on the true posterior mode \hat{f}^c , and with a variance matching a second-order Taylor expansion of p at that point (see Figure 3.4). Approximating the posterior with a single multivariate Gaussian allows us to conveniently reuse it as prior in subsequent Bayesian updates which is important for online and active learning from user interactions (Saar-Tsechansky and Provost, 2004). A limitation of Laplace’s method is that the approximation becomes poor if the true posterior is strongly multi-modal or skewed. However, prior work has found no significant impact of this limitation in the preference learning context, e.g., (Chu and Ghahramani, 2005).

Our development of Laplace inference in *GTM* proceeds in two steps. First, we describe an efficient procedure for finding the posterior mode \hat{f} (Algorithm 1). We then describe how the posterior variance and predictions for new trade-offs t_* can be computed (Algorithm 2). Additional mathematical details are provided in Appendix B.1.

The mode \hat{f} of the posterior is the maximizer of the log posterior $\log p(f^c|C) \propto \log p(C|f^c) + \log p(f^c)$ (see Equation (3.2)) which can be found by setting the first derivative of $\log p(f^c|C)$ to zero and solving for f^c . Because the Probit likelihood is log concave, there exists a unique maximum \hat{f} , which we obtain iteratively by using the Newton-Raphson method (Press et al., 2007) with the update step:

$$f^{new} = (K^{-1} + W)^{-1} \underbrace{(Wf + \nabla \log p(C|f))}_b \quad (3.7)$$

$$= K(b - L(I + L^T KL)^{-1} L^T Kb) \quad (3.8)$$

until f converges. The matrix W in Equation (3.7) denotes the negative Hessian of the log likelihood, $W = -\nabla\nabla \log p(C|f^c)$, a sparse matrix consisting of $n_c \times n_c$ diagonal sub-matrices of size $n_T \times n_T$.¹⁰ The sparsity of W allows us to compute its Cholesky decomposition $W = LL^T$, which we use instead of W in Equation (3.8). This alternative representation eliminates the numerically unstable K^{-1} and the unwieldy inverse of the first factor in Equation (3.7). All matrices in Equation (3.8) are of size $(n_c n_T) \times (n_c n_T)$ and therefore usually large. However, L has at most $\frac{n_T n_c (n_c - 1)}{2}$ non-zero elements (less if not all possible trade-offs from T are observed), and we never have to generate K explicitly (see Section 3.4.1).

Algorithm 1 Laplace mode finding

```

1: function LAPLACEMODE(covariance matrix  $K$ , choices  $C$ , user characteristics  $\Gamma$ )
2:    $f = \mathbf{0}$ 
3:   repeat
4:      $W \leftarrow -\nabla\nabla \log p(C|f)$ 
5:      $L \leftarrow \text{CHOLESKY}(W)$ 
6:      $b \leftarrow Wf + \nabla \log p(C|f)$ 
7:      $a \leftarrow b - L(I + L^T K L)^{-1} L^T K b$  ▷ using conjugate gradients
8:      $f \leftarrow Ka$ 
9:   until  $f$  converges
10:  return posterior mode  $\hat{f}$ 
11: end function

```

Using Equation (3.8), we can efficiently compute the posterior mode by following the steps outlined in Algorithm 1. The operations in lines 6 through 8 are all matrix-vector operations which generate vectors as intermediate results. Importantly, because of K 's Kronecker structure, and because L consists only of diagonal sub-matrices, multiplications with K and L have linear time and space complexity. Rather than calculating the inverse in line 7 explicitly, we use conjugate gradients (Press et al., 2007) to solve the system $(I + L^T K L)x = L^T K b$ by repeatedly multiplying the parenthesized term with candidates for x , as in (Cunningham et al., 2008). This operation is fast because the required multiplications can be carried out in linear time.

We next compute the variance $V_q(f)$ of the approximate posterior q , which can be written as (Rasmussen and Williams, 2006):

$$V_q(f) = \text{diag}(K) - \text{diag}(KL(I + L^T K L)^{-1} L^T K) \quad (3.9)$$

The computations in Equation (3.9) involve full matrix operations, and are therefore more

¹⁰ W is computed using Equation (B.1) in Appendix B.1.1. There, we also give other computational details regarding our Probit likelihood.

expensive than the matrix-vector operations used for mode-finding. However, we can limit the computations to points of interest t_* only, which reduces the number of rows in K being considered. To further reduce the size of the involved matrices, we approximate K via a low-rank decomposition with exact diagonal given by:

$$K \approx QSQ^T + \Lambda, \text{ where } \Lambda = \text{diag}(K) - \text{diag}(QSQ^T) \quad (3.10)$$

The decomposition can be efficiently computed when K has Kronecker structure (see Section 3.4.1). The matrix S in Equation (3.10) is a diagonal matrix with the n_e largest Eigenvalues of K on its main diagonal. Q contains the corresponding Eigenvectors, and it has the same number of rows as K but only n_e columns. Λ is a diagonal matrix of the same size as K , making the low-rank approximation of K exact on the diagonal (Quiñonero-Candela and Rasmussen, 2005, Vanhatalo et al., 2010). The number of Eigenvalues n_e in the approximation is a user-defined input and it can be used to balance computing time against accuracy of the approximated posterior variance. As we will show below, even choices of small numbers of Eigenvalues n_e often yield posterior variances close to those obtained with the full matrix K . Under this low-rank approximation, Equation (3.9) can be re-written as:

$$\begin{aligned} V_q(f) &\approx \text{diag}(K) - \text{diag}(KL(I + L^T(QSQ^T + \Lambda)L)^{-1}L^TK) \\ &= \text{diag}(K) - \text{diag}(K\Pi K) + \text{diag}(K\Pi Q(\underbrace{S^{-1} + Q^T\Pi Q}_P)^{-1}Q^T\Pi K) \end{aligned} \quad (3.11)$$

where P is a small matrix of size $n_e \times n_e$, and where $\Pi = L(I + L^T\Lambda L)^{-1}L^T$ can be computed efficiently, because L is sparse and Λ is diagonal. Π itself is also sparse, consisting of $n_c \times n_c$ diagonal blocks like W . Because K has Kronecker structure, the first two terms in Equation (3.11) can be computed efficiently and without resorting to approximations. We address the computation of the third term next.

In Algorithm 2, we first calculate the Cholesky factor C of P (line 5), which we then use in solving¹¹ the system ΠQC^{-1} . The product V in line 6 is equivalent to n_e matrix-vector products with a Kronecker matrix and is computationally inexpensive when n_e is sufficiently small. In line 7, we exploit the symmetry of the third term in Equation (3.11), and the fact that only its diagonal is needed, to reduce calculations to an efficient element-wise product of the smaller V . Finally, in line 9, we use the posterior variances to calculate the predictive probabilities p_* at the trade-off points T_* using Equation (3.5).

Figure 3.5 illustrates the output of Algorithms 1 and 2 for the choices of a single user,

¹¹FORWARD SOLVE denotes the operation that solves the linear system $Ax = b$ for x . BACKWARD SOLVE similarly solves $xA = b$.

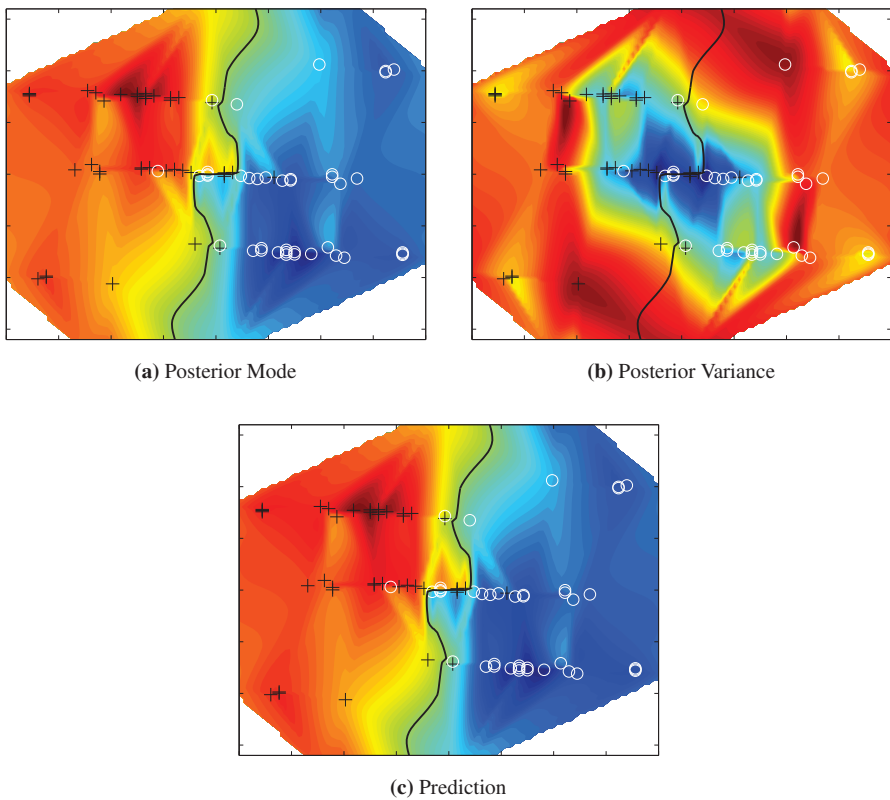


Figure 3.5. Outputs of Algorithms 1 and 2 for a Single User from a Popular Preference Benchmark Dataset. Observed choices are represented by black pluses (favorable trade-offs) and white circles (unfavorable trade-offs); red colors represent higher values; bold lines represent the boundaries where $E[f_u] = 0$ (Panels (a), (b)), or $p(y = +1|f) = 0.5$ (Panel (c)). The diamond-like shape of the plot results from mapping the four-dimensional trade-off space to two dimensions using its first two principal components.

Algorithm 2 Laplace prediction

```

1: function LAPLACEPREDICT(covariance matrix  $K$ , choices  $C$ , user characteristics  $\Gamma$ ,
   posterior mode  $\hat{f}$ , trade-offs  $T_*$ , # Eigenvalues  $n_e$ , Cholesky factor  $L$ )
2:
3:    $QSQ^T + \Lambda \leftarrow \text{LOWRANKAPPROXIMATION}(K, n_e)$  ▷ Equation (3.10)
4:    $\Pi \leftarrow L \cdot \text{FORWARDSOLVE}(I + L^T \Lambda L, L^T)$ 
5:    $C \leftarrow \text{CHOLESKY}(S^{-1} + Q^T \Pi Q)$ 
6:    $V = K_* \cdot \text{BACKWARDSOLVE}(\Pi Q, C)$ 
7:    $v_* \leftarrow \text{diag}(K_*) - (K \circ K) \cdot \text{diag}(\Pi)|_* + \sum_j [V \circ V]_{i,j}$  ▷ Equation (3.11)
8:
9:    $p_* \leftarrow \Phi\left(\frac{\hat{f}_*}{\sqrt{1+v_*}}\right)$  ▷ Equation (3.5)
10: return posterior variances  $v_*$ , predictive probabilities  $p_* = p(y = +1|f, T_*)$ 
11: end function

```

using data from a popular preference benchmark dataset (Kamishima and Akaho, 2009). Panel (a) shows the posterior mode $\hat{f}_u = E[f_u]$, which is expectedly high in regions of the trade-off space perceived as favorable, and low otherwise. The bold line indicates the zero boundary $\hat{f}_u = 0$, and it is sufficient as a predictor of future choices when predictive certainty estimates are not required. Importantly, it can be computed using only Algorithm 1 and is therefore very fast. However, the distinguishing feature of our probabilistic approach is the variance estimates shown in Panel (b). Here, the algorithm correctly identifies a region at the center of the panel where the decision boundary already follows a closely determined course to match earlier observations (blue coloring, low variance). If additional observations were to be acquired for the purpose of improving predictions, they should be located in the upper or lower regions of the decision boundary instead, where less evidence is presently available (red coloring, high variance). Panel (c) shows the combination of both outputs into predictive probabilities $p(y = +1|f)$. While the decision boundary at $p(y = +1|f) = 0.5$ is the same as the one in Panel (a), this panel additionally incorporates predictive variances by shrinking the predictive probabilities towards indifference ($p = 0.5$) in high-variance regions (see Equation (3.5)). Consequently, the corridor in which *GTM* is indifferent (yellow/green coloring, intermediate probabilities) is narrower in areas with extensive evidence from the data, and wider towards the edges of the panel.

3.4.3 Learning User Characteristics

So far we have assumed that the user characteristics $\Gamma = [\gamma_u^c]_{u,c}$ are known. We now relax this assumption by learning user characteristics from the data. Recall from Section 3.3 that γ_u^c denotes the fraction of user u 's behavior explained by characteristic c , that is, $f_{u,t} = \sum_c \gamma_u^c \cdot f_t^c$

with $\sum_c \gamma_u^c = 1$. In a full Bayesian treatment, we would consider Γ another latent quantity of interest, and infer its posterior distribution. Previous work has addressed similar challenges by either imposing a multinomial or a Dirichlet process prior on Γ (Houlsby et al., 2012, Abbasnejad et al., 2013). However, these approaches are conceptually complex and computationally expensive and therefore conflict with our key design objectives towards wide adoption in practice. Instead, we employ a simple and fast iterative scheme that calculates point estimates of Γ .

Algorithm 3 Learning user characteristics

```

1: function LEARNUSERCHARS(covariance matrix  $K$ , choices  $C$ , # characteristics  $n_c$ )
2:    $\Gamma \leftarrow$  random user characteristics
3:   repeat
4:      $E[\hat{f}^c] \leftarrow \text{LAPLACEMODE}(K, C, \Gamma)$ 
5:     for  $c = 1 : n_c$  do
6:        $\Gamma_{*,c} \leftarrow \prod_i \Phi(y_i \cdot E[\hat{f}_{t_i}^c]) \triangleright$  approx. to Eq. (3.5):  $p(C|f^c) = \Phi\left(\frac{y_i \cdot E[f_{u,t_i}]}{\sqrt{1 + \text{Var}[f_{u,t_i}]}}\right)$ 
7:     end for
8:      $\text{NORMALIZEROWS}(\Gamma)$ 
9:   until  $\Gamma$  converges
10: return user characteristics  $\Gamma$ 
11: end function

```

Our Algorithm 3 is an EM-type algorithm (Dempster et al., 1977) that alternates between optimizing \hat{f}^c given user characteristics Γ (line 4), and inferring new user characteristics Γ given \hat{f}^c (line 6). For each user, the resulting user characteristics are re-scaled so that they add to one in line 8. The process is repeated until Γ converges.

The Γ update in line 6 computes an approximation to the likelihood that characteristic c alone generated the observed choices. Each iteration of the surrounding loop calculates one column of the Γ matrix, corresponding to one characteristic. The intuition behind this update is that users should have higher γ_u^c values for those functions f^c that have generated their choices with a higher probability. The update is a reduced version of Equation (3.5). In particular, the reduced version does not include the denominator because it is expensive to compute while merely shrinking high-variance estimates to the point of uncertainty (see Section 3.4.2). Finally, while the number of user characteristics n_c has to be set manually, we find that, consistent with prior work, our method is insensitive to the choice of this parameter when it is not excessively small, e.g., (Houlsby et al., 2012).

3.4.4 Learning Hyperparameters

We finally consider learning the hyperparameters $\theta = \{l_d\}$. As in the case of Γ , a full Bayesian treatment of these parameters is prohibitively expensive. Prior work has often resorted to either gradient-based optimization of the marginal likelihood Z , e.g., (Chu and Ghahramani, 2005), or to heuristics, e.g., (Stachniss et al., 2009) to learn them from the data. In the experiments that follow, we employ a heuristic and set the length-scales to the median distance between trade-offs t . This approach has been found in prior work to be a computationally fast heuristic that yields consistently good empirical results (Houlsby et al., 2012).

3.5 Empirical Evaluation

To assess *GTM*'s predictive and performance characteristics, we compared it to three other recent GP models, and to the well-established Mixed Logit model. Our evaluation employs two real-world datasets from the literature, as well as a novel stated preference dataset collected specifically for this work. Our *GTM* implementation is based on a standard GP toolbox (Vanhatalo et al., 2013), and we make both implementation and data publicly available at <https://bitbucket.org/gtmanon/gtmanon>.

3.5.1 Datasets and Benchmark Methods

We ran *GTM* and its benchmarks on three preference datasets collected from human decision-makers. Recall, that a key motivation for this work is the need for computationally fast and scalable preference modeling techniques for use in contemporary applications with millions of users and many alternatives to choose from. A prime application of significant global importance is the modeling of electricity tariff choices of smart grid consumers. In future smart grids, tariffs may be revised frequently to reflect changes in the cost and availability of renewable energy, and therefore tariff choice is anticipated to become a near-continuous process in which both retailers and customers will rely on automated, data-driven decision support. The ability to predict and act on tariff choices quickly and with adequate accuracy is therefore an important challenge. To evaluate our approach in this setting, we employed real data on electricity tariffs offered in the Texas retail electricity markets, which is one of the most advanced retail markets in the United States, and which provides daily information on available tariffs (see <http://www.powertochoose.org>). Using the Amazon Mechanical Turk crowdsourcing platform, we acquired data on American participants' choices between pairs of tariffs, randomly drawn from a set of 261 tariffs offered in Austin, Texas in February

| Dataset | Instances | Users | Trade-Offs Stated Preferences | Orig. Dim. | Sel. Dim. | Grid Size |
|-----------|-----------|-------|----------------------------------|------------|-----------|-----------|
| Tariffs | 261 | 61 | 610 | 9 | 5 | 12288 |
| Cars | 10 | 53 | 2362 | 5 | 4 | 216 |
| Elections | 8 | 264 | 7392 | 20 | 8 | 30375 |

Table 3.2. Characteristics of the Datasets Used in this Study. **Instances**, **Users**, and **Trade-Offs** refer to the number of elements in X , U , and T , respectively. **Orig. Dim.** and **Sel. Dim.** refer to the number of trade-off dimensions (the size of each t) before and after feature selection. And **Grid Size** refers to the number $|T|$ of points on *GTM*’s grid.

2013. Appendix B.2 provides complete details on the data collection process for the Tariffs dataset, including an example tariff choice (Table B.1). The Tariffs dataset contains many instances (tariffs) but relatively few observed choices per user (see Table 3.2). The choices of different users likely correspond to different instances and are hence sparsely distributed and difficult to learn from. This is an important property of the Tariffs dataset which is common in many real world applications, but which is not reflected in prior benchmark datasets.

We complemented our evaluations with two prior benchmark datasets. Specifically, we used the **Cars** dataset collected by Abbasnejad et al. (2013), which contains stated preferences for automobile purchases, and on the **Elections** dataset prepared by Houlshby et al. (2012), which captures revealed voters’ preferences over eight political parties in 650 constituencies in the United Kingdom. A summary of the key characteristics of these datasets is presented in Table 3.2. The number of instances in all datasets is such that even computationally intensive methods can be trained in a reasonable time. However, we will closely examine the scalability properties of each method below, because they are key to the feasibility of the approaches and their subsequent adoption in practice.

We applied greedy forward feature selection to reduce the dimensionality of the datasets to a subset of predictive features without a significant loss in information content. The quality criterion guiding the feature selection was the predictive accuracy achieved by our most accurate benchmark method (Birlutiu et al., 2012). The resulting accuracies after feature selection were essentially the same as those reported on the complete feature set in Houlshby et al. (2012). Importantly, that the information content was maintained by the feature selection procedure reaffirms prior findings which *GTM* exploits, namely that only a subset of relevant dimensions effectively informs human choices. Using the most accurate benchmark as quality criterion also ensured that the procedure was not biased in favor of *GTM*. Finally, *GTM* was applied to versions of the datasets in which the continuous attributes were discretized to between 5 and 25 levels, with the objective of minimizing information loss while

keeping the resulting grid size manageable.¹² We employed the Natural Breaks algorithm of Jenks and Caspall (1971) to identify bins for discretization. Natural Breaks is a univariate variation of the k-means algorithm which selects bin boundaries such that within-bin variances are minimized while between-bin variances are maximized. All benchmark methods ran on the original, non-discretized datasets.

We compared *GTM* to four benchmarks: a computationally fast and conceptually simple method, **GP classification**, whose computational performance *GTM* aims to match; two methods with state-of-the-art predictive accuracies but higher computational and conceptual demands, **Hierarchical GP** (Birlutiu et al., 2012) and **Collaborative GP** (Houlsby et al., 2012); and the popular parametric **Mixed Logit** model. These benchmarks reflect different trade-offs between computational efficiency and scalability, accuracy, and conceptual simplicity.

Specifically, the **Hierarchical GP** model (Birlutiu et al., 2012) is based on a semi-parametric Bayesian formulation that builds on the framework proposed in Bradley and Terry (1952). The authors estimate the distribution of parameters w of the user-specific utility functions $f_u(x) = w_u^T \phi(x)$. These parameters are drawn from a hierarchical Gaussian prior that couples similar users (hence our choice of name for the approach). An EM-type algorithm is used for learning, which iteratively refines the parameters of the hierarchical prior. The Hierarchical GP model is conceptually simple and known to offer state-of-the-art accuracy. However, inference in it is computationally expensive.

The **Collaborative GP** method of Houlsby et al. (2012) also builds on Bradley and Terry (1952). Like *GTM*, it estimates latent characteristics $h^c(x^i, x^j)$ that quantify the utility derived from the trade-off between x^i and x^j , and that users possess to different degrees. A key distinction is that Collaborative GP operates on pairs of alternative instances (x^i, x^j) instead of the associated trade-offs $t = \tau(x^i, x^j)$, and it estimates instance utilities rather than trade-off evaluations. This makes inference in the model more demanding, and the authors employ a combination of Expectation Propagation and Variational Bayes to address this challenge. Their design choice yields comparable accuracies to the Hierarchical GP method at lower computational cost, but it makes inference in the model conceptually much more demanding.

In the limit of a single latent characteristic, Collaborative GP reduces to regular **GP classification** with a specific preference kernel (Rasmussen and Williams, 2006, Houlsby et al., 2012). Inference in this model is fast and conceptually simple, and as such it constitutes the strongest computational benchmark for *GTM*. While it is desirable to achieve GP classification's computational performance, its computational advantage comes at the cost of limited flexibility to adapt to complex data and at a significant loss in predictive accuracy relative to

¹²On our hardware, we restricted overall grid sizes to $10^4 - 10^5$ points.

Collaborative GP. Hence, it would be desirable to match GP classification’s computational scalability while maintaining high predictive accuracy.

The final benchmark is the well-established parametric **Mixed Logit** model which estimates $f_u^i = w_u x_u^i + \epsilon_u^i$ where ϵ_u^i is extreme-value distributed, and the w_u are drawn from a hierarchical prior. Like the other benchmarks, Mixed Logit accommodates random variations in taste among users. This makes inference more difficult than in the regular Logit model, a challenge which is addressed using a computationally expensive sampling procedure. Moreover, in comparison to the non-parametric models, Mixed Logit is significantly less flexible to adapt to the data. In our experiments we used the implementation of Train (2003).

In summary, in our empirical evaluations we examine *GTM*’s ability to match the scalability of GP classification. Simultaneously, the state-of-the-art predictive accuracies exhibited by the Hierarchical GP and Collaborative GP methods will allow us to assess the reduction in predictive accuracy that *GTM*’s computational benefits entail.

3.5.2 Model Scalability and Predictive Accuracy

We evaluated *GTM* and its benchmarks on 20 random splits of the data into training and test sets.¹³ In all *GTM* runs, the number of characteristics was set to $n_c = 10$, and we used the Eigenvectors corresponding to the $n_e = 100$ largest Eigenvalues in the sparse approximation.

Figure 3.6 shows the model training time required by *GTM* and its benchmarks on increasing training set sizes for each dataset. The learning curves show average running times over 20 random experiments, and the error bars reflect 90% confidence intervals. In all plots, training times correspond to running Algorithms 1 through 3. Our algorithm development aims to offer a new benchmark for preference learning in real-world applications with large numbers of users and choices. Fast and *consistent* training times are critical in these applications, because new observations must quickly be reflected in the preference model, and in the decision-making processes that build on it. Training times must therefore scale gracefully with the number of training choices.

As shown in Figure 3.6, *GTM* achieves two important goals. First, it trains significantly faster than the Hierarchical GP, Collaborative GP, or Mixed Logit models. Second, *GTM*’s training times barely increase with the size of the training set. *GTM*’s training efficiency matches that of the relatively simple GP classification model. However, as we will see below, the simplicity offered by GP classification limits its flexibility to adapt to complexities in the data, which undermines its predictive accuracy under some conditions and results in

¹³i.e., repeated-sampling cross-validation with varying training sizes. The horizontal axis in the following figures indicates the different training/test data splits.

inconsistent performance.

GTM's fast training times and superior scalability as a function of the number of training observations follow directly from our design which employs Kronecker-structured covariance matrices. Importantly, once a given grid size is set for inference, new observations only entail modest additional training time through additional likelihood terms. In contrast, in other methods, additional observations increase the size of covariance matrices which undermines their scalability. The added (typically cubic) costs of matrix operations are the primary factor for their poor scalability in the number of observations.

While *GTM* offers computational efficiency and scalability comparable to that of the simple GP classification approach, we now explore whether this efficiency entails the same inconsistent predictive accuracy exhibited by GP classification. Figure 3.7 presents the predictive accuracies (captured by the proportion of correctly predicted test choices) corresponding to the training times in Figure 3.6. Interestingly, we find that *GTM*'s superior speed and scalability do not result in lower accuracy on the Tariffs dataset. Actually, *GTM* instead exhibits higher predictive accuracy than all other methods. It is useful to recall here that, similar to many real-world choice situations, the Tariffs dataset contains many instances but relatively few observations (see Table 3.2). We conjecture that estimating the f^c , our focus in the development of *GTM*, is more critical in such cases than determining user characteristics Γ . Furthermore, *GTM* yields predictive accuracies that are comparable to the most accurate and computational intensive methods on the other datasets, and that are consistently good across domains. By comparison, GP classification's high training speed and scalability come at the cost of inconsistent predictive accuracies. GP classification yields particularly poor predictions on the Elections dataset, and it is unable to benefit significantly from additional training data.

The Hierarchical and Collaborative GP methods deliver almost perfect predictions on the Cars and Elections datasets, yet, as shown in Figure 3.6, at the cost of slow training times and poor scalability in the number of observations. Key to this discussion is that all non-parametric methods, including *GTM*, can exploit additional training observations to automatically adapt the model to the inherent complexity in the data. Having additional training data allows these methods to capture more predictive structure in the data, as reflected by the inclining accuracy curves (see, in particular, Figure 3.7(b)). In sharp contrast, the parametric Mixed Logit fails to benefit from additional data because its fixed number of parameters underfits larger training sets. A related effect can be observed in GP classification's performance on the Elections dataset. On this revealed preference dataset, the sophisticated GP methods benefit substantially from additional training data early on in the learning curve as reflected by the rate of performance improvement. As shown, once a representative train-

ing sample is available, these methods are *able* to exploit more observations to capture the heterogeneity in the data. GP classification benefits somewhat from additional training data, but to a lesser degree. While its single latent characteristic yields a significant speed-up in computation, it also undermines its flexibility to capture all the heterogeneity inherent in the Elections data. In contrast, *GTM* shares the capacity of other non-parametric methods to adapt to the complexity in the data and to improve its predictive performance with more observations. However, by design, and unlike other non-parametric methods, *GTM* also scales gracefully as more data becomes available.

In summary, *GTM* strikes a new balance which combines the scalability of GP classification with a conceptually simple approach, and at the same time retains the modeling flexibility and expressiveness of more complicated and expensive non-parametric GP approaches. *GTM* is highly scalable with respect to the number of users and observations in consumer choice setting. It also offers both good and consistent predictive performance which is particularly pronounced in sparse domains commonly encountered in preference settings with a large number of users and instances. *GTM*'s predictive accuracy closely follows that of its most accurate benchmarks, and it is significantly more accurate than Mixed Logit. Overall, its scalability along with good and consistent predictive performance establish *GTM* as a new benchmark that is a likely method of choice in large-scale applications involving many users and observations, such as smart grid and healthcare choice modeling.

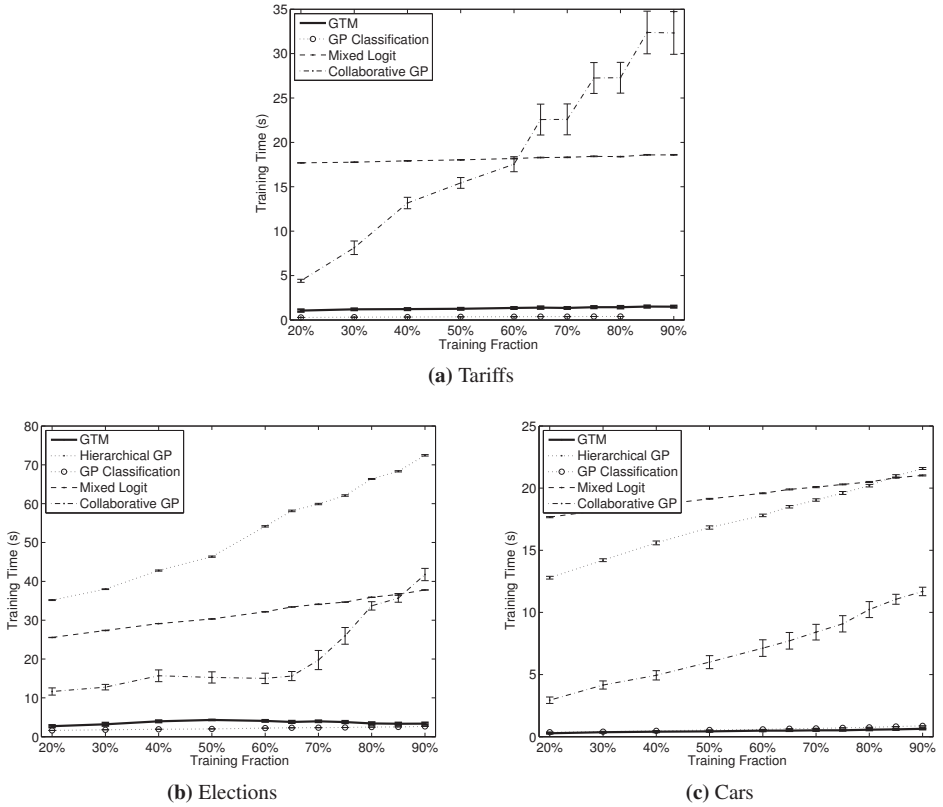


Figure 3.6. Training Time. Training times for *GTM* and its benchmarks in seconds. The horizontal axis indicates the fraction of data used for training the model, all remaining data is used for testing. Error bars show 90% confidence intervals based on 20 repetitions with different random splits between training and test data. The Hierarchical GP implementation available to us failed to predict on the Tariffs dataset due to numerical errors.

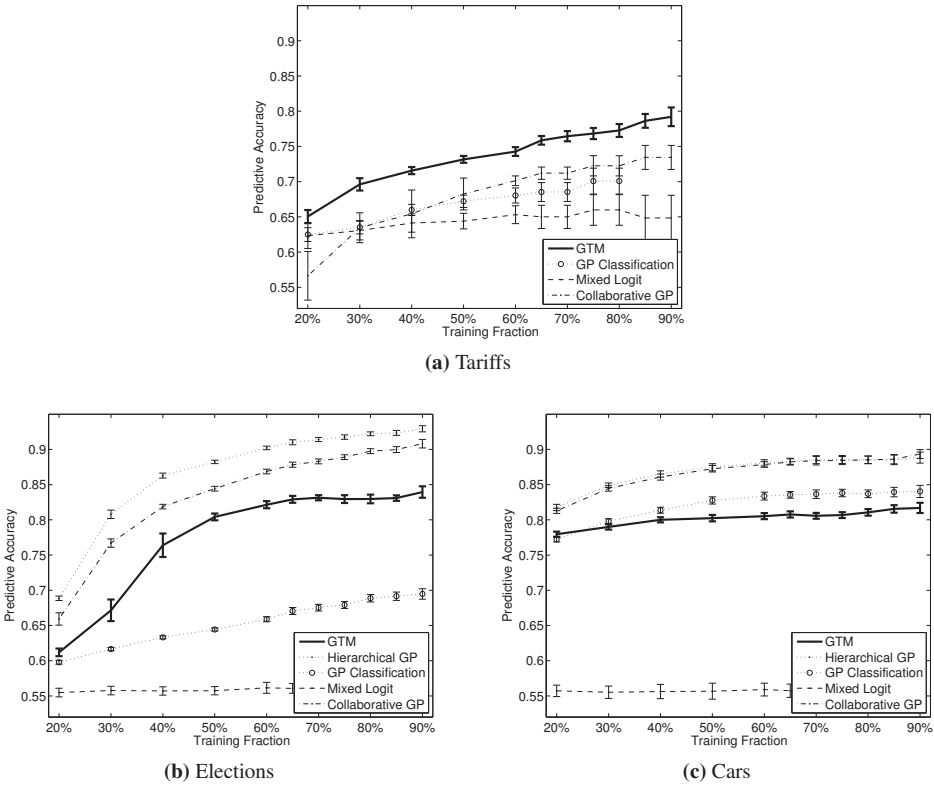


Figure 3.7. Predictive Accuracy. Accuracy of predictions on held-out test data corresponding to the training time measurements in Figure 3.6. The horizontal axis indicates the fraction of data used for training the model, all remaining data is used for testing. The vertical axis indicates the fraction of correctly predicted choices on held-out test data.

3.5.3 Dimensionality Characteristics

Our design choices for *GTM* aim to produce state-of-the-art scalability in the number of observations to accommodate real-world applications with millions of observed choices. Our solution relies on prior findings that human choices are determined by a small number of dimensions. Based on these findings, *GTM* has been designed to provide superior scalability for learning and inference when trade-offs are characterized by a small number of dimensions. Our experiments also demonstrate that dimensionality reduction incurs only a modest loss in predictive accuracy compared to prior results. The trade-off inherent in *GTM*'s ability to offer state-of-the-art scalability in the number of observations and consistently good predictive performance, is typical of structured GP methods: *GTM* is fast and highly scalable with respect to the number of observations for low-dimensional settings, but it is unsuitable for use in higher dimensions as this yields exponential growth in its grid size.

To demonstrate the implications of this trade-off, we studied the performance of *GTM* and the two fastest benchmarks, GP classification and Collaborative GP, on synthetic choice datasets for which we can directly control the number of users (n_U) and dimensions (d_X). Specifically, for each user, we randomly constructed a utility plane in a d_X -dimensional instance space from which we read utilities for $n_X = 15$ randomly drawn instances. These instance utilities were distorted with Gaussian noise, and then used to compute each user's choices between all $\frac{15 \cdot 14}{2} = 105$ instance pairs. Eighty percent of these choices were used to train the model while the remaining 20% were held out for model evaluation. Note, that our synthetic generation procedure closely follows the key assumption made by the two benchmark methods, namely that users make choices based on their predetermined, latent utility functions. Our synthetic generation procedure should therefore work in the favor of these methods.

Figure 3.8 shows the resulting training times for several dimensionalities (panels) and levels of discretization (three *GTM* lines per plot). *GTM*'s computational costs are dominated by the fixed cost associated with a given grid size. In particular, because *GTM*'s grid grows exponentially in the number of dimensions, this fixed cost outgrows the variable cost of other methods as the data's dimensionality increases (see Panel (c) for 9 dimensions and 7 levels). At the same time, as shown in Figure 3.8, *GTM*'s training curves are relatively flat, which means that it scales better for large numbers of users and choices in the consumer choice settings for which it is designed.

3.5.4 Sparse Approximation Quality

Another parameter affecting the overall computational cost of *GTM* is the number of Eigenvectors n_e used in the sparse approximation in Algorithm 2. In our experiments, we used $n_e = 100$ throughout, and we now illustrate that *GTM*'s output is relatively unaffected by this choice as long as n_e is not excessively low. Note, that n_e has no bearing on predictive accuracy as the sparse approximation is only used in the posterior variance computation. The posterior mode, and therefore also the iterative procedure for learning user characteristics Γ (Algorithm 3), are unaffected by n_e .

Figure 3.9 depicts the posterior variance for the first user from a popular preference benchmark dataset, and for varying numbers of Eigenvectors. Note, that the general shape of the posterior variance is similar in all three panels, which indicates that our sparse Laplace approach delivers reasonable results starting from small n_e values. Differences between panels are primarily limited to the step from $n_e = 10$ (Panel (a)) to $n_e = 100$ (Panel (b)). In Panel (b), the low-variance area at the center of the panel is noticeably larger than in Panel (a). Surrounding areas similarly shift to lower variances. The subsequent step to $n_e = 1000$ (Panel (c)) entails almost no further change in posterior variance.

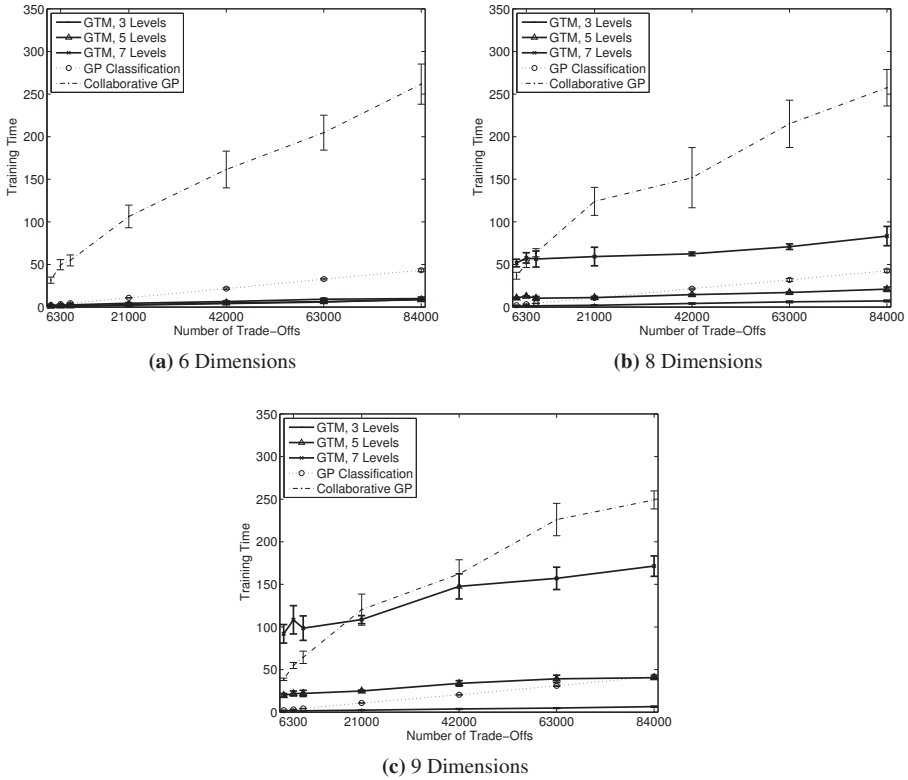


Figure 3.8. Scalability in the Number of Dimensions. Training times for *GTM* at various levels of discretization, and its two fastest benchmarks, GP Classification and Collaborative GP. Experiments are based of synthetic data of various dimensionality. Half of the dimensions were assumed to be binary, the other half continuous. Continuous dimensions were discretized to the indicated number of levels for *GTM* only. Error bars show 90% confidence intervals based on ten repetitions with different random splits between training and test data.

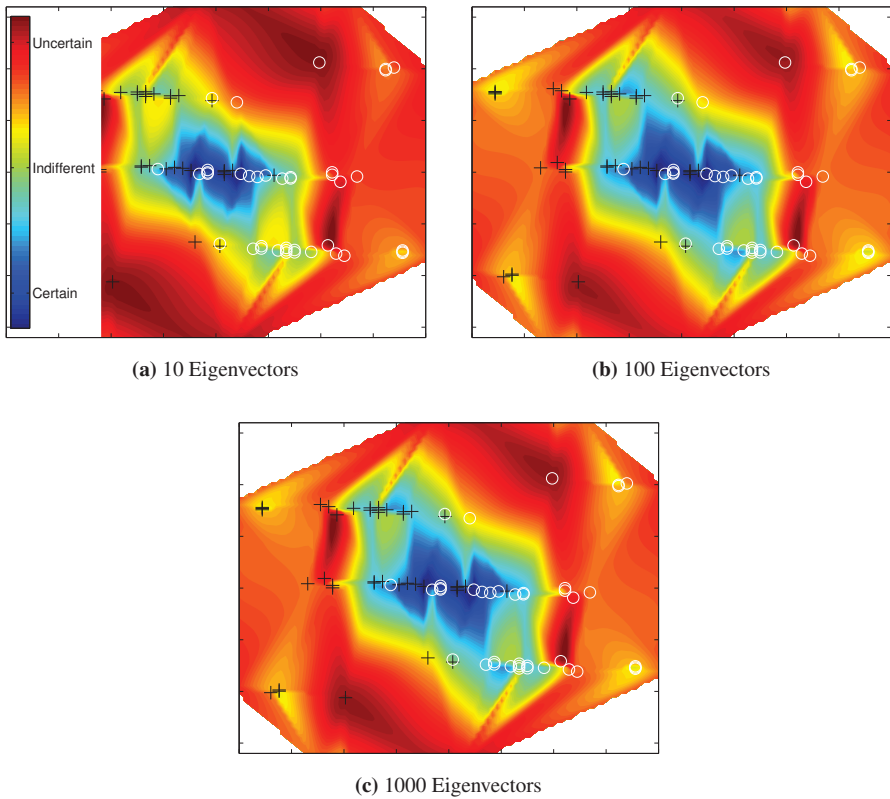


Figure 3.9. Posterior Variance for Different Numbers of Eigenvectors (n_e) in the Low-Rank Approximation. Panels show increasingly finer posterior variance estimates for a single user from a popular preference benchmark dataset. Between panels (a) and (b), the low-variance area in the center expands, and the adjacent regions shift towards lower variance, reflecting the better estimate. The addition of more Eigenvalues in Panel (c) has no noticeable effect, however.

3.6 Discussion and Conclusions

The *GTM* preference model we develop here is a new benchmark that strikes a novel balance between computational scalability and predictive performance, and it is designed to facilitate the use of preference models for autonomous decision-making in consumer choice settings. *GTM* provides state-of-the-art speed and scalability which allows it to accommodate data on millions of users and observations. These properties are particularly critical in important emerging applications, including smart electric grids and complex B2B marketplaces where preference models must be learned in real-time, based on a large number of users and observations. *GTM* provides principled probabilistic uncertainty estimates that are fundamental for automated, data-driven decisions. Finally, its conceptual simplicity makes *GTM* easier to understand and validate for practitioners than existing state-of-the-art, non-parametric choice models.

These advantages are possible by exploiting common characteristics of consumer choice settings. Specifically, *GTM* is designed for settings in which the number of users, instances, and observed choices is very large and thus excellent scalability is critical. In addition, given users approximate when evaluating alternatives to avoid information overload (Caussade et al., 2005), *GTM* exploits settings in which trade-offs can be captured by a small number of attributes with a few levels for each attribute. In our empirical evaluations we find that exploiting these properties allows *GTM* to offer order-of-magnitude performance improvements, making it possible to deploy preference modeling in a wide variety of contemporary consumer choice domains. Our empirical evaluations also show that the computational benefits entail only a modest reduction in predictive accuracy. *GTM* is significantly more accurate than the traditional Mixed Logit model, and its predictive performance compares favorably with conceptually and computationally more demanding non-parametric GP approaches that scale poorly and are infeasible in applications with millions of users.

We also show that for settings where both the dimensionality and the number of observations is high, standard GP classification provides similarly fast predictions as *GTM* in lower-dimensional settings. However, this capability comes at the cost of inconsistent predictive accuracies because GP classification is limited in capturing complex patterns. Overall, *GTM*'s state-of-the-art scalability and its consistently good predictive performance has not been possible with existing approaches, and it thus constitutes a new benchmark for preference modeling with a large number of users, instances, and observed choices.

The focus of the research we present here has been on modeling fast learning of probabilistic trade-off evaluations f^c that characterize segments of the user population. We solved the related problem of learning what combination of these evaluations describes each user through a simple, yet effective iterative scheme. We find that existing alternatives to this




Figure 3.10. Summary of Empirical Results. Darker colors indicate higher predictive accuracy. *GTM* provides high predictive accuracy and higher scalability than existing methods for consumer choice settings with few dimensions and attribute levels. In higher-dimensional settings, the Hierarchical and Collaborative GP models are more efficient, but their scalability with respect to the number of users and choices is limited. GP classification scales to high dimensions and large numbers of observations, but its predictive accuracy is inconsistent across datasets due to its limited expressive power.

simple iterative scheme entail significantly higher computational and conceptual costs, making them impractical for the settings we consider. It would be valuable for future work to explore alternatives that learn the number of characteristics n_c from the data at a reasonable cost.

We also focused on learning from pairwise choices of the form “User u prefers alternative a to alternative b ,” which are objective and cognitively less demanding for the user, but also more difficult to learn from than ratings. However, unlike some existing approaches, the natural separation between model and observations inherent in Bayesian modeling allows us to adapt *GTM* to learn from a variety of other data types in addition to pairwise choices. In particular, Jensen and Nielsen (2014) provide likelihood models for ordinal ratings that are compatible with the framework underlying *GTM*, and that would allow *GTM* to learn from heterogeneous observations containing pairwise choices and ratings simultaneously. This flexibility of our framework further enhances its applied properties.

In machine learning terms, *GTM* corresponds to a collaborative classification approach for preference learning, a common interpretation of pairwise choice settings. Importantly, the contributions presented here towards efficient, scalable, and conceptually simple inference also generalize to other important classification problems such as those arising in credit scoring, quality assurance, and other impactful business challenges. As such, *GTM* contributes to the foundations of next-generation information systems in a broad range of domains, where its reliable and consistent computational and predictive performance make it suitable for autonomous decision-making on users’ behalf.

Chapter 4

Competitive Benchmarking of Electricity Brokers

4.1 Introduction

In a significant change to its strategic research policy, the European Commission announced in 2010 that it would shift large parts of its research resources to addressing societal challenges in Europe and beyond (European Commission, 2011). Instead of working on projects that are neatly delineated by disciplinary borders, researchers are increasingly encouraged to work on solutions to overarching societal challenges such as health and demographic change, climate action and resource efficiency, and clean and efficient energy. This policy shift is not just a formal departure from an earlier emphasis on individual research areas or a peculiar European phenomenon. It is a symptom of the deeper need for comprehensive solutions to the complex challenges of our time, and it is a testament to the fact that we need more than a loose collection of clever technological advances to address these challenges.

Consider the challenge of supplying clean and efficient energy at societal scale. In the United States, more than 40% of all energy is supplied through the electric grid, a threefold increase over the 10% supplied in electric form in 1940 (US Energy Information Administration, 2013a,b). Electrification and its counterpart digitization have spurred substantial innovation and economic growth, as witnessed by projected global information technology (IT) investments of around four trillion dollars annually over the next five years (Gartner Group, 2013). But the boon of electrification also has its downsides.

Most American electricity is generated in large, central power plants from fossil (68%) or nuclear (20%) fuels that have been heavily criticized for their environmental impact (US

Energy Information Administration, 2013a). The average US power plant converts only one third of its primary fuel into usable electricity, while 6% of generated electricity is lost to the aging power lines that connect generators with consumers (US Department of Energy, 2003). Required investments for maintaining and updating these lines have been held back by budget cuts and lengthy approval processes, and the consequences of these delays are starting to show as disruptions like the 2003 Northeastern blackout are leaving millions of households and businesses without power (Fox-Penner, 2005).

These challenges are even greater beyond the highly developed economies. The second most populous country in the world, India generates close to 80% of its electricity from fossil fuels, more than a fifth of which is lost to its unreliable network infrastructure (Worldbank, 2013). In 2012, the Indian electricity network became the scene of the worst blackout in power system history when a cascading fault left more than 600 million people without electricity and triggered breakdowns in water and health services (Romero, 2012). At 5.2% of American per capita generation, India's power system may be small in relative terms (Worldbank, 2013), but its influence on the country's economic development is significant and its problems will continue to exacerbate as Indian standards of living advance.

Generation from renewable or decentralized sources could alleviate these problems, but they are difficult to integrate when consumers have little information about their electricity usage and few incentives to invest in smart appliances that adapt to the changing availability of wind and sun. A less predictable supply of clean electricity must be counteracted by more intelligently managed demand, but the right combination of technology, information, and policies to control the intricate interplay between individual behaviors and physical constraints remains elusive. The only universally accepted fact to date is that, to tackle a challenge of this scale and scope, an interdisciplinary effort is required that bundles the ingenuity of some of the sharpest minds in academia, industry, and policy.

Societal challenges like sustainable energy provisioning are "wicked problems" (Rittel and Webber, 1973), that arise from complex socio-technical systems where numerous social, technical, economic, and political factors interact. The overall behavior of such systems cannot be explained by considering each of its parts in isolation which makes it difficult to design targeted interventions that correct perceived misbehaviors of the system (Kling, 2007). Worse yet, even where promising interventions are known, the prohibitive cost of potential social negatives makes it impossible to thoroughly evaluate candidate interventions realistically and at scale. To continue our example, even though the sustainable energy challenge is, to all appearances, an information and coordination problem, legislators are understandably wary of funding large-scale IT deployments without reliable guidance on their long-run effects.

Providing such guidance should be familiar ground for the Information Systems (IS) discipline with its rich tradition of studying and resolving socio-technical challenges for which “solutions cannot be deduced from scientific principles alone” (Hevner and Chatterjee, 2010). Events like electrical blackouts or recent financial market flash crashes have left the public wondering whether “we may be becoming critically dependent on [large-scale IT systems] that we simply do not understand” (Cliff and Northrop, 2012). But even though the IS discipline seems well-positioned to engage in these debates, its impact on resolving societal challenges has remained limited (Straub and Ang, 2011, Lucas Jr et al., 2013). Although there are various reasons for this, we consider two reasons to be particularly important.

Research Method Scalability: IS research has historically favored the individual, group, organization, and market levels of inquiry over the societal level (Sidorova et al., 2006). We suspect that the limited scalability of the single-investigator model (Ibid.) of IS research plays a significant role in this.

Insights and Solutions: Decision-makers in politics and industry are increasingly looking for solutions in addition to mere insights (Aken, 2004). Researchers must expand their vision to “inventing new systems that address information needs not covered by current systems. [They] must not only be observers and historians of technology, [but] make technological contributions” (Nunamaker and Briggs, 2012).

One of the many IS innovations needed in the sustainable energy context are systems that broker spontaneous, self-sufficient coalitions of appliances, solar-panels, and electric vehicles to better utilize local infrastructure without inconveniencing users.¹ But delivering this and other innovations despite the complexity that governs the underlying system requires research methods that scale with the challenge, and that deliver solutions in addition to insights. We aim to make three contributions to this end:

First, we characterize the difficulties that wicked problems of societal scale pose to IS researchers. We contend that several obstacles limit the ability of current research methods to tackle problems of essential complexity that are large in scale and scope, that are currently unrealized, that progress at a rapid pace, and for which the social costs of erroneous interventions are prohibitive.

Second, we propose **Competitive Benchmarking**, a novel IS research method that builds on the competitive research approach pioneered by the Trading Agents community (Greenwald and Stone, 2001, Wellman, 2011, Ketter and Symeonidis, 2012), and that specifically

¹This is a special case of a more general idea known as Microgrids or Virtual Power Plants (Pudjianto et al., 2008).

addresses these obstacles. Our method emphasizes the importance of rich problem representations that are jointly developed among stakeholders and researchers, and it leads to actionable research results complete with comprehensive supporting data. Competitive Benchmarking supports both behavioral IS research (insights) and design science research (solutions).

Third, we evaluate the efficacy of Competitive Benchmarking using a case study on the **Power Trading Agent Competition** for research on sustainable energy systems (Power TAC, Ketter et al. 2013a). Power TAC challenges researchers to design Brokers, a novel type of information system that can play a pivotal role in modern coordination mechanisms for sustainable energy systems specifically, and other smart market environments more generally (Bichler et al., 2010).

To date, Power TAC has brought together more than a dozen research groups from various academic disciplines, and stakeholders from utilities to customer lobby groups to competitively design, evaluate, and improve Brokers. We conclude this article with preliminary empirical evidence of the benefits of CB as a research methods, and evidence of the benefits that competitively designed Brokers have to offer towards resolving the “grand challenge” of providing affordable, reliable, and sustainable energy for the twenty-first century (Masoud Amin and Wollenberg, 2005).

4.2 Information Systems Research for Societal Challenges

We set the scene for Competitive Benchmarking by first considering the difficulties that societal challenges pose to IS researchers. Two fundamental types of scientific inquiry can be distinguished in the IS discipline, both of which are important in resolving societal challenges: behavioral research and design science research (March and Smith, 1995, Walls et al., 1992). The research framework of Hevner et al. (2004) depicted in Figure 4.1 illustrates the interaction between the two.

An IS research effort might start with the realization that IT can improve the effectiveness or efficiency of a particular socio-technical system, such as an organization’s use of IT, or that of a whole society ❶. If the goal of the research effort is to describe or explain phenomena occurring within the system, researchers develop and justify new descriptive or explanatory theories, whereas if the goal is to improve the system, they build and evaluate artifacts and corresponding prescriptive design theories ❷ – ❹. The outcomes of these efforts are both applied to the original system ❺, and added to the scientific knowledge base for future use ❻.

Note that both types of scientific inquiry are interrelated. Descriptive and explanatory theories provide the understanding needed for designing effective artifacts, whereas artifacts

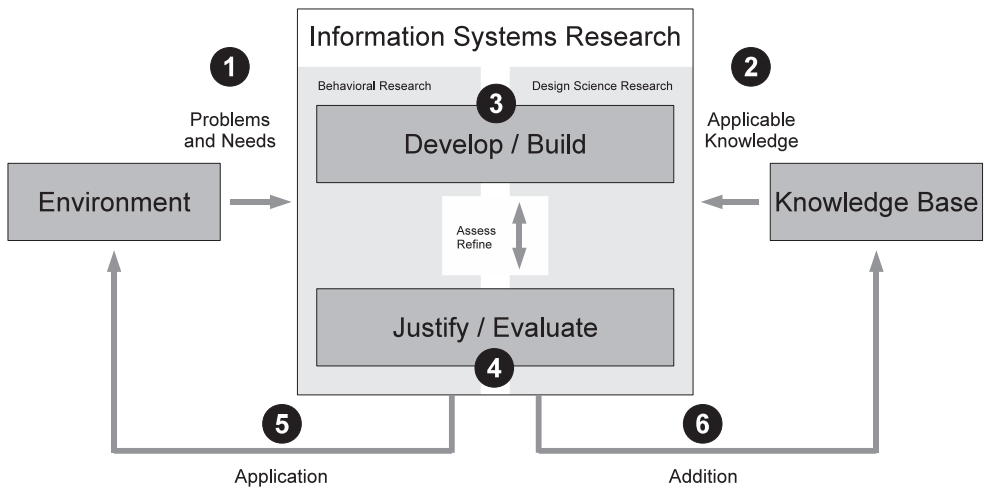


Figure 4.1. IS research encompasses behavioral and design science research. While the former seeks truth in the form of justified theories, the latter seeks to create utility through designed artifacts. The depicted framework is adapted from Hevner et al. (2004). Circled numbers are referenced from the main text.

embedded in context are the subject of new theories. In the remainder of this article, we will illustrate many of our arguments using design science examples, which, in our opinion, holds the greatest need and the greatest opportunity for advancing the impact of IS on societal challenges. But our arguments hold true for behavioral research as well, and we will highlight several such instances below.

This general research framework applies to societal challenges as well as to challenges of smaller scale. However, a number of issues arise in each step of the framework when applying it at the societal level. We discuss these issues below (see Table 4.1 for a summary).

Defining Problems and Needs ①

Societal challenges exceed the capacity of individual research groups to **interact with all stakeholders** to build and maintain an understanding of an unfolding challenge (Arias et al., 2000). For example, a research group attempting to design IT-based interventions to climate change would have to interact with meteorologists, geologists, politicians, chemists, economists, sociologists, and many other stakeholders to develop an understanding of climate change and its expected societal impact. But even if time and resources were unlimited, societal challenges like climate change would **defy the comprehensive formalization** of the challenge itself and the objectives for possible interventions. This is a direct consequence

| Characteristic | Methodological Challenges |
|--|--|
| Cost of Social Negatives: Failures of real-world interventions, even at small scale, entail prohibitive costs | <ul style="list-style-type: none"> • Reduce social negatives through high external validity evaluations • Produce rigorous prescriptive design theories, not singular designs |
| Unrealized Challenges: Many challenges are only anticipated, solutions must be preemptive | <ul style="list-style-type: none"> • Produce solutions in addition to descriptive and explanatory insights • Demonstrate viability of candidate interventions for expensive real-world evaluation |
| Rapid Pace: Real world progresses quickly and unpredictably | <ul style="list-style-type: none"> • Avoid wasteful duplication in developing a joint understanding of the challenge • Maintain an up-to-date understanding of the challenge • Find the right balance between rigor and relevance • Benchmark alternative interventions swiftly • Disseminate results in a timely manner |
| Scale and Scope: Societal challenges have vastly broader scales and scopes (societal, global) than most traditional IS research (individual, group, organization, market) | <ul style="list-style-type: none"> • Interact with all stakeholders • Understand the problem and find solutions, effectively coordinate many smaller research groups • Evaluate candidate interventions swiftly, rigorously, and with high external validity |
| Essential Complexity: Increasing use of IS, (smart) markets, and other social forms of organization creates essential complexity | <ul style="list-style-type: none"> • Explore broad solution space • Produce comparable artifacts based on a shared paradigm • Comprehensively formalize the problem and solution quality criteria, quickly converge on a paradigm • Evaluate strategic interactions of evolving candidate artifacts, relate locally optimized artifact performance to overall system performance, deal with asymmetrically dominated options |

Table 4.1. Summary of key characteristics that societal challenges pose to IS researchers.

of the essential complexity governing the systems that these challenges emerge from (von Hayek, 1989). In the climate change example, an intervention might aim to protect biodiversity, mitigate the short-run impact on the global food supply, or maintain economic growth. Each of these objectives gives rise to a different set of interventions and to a different delineation of the challenge. In other words, the definition of the challenge, the vocabulary used to describe it, and the questions that researchers ask about it all become a crucial part of the challenge itself (Rittel and Webber, 1973).

Two conventional responses to these issues have been to either work on a small sub-

set of the challenge, or to establish large, centrally composed and hierarchically organized research consortia (Hey et al., 2009). But by focussing on small subproblems, researchers ignore essential facets of the challenge, create candidate interventions that are incomparable to interventions for adjacent subproblems, and ignore important system-level consequences. And centrally composed and hierarchically organized research consortia forgo the opportunity of leveraging the diversity of various research groups for understanding the problem from a wide range of angles. Large consortia also tend to move more slowly than the rapidly evolving challenges they aim to address (Moss et al., 2010). Unsurprisingly therefore, practitioners find “science [to be] lagging behind the commercial world in the ability to infer meaning from data and take action based on that meaning” (Hey et al., 2009).

We argue that methodological advances are needed to support interdisciplinary communities of stakeholders and researchers in jointly developing (1) problem definitions and models of societal challenges, (2) shared vocabularies, and (3) lists of important research questions. Loosely following Kuhn (1996), we refer to this triplet as a scientific **paradigm**. Any method fit for this purpose must effectively distribute the limited capacity of individual research groups by facilitating a separation of concerns between them. It must avoid the wasteful duplication inherent in defining small, disparate sub-challenges, and it must be able to respond quickly to evolutions in the current understanding of the challenge. Richer, more malleable problem representations beyond printed articles are needed that can be created and updated with ease. The paradigm itself must become a device through which researchers and stakeholders jointly coordinate their efforts, and new incentives are needed for researchers to invest in its development.

Using the Knowledge Base and Building Artifacts ② ③

The scale and complexity of societal problems comes paired with a vast number of possible interventions. In our climate change example, these interventions might include organizational redesigns, legislation, economic incentives, deployment of technology, or a combination thereof. Research on societal challenges must consider a **broad range of diverse candidate interventions** based on experiences of researchers from various disciplines to understand the nature of *good* interventions in the absence of a unique quality criterion (Pries-Heje and Baskerville, 2008, Collins et al., 2009). In the case of technological interventions, studying a broad range of candidate artifacts is particularly important, because the effects of strategic interactions between artifacts can easily dominate the performance of artifacts studied in isolation (Hanusch and Pyka, 2007).

Quickly generating diverse candidate interventions presents current scientific methods with difficulties. We already discussed shared paradigms as a critical precondition to compa-

rability. In contrast, the current norm are disparate candidate interventions based on different problem definitions that hamper a quick cycle of comparison and improvement.

We argue that methodological advances are needed to foster interdisciplinary communities of researchers working from a shared paradigm. This will require new forms of coordination between a multitude of smaller research groups, and a mindset that favors a peer-reviewed and community-owned paradigm over after-the-fact comparisons of results based on disparate problem definitions.

Evaluating IS Artifacts ④

Interventions in complex systems must be **evaluated at the system level** where strategic interaction effects can be observed. This is particularly difficult for societies, where the increasing use of markets and other social forms of organization has vastly increased the number and diversity of interactions (Bichler et al., 2010). Consider the case of the global financial markets with their continuously evolving setup. These markets

involve or acquire significant degrees of variability in components and heterogeneity of constituent systems ... For this reason traditional engineering techniques, which are predicated on very different assumptions, cannot necessarily be trusted to deliver acceptable solutions. ... [N]ew approaches are required: new engineering tools and techniques, new management perspectives and practice (Cliff and Northrop, 2012).

Analytical methods may provide important insights in stylized settings, but they are necessarily limited when it comes to evaluating complex system interventions. Real world evaluations such as field experiments, on the other hand, are problematic because of the prohibitive cost of well-intended interventions gone awry, that is, the **cost of social negatives**. Pilot evaluations could alleviate these consequences, but they are expensive and their realism is bounded by a homogenous, small-scale setup where one consortium controls the entire pilot. Finally, many important societal challenges like climate change, aging societies, and depleted carbon-based energy sources have **not fully materialized** yet, rendering real world evaluations simply impossible.

IS researchers have extensive experience with system level evaluations that often include strategizing actors and artifacts, e.g., (Bapna et al., 2004, Wang and Benbasat, 2005). But because of the **vast number of interactions, decentrally evolving artifacts**, and the **evolving web of interactions** between them, interventions in societal challenges are particularly difficult to evaluate. Research must anticipate and preempt societal challenges instead of

studying them in retrospect, but it is unclear how researchers can cater to unrealized futuristic needs while meeting standards of academic rigor today.

We argue that methodological advances are needed in system-level evaluations of de-centrally evolving artifacts and their strategic interactions for currently unrealized problems of societal scale. Evaluation facilities must provide detailed, comparable data on artifact performance and evolution, and balance swift evaluation against the risk of incurring social negatives. We should emphasize that we do not attempt to prescribe a single best tradeoff between rigor and relevance. Instead, we see this tradeoff as a conscious choice, jointly made by researchers and stakeholders during the definition of their paradigm.

Communicating with the Environment and the Knowledge Base 5 6

Producing research results on societal challenges that are both impactful and rigorous is difficult for at least three reasons. First, stakeholders expect researchers to proactively provide **solutions in addition to mere insights**. Policy makers, for instance, seek concrete guidance on the technologies, rules, and institutions of future energy infrastructures (Kassakian and Schmalensee, 2011). Second, due to the scale and complexity of societal challenges it is difficult to communicate the problem and possible interventions, and to **convince stakeholders of the viability of interventions** for further evaluation in the real world. And finally, the established scientific **publication cycle cannot keep up with the pace of societal challenges**, which reduces the timeliness of research results and their potential impact.²

We argue that methodological advances are needed that encourage researchers to produce tangible representations of their results in addition to textual descriptions. These representations must be based on a credible, peer-reviewed paradigm, invite further experimentation by researchers or practitioners, be readily comparable to alternatives, and come with detailed performance records in the form of curated experimental data. By working from a shared paradigm, and by making data and designed artifacts first-class citizens of the scientific process, frictions in building on other researchers' results must be reduced, and the credibility and concreteness of results must be increased.

Towards IS Research for Societal Challenges

To summarize, while nothing limits the application of existing research frameworks to societal challenges per se, we find several obstacles when applying them at the societal level. In this section, we have identified and described these obstacles, and have put forward several required advances.

²Our focus in this article is on IS research, but we conjecture that this is also true for other disciplines.

4.3 Competitive Benchmarking

We next propose Competitive Benchmarking (CB), a novel IS research method that aims to remove these obstacles. At the heart of CB is a new **separation of concerns** around rich representations of scientific paradigms and research results. CB enables scalable interdisciplinary research communities in which coordination and peer review are shifted to the earliest possible time. The return on this up-front investment comes in the form of comparable, actionable research results, and timely dissemination.

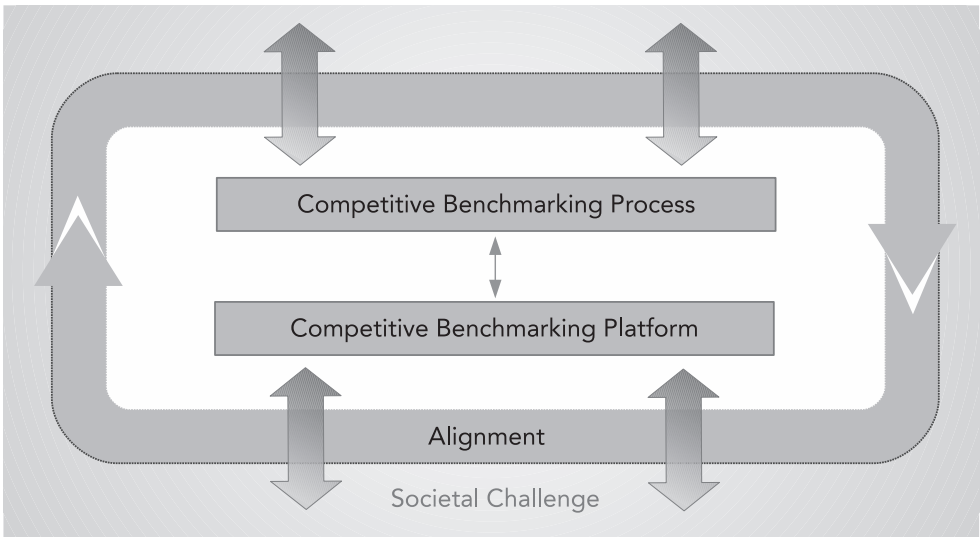


Figure 4.2. Competitive Benchmarking (CB) consists of three related core elements: the PROCESS, the PLATFORM, and a continuous process of ALIGNMENT. The boundaries between these core elements enable the new separation of concerns within CB.

Figure 4.2 illustrates CB's three core elements:

1. CB ALIGNMENT³ refers to a continuous synchronization process between a scientific paradigm and a societal challenge, and it provides for the timely dissemination of late-breaking results.
2. The CB PLATFORM is the medium in which researchers and stakeholders represent an evolving scientific paradigm, and it provides the infrastructure for the PROCESS.
3. The CB PROCESS is where independent researchers iteratively build novel theories

³We use small capitals to distinguish ALIGNMENT, PROCESS, and PLATFORM as defined in Competitive Benchmarking from their usual interpretations.

and design artifacts, while benchmarking and improving their work in direct sight of each other.

We now elaborate on these core elements and describe where CB departs from conventional IS research. In the following sections, we instantiate CB to the sustainable energy challenge, give empirical evidence for its efficacy, and discuss connections to related work.

4.3.1 Competitive Benchmarking ALIGNMENT

No single research group is likely to understand the full extent of a societal challenge, and we therefore propose a shared scientific paradigm, established through a community-based process. This paradigm must be updated continuously as technologies, regulations, or objectives change (**synchronization** function). And research results must be disseminated in a targeted and timely fashion in order to have impact (**dissemination** function). In CB, these two functions are realized through a continuous ALIGNMENT process, depicted in Figure 4.3.

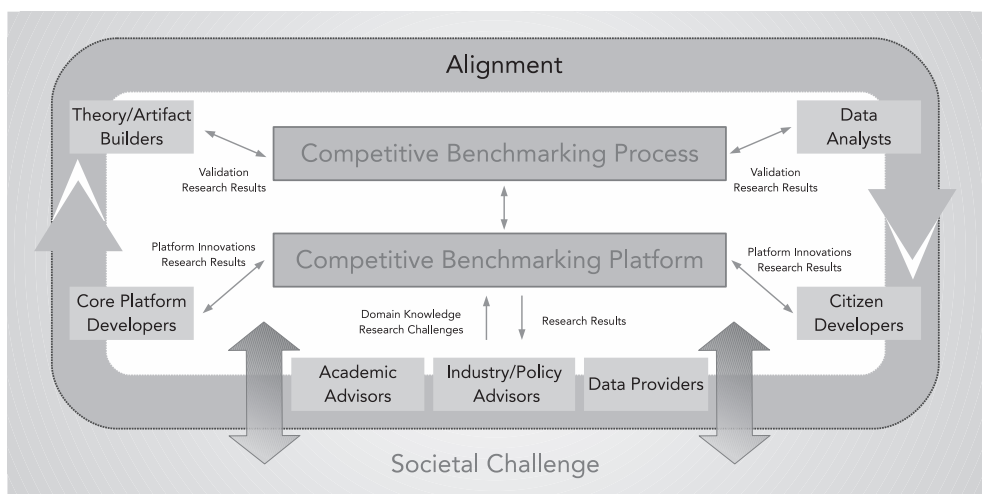


Figure 4.3. Alignment is a continuous, community-based synchronization process between a scientific paradigm and a societal challenge.

Let us first consider **synchronization**. Establishing and maintaining an accurate model of a societal challenge is an important precondition for research that generates useful theories and artifacts, and that offers reliable policy guidance (Pyka and Fagiolo, 2007). Neither the idea of continuous analysis nor the methods CB researchers use to this end differ from conventional research and we will therefore not discuss them further (see, e.g., Gray, 2004; Majchrzak and Markus, 2013). ALIGNMENT's distinguishing feature is that it encourages the

establishment of one shared, peer-reviewed paradigm early on, to increase the speed, effectiveness, and credibility of the research efforts that follow.

The basic idea is to replace the single-investigator model and its numerous smaller, incompatible problem definitions with a social learning process that is better suited for gathering dispersed, often tacit stakeholder knowledge. The resulting paradigm is continuously updated and represented in a software-based CB PLATFORM, a choice of medium that we discuss in detail below. In practice, community-based paradigm development requires initial investments from a core community of dedicated researchers. Once a critical mass of groundwork has been laid, its benefits become evident and a virtuous cycle of peer-review, incremental refinement, and increase in paradigm value sets in. As researchers from diverse backgrounds begin adopting and contributing to the paradigm, they increase the community's capacity for understanding the challenge, improve the coverage and detail of the paradigm, challenge prior assumptions, and provide additional validation in the process.

Besides building the paradigm, maintaining its connection to the societal challenge under study is equally important. In our own CB efforts, we institutionalize this idea through industry and policy **advisory boards** that meet regularly to provide guidance on the most important aspects of the challenge. The upshot is an intellectual capital base with high managerial and societal relevance, that each researcher has an interest to invest in, and that benefits the entire community by providing one high-quality research infrastructure.

We hasten to add that the goal of ALIGNMENT is not to establish one universally accepted world-view, nor to socialize the scientific process. As we shall see below, CB encourages a type of intense, competitive innovation in which individual achievements are promoted rather than attenuated. But for such competitive innovation to be effective, researchers must start from compatible assumptions and distribute their limited time judiciously. ALIGNMENT provides upfront coordination and open dispute resolution *before* major research efforts are undertaken. It avoids duplicate work during the problem definition phase, it promotes research results that are comparable after the fact, and it leads to a greater confidence that the community's efforts flow into the highest-value research questions.

The results of these efforts must be communicated in a targeted and timely fashion to have impact and to accelerate progress (Garvin, 1993). CB supports the timely communication of results through the **dissemination** function of ALIGNMENT. Clearly, the community of stakeholders and researchers involved in CB is a natural starting point for dissemination, with a vested interest in results guided by their own ideas. But the dissemination function adds at least two other novel and important benefits.

First, by combining a peer-reviewed paradigm with a swift but rigorous PROCESS, CB offers an alternative to the protracted ex-post review of assumptions and results that is the

current scientific norm. Because a significant share of review is performed up-front at the paradigm level by numerous independent researchers and stakeholders, ALIGNMENT is “a way of effecting ... validation. The interaction between the modeler and the client in mutually understanding the model and the process establishes the model’s significance; that is, its warranty” (Kleindorfer et al., 1998). Individual researchers then develop new theories and artifacts based on the pre-validated paradigm, which are ultimately evaluated by an independent party during the public CB PROCESS. There, theories and artifacts have to perform well under demanding conditions that are partly determined by the evaluators, and partly by other researchers’ designs. Fine-grained protocols of these evaluations are made publicly available to support their credibility. Overall, this procedure greatly reduces the need for ex-post scrutiny and time to disseminate.

Second, because ALIGNMENT is problem-centric and continuously seeks to identify the next most important insights and solutions, it reduces the risk of addressing outdated problems. It thereby generalizes the idea of applicability checks (Rosemann and Vessey, 2008) to a continuous process that guides a research community.

The prohibitive cost of potential social negatives will make decision-makers in industry and policy understandably sceptical of trusting just any result. A diligently executed process of ALIGNMENT leads to an improved rapport with these stakeholders and adds credibility to research results obtained through CB. Combined with timely, tangible results in the form of data and executable artifacts, this creates attractive opportunities for high-impact dissemination.

4.3.2 Competitive Benchmarking PLATFORM

The PLATFORM is the central point of coordination for CB participants. It is the malleable, executable **representation of the shared paradigm** created and updated during ALIGNMENT, and it provides a **scientific toolset** to the PROCESS (see Figure 4.4).

Given the central role of the paradigm within CB, the medium used to **represent** it is important. The current medium (natural language) has three significant shortcomings: it has no safeguards against imprecisions and inconsistencies, it is difficult to update as the challenge evolves, and it must be translated into other media to become actionable. Formal representations address the first concern, but they are limited in terms of problem sizes they can address.

CB instead promotes the use of software-based PLATFORMs to represent the paradigm, which leverage the great strides that software engineering has made in understanding and representing complexity. These started with the realization that modeling complex socio-technical systems should be an iterative, social learning process. Related progress in com-

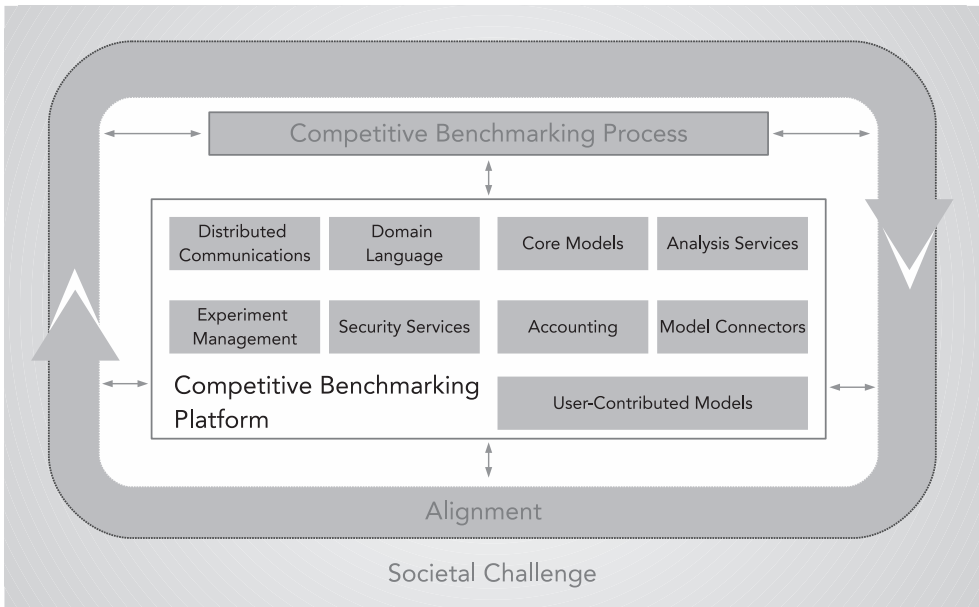


Figure 4.4. The PLATFORM is a community-developed, executable representation of the paradigm. It is the central point of coordination between researchers and stakeholders, and it provides a toolset for empirical science.

puter language theory has bred a generation of highly expressive, problem-centric languages that put stakeholder needs before machine considerations (Meyer, 1998). And advances in program design and architecture have made software extensible and adaptive to changing environments. The upshot is a proven, scalable, and social approach to capturing complexity (Baetjer, 1997).

The advantages of software-based paradigm representations come at a greater cost of initially describing the problem, which must therefore be spread over several research groups. We should also note that certain technical qualities of software-based representations may require advanced software engineering skills, a point we revisit in the discussion. Among these qualities are a **clear design** that makes it easy for other researchers to use and extend the paradigm, good **readability** and thoroughly documented assumptions, a **modular architecture** that enables specialist contributions in clearly delineated areas, and a **licensing** model that encourages free redistribution and extension.

The second PLATFORM function is that of a **toolset for empirical science**. Because the PLATFORM encodes a shared understanding of a societal challenge, research results and tools derived from it will be comparable and technically compatible. For the purpose of theory validation, PLATFORM data can be related to data obtained from studies under different

environmental conditions, or reproduced under identical environmental circumstances (Tsfatsion, 2006, Pyka and Fagiolo, 2007). Designed artifacts can readily be benchmarked against artifacts from other research groups. And ecosystems of scientific tools can be built around the PLATFORM to aid researchers in routine tasks such as data screening, reporting, and distributed experiment management.

We should emphasize that the presence of an executable representation of the paradigm also means that fully executable interventions like dynamic decision rules, economic mechanisms, or IS artifacts can be evaluated using the PLATFORM as a testbed. These interventions are tangible and interesting to study for practitioners and researchers alike.

4.3.3 Competitive Benchmarking PROCESS

The iterative PROCESS is where researchers iteratively build novel theories and design artifacts, while continually improving their work in direct sight of each other. It also encompasses the competitive, benchmarking-based elements from which CB derives its name.

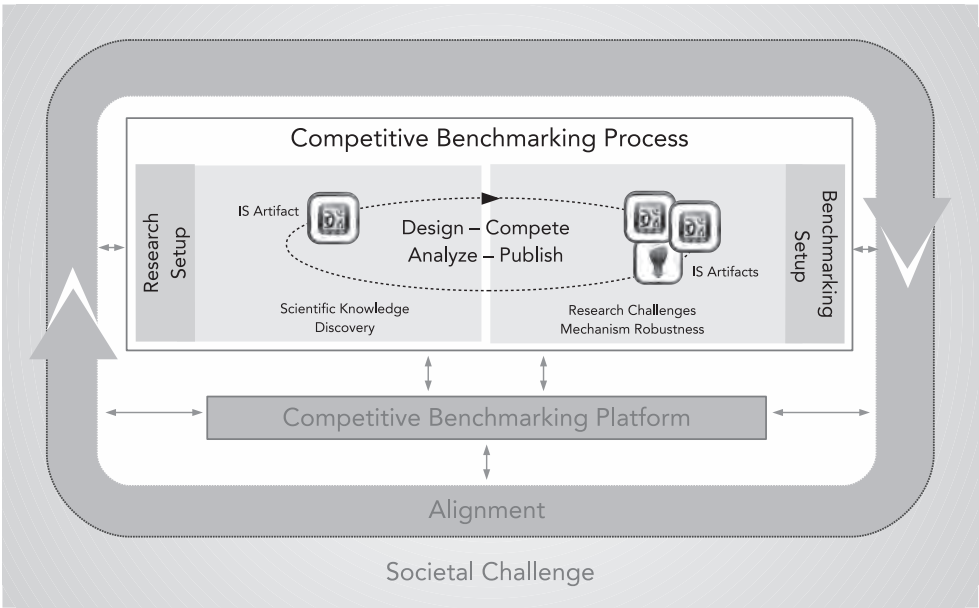


Figure 4.5. The PROCESS consists of four phases in which researchers subsequently (1) design new theories and artifacts, (2) pit them against other researchers’ work, (3) analyze the quality of their ideas, and (4) improve iteratively. The PROCESS leverages the aligned PLATFORM, and it produces theories and executable artifacts with comprehensive supporting data for quick dissemination.

Any effective research method is a structured approach to exploring and learning about phenomena (descriptive and explanatory research) and solution spaces (design science research). Researchers create new theories and designs, evaluate their realism and usefulness, learn from experience, and iterate to improve their work (see Figure 4.1). This structured form of learning and improvement is related to benchmarking in that it requires skills in “systematic problem solving, experimentation with new approaches, learning from ... own experience and past history, learning from the experiences and best practices of others, and transferring knowledge quickly and efficiently.” Its best practitioners “[rely] on the scientific method, rather than guesswork, for diagnosing problems” and “[insist] on data, rather than assumptions, as background for decision making” (Garvin, 1993).

Figure 4.5 illustrates how CB incorporates the notion of benchmarking in its PROCESS. Suppose a community of researchers and stakeholders is interested in understanding the effects that different transaction tax regimes have on the trading behavior of commercial banks and in stability implications for global financial markets. Starting from these goals, they engage in ALIGNMENT and model the behaviors of private and institutional investors, a market infrastructure, central banks, etc. until they agree on having captured the most salient features of the challenge. The result of this work is an aligned PLATFORM on which the PROCESS proceeds iteratively, in cycles consisting of four phases:

Design: Several research groups design trading strategies against the PLATFORM, which can be based on both *ad-hoc* designs or on sound kernel theories, as long as they remain within the agreed-upon paradigm (see Section 4.2).⁴ Strategies can even involve human traders, which opens interesting avenues for work on behavioral theories (Babb et al., 1966, Collins et al., 2009, 2010). Researchers repeatedly evaluate their strategies against the PLATFORM in research setup to detect and remove weaknesses.

Compete: Participants then pit their artifacts against each other in a formal tournament where strategic interactions and system-level properties can be observed. An independent party determines the tournament schedule, including the pairing of artifacts and environmental conditions (e.g., various tax levels or trading intensities). CB is inherently meritocratic and good performance in a strong field of competitors is reward and incentive for further improvement.

Analyze: The tournament outcome is a ranking of strategies, together with fine-grained data on artifact *and* system-level behavior made publicly available. The PLATFORM and its accompanying scientific tools promote credible analyses that can be produced quickly

⁴Note, that this does not preclude artifacts from exploiting loopholes within the PLATFORM. One of the benefits of CB is the discovery of unintended loopholes through a wide array of creative artifacts.

and distributed along with the underlying data.

Publish: The insights gleaned from these analyses are disseminated to researchers and stakeholders. Analyses can, for example, pinpoint drivers of artifact performance that research groups can use to direct their future efforts, e.g., (Jordan et al., 2007, Ketter et al., 2013c). Researchers also make executable versions of their tournament artifacts available for study, which is an attractive additional channel for distributing tangible, actionable results.

This example is mostly concerned with artifact design, and system-level artifact evaluations are indeed among CB's key benefits, but the PROCESS equally supports other types of scientific inquiry that close the IS research cycle described by Hevner et al. (2004). Most importantly, a PLATFORM together with a fixed set of high-performing artifacts can be used as a conventional Agent-based Virtual World (ABVW) to perform controlled experiments in pursuit of descriptive or explanatory theories (Ketter et al., 2008, Chaturvedi et al., 2011). These theories can then be used by artifact designers to improve their designs. Table 4.2 provides an overview of the supported research types.

CB's PROCESS contains four novelties that aim to improve the capacity of IS research for tackling societal challenges. Most importantly, it adds **naturalistic dynamics** to artifact validations. In our example, researchers cannot hope to experiment with real tax regimes and must therefore resort to working against a model of the challenge. However, one particularly important facet of real-world evaluation *can* be brought into the laboratory: the competitive co-evolution of artifacts. Like firms and individuals in the real world, CB participants constantly seek to improve their designs by adapting to the behavior of the environment and of others in a type of Emergent Knowledge Process (EKP, Markus et al. 2002). The ensuing dynamics provide a unique tradeoff between artificial and naturalistic elements for high-risk evaluations in complex economic environments, see Figure 4.6.⁵

Second, the aligned PLATFORM is **validated by other researchers and stakeholders**, and evaluation conditions are determined by an independent party. That artifacts and theories must perform well under many different circumstances in a realistic environment increases **external validity** and researchers' confidence in the absence of unanticipated social negatives.

Third, community-based ALIGNMENT and PLATFORM development spreads the effort of understanding and modeling a challenge across many researchers to **increase scientific**

⁵An alternative view on this is based on the "increasing recognition of the mutable nature of these artifacts. That is, they are artifacts that are in an almost constant state of change" (Gregor and Jones, 2007). Designs, in the context of CB, are by definition "evolutionary trajectories," not static blueprints, and an important benefit of CB is the ability to generate such trajectories realistically, and to study their development over time.

| Research Type | Research Setup | Examples |
|--------------------------|---|--|
| Artifact Design | <ol style="list-style-type: none"> 1. Use PLATFORM for distributed artifact design 2. Benchmark and improve artifacts iteratively | <ul style="list-style-type: none"> • Trading strategies • Dynamic pricing • Brokers |
| Controlled Experiments | <ol style="list-style-type: none"> 1. Hold set of high-performing artifacts constant 2. Execute artifacts against PLATFORM while varying environmental parameters 3. Measure resulting system-level properties | <ul style="list-style-type: none"> • Social welfare studies • Distribution studies • Concentration and competitiveness measures |
| Falsification Studies | <ol style="list-style-type: none"> 1. Vary set of high-performing artifacts 2. Execute artifacts against the PLATFORM 3. Assess stability of mechanism or theory | <ul style="list-style-type: none"> • Market mechanisms • Circuit breakers |
| Mixed Initiative Studies | <ol style="list-style-type: none"> 1. Vary set of high-performing artifacts 2. Task human participants 3. Execute artifacts against PLATFORM 4. Assess human or artifact performance | <ul style="list-style-type: none"> • Decision support systems • User interfaces |

Table 4.2. CB supports descriptive and explanatory research (insights) as well as design science research (solutions). The continuous evolution of artifacts in the PROCESS yields diverse, high-performing artifacts that can be studied towards descriptive or explanatory theories.

cycle speed later. The initial investment amortizes as researchers gain the ability to rapidly test artifacts and theories without the frictions of first finding compatible benchmarks. The publicly evaluated artifacts and theories can then be swiftly disseminated. The evaluation data also can be used to derive **rigorous design theories**, which is an important step in reconciling the need for scientific rigor with leveraging the creativity of pragmatic designs. It may not even be known why a particular artifact works at the time of evaluation, but the availability of evaluation data allows the community to discover theoretical principles behind its working later on.⁶

And finally, the comprehensive data generated in the PROCESS provides **clear visibility of the progress** that designers make in improving their artifacts, which also gives a measure of the benefits of CB as a research method (Venable and Baskerville, 2012). When progress

⁶A similar separation of concerns led Johannes Kepler to discover the laws of planetary motion from recordings in the notebooks of Tycho Brahe (Hey et al., 2009). We speculate that the lack of comparability between artifacts causes this separation of concerns to be virtually absent from design research today.

| | Ex Ante | Ex Post |
|--------------|--|--|
| Naturalistic | <div>Competitive Benchmarking</div> <div>Action Research</div> <div>Focus Groups</div> | <div>Ethnography</div> <div>Phenomenology</div> <div>Survey</div> <div>Participant Observation</div> <div>Case Study</div> <div>Action Research</div> <div>Focus Group</div> |
| Artificial | <div>Lab Experiment</div> <div>Computer Simulation</div> <div>Mathematical or Logical Proof</div> <div>Criteria-Based Evaluation</div> | <div>Lab Experiment</div> <div>Role Playing</div> <div>Simulation</div> <div>Computer Simulation</div> <div>Field Experiment</div> <div>Mathematical or Logical Proof</div> |

Figure 4.6. The CB PROCESS brings competitive coevolution of artifacts, one particularly important facet of naturalistic evaluation, into the laboratory. The depicted classification is adapted from Venable et al. (2012). The dotted line indicates the partly artificial, partly naturalistic character of evaluation in CB.

tapers off, the community may also decide to call its advisory board for new challenges.

4.3.4 Interaction effects of ALIGNMENT, PLATFORM, and PROCESS

We should emphasize that CB does not attempt to replace the existing process of scientific knowledge discovery. It rather aims to remove several common obstacles, and it adds a structured approach to benchmarking which, in our opinion, is insufficiently represented in current IS research practices. One of the resulting benefits for IS research on societal challenges is a clear separation of concerns between various stakeholder and researcher groups around the PLATFORM, which ultimately leads to better scalability, and which we summarize in Table 4.3.

The improvement in scalability stems partly from reducing the waste and redundancy inherent in incomparable research results, and partly from redistributing efforts between individuals and the community. In particular, the early coordination during ALIGNMENT enables the reuse of domain knowledge obtained from stakeholders, and of the scientific toolset provided by the PLATFORM. Figure 4.7 illustrates how individual effort is supplemented by

| Separation of concerns between ... | Enables ... |
|--|---|
| Stakeholders and Researchers | <ul style="list-style-type: none">• Researchers to effectively learn about the challenge• Stakeholders to learn about new research insights and solutions in a timely fashion |
| Researchers from different disciplines or with different expertise | <ul style="list-style-type: none">• Scalable, expert model-building and concurrent work on one joint problem definition. For example, a battery expert might build realistic models of e-vehicle charging behavior to be used by an economist in the design of market mechanisms.• Competitive design. For example, a machine learning (ML) expert and an operations research (OR) expert might design alternative solutions to a given problem. The shared use of a PLATFORM ensures that their artifacts remain technically compatible and comparable. |
| Theory/Artifact Designers and Data Scientists | <ul style="list-style-type: none">• Independent data analysis and validation. PROCESSES generate publicly available data for analysis. An economist could, e.g., analyze the welfare effects of deploying the ML- and OR-based artifacts described above. |
| Academic Researchers and Pragmatic Designers | <ul style="list-style-type: none">• Leveraging the creativity of pragmatic designers (Hevner and Chatterjee, 2010). CB imposes very few constraints on the theoretic underpinnings of designed artifacts. Practitioners can contribute high performing ad-hoc artifacts that are then further analyzed by academic researchers.• Effective industry cooperations. Industrial designers can contribute artifacts that are rigorously evaluated according to the standards of design theories. |

Table 4.3. CB’s three core elements facilitate an effective collaboration between various groups of contributors. This leads to better scalability in the challenge size, and in the number of independent contributors.

community effort in defining the problem and in evaluating and communicating results. The upshot is more time spent on the value-generating core activities of theory development and artifact building for each individual researcher.

The next section evaluates the efficacy of CB in a case study on design for sustainable energy systems that we have been conducting together with a global community of researchers over the past four years.

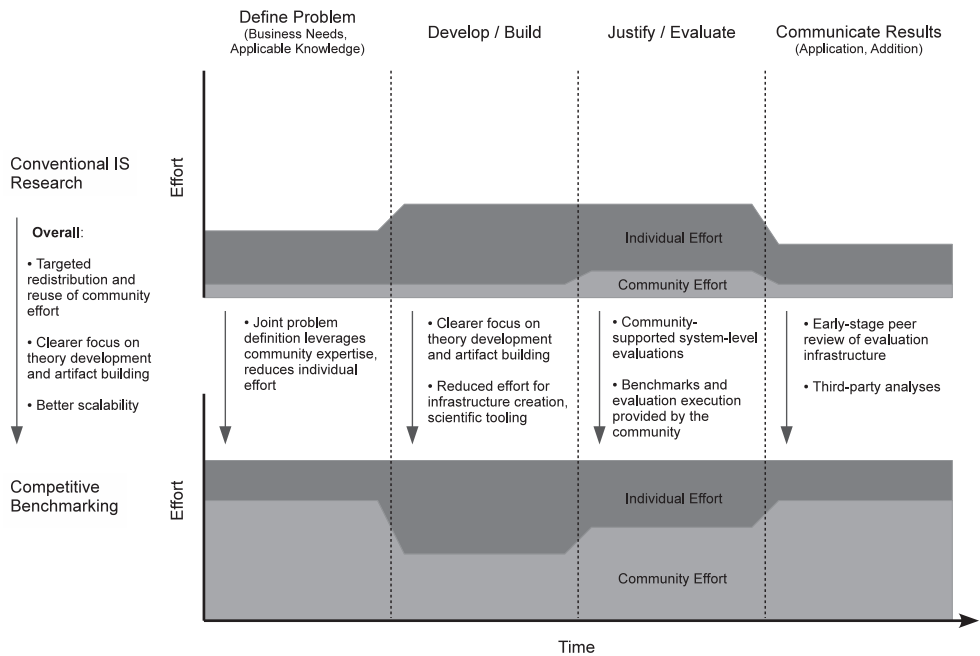


Figure 4.7. Competitive Benchmarking leads to redistributing efforts between individual researchers and the community. Early joint problem definition and peer review lead to conceptually compatible theories and artifacts that can be re-used for evaluation and benchmarking purposes. Clear separations of concerns enable, e.g., the development of system-level analyses by independent third-party data scientists.

4.4 The Power TAC Case: Competitive Benchmarking for Sustainable Energy Systems

From relatively modest beginnings 130 years ago, electricity has quickly revolutionized the way we consume energy and transport information. But next to substantial economic benefits, the electric revolution has also brought about environmental and reliability concerns. The drivers behind these negatives are numerous and complex, but one important underlying theme is the mismatch between increasing demands for volume, sustainability, and affordability on one hand, and hierarchical control structures that are essentially unchanged from electricity's early days on the other.

Modernizing these control structures is an extremely challenging proposition. The **Smart Grid** of the future will have to (a) efficiently allocate electricity among hundreds of millions of users with unique preferences, (b) integrate production from renewable and decentral-

ized power sources like electric vehicle fleets, (c) respect complicated constraints imposed by power flow physics, privacy concerns, and several layers of regulation, and (d) uphold real-time control under uncertainty, all the while ensuring a smooth transition from the operational grid of today. IS scholars can make substantial contributions to this grand societal challenge by “integrating new information and communications technologies, combining them with active support from electricity consumers, and leveraging the optimizing power of markets” (Coll-Mayor et al., 2007).

The scale and complexity of the challenge, and the interrelated advances required in theory and artifact design prompted a global community of researchers to address it through the **Power Trading Agent Competition (Power TAC)** (Ketter et al. 2013b, 2013c, see also www.powertac.org), a Competitive Benchmarking effort that we describe next, together with empirical evidence for its efficacy.

4.4.1 Power TAC ALIGNMENT

At the beginning of the Power TAC project stood an initial ALIGNMENT during which a core group of researchers performed stakeholder interviews and surveyed the literatures on power systems and smart grids. Key stakeholders were identified in utility companies, network infrastructure providers, communication electronics manufacturers, electricity cooperatives, and electricity customer lobby groups. These stakeholders were interviewed repeatedly, and many joined an advisory board which now institutionalizes Power TAC’s ongoing ALIGNMENT. The board meets twice yearly to provide researchers with industry insights, to ensure that important challenges are being tackled, and to disseminate the latest research results.

After several ALIGNMENT iterations, Power TAC began to attract outside researchers interested in leveraging the publicly available PLATFORM for their own work. Several groups contributed specialized knowledge that improved its realism in areas where no other community member possessed the requisite expertise or resources, e.g., customer modeling (Reddy and Veloso, 2012) and balancing (de Weerd et al., 2011). In exchange, the contributors could study their models in a rich, realistic environment that they could not have created otherwise, including a dedicated community that validated and critiqued their models. Other groups created experimental tools for and third-party analyses of Power TAC (Babic and Podobnik, 2013, Kahlen et al., 2012), compared the PLATFORM against real-world behaviors (Nanoha, 2013), and designed and evaluated artifacts, e.g., (Peters et al., 2013c, Kuate et al., 2013). Importantly, many of these new participants had the technical expertise but no prior domain knowledge or interest in contributing to the sustainable energy challenge. It was the availability of a community-supported, executable model of a real-world challenge and a predefined list of important research questions that triggered them to apply their

diverse technical skills to sustainable energy. Conversely, researchers and external stakeholders with energy domain knowledge benefited from the innovative contributions of these technical experts.

Our case study illustrates how ALIGNMENT provides scalability to communities of researchers coordinating through a shared paradigm. Establishing and maintaining this paradigm regularly requires incisive modeling decisions from the community. But through ongoing ALIGNMENT, these decisions can be made early, thereby keeping subsequent research results technically and conceptually comparable. For example, Power TAC currently:

- models the electric **distribution system** but not the transmission system, because while controlling the latter is well understood, much scientific guidance is needed on making the former “smarter” (EPRI - Electric Power Research Institute, 2011).⁷
- models the **economic aspects** of the smart grid but not the physical power flows, because of an urgent need for insights on how a combination of IT and economic forces can incentivize sustainable electricity consumption (Watson et al., 2010).
- models **retail electricity tariffs**, but not bilateral price negotiations with commercial customers, because end users “can provide remarkable local intelligence ... [but] any technology is doomed to fail if the involved users do not like or understand it” (Palensky and Dietrich, 2011).

These ALIGNMENT results are continuously translated into the executable and peer-reviewed Power TAC PLATFORM.

4.4.2 The Power TAC PLATFORM

The PLATFORM models a competitive retail power market in a medium-sized city, in which consumers and small-scale producers may choose from among a set of alternative electricity providers, represented by competing **Brokers**. Brokers are autonomous software agents, built by individual research groups. The remainder of the paradigm (see Section 4.2) is modeled by the PLATFORM (see Figure 4.8).

Brokers offer electricity tariffs (also known as plans or rates) to household and business customers through a **retail market**. Some customers are equipped with solar panels and wind turbines, which produce and consume power, and many own demand-side management capabilities such as remotely controllable heat pumps or water heaters. All customers are

⁷The distribution system is responsible for providing regional electricity to commercial and residential end-customers. The transmission system is where large-scale generators like wind farms and coal power plants feed in high-voltage electricity for long range transmission.

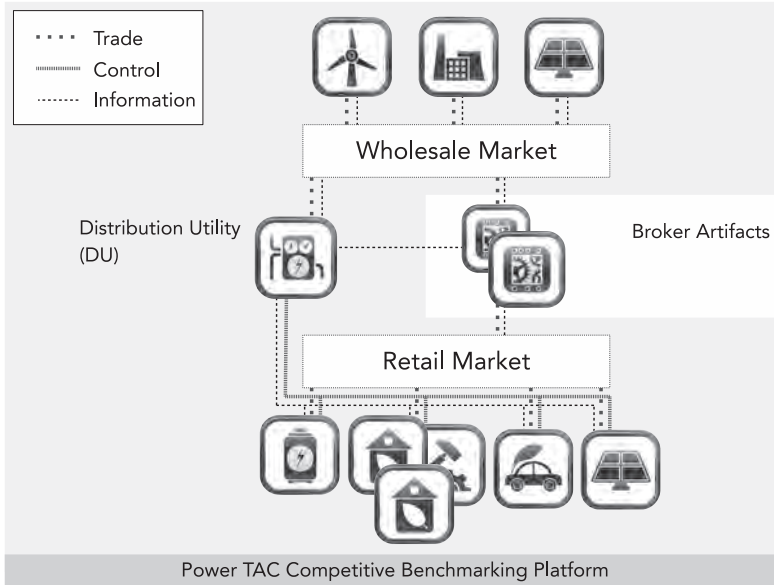


Figure 4.8. Main elements of the Power TAC scenario. Brokers are autonomous software agents built by individual research groups. The remainder of the scenario is modeled by the PLATFORM

equipped with smart meters from which consumption and production is reported every hour. Customers are sensitive to price changes, weather conditions, and calendar factors such as day of week and hour of day, and they have a range of preferences over tariff terms. For example, some are willing to subscribe to variable-rate tariffs if they have the opportunity to save by adjusting their power usage, while others are willing to pay higher prices for the simplicity of fixed-rate or time-of-use tariffs. Many of these models are contributions from the user community, e.g., (Gottwalt et al., 2011, Reddy and Veloso, 2012). Brokers buy and sell energy either from retail customers, or in the day-ahead **wholesale market**, where utility-scale power suppliers sell their output. These suppliers represent different price points and lead-time requirements, e.g., fossil and nuclear power plants, gas turbines, and wind parks.

The **Distribution Utility (DU)** models a regulated monopoly that owns and operates the physical facilities (feeder lines, transformers, etc.) and is responsible for real-time balancing of supply and demand within the distribution network.⁸ It does this primarily by operating in the balancing market, the real-time facet of the wholesale market, and by exercising demand and supply controls provided by Brokers. The associated costs are allocated to imbalanced

⁸In the real world, balancing responsibility is typically handled at the transmission level; the simulation implements a generalization of proposals to move some balancing responsibility to the distribution level (Strbac, 2008).

Brokers. Given a portfolio of customers, Brokers compete in the wholesale market to minimize the cost of power they deliver to their consuming customers, and to maximize the value of power delivered to them by their producing customers.

The Power TAC PLATFORM as described here has quickly evolved into the most comprehensive economic simulation for smart distribution networks worldwide. Its source code is licensed under a research and business-friendly Apache license and can be downloaded for free from <https://github.com/powertac>. Use and modification of the PLATFORM are not predicated on participating in the PROCESS, but several important CB benefits can only be reaped by actively engaging with the Power TAC community in this way.

4.4.3 The Power TAC PROCESS

Cornerstones of the CB PROCESS are annual championships, and pilots that provide additional informal benchmarking opportunities. To date, pilots have been held at IJCAI 2011 in Barcelona, at AAMAS 2012 in Valencia, and at IEEE SG-TEP 2012 in Nuremberg. The first two official Power TAC championships were held at AAAI 2013 in Bellevue, WA, and at AAMAS 2014 in Paris, France.⁹

All tournaments consisted of qualifying rounds in which Brokers were screened for technical flaws, followed by final rounds in which varying combinations of three, five, and eight Brokers competed. Table 4.4 shows all finalists of the 2012 Nuremberg pilot and the first two Power TAC championship. These Brokers were designed by researchers with expertise in Artificial Intelligence, Electrical Engineering, Information Systems, Machine Learning, and other areas, and their heterogeneous design approaches have contributed to a rich repository of design ideas, executable artifacts, and artifact performance data that we analyzed.

Figure 4.9 gives a high-level view of Brokers' performances at different levels of competition in the 2012-2014 tournaments.¹⁰ The figure supports several preliminary conclusions:

Heterogeneity Matters: Heterogeneity in designs led to significant performance differences.

The best Broker designs outperformed competitors by more than an order of magnitude. Attracting heterogeneous designs, benchmarking them against each other, and understanding the theoretical basis of their functioning is therefore critically important.

Breakthrough Innovation: The two best Broker designs in 2013 had only submitted an

⁹AAAI = Conference of the Association for the Advancement of Artificial Intelligence; AAMAS = International Conference on Autonomous Agents and Multiagent Systems; IJCAI = International Joint Conference on Artificial Intelligence; SG-TEP = IEEE Conference on Smart Grid Technology, Economics, and Policies.

¹⁰For simplicity, we refer to the 2012 Nuremberg pilot and the 2013/2014 Power TAC championship as the 2012-2014 tournaments henceforth.

| Broker | Institute | Country |
|----------------|--|-------------|
| AgentUDE | University Duisburg-Essen | Germany |
| AstonTAC | Aston University Birmingham | UK |
| coldbroker | INAOE, Natl. Institute for Astrophysics, Optics, and Electronics | Mexico |
| CrocodileAgent | University of Zagreb | Croatia |
| cwiBroker | CWI, Natl. Research Institute for Mathematics and Computer Science | Netherlands |
| Incumbent | Incumbent Monopoly, provided by Power TAC | - |
| LARGE | Erasmus University Rotterdam | Netherlands |
| Maxon | Westfälische Hochschule | Germany |
| Mertacor | Aristotle University Thessaloniki | Greece |
| MinerTA | University of Texas at El Paso | USA |
| MLLBroker | University of Freiburg | Germany |
| SotonPower | University of Southampton | UK |
| TacTex | University of Texas at Austin | USA |

Table 4.4. Participants in the 2012-2014 Power TAC finals. The list excludes several other participating groups who did not qualify for the final rounds.




Figure 4.9. Average profit share of brokers in the 2012-2014 tournaments. Larger circles indicate a higher percentage of profits captured on average. Darker circles indicate a higher certainty, that is, a lower standard deviation of average profit shares. CrocodileAgent participated in the 2012 pilot, but all of its games were affected by a technical flaw and are therefore excluded from the analysis. The Incumbent broker is provided by the Power TAC PLATFORM and models the incumbent monopoly.

early-stage prototype in 2012 (TacTex) or no entry at all (cwiBroker). The 2014 winner AgentUDE similarly was a new entrant. Nevertheless, these designed surpassed

relatively more established designs from previous years, which illustrates the powerful effect of competitive innovation.

Game Context Matters: Performance depended on the level of competitiveness and other environmental conditions. AstonTAC’s design, for example, lost ground to several competitors in three-broker games in 2013, but continued to perform well in highly competitive eight-broker games. Repeatable games under different environmental conditions therefore play an important role in fully understanding system dynamics and reducing the risk of incurring social negatives.

Note, that it is only through Power TAC’s shared scientific paradigm, the common PLATFORM, and the public nature of the PROCESS that we can freely access, understand, and compare such a broad variety of creative designs towards an understanding of the Broker artifacts’ role in the sustainable energy challenge.

4.4.4 Case Study Evaluation

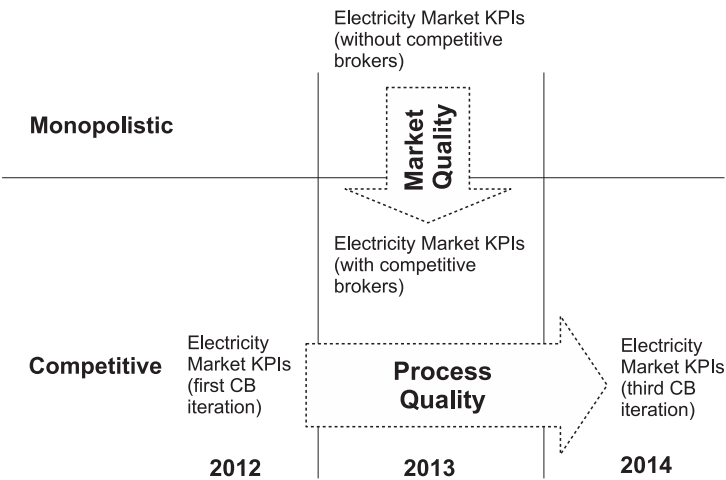


Figure 4.10. Two complementary viewpoints on the qualities of CB: PROCESS Quality is defined as improvements in several key electricity market characteristics over subsequent iterations of the PROCESS. Market Quality is defined to improvements in these characteristics through the introduction of competitive brokers.

Towards this deeper understanding, we now analyze Power TAC’s CB PROCESS from two complementary viewpoints (see Figure 4.10):

PROCESS Quality: Our key objective in this case study is to understand the efficacy of CB as a research method. To this end, we quantify the progress made by broker designers in the 2012-2014 timeframe by evaluating its impact on a set of key market characteristics defined below. Because the development of these market characteristics over time is driven by the progress of the PROCESS, we refer to this view as PROCESS Quality.

Market Quality: As a secondary objective, the case study provides an evaluation of the potential of competitively designed brokers for addressing the sustainable electricity challenge. The Power TAC PLATFORM allows us to contrast the performance of competitive and monopolistic retail electricity markets by re-running all competitive games *without* brokers. Because the change between the competitive and monopolistic scenario lies in the introduction of competitive markets, we refer to this view as Market Quality.

Our results are based on data from the three most recent Power TAC tournaments and are therefore preliminary. In the spirit of CB, the complete data set together with scripts used for the statistical analyses is publicly available at <https://bitbucket.org/journalanon/pla/src>, so that readers can reproduce and extend these result as more data becomes available.

PROCESS Quality

We compared the results of 48 final games played during the 2012 tournament to an equal number of final games played in 2013. The latter were selected from a total of 59 games, and paired with the 2012 games such that pairs were maximally similar.¹¹ Using the same procedure, we paired the 59 games played in 2013 with 59 out of 67 games played in the 2014 championship. Together, these games represent about 30 years of simulated artifact behavior at hourly resolution.

Table 4.5 shows the key performance indicators (KPIs) that we used in our comparison. They can be categorized as follows:

Retail Market Behavior: Broker behaviors and overall distribution system performance critically depend on Brokers' interactions with retail customers. We tracked realized

¹¹The 2012 tournament used a public beta version of Power TAC and for our analyses we excluded games that were affected by a technical flaw in this beta version. As an example for the similarity pairing procedure, a three-broker game between AstonTAC, Mertacor, and TacTex played in 2012 would be paired either with a game between the same brokers in 2013, or otherwise with a similar alternative.

prices, price variability, overall energy traded, and retail market competitiveness as indicated by the Herfindahl-Hirschman concentration index (HHI).¹²

Sourcing Behavior: Brokers cover retail consumption through production commitments procured in the wholesale market, or through small-scale production from retail customers. We tracked the percentage of retail consumption covered and the development of corresponding profit margins for each of these sources.

Balancing Behavior: In case of residual imbalances, the balancing market acts as the energy source and sink of last resort. We tracked the development of balancing energy required by Brokers as a percentage of their overall energy trading business.

¹²The HHI is defined as $HHI = \sum_{n=1}^N s_n^2$ where N denotes the total number of firms in a market, and s_n the market share of firm n in percent. Possible index values range from 0 (perfect competition) to 10,000 (monopoly).

| Metric | 2012 → 2013 | | | | 2013 → 2014 | | | | Target | |
|-------------------------------------|--------------------|-----------|-----------|-----------|-------------|-----------|-----------|------|--------|--|
| | 3 Brokers | 5 Brokers | All Games | 3 Brokers | 5 Brokers | 8 Brokers | All Games | 2013 | 2014 | |
| Retail Price [\$ / kWh] | Sale -0.03* | -0.03* | -0.03* | -0.02 | 0* | 0 | 0 | ↗ | ✓ | |
| | Purchase -0.02* | -0.01* | -0.01* | -0.01* | -0.02* | -0.03* | -0.02* | ↗ | ✗ | |
| Retail Price Variability [\$ / kWh] | Sale -0.023* | -0.02* | -0.022* | 0.001 | -0.001* | 0 | 0 | ↗ | ✓ | |
| | Purchase -0.001 | -0.003 | -0.003 | -0.002 | 0.002* | 0.001* | 0.004* | ↗ | ✗ | |
| Retail Volume Change | Sale 1.12* | 1.32* | 1.2* | 0.99* | 1.03* | 1.01* | 1.02* | ↗ | ✗ | |
| | Purchase 1.58* | 1.93* | 1.71* | 1.01* | 1.03* | 1.02* | 1.02* | ↗ | ✓ | |
| Herfindahl-Hirschman index | Sale -2428* | -1268* | -1993* | 678 | 1364* | 1257 | 1118* | ↗ | ✓ | |
| | Purchase -367 | -732 | -504* | 1650* | 2508* | 4885* | 2499* | ↗ | ✗ | |
| Wholesale Sourcing | -0.49* | -0.6* | -0.57* | -0.06 | -0.09* | -0.08 | -0.08* | ↗ | ✓ | |
| Wholesale Margin | -0.08 | 0.17 | 0.05 | 0.83 | 0.1 | 0.29 | 0.37 | ↗ | ✗ | |
| Retail Sourcing | 0.05* | 0.05* | 0.05* | 0 | 0 | 0 | 0 | ↗ | ✗ | |
| Retail Margin | -3.91* | -2.98* | -3.48* | 2.91 | 1.46* | 1.85* | 2* | ↗ | ✗ | |
| Balancing Ratio | Sale -1.39 | -0.55 | -0.89 | -2.03 | -0.71 | -0.21 | -1.1* | ↗ | ✓ | |
| | Purchase -0.12* | -0.17* | -0.14* | -0.03* | -0.05* | -0.06 | -0.04* | ↗ | ✓ | |

Table 4.5. Quality indicators for the Power TAC CB PROCESS as mean difference between the 2012/2013, and 2013/2014 tournaments, interpreted from the Broker's perspective. For example, a value of -0.03 for retail prices in sell direction means that Brokers earned 3 cents less per kWh on average when selling electricity to retail customers. Targets are given at the system level and with an emphasis on sustainability objectives. Stars indicate significance at the 95% level, circled numbers are referenced from the main text.

The *targeted* change in each of these KPIs depends on perspective, and the directions given in Table 4.5 reflect Power TAC's objective of contributing to the resolution of the sustainable energy challenge. As the PROCESS progresses, we hope to see highly competitive ① retail markets with low ② and stable ③ electricity prices for consumers. Overall electricity consumption ④ should remain constant or decline, electricity should increasingly come from renewable sources ⑤, and brokers should rely less on balancing power as their forecasting abilities improve ⑥. Conversely, small-scale producers should benefit from high ⑦ and stable ⑧ prices for their willingness to produce locally from renewable sources.

The results of the first PROCESS iterations have met many of these objectives. Most importantly, Broker designs have matured, as suggested by the reduced need for balancing power ⑥, and retail markets have become significantly more competitive ①.¹³ These findings are further confirmed in Figure 4.11 which illustrates Brokers' retail market strategies for the sale (left) and purchase (right) of energy. The left panels show a remarkable reduction in overall price levels ② for consumers relative to the incumbent monopoly, as well as a reduction in price *differences* between Brokers. The right panels additionally shows that Brokers now routinely include small-scale production in their sourcing strategies, which suggests that participants have been probing into more advanced design options for the most recent iteration.

But our analysis also highlights several behaviors that have developed in unintended ways. Using the data produced in the PROCESS, we can identify these developments and suggest remedies. First, lower retail consumption prices have naturally lead to higher electricity consumption ④. While some of this additional consumption comes from renewable sources, we see the need for incentives that encourage consumer energy efficient behavior even when prices decline. Brokers could, for example, be rewarded for selling *smarter* through distribution charges coupled to a steadier, more efficient utilization of the distribution infrastructure. Second, the desired increase in price levels for small-scale producers has not materialized ⑦. Clearly, additional buy-side competition in that market should lead to higher prices, but Brokers have simultaneously improved their pricing strategies. The upshot is a slightly reduced price level for small-scale producers in later years. And finally, while Brokers now buy less energy in the balancing market, we observe a new tendency to sell oversupply on short notice ⑥. While this need not be a problem per se, it suggests that we need ways of absorbing such excess energy, such as cold storage facilities that double as thermal energy storage.

¹³The trend towards more competitiveness has partly been reversed in 2014 with AgentUDE, cwiBroker, and TacTex increasingly dominating the field, see also Figure 4.9.




Figure 4.11. Brokers’ strategic retail market positioning. Each panel shows the relationship between volume captured and average price paid. Left: Market for energy consumption (Brokers’ sales). Right: Market for small-scale production (Brokers’ purchases). Lighter labels indicate higher variations in market shares between games.

Market Quality

We re-ran the 59 games of the 2013 finals under identical environmental circumstances but with only the incumbent monopoly, and without competitive Brokers. That is, we conducted 59 alternative simulations that only differed in the absence of a competitive retail market and competitively designed brokers. This procedure yielded 59 pairs of games for which we analyzed several important societal indicators. Most of the indicators from Table 4.6 were already introduced in the PROCESS Quality analysis. The key difference here is that we study their development between *alternate realities* that differ only in the presence of the Broker artifacts, not between multiple iterations of the PROCESS.

The targeted directions of change are again viewed from a societal perspective. In particular, we aim for competitive markets ❶ with low ❷ and stable ❸ consumer prices after the introduction of Brokers. Because the incumbent monopoly already offers stable (albeit high) prices, retail price variability can at best be expected not to increase by much. Balancing ratios should decline ❹ as Brokers improve their forecasting abilities in an effort to avoid high balancing costs. One additional indicator of importance is the **Load Factor**, the

| Metric | | 3 Brokers | 5 Brokers | 8 Brokers | All Games | Target | |
|--------------------------|----------|-----------|-----------|-----------|-----------|--------|-----|
| Retail Price | Sale | -0.42* | -0.44* | -0.45* | -0.43* | ↘ | ✓ ② |
| | Purchase | -0.01* | -0.02* | -0.04* | -0.02* | ↗ | ✗ ⑥ |
| Retail Price Variability | Sale | 0.02* | 0.017* | 0.018* | 0.018* | → | ✗ ③ |
| | Purchase | 0.002* | 0.005* | 0.002* | 0.004* | → | ✓ |
| Retail Volume Change | Sale | 1.4* | 1.41* | 1.4* | 1.4* | ↘ | ✗ ⑦ |
| | Purchase | 1.02* | 1.02* | 1.02* | 1.02* | ↗ | ✓ ⑧ |
| HHI | Sale | -3050* | -4953* | -6802* | -4433* | ↘ | ✓ ① |
| | Purchase | -216 | -1132* | -27* | -746* | ↘ | ✓ |
| Balancing Ratio | Sale | 0.02* | 0.07* | -0.01 | 0.05* | ↘ | ✗ ④ |
| | Purchase | 0.02* | 0.07* | -0.01 | 0.05* | ↘ | ✗ ④ |
| Load Factor | | 0 | -0.01 | -0.01 | -0.01 | ↗ | ✗ ⑤ |

Table 4.6. Quality indicators for Power TAC's competitive retail electricity market as mean difference between competitive 2013 tournament games and corresponding monopolistic games, interpreted from the Broker's perspective. Stars indicate significance at the 95% level, circled numbers are referenced from the main text, checkmarks indicate targets that have been met.

ratio between the average and the maximum load on the distribution infrastructure. High load factors ⑤ are desirable in that they indicate a steadier, more efficient use of the expensive infrastructure which in turn leads to lower environmental impact (e.g., fewer power line corridors) and reduced consumer prices.

Table 4.6 shows an ambivalent impact of introducing competitive retail Brokers compared to the monopolistic situation. As before, we observe large and significant increases in competitiveness ① and decreases in retail prices ②. We also note a slight increase in the volume of energy sourced from small-scale renewables ⑥. But these social positives come with several negatives that require careful consideration. Through competition, retail markets become significantly more *turbulent* as reflected in higher retail price variability ③, the need for additional balancing power ④, and slightly lower load factors ⑤. Parts of the favorable competitive retail prices are, in other words, socialized by Brokers that put higher stress on consumers and the physical infrastructure in an effort to offer the lowest possible prices while remaining competitive.

4.4.5 Interim Balance

These preliminary results are ambivalent, but expected and desirable from a CB perspective. They are expected in that designers have made meaningful, high-return design choices to obtain control over their complex environment. After only two iterations, Brokers autonomously set competitive prices, forecast demand with acceptable accuracy, and make explicit use of most strategic options such as sourcing from small-scale production. As the field of competitors grows more sophisticated, designers will also need to consider more subtle effects arising from complex interactions like those in the balancing market. CB's iterative nature continually raises the bar for designers. In the future, we expect to see more nuanced strategies that incorporate ideas with successively lower return on design investment.

The results are desirable in that they raise the right questions. The analyses we presented above are objective in the sense that we provide no value judgement on the relative importance of individual targets. It is unclear whether all targets can be achieved simultaneously and it is the subject of politics to prioritize between them. Should society care about higher stress on distribution infrastructures? Should the bulk of the welfare effects from price reductions go to consumers, or should parts of these benefits be used to incentivize more efficient energy use and investments in renewables? And should consumers be subjected to more pronounced price swings, or are additional policies necessary to dampen these effects? Competitive Benchmarking as instantiated in Power TAC invites these questions, and provides the means to study the effect of alternate answers on businesses and consumers.

4.5 Discussion

Throughout this article we have portrayed Competitive Benchmarking as an effective combination of existing tools and techniques, integrated into a coherent method for IS research on societal challenges. In this section, we discuss connections with several influential streams of work, and we describe a set of best practices for implementing CB based on our own experiences with Power TAC.

4.5.1 Related Work

Benchmarking has long been recognized as an important tool for improving products and organizational performance (Hindle, 2008). Walter Chrysler regularly bought and disassembled new Oldsmobiles to better understand his competition (Shetty, 1993), and Ford engineers allegedly anatomized some fifty German and Japanese cars before embarking on the construction of the popular Ford Taurus (Mittelstaedt, 1992). But the key event that popular-

ized benchmarking as a distinct concept among management practitioners and scholars was Xerox Corp.'s benchmarking-driven turnaround in the late seventies (Garvin, 1993). Today, a wide range of activities are recognized as benchmarking, ranging from informal comparisons within corporate boundaries to highly structured analyses of competitive postures across industries.

Competitive Benchmarking as we define it here is rooted in the competitive research approach pioneered by the **Trading Agents** community (Greenwald and Stone, 2001, Wellman, 2011, Ketter and Symeonidis, 2012) which aims to deploy techniques from Artificial Intelligence and other computational disciplines to trading applications. Trading Agent Competitions (TAC) challenge researchers to devise software agents for complex, uncertain environments like supply chains (Arunachalam and Sadeh, 2005) and advertisement auctions (Jordan and Wellman, 2010), and to benchmark and improve them iteratively. This practice has been found to foster creativity, improve learning, and facilitate innovation based on deep introspection (Garvin, 1993, Shetty, 1993, Drew, 1997).

The key difference between CB and TAC is our emphasis on real-world **ALIGNMENT**. Theories developed through CB must be representative of real-world dynamics to the degree that they can be used for policy guidance. And artifacts designed through CB must meet the usefulness criterion of IS design science. That is, they must address an important and relevant business problem, and their utility, quality, and efficacy must be clearly demonstrated (Hevner et al., 2004). While previous TACs have been inspired by business settings, their focus has been on stylized decision problems among autonomous software agents, and system-level consequences of interventions played no significant role in them. CB also improves over TAC by providing human-system interaction facilities that can be used in training human decision-makers and in decision support studies. Such facilities are valuable in complex environments like financial markets, where training based on historical data streams “cannot readily model market impact [and offers] essentially nothing toward understanding the current or future overall system-level dynamics ... [it] can tell you what happened, but not what might happen next, nor what might have happened instead” (Cliff and Northrop, 2012).

A competitive element is also present in **Research Competitions** like those organized by Netflix (Bell and Koren, 2007) and Kaggle (<http://www.kaggle.com>). Research competitions encourage participants to develop solutions for data mining, forecasting, and optimization problems ranging from chess predictions to disease spread analyses. Like CB, they attract diverse communities of experts from various technical backgrounds, and the rich repositories of resulting artifacts can be used to explore broad solution spaces and derive rigorous design theories.

But while research competitions leverage iteration and benchmarking, they forgo the benefits of collaborative analysis, learning, and improvement that are central to CB. Furthermore, they are based on static datasets instead of dynamic PLATFORMs, because they aim to address *complicated* instead of complex problems. Participants are limited to deploying promising techniques to prefabricated datasets provided by a self-interested sponsor, whereas identification and modeling issues remain out of their scope. Research competitions are therefore limited in their ability to produce insights and solutions for societal challenges for which the problem definition itself constitutes a significant hurdle (Wagstaff, 2012).

By contrast, work on **Agent-based Computational Economics** (ACE, Tesfatsion, 2006) and **Agent-based Virtual Worlds** (ABVW, Chaturvedi et al. 2011) foregrounds these modeling aspects in an effort to derive possible futures of high-complexity environments, and the paths to these futures, based on realistic assumptions. Creating agent-based models that faithfully capture interesting aspects of real-world phenomena is difficult, because the represented phenomena are often vague, unstructured, and perennially changing. ABVW research therefore promotes the use of simulation platforms on which user-contributed content can be executed, so that users become *citizen developers* who contribute to the richness and validation of the models.

CB PLATFORMs are Virtual Worlds by definition, and design guidelines like the involvement of citizen developers are important in their construction.¹⁴ But in contrast to work on ABVW, we *make use* of PLATFORMs as one of several components in an overarching method for IS research on societal challenges to alleviate the problem that:

[a]nalytical methods give elegant closed-form solutions to narrow, but well-defined, problems; empirical methods allow researchers to test theories at different levels of analyses; and computational methods allow researchers to build high fidelity simulations. However, none of these methods are particularly effective for studying large-scale problems (Chaturvedi et al., 2011, p.682).

Beyond Virtual Worlds, CB therefore adds the novel notion that software-based PLATFORMs can be used as the medium for capturing a community-created scientific paradigm, and as the infrastructure for a new type of competitive research process. The iterative, competitive nature of the PROCESS is essential in the context of societal challenges, because it brings the competitive co-evolution of artifacts into the laboratory, as well as the environmental complexity captured by regular ABVWs.

Bringing elements of real-world evaluations into the laboratory is also prominent in the use of **Serious Games** for artifact evaluation (Lang et al., 2009) where participants engage in

¹⁴More specifically, CB platforms are so-called Mirror Worlds, one of the two subtypes of ABVW (Chaturvedi et al., 2011).

games that incorporate the artifact under study, e.g., a particular market mechanism. Similar to a CB PROCESS, these participants can evaluate the artifact more realistically than an isolated research group, since their diverse, creative behaviors will better pinpoint unintended design flaws. But Serious Games focus on human evaluations of a single artifact, whereas CB studies the competitive co-evolution of artifacts in complex environments. Moreover, unlike CB, Serious Games provide no tools for handling the scale and complexity inherent in research on societal challenges.

Finally, we should point out that Power TAC fills several recently proposed **IS research agendas** on energy and sustainability with life (Bichler et al., 2010, Melville, 2010, Watson et al., 2010). We provide more detail on this connection, and on Power TAC's coverage of the three research agendas in the appendix.

4.5.2 Best Practices

Power TAC is the first comprehensive implementation of the ideas presented in this article, and initiating and organizing this CB effort together with a global community of dedicated researchers has been truly insightful. In this section, we present a set of best practices that we have collected so far, and that can inform future CB researchers.

CB's reliance on software-based representations is perhaps the most visible departure from established IS research methods. These representations have a fruitful tradition in the software engineering field, and agent-based models in particular have long been accepted into the scientific mainstream (Gilbert and Troitzsch, 2005). But a scientific paradigm must be readily understandable and modifiable before it is accepted by a community. Purely code-based representations become problematic as the diversity of the community grows and we must therefore emphasize the important distinction between **software** (which includes logical representations like decision trees and process charts) on the one hand, and **program code** on the other. For a PLATFORM to successfully represent a scientific paradigm requires high-quality software, not just high-quality program code. We contest that PLATFORMs that consist only of the latter are bound to fail in addressing interdisciplinary challenges of societal scale, and we recommend a clear distinction between logical models of the paradigm, and their translation into machine-executable statements. While the former are of critical importance to CB, the latter can be created by well-trained engineers that do not necessarily have to be involved in the core research effort.

The intention behind our recommendation is not to downplay the difficulty or importance of actually building a PLATFORM. In CB, like in software projects more generally, it is easy to underestimate the effort and skill needed for the implementation of an idea. A PLATFORM must provide a sound foundation on which the PROCESS can proceed, which

in turn requires technical qualities like scalability, security, and extensibility. Moreover, the distributed nature of ALIGNMENT requires an understanding of concepts like version and configuration management that are well-understood by professional software engineers but not often by amateur programmers. In Power TAC, many of the initial software engineering tasks were done by experienced programmers that are part of our research community. But since these initial days, we have gradually hired several full-time software engineers who are now responsible for maintaining the technical foundation of Power TAC's PLATFORM. Separating the logical problem definition from its translation into code is another important separation of concerns and we recommend this also because it frees up additional research capacity.

Moving to the conceptual level, one critique we sometimes encounter is that formal representations like those built on by CB require a full understanding of the dynamics of each part of the system, some of which may be unknown for societal challenges that have yet to unfold. For example, it is largely unclear how electricity consumers would react to real-time prices, simply because such pricing schemes are technically unfeasible or forbidden in most states. While we consider the critique legitimate, we feel that forcing modelers to make all assumptions explicit is a benefit of agent-based modeling and, by extension, CB. Perhaps a more fitting critique is that CB requires the right choice of abstraction. We might counter that artificial systems such as markets or other social forms of organization exhibit

properties that make them particularly susceptible to simulation via simplified models. ... [T]he possibility of building a mathematical theory of a system or of simulating that system does not depend on having an adequate micro theory of the natural laws that govern the system components. Such a micro theory might indeed be simply irrelevant (Simon, 1996, p.19).

But the combination of naturalistic and simulation-based elements in CB indisputably inherits strengths and weaknesses from each that researchers must carefully consider before embarking on a CB effort (e.g., North and Macal 2007). In particular, skillful modeling remains as critical in CB as it is in any other simulation effort. One common modeling pitfall is the temptation to *boil the ocean*, that is, the attempt to capture every possible detail. ALIGNMENT and PROCESS are purposely iterative and allow researchers to start small and gradually increase the level of sophistication. We encourage CB researchers to make use of these facilities and to resist the temptation of *big design up-front*.

Our previous recommendation comes with one caveat: PLATFORMs must provide a certain level of realism before they attract a research community. This should not keep CB initiators from iteratively building their understanding of a societal challenge and the corresponding PLATFORM. But they must brace themselves for an initial investment of time and

resources for which current academic incentive systems offer little reward. In the case of Power TAC, it took approximately two years for the PLATFORM to become attractive and stable enough for other researchers to build upon it. We advise future CB researchers to carefully plan this period, and to ensure that resources are available for its duration.

Finally, we should remark that some challenges stand to gain more from CB research than others. In our opinion, these are the challenges that are large in scale and scope, characterized by essential complexity and prohibitive costs of potential social negatives, and that require interrelated advances in theory development and design. Many important CB principles certainly carry over to, for example, common design problems like those currently tackled in research competitions. But CB's characteristic up-front investments in ALIGNMENT and PLATFORM development offer the highest benefits in situations where the swift, community-based development of a shared paradigm, the strategic co-evolution of artifacts, and system-level evaluations matter.

4.6 Conclusions

Many important challenges of our time transcend individuals, organizations, and markets, which have been the traditional focus areas of IS research. Grand challenges like sustainable energy, climate change, and financial market stability can only be fully understood at the societal or global level, and they require interrelated advances in theory development and design that are best provided by interdisciplinary research communities (European Commission, 2011). IS innovations have fueled these challenges through their enabling role in globalization, and they should play a similarly important part in their resolution.

Any intervention in complex social systems requires careful consideration of system-level consequences including potential social negatives (Rittel and Webber, 1973) which is particularly difficult in the case of societal challenges. We argue that the single-investigator research model of IS research is limited in its ability to scale to the societal level, and to deliver proactive solutions in addition to reactive insights. Competitive Benchmarking effectively addresses these limitations through a coherent combination of ideas from Benchmarking, Trading Agents research, Agent-Based Computational Economics, Agent-Based Virtual Worlds, and several other fields. CB scales to large, interdisciplinary communities of researchers, and it encompasses both behavioral research (insights) and design science (solutions).

At the heart of CB is the notion of a community-created problem definition that we call a scientific paradigm, i.e., a triplet consisting of (1) problem definitions and models of societal challenges, (2) shared vocabularies, and (3) lists of important research questions

(see Section 4.2). This paradigm is captured using malleable, executable software-based representations. It is important that these representations are malleable, because they are maintained and validated by a distributed research community through ALIGNMENT. And it is important that they are executable, because they serve as the foundation for a competitive research PROCESS in which artifacts and theories are created, benchmarked, and iteratively improved. This PROCESS starts from one set of peer-reviewed assumptions which breeds artifacts and theories that are readily comparable after the fact, and it reduces the need for protracted ex-post scrutiny and increases scientific cycle speed.

The key to scalability within CB lies partly in reducing waste from multiple, incomparable paradigms and results, and partly in better focusing individual efforts on theory and artifact creation. CB does not aim to socialize the scientific process, but it distributes the community's joint efforts to a greater effect, and it promotes intense scientific competition in areas where individual achievements matter.

Power TAC is one concrete instance of CB that aims to address the sustainable energy challenge and we have been conducting it together with a global community of researchers for over four years. To date, this community has created a diverse set of candidate designs for a novel class of IS artifacts that we call Brokers, and that can contribute to sustainability objectives like better integration of renewable energy sources. By turning worthy causes into viable business models, Brokers provide “an opportunity to create shared value – that is, a meaningful benefit for society that is also valuable to the business” (Porter and Kramer, 2006). Brokers form the foundation for a series of radically new IS-driven, customer-centric business models based on personalized services, and they afford customer participation where it is currently unfeasible. Power TAC has also spawned work on the economic benefits that competitively designed Brokers can offer towards solving the “grand challenge” of providing affordable, reliable, and sustainable energy (Massoud Amin and Wollenberg, 2005). Just how grand this challenge is may still be the subject of heated debates. But significant upsides can be achieved in the form of sustainable energy systems with reduced environmental footprints and improved economic productivity that critically depend on IS innovations.

Like the sustainable energy challenge itself, Power TAC continually evolves through ALIGNMENT. Potential new research questions that are currently under consideration include the costs and benefits of large-scale storage and locational marginal pricing (Stoft, 2002), the role of customer-side automation and security considerations, and demand-based distribution charges that incentivize more efficient usage of the existing distribution infrastructures.

An interesting methodological question is how rigorous design theories can be derived

from the comprehensive data recorded by Power TAC. Although a lively debate has been held on what constitutes a proper design theory, e.g., (Walls et al., 1992, Gregor and Jones, 2007, Venable et al., 2012), it is less clear how such a theory is best constructed starting from raw observational data. We suspect this is because artifact performance data like those generated by Power TAC's PROCESS were previously simply unavailable. Within the Trading Agents community similar questions have been answered using descriptive analyses (Ketter et al., 2013c), formal statistical or information-theoretical methods (Andrews et al., 2009), and empirical game theory (Jordan et al., 2007), and we are currently evaluating their benefits for the derivation of principled IS design theories.

Interestingly, CB as a research method is itself a designed artifact, and the Power TAC process quality analysis we presented in this article is one example of how CB's comprehensive data record reaches beyond analyses of theories and designed artifacts, into the assessment of the method itself. This has important ramifications as this new level of visibility allows the Power TAC community to purposely control the degree of novelty and challenge admitted into the PROCESS, and it provides a sound measurement of the rates of insight and innovation delivered (Venable and Baskerville, 2012). Based on our results to date, Power TAC is making good progress on both counts, and we propose that IS research on other societal challenges stands to gain from the benefits of Competitive Benchmarking as well.

Chapter 5

Conclusions

The preceding work aimed to advance the development of IS design theories for brokers. In Section 5.1 we summarize our main findings, and their implications for theory and practice. In Section 5.2, we discuss areas for future research.

5.1 Summary of Main Findings and Implications

We presented our contributions in three separate studies that covered different facets of the broker design challenge (see Figure 5.1).

Chapter 2 – Broker Design

In our first study, we developed a top-level design theory for brokers that learn from experience, and that behave effectively in their ever-changing retail electricity environments. Our work builds on the theory of Reinforcement Learning (Sutton and Barto, 1998), which we combined with regularization and feature selection techniques. We found that different combinations of these design elements perform well under different environmental circumstances, and we provided guidance on when to use which combination.

Specifically, we found that simple manual and regularized approaches performed well in static environments, but their performance deteriorated quickly as environmental conditions changed. The regularized approach generalized slightly better and required significantly less domain knowledge in the design process. Both approaches are computationally efficient and can be applied without a model of a target environment. When an environmental model is available, feature selection techniques can be leveraged to obtain significantly better performance in the target environment and beyond. The application of feature selection techniques




Figure 5.1. Areas of contribution for the three studies that comprise the dissertation, repeated here from Figure 1.1 for convenience

led to strategies that generalized better than, e.g., the regularized approach. Generalization is desirable because of potential shifts in a smart market environment, but also to effectively accommodate potential mismatches between the environmental model used for feature selection and the actual target environment. Feature selection requires significant computation before executing the strategy, but has little impact on the derived strategies' runtime requirements. In our experiments, we found that greedy feature selection generally delivers high-quality results at low computational costs, and we therefore recommend it over other feature selection techniques (e.g., using genetic algorithms) for our application. We summarize the resulting theory for design and action in Table 5.1.

The **theoretical contribution** of this study lies in advancing the state-of-the-art in the area of retail electricity broker design. Furthermore, we aimed to provide guidance on this nascent topic area to future broker designers. Other RL-based design work has indeed referenced our study since it was originally published, e.g., (Babic and Podobnik, 2014, Reddy, 2013, Urieli and Stone, 2014, Valogianni et al., 2014a). Its **managerial and societal importance** lie in providing evidence for the viability of a decentralized distribution system control scheme built on smart markets and economic incentives. Furthermore, intelligent intermediation is one possible business model for utility companies in future Smart Grids.

| Theory component | Instantiation |
|-------------------------|--|
| Means of representation | Mathematical formulae, computer programs, words |
| Primary constructs | Broker, tariffs and subscriptions, small-scale consumers and producers, retail market, wholesale market |
| Scope | Brokers in retail electricity markets |
| Causal explanations | Provided by the Reinforcement Learning framework (Sutton and Barto, 1998) |
| Testable propositions | <ul style="list-style-type: none"> • Profitable autonomous electricity brokers for Smart Electricity Markets can be constructed using the Reinforcement Learning framework. • Regularization and feature selection techniques lead to improved control performance and better generalizability when the characteristics of the consumer and producer population are volatile. • Greedy feature selection generally delivers high-quality results at lower computational costs than other considered feature selection techniques. |

Table 5.1. Components of our broker design theory. A theory for design and action is a prescriptive theory that informs the choices of artifact designers. As such, its principal components are constructs and relationships between those constructs that are prescriptive in nature, backed by causal explanations, and empirically testable.

We contend that the insights gleaned from this work also generalize to domains other than electricity. For example, other work has recently explored the potential of concurrent RL for controlling general customer interactions through web and e-mail channels (Silver et al., 2013). Domains such as traffic congestion pricing, where real-time visibility and controllability are now becoming feasible, or complex B2B marketplaces (Lu et al., 2013) also stand to benefit from these advances.

Chapter 3 – Preference Learning

The idiosyncratic preferences of their customers are a particular challenge for brokers. Brokers must be capable of reasoning about customer responses to changes in tariff terms to incentivize behavioral change. We therefore studied in greater detail the problem of learning from past customer choices to improve future decisions in our second study.

In the study, we presented *GTM*, a Bayesian nonparametric preference model that leverages characteristics of consumer choice settings (many consumers and observed choices; choices described by a small number of attributes; and consumers with limited cognitive

resources for evaluating trade-offs) to efficiently learn from noisy observations. To assess *GTM*'s predictive and performance characteristics, we compared it against two recent nonparametric GP models, one semi-parametric GP model, and the well-established Mixed Logit model. We found that *GTM*'s predictions are of comparable quality to those made by state-of-the-art nonparametric choice models. However, *GTM* scales much better in consumer choice situations than its benchmarks. We summarize the resulting theory for design and action in Table 5.2.

| Theory component | Instantiation |
|-------------------------|---|
| Means of representation | Mathematical formulae, computer programs, words |
| Primary constructs | Users, alternatives, choices |
| Scope | Discrete choice problems |
| Causal explanations | Provided by the Gaussian process framework (Rasmussen and Williams, 2006) |
| Testable propositions | <ul style="list-style-type: none"> The Gaussian process trade-off model (<i>GTM</i>) predicts future choices with an accuracy that is at par with existing state-of-the-art models, but at a significantly lower computational cost. |

Table 5.2. Components of our design theory for discrete choice prediction.

The **theoretical contribution** of this study lies in combining elements of sparse and structured Gaussian process models with Laplace inference, and in showing the effectiveness of this combination for consumer preference learning tasks. Its **managerial importance** lies in the development of a practical, conceptually simple method for predicting consumer choices that nevertheless provides a principled Bayesian treatment of predictive uncertainty. Uncertainty estimates are an important ingredient for decision-making tasks such as that faced by electricity brokers.

We intentionally separated this second study from the broker design work of Chapter 2, because it is a general purpose preference learning method that transcends the broker problem. In our empirical evaluation, we also assessed the performance of our method on other preference learning tasks (political elections, automobile purchases), and it unsurprisingly performs well as long as the choice task matches the characteristics outlined above. Moreover, the technical foundations developed in this work fully generalize to approximate posterior inference using Laplace's method in other machine learning tasks such as classification.

Chapter 4 – Competitive Benchmarking

In our final study, we proposed and evaluated Competitive Benchmarking (CB), a novel research method for accelerating progress on IS design tasks where the underlying real-world phenomena evolve rapidly, and where the social cost of failure is high. The first instance of Competitive Benchmarking is Power TAC, a research effort through which a global community of researchers now addresses the broker design challenge. Their work has resulted in a rich repository of fine-grained data that are interesting to practitioners and policy makers alike. Analyzing the 2012 through 2014 Power TAC data, we find preliminary evidence of the benefits of CB as a research method (process quality), and of the system-level effects of introducing brokers into retail electricity markets (market quality).

With respect to **process quality**, we find that the heterogeneous broker designs attracted by Power TAC indeed led to large and significant performance differences. The best broker designs outperformed competitors by more than an order of magnitude. Several outstanding designs were new Power TAC entrants who surpassed relatively more established designs, which illustrates the powerful effect of competitive innovation. After only two CB iterations, brokers now autonomously set competitive prices, forecast demand with acceptable accuracy, and make explicit use of most strategic options such as sourcing from small-scale production.

Our **market quality** analyses showed an ambivalent impact of introducing competitive retail brokers. We observed large and significant increases in competitiveness, decreases in retail price, and a slight increase in the volume of energy sourced from small-scale renewables. But these social positives came with several negatives that require careful consideration. Through competition, retail markets become significantly more turbulent as reflected in higher retail price variability, the need for additional balancing power, and slightly lower load factors. Parts of the favorable competitive retail prices are, in other words, socialized by brokers that put higher stress on consumers and the physical infrastructure in an effort to offer the lowest possible prices while remaining competitive. These analyses are objective in the sense that we provide no value judgement on the relative importance of individual targets. It is unclear whether all targets can be achieved simultaneously, and it is the subject of politics to prioritize between them. Competitive Benchmarking as instantiated in Power TAC invites these discussions, and it provides the means to study the effect of alternate outcomes on businesses and consumers.

The **theoretical contribution** of this study lies in formalizing a novel research method that can help interdisciplinary research communities tackle complex challenges of societal scale. The key to swift innovation and scalability within CB lies partly in reducing waste, and partly in better focusing individual efforts on theory and artifact creation. CB distributes

a community's joint efforts to greater effect, and it promotes intense scientific competition in areas where individual achievements matter. Its primary **societal and managerial benefits** lie in its ability to address challenges that evolve rapidly, and that require solutions in addition to mere insights, a combination that otherwise presents IS researchers with difficulties (Hey et al., 2009). The CB process yields tangible research results, and it facilitates a new separation of concerns where practitioners can contribute design ideas that are then rigorously evaluated to according scientific standards.

In the study, we carefully separated the formalization of the CB research method from its instantiation to the sustainable electricity challenge (Power TAC). We expect that CB immediately generalizes to other societal challenges of essential complexity that are large in scale and scope, that are currently unrealized, that progress at a rapid pace, and for which the social costs of erroneous interventions are prohibitive.

5.2 Limitations and Future Research

The development of design theories is an ongoing process. Much like our studies aimed to improve over previous work, we hope and expect that future work will improve over ours. Improvements in our reference theories (Reinforcement Learning, Gaussian processes, etc.) are one possible source of progress. But even within their current state-of-the-art, several areas for future research remain, and we now discuss them in turn.

First, all of our studies are based on environmental models that can be extended to cover other important real-world phenomena. For example, the customer interactions we modeled in Chapter 2 were limited to fixed tariff offerings, and a model of customer choice driven by real-world data, but idealized nonetheless. Future work could explore the performance of our broker design in increasingly sophisticated retail electricity markets with advanced tariff structures, storage devices, and customers who adopt electronic agents of their own. The preference learning study in Chapter 3 made use of stated preferences from human decision-makers, and it would be interesting to assess our model's performance on actual real-time tariff choices. Chapter 4 employed what is perhaps the most sophisticated Smart Grid economics model to date, fueled by the contributions of the global Power TAC community. But even within Power TAC's sophisticated environmental model, many possibilities for extensions (e.g., physical network characteristics) remain.

Second, while the analyses in Chapter 4 give a good first impression of the benefits of CB, and of the societal ramifications of competitive electricity brokerage, they are based on preliminary data. More Power TAC iterations will generate more data, and provide more conclusive evidence. CB's iterative nature continually raises the bar for designers. In the

future, we expect to see more nuanced strategies that incorporate ideas with successively lower return on design investment.

Third, more work is required at the seams between the three studies. For example, while we presented comprehensive empirical evaluations of our preference model in Chapter 3, these evaluations were conducted on static tariff choice datasets. More work remains to integrate our *GTM* preference model with the dynamic decision-making scenario of Chapter 2. The work of Chapter 2, in turn, predates the first Power TAC tournament held in July 2013, and it would be interesting to assess our broker's performance in the broad field of sophisticated benchmark strategies that Power TAC now offers.

Finally, the models in Chapters 2 and 3 should make the leap into Smart Grid pilot projects. For example, Power TAC is now being used as part of the Cassandra platform (www.cassandra-fp7.eu) for strategic decision-making. First versions of this platform have been deployed in pilot projects in Coventry (UK), Milan (Italy), and Lulea (Sweden). Customers in participating pilot buildings receive real-time feedback on their electricity consumption, see the impact of their own decisions on the overall power system, and coordinate through social networks to reduce their impact as a community. The performance of autonomous electricity brokers, including the automatic learning of customer preferences, should similarly be evaluated under such real-world conditions.

Summary

The shift towards sustainable power systems is one of the grand challenges of the twenty-first century. Decentralized production from renewable sources, electric mobility, and related advances are at odds with traditional power systems where central large-scale generation of electricity follows inelastic consumer demand. Smart Markets and intelligent Information Systems (IS) could alleviate these issues by providing new forms of coordination that leverage real-time consumption information and prices to incentivize consumer behaviors that remain within the grid's operational bounds. But the best design for these systems, and the societal implications of different design choices are largely unclear. This dissertation makes three contributions to the debate:

First, we propose and evaluate a design theory for Brokers, a novel class of IS-based intermediaries in retail electricity markets that provide participants with additional information and fine-grained economic incentives. Endowing Brokers with all possibly successful behaviors at design time is futile, because these behaviors are impossible to enumerate, and quickly become obsolete in Brokers' ever-changing retail electricity environments. Instead, our designs employ Reinforcement Learning (RL) to *learn* effective behaviors by observing their environment, taking actions, monitoring the long-run consequences that these actions entail, and updating their behavior accordingly. We propose and evaluate several RL-based designs, and we identify the most effective design conditional on prevailing market conditions. Our simulation-based evaluations are based on data from the Ontario wholesale market, a complete micro-level model of appliance usage in private households, and several benchmark designs for Brokers proposed in the literature. We find that our designs outperform these earlier works by a significant margin.

Second, we study in greater detail the Broker's core problem of learning from past customer choices as a basis for future decisions. To incentivize favorable behaviors, Brokers must be capable of reasoning about customer responses to changes in tariff terms, even if the prior experiences they can draw from are limited. Furthermore, observed choices are usually subject to behavioral inconsistencies, such as inertia. A Broker preference model

must respect these inconsistencies, because the uncertainty arising from them is a crucial ingredient for autonomous decision-making: Brokers should only make high-value decisions autonomously if past evidence suggests that they will be correct with high probability, and prompt their users for additional information otherwise. To address these requirements, we propose a non-parametric Bayesian preference model based on Gaussian processes that learns from limited data, and that quantifies the certainty of its predictions as input to the Broker's autonomous decision-making task. Probabilistic inference in non-parametric Bayesian models is often computationally expensive, but by using advances in sparse and structured Gaussian processes, we are able to reduce the costs of inference substantially. We evaluate our model on several real-world choice datasets, including an electricity tariff choice dataset that we collected specifically for this study on a crowdsourcing platform. We find that our model is competitive with state-of-the-art approaches in terms of predictive accuracy, that it is significantly faster, and that it scales better to large customer populations.

Finally, we propose and study Competitive Benchmarking (CB), a novel research method for effective IS artifact design in complex environments like power systems, where the social cost of failure is prohibitive. Traditional design science studies are well-suited for initially identifying and studying individual designs, but their homogeneous setup limits their ability to quickly detect promising alternatives and possible social negatives. CB challenges researchers to devise alternative designs that are regularly pitted against each other in competitions to foster a swift cycle of innovation. The first instantiation of CB is the Power Trading Agent Competition (Power TAC), in which more than a dozen research groups from four different continents now jointly devise, benchmark, and improve Broker designs. Using fine-grained records of this community's results, we quantify performance differences between alternative Broker designs, and between subsequent iterations of the same designs to give preliminary empirical evidence of CB's efficacy as a research method.

The results reported in this dissertation provide guidance on IS design choices for sustainable electricity systems, and they contribute to the foundations for new Smart Grid business models.

Samenvatting

De verschuiving naar duurzame energiesystemen is een van de grote uitdagingen van de eenentwintigste eeuw. Decentrale productie uit duurzame bronnen, elektrische mobiliteit en hieruit volgende nieuwe ontwikkelingen staan op gespannen voet met traditionele energiesystemen waar de centrale productie de inelastische vraag van consumenten volgt. Smart Markets en intelligente informatiesystemen (IS) zouden deze problemen kunnen verlichten door te voorzien in nieuwe vormen van coördinatie die, door gebruik te maken van real-time consumptiegegevens en prijssignalen, consumenten belonen voor gedrag waarmee consumenten binnen de operationele grenzen van het system blijven. Echter, het beste design van zulke systemen en de maatschappelijke gevolgen van het gebruik ervan zijn onduidelijk. Dit proefschrift levert een bijdrage aan deze discussie in de vorm van drie studies.

De eerste studie is gericht op de introductie en evaluatie van een designtheorie voor Brokers bestaand uit een nieuwe klasse van IS-gebaseerde intermediairs voor de retailmarkt van elektriciteit, die klanten van aanvullende informatie en economische prikkels voorzien. Het feit dat de retailmarkt voor elektriciteit sterk dynamisch is maakt het onmogelijk om Brokers vooraf van alle mogelijke succesvolle gedragingen te voorzien. Daarom maken onze designs gebruik van Reinforcement Learning (RL) waarmee effectieve gedragingen geleerd kunnen worden door waarneming van de omgeving, het nemen van beslissingen, het observeren van de gevolgen en het (op basis hiervan) bijstellen van het gedrag. We introduceren en evalueren meerdere RL-gebaseerde designs waarna het meest effectieve design wordt geïdentificeerd afhankelijk van de heersende marktcondities. Onze simulatie-gebaseerde evaluatie maakt gebruik van gegevens uit de groothandel in elektriciteit in Ontario, van een gedetailleerd model van het gebruik van elektrische apparaten in huishoudens en van verschillende benchmark designs uit de wetenschappelijke literatuur. Uit onze studie blijkt dat de resulterende nieuwe designs aanzienlijk betere prestaties opleveren dan bestaande designs.

In de tweede studie focussen we op het centrale probleem van de Broker: het nemen van beslissingen op basis van klantpreferenties uit het verleden. Om prikkels te kunnen aanbieden waarmee gewenst gedrag bij klanten gestimuleerd wordt, moeten Brokers in staat

zijn de reacties van klanten op veranderende economische voorwaarden (zoals prijsveranderingen) te kunnen beredeneren. Echter, gegevens hierover uit het verleden zijn vaak beperkt en bevatten bovendien gedragsmatige inconsistenties zoals inertie. Een Broker preferentiemodel moet met inconsistenties rekening houden omdat de onzekerheid die hieruit voortvloeit een belangrijke factor is in het nemen van autonome beslissingen. Brokers zullen hoogwaardige beslissingen enkel autonoom moeten nemen als er voldoende indicatie voor hun juistheid bestaat, en de gebruiker om aanvullende informatie moeten vragen als dit niet zo is. Om aan deze eisen te voldoen stellen we een non-parametrisch Bayesiaans model gebaseerd op Gauss processen (GP) voor dat uit beperkte gegevens kan leren en dat onzekerheden als basis voor autonome besluitvormingen berekent. Het berekenen van waarschijnlijkheden in zulke modellen is vaak kostbaar in termen van rekenkracht. We maken daarom gebruik van nieuwe Sparse GP and Structured GP technieken om deze kosten aanzienlijk te verlagen. We evalueren ons model op meerdere realistische datasets, onder andere een set met keuzes van elektriciteitstarieven die via een crowdsourcing platform special voor deze studie werd verzameld. De voorspellingen van ons model zijn vergelijkbaar met die van andere state-of-the-art modellen in termen van precisie, maar ons model is substantieel sneller en schaaft beter voor populaties met grote aantallen klanten.

In de derde en laatste studie stellen we Competitive Benchmarking (CB) voor als nieuwe methode voor onderzoek naar effectieve IS designs in complexe omgevingen zoals energiesystemen waar de maatschappelijke kosten van mislukking groot zijn. Traditionele design science studies zijn geschikt voor het identificeren en bestuderen van enkele designs. Echter, hun homogene natuur beperkt hun vermogen snel alternatieven te vinden en negatieve maatschappelijke gevolgen duidelijk te maken. CB daagt onderzoekers uit alternatieve designs te ontwerpen en die in regelmatige wedstrijden tegen elkaar te laten wedijveren wat tot snellere innovatie leidt. De eerste CB instantie is de Power Trading Agent Competition (Power TAC), waarin ruim een dozijn onderzoeksgroepen uit vier continenten Broker designs ontwerpen, vergelijken en verbeteren. Door gebruik te maken van gedetailleerde verslagen van hun werk meten we prestatieverschillen tussen designs, en tussen opeenvolgende CB iteraties. Dit heeft geleid tot een voorlopig bewijs van de effectiviteit van CB als onderzoeksmethode.

De resultaten uit dit proefschrift dragen bij aan betere designs voor duurzame elektriciteitssystemen en aan de fundering van nieuwe Smart Grid bedrijfsmodellen.

Bibliography

Ehsan Abbasnejad, Scott Sanner, Edwin Bonilla, and Pascal Poupart. Learning community-based preferences via Dirichlet process mixtures of Gaussian processes. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, pages 1213–1219. AAAI Press, 2013.

Gediminas Adomavicius and Alexander Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, pages 734–749, 2005.

Gediminas Adomavicius, Alok Gupta, and Dimitri Zhdanov. Designing intelligent software agents for auctions with limited information feedback. *Information Systems Research*, 20(4):507, 2009.

Joan E. van Aken. Management research based on the paradigm of the design sciences: the quest for field-tested and grounded technological rules. *Journal of Management Studies*, 41(2):219–246, 2004.

Greg M. Allenby and Peter E. Rossi. Marketing models of consumer heterogeneity. *Journal of Econometrics*, 89(1):57–78, 1998.

John A. Anderson. Electricity restructuring: a review of efforts around the world and the consumer response. *The Electricity Journal*, 22(3):70–86, 2009.

James Andrews, Michael Benisch, Alberto Sardinha, and Norman Sadeh. Using information gain to analyze and fine tune the performance of supply chain trading agents. In *Agent-Mediated Electronic Commerce and Trading Agent Design and Analysis*, pages 182–199. Springer, 2009.

Ernesto Arias, Hal Eden, Gerhard Fischer, Andrew Gorman, and Eric Scharff. Transcending the individual human mind – creating shared understanding through col-

- laborative design. *ACM Transactions on Computer-Human Interaction*, 7(1):84–113, 2000.
- Raghu Arunachalam and Norman Sadeh. The supply chain trading agent competition. *Electronic Commerce Research and Applications*, 4:63–81, 2005.
- Pelin Atahan and Sumit Sarkar. Accelerated learning of user profiles. *Management Science*, 57(2):215–239, 2011.
- Robert Axelrod. *The evolution of cooperation*. Basic Books, 2006.
- Emerson M. Babb, M.A. Leslie, and M.D. Van Slyke. The potential of business-gaming methods in research. *The Journal of Business*, 39(4):465–472, 1966.
- Jurica Babic and Vedran Podobnik. An analysis of Power TAC 2013 Trial. In *Trading Agent Design and Analysis Workshops*, 2013.
- Jurica Babic and Vedran Podobnik. Adaptive bidding for electricity wholesale markets in a smart grid. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS14)*, Paris, France, May 2014.
- Howard Baetjer. *Software as capital: An economic perspective on software engineering*. IEEE Computer Society Press, 1997.
- Elizabeth White Baker. Relational model bases: A technical approach to real-time business intelligence and decision making. *Communications of the Association for Information Systems*, 33(1):23, 2013.
- Ravi Bapna, Paulo Goes, Alok Gupta, and Yiwei Jin. User heterogeneity and its impact on electronic auction market design: An empirical exploration. *Management Information Systems Quarterly*, pages 21–43, 2004.
- Robert Bell and Yehuda Koren. Lessons from the Netflix prize challenge. *ACM SIGKDD Explorations Newsletter*, 9(2):75–79, 2007.
- Martin Bichler, Alok Gupta, and Wolfgang Ketter. Designing smart markets. *Information Systems Research*, 21(4):688–699, 2010.
- Adriana Birlutiu, Perry Groot, and Tom Heskes. Efficiently learning the preferences of people. *Machine Learning*, pages 1–28, 2012.
- Christopher M. Bishop. Training with noise is equivalent to Tikhonov regularization. *Neural Computation*, 7(1):108–116, 1995.

- Christopher M. Bishop. *Pattern recognition and machine learning*, volume 4. Springer New York, 2006.
- Avrim L. Blum and Pat Langley. Selection of relevant features and examples in machine learning. *Artificial Intelligence*, 97(1):245–271, 1997.
- Severin Borenstein. The trouble with electricity markets: Understanding California’s restructuring disaster. *Journal of Economic Perspectives*, 16(1):191–211, 2002.
- Craig Boutilier, Ronen I. Brafman, Carmel Domshlak, Holger H. Hoos, and David Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research*, pages 135–191, 2004.
- Ralph Allan Bradley and Milton E. Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, pages 324–345, 1952.
- Steven J. Bradtke and Andrew G. Barto. Linear least-squares algorithms for temporal difference learning. *Machine Learning*, 22(1):33–57, 1996.
- Ronen Brafman and Carmel Domshlak. Preference handling - an introductory tutorial. *AI Magazine*, 30(1):58–86, 2009.
- Lucian Busoniu, Robert Babuska, Bart De Schutter, and Damien Ernst. *Reinforcement Learning and Dynamic Programming using function approximators*. CRC, 2010.
- Lisa Cameron and Peter Cramton. The role of the ISO in US electricity markets: A review of restructuring in California and PJM. *The Electricity Journal*, 12(3):71–81, 1999.
- Sebastián Caussade, Juan Ortúzar, Luis I. Rizzi, and David A. Hensher. Assessing the influence of design dimensions on stated choice experiment estimates. *Transportation Research Part B: Methodological*, 39(7):621–640, 2005.
- Alok R. Chaturvedi, Daniel R. Dolk, and Paul L. Drnevlch. Design principles for virtual worlds. *Management Information Systems Quarterly*, 35(3):673–684, 2011.
- Wei Chu and Zoubin Ghahramani. Preference learning with Gaussian processes. In *Proceedings of the 22nd International Conference on Machine Learning*, pages 137–144. ACM, 2005.

Dave Cliff and Linda Northrop. The global financial markets: an ultra-large-scale systems perspective. In *Large-Scale Complex IT Systems. Development, Operation and Management*, pages 29–70. Springer, 2012.

Debora Coll-Mayor, Mia Paget, and Eric Lightner. Future intelligent power grids: Analysis of the vision in the European Union and the United States. *Energy Policy*, 35(4):2453–2465, 2007.

John Collins, Wolfgang Ketter, and Maria Gini. Flexible decision control in an autonomous trading agent. *Electronic Commerce Research and Applications*, 8(2):91–105, 2009. ISSN 1567-4223.

John Collins, Wolfgang Ketter, and Maria Gini. Flexible decision support in dynamic interorganizational networks. *European Journal of Information Systems*, 19(3):436–448, August 2010.

Antonio J. Conejo, Javier Contreras, and Miguel A. Plazas. Forecasting electricity prices for a day-ahead pool-based electric energy market. *International Journal of Forecasting*, 21(3):435–462, 2005.

Peter Cramton. Electricity market design: The good, the bad, and the ugly. In *Proceedings of the 36th Annual Hawaii International Conference on System Sciences*. IEEE, 2003.

Dapeng Cui and David Curry. Prediction in marketing using the Support Vector Machine. *Marketing Science*, 24(4):595–615, 2005.

John P. Cunningham, Krishna V. Shenoy, and Maneesh Sahani. Fast Gaussian process methods for point process intensity estimation. In *International Conference on Machine Learning*, pages 192–199. ACM, 2008.

Rajdeep K. Dash, Nicholas R. Jennings, and David C. Parkes. Computational-mechanism design: A call to arms. *IEEE Intelligent Systems*, 18(6):40–47, 2003.

Thomas H. Davenport and Jeanne G. Harris. Automated decision making comes of age. *MIT Sloan Management Review*, 46(4):83, 2005.

Kenneth De Jong. Learning with Genetic Algorithms: An overview. *Machine Learning*, 3(2):121–138, 1988.

Mathijs de Weerd, Wolfgang Ketter, and John Collins. A theoretical analysis of pricing mechanisms and broker’s decisions for real-time balancing in sustainable regional

electricity markets. In *Conference on Information Systems and Technology*, pages 1–17, Charlotte, November 2011.

David Dean, Sebastian DiGrande, Dominic Field, Andreas Lundmark, James ODay, John Pineda, and Paul Zwillenberg. Boston Consulting Group report – The Internet Economy in the G-20: The \$4.2 trillion growth opportunity, March 2012. URL <http://www.bcg.com/documents/file100409.pdf>.

William H. DeLone and Ephraim R. McLean. The DeLone and McLean model of information systems success: a ten-year update. *Journal of Management Information Systems*, 19(4):9–30, 2003.

Arthur P. Dempster, Nan M. Laird, and Donald B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 1–38, 1977.

US Department of Energy. 2010 Smart grid system report. Technical report, US Department of Energy, February 2012.

Stephen A.W. Drew. From knowledge to action: the impact of benchmarking on organizational performance. *Long range planning*, 30(3):427–441, 1997.

EPRI - Electric Power Research Institute. Estimating the costs and benefits of the smart grid - a preliminary estimate of the investment requirements and the resultant benefits of a fully functioning smart grid, March 2011.

EPRI - Electric Power Research Institute. Understanding electric utility customers - Summary report, October 2012.

ETPSG. European Technology Platform Smart Grids: Strategic deployment document for Europe’s electricity networks of the future, 04 2010.

European Commission. Horizon 2020 - The Framework Programme for Research and Innovation, November 2011. URL <http://ec.europa.eu/research/horizon2020>.

Eurostat. Electricity consumption of households for EU 27, 2011.

Theodoros Evgeniou, Constantinos Boussios, and Giorgos Zacharia. Generalized robust conjoint estimation. *Marketing Science*, 24(3):415–429, 2005.

- Theodoros Evgeniou, Massimiliano Pontil, and Olivier Toubia. A convex optimization approach to modeling consumer heterogeneity in conjoint estimation. *Marketing Science*, 26(6):805–818, 2007.
- Peter Fox-Penner. Rethinking the grid: Avoiding more blackouts and modernizing the power grid will be harder than you think. *The Electricity Journal*, 18(2):28–42, 2005.
- Johannes Fürnkranz and Eyke Hüllermeier. *Preference learning*. Springer, 2011.
- Gartner Group. IT Spending Report, Q4 update, 2013.
- David A. Garvin. Building a learning organization. *Harvard Business Review*, pages 78–91, 1993.
- Mohammad Ghavamzadeh, Alessandro Lazaric, Odalric-Ambrym Maillard, and Remi Munos. LSTD with random projections. In *Advances in Neural Information Processing Systems*, pages 721–729, 2010.
- Nigel Gilbert and Klaus Troitzsch. *Simulation for the social scientist*. McGraw-Hill International, 2005.
- Elad Gilboa, Yunus Saatci, and John P. Cunningham. Scaling multidimensional inference for structured Gaussian processes. *Transactions on Pattern Analysis and Machine Intelligence*, 2013.
- Itzhak Gilboa and David Schmeidler. Case-based decision theory. *The Quarterly Journal of Economics*, 110(3):605–639, 1995.
- Sebastian Gottwalt, Wolfgang Ketter, Carsten Block, John Collins, and Christof Weinhardt. Demand side management - a simulation of household behavior under variable prices. *Energy Policy*, 39:8163–8174, 2011.
- David E. Gray. *Doing research in the real world*. Sage Publications, 2004.
- William Greene. *Econometric Analysis*. Prentice Hall, 7th edition, 2012.
- Amy Greenwald and Peter Stone. Autonomous bidding agents in the Trading Agent Competition. *IEEE Internet Computing*, (2):52–60, 2001.
- Shirley Gregor. The nature of theory in information systems. *Management Information Systems Quarterly*, 30(3):611, 2006.

- Shirley Gregor and David Jones. The anatomy of a design theory. *Journal of the Association for Information Systems*, 8(5):312–335, 2007.
- Shengbo Guo, Scott Sanner, and Edwin V. Bonilla. Gaussian process preference elicitation. In *Advances in Neural Information Processing Systems*, pages 262–270, 2010.
- Isabelle Guyon and André Elisseeff. An introduction to variable and feature selection. *The Journal of Machine Learning Research*, 3:1157–1182, 2003.
- Horst Hanusch and Andreas Pyka. Principles of neo-schumpeterian economics. *Cambridge Journal of Economics*, 31(2):275–289, 2007.
- David A. Hensher. How do respondents process stated choice experiments? Attribute consideration under varying information load. *Journal of Applied Econometrics*, 21(6):861–878, 2006.
- Karen Herter, Patrick McAuliffe, and Arthur Rosenfeld. An exploratory analysis of California residential customer response to critical peak pricing of electricity. *Energy*, 32(1):25–34, 2007.
- Alan Hevner and Samir Chatterjee. *Design Research in Information Systems – Theory and Practice*. Springer, 2010.
- Alan R. Hevner, Salvatore T. March, Jinsoo Park, and Sudha Ram. Design Science in Information Systems research. *Management Information Systems Quarterly*, 28(1):75–105, 2004.
- Anthony J.G. Hey, Stewart Tansley, Kristin Michele Tolle, et al., editors. *The fourth paradigm: data-intensive scientific discovery*. Microsoft Research, 2009.
- Tim Hindle. *The Economist guide to management ideas and gurus*, volume 42. John Wiley & Sons, 2008.
- Andy Hoar, Patti Freeman Evans, Carrie Johnson, Douglas Roberge, and Susan Huynh. Forrester Report: Key trends in B2B eCommerce for 2013, October 2012. URL <http://www.forrester.com/Key+Trends+In+B2B+eCommerce+For+2013/fulltext/-/E-RES82102>.
- Neil Houlsby, Ferenc Huszar, Zoubin Ghahramani, and Jose M. Hernández-Lobato. Collaborative Gaussian processes for preference learning. In F. Pereira, C.J.C. Burges, L. Bottou, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 25, pages 2096–2104. 2012.

- Zan Huang, Huimin Zhao, and Dan Zhu. Two new prediction-driven approaches to discrete choice prediction. *ACM Transactions on Management Information Systems*, 3(2):9:1–9:32, July 2012.
- Hussein Ibrahim, Adrian Ilinca, and Jean Perron. Energy storage systems – characteristics and comparisons. *Renewable and Sustainable Energy Reviews*, 12(5):1221–1250, 2008.
- George F. Jenks and Fred C. Caspall. Error on choroplethic maps: definition, measurement, reduction. *Annals of the Association of American Geographers*, 61(2): 217–244, 1971.
- Bjorn Sand Jensen and Jend Brehm Nielsen. Pairwise judgements and absolute ratings with Gaussian process priors. Technical report, 2014.
- Patrick R. Jordan and Michael P. Wellman. Designing an Ad Auctions game for the Trading Agent Competition. In *Agent-Mediated Electronic Commerce. Designing Trading Strategies and Mechanisms for Electronic Markets*, pages 147–162. Springer, 2010.
- Patrick R. Jordan, Christopher Kiekintveld, and Michael P. Wellman. Empirical game-theoretic analysis of the TAC supply chain game. In *Seventh International Conference on Autonomous Agents and Multi-Agent Systems*, pages 1188–1195, May 2007.
- Paul L Joskow. Creating a smarter US electricity grid. *The Journal of Economic Perspectives*, 26(1):29–47, 2012.
- Micha Kahlen, Konstantina Valogianni, Wolfgang Ketter, and Jan van Dalen. A profitable business model for electric vehicle fleet owners. In *IEEE Conference on Smart Grid - Technologies, Economics, and Policies*, pages 1–5, Nuremberg, December 2012. IEEE.
- Micha Kahlen, Wolfgang Ketter, and Jan 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.
- Aijt Kambil and Eric van Heck. *Making Markets: How Firms Can Design and Profit from Online Auctions and Exchanges*. Harvard Business School Press, 2002.
- Toshihiro Kamishima and Shotaro Akaho. Efficient clustering for orders. In *Mining complex data*, pages 261–279. Springer, 2009.

John G. Kassakian and Richard Schmalensee. The future of the electric grid: An interdisciplinary mit study. Technical report, Massachusetts Institute of Technology, 2011. ISBN 978-0-9828008-6-7.

Wolfgang Ketter and Andreas Symeonidis. Competitive benchmarking: Lessons learned from the Trading Agent Competition. *AI Magazine*, 33(1), 2012.

Wolfgang Ketter, John Collins, and Maria Gini. A survey of agent designs for TAC SCM. In *Workshop for Trading Agent Design and Analysis*, 2008.

Wolfgang Ketter, John Collins, Maria Gini, Alok Gupta, and Paul Schrater. Real-time tactical and strategic sales management for intelligent agents guided by economic regimes. *Information Systems Research*, 23:1263–1283, December 2012a.

Wolfgang Ketter, John Collins, Prashant Reddy, and Mathijs de Weerd. The 2012 Power Trading Agent Competition. Technical Report ERS-2012-010-LIS, RSM Erasmus University, Rotterdam, The Netherlands, 2012b. URL <http://ssrn.com/paper=2144644>.

Wolfgang Ketter, John Collins, and Prashant Reddy. Power TAC: A competitive economic simulation of the smart grid. *Energy Economics*, 39:262–270, 2013a.

Wolfgang Ketter, John Collins, Prashant Reddy, and Mathijs de Weerd. The 2013 Power Trading Agent Competition. Technical Report ERS-2013-006-LIS, RSM Erasmus University, Rotterdam, The Netherlands, 2013b. URL <http://ssrn.com/paper=http://ssrn.com/abstract=2268852>.

Wolfgang Ketter, Markus Peters, and John 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 2013c.

Wolfgang Ketter, John Collins, Prashant Reddy, and Mathijs de Weerd. The 2014 Power Trading Agent Competition. Technical Report ERS-2014-004-LIS, RSM Erasmus University, Rotterdam, The Netherlands, 2014. URL http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2411847.

Daniel Kirschen and Goran Strbac. *Fundamentals of Power System Economics*. John Wiley & Sons, 2005.

- George B. Kleindorfer, Liam O'Neill, and Ram Ganeshan. Validation in simulation: Various positions in the philosophy of science. *Management Science*, 44(8):1087–1099, 1998.
- Rob Kling. What is social informatics and why does it matter? *The Information Society: An International Journal*, 23(4):205–220, 2007.
- Jens Kober, J. Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013.
- Ron Kohavi, Llew Mason, Rajesh Parekh, and Zijian Zheng. Lessons and challenges from mining retail e-commerce data. *Machine Learning*, 57(1-2):83–113, 2004.
- Daphne Koller and Nir Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, 2009.
- J. Zico Kolter and Andrew Y. Ng. Regularization and feature selection in least-squares temporal difference learning. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 521–528. ACM, 2009.
- Rodrigue T. Kuate, Minghua He, Maria Chli, and Hai H. Wang. An intelligent broker agent for energy trading: an MDP approach. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, pages 234–240. AAAI Press, 2013.
- Thomas S. Kuhn. *The structure of scientific revolutions*. University of Chicago press, third edition, 1996.
- Michail G. Lagoudakis and Ronald Parr. Least-squares policy iteration. *The Journal of Machine Learning Research*, 4:1107–1149, 2003.
- Fabian Lang, Tim Püschel, and Dirk Neumann. Serious gaming for the evaluation of market mechanisms. In *Proceedings of the International Conference on Information Systems (ICIS 2009)*, 2009.
- Sarah Lichtenstein and Paul Slovic. *The construction of preference*. Cambridge University Press, 2006.
- Gunar E. Liepins and MR Hilliard. Genetic Algorithms: Foundations and applications. *Annals of Operations Research*, 21(1):31–57, 1989.

- Nathan N. Liu, Min Zhao, and Qiang Yang. Probabilistic latent preference analysis for collaborative filtering. In *ACM Conference on Information and Knowledge Management*, pages 759–766. ACM, 2009.
- Manuel Loth, Manuel Davy, and Philippe Preux. Sparse temporal difference learning using LASSO. In *IEEE International Symposium on Approximate Dynamic Programming and Reinforcement Learning*, pages 352–359. IEEE, 2007.
- Yixin Lu, Alok Gupta, Wolfgang Ketter, and Eric van Heck. Designing intelligent software agents for B2B sequential dutch auctions: A structural econometric approach. In *International Conference on Information Systems*, Milan, Italy, 2013.
- Henry C. Lucas Jr, Ritu Agarwal, Omar A. El Sawy, and Bruce Weber. Impactful research on transformational information technology: An opportunity to inform new audiences. *Management Information Systems Quarterly*, 37(2):371–382, 2013.
- David J.C. MacKay. Bayesian methods for backpropagation networks. In *Models of neural networks III*, pages 211–254. Springer, 1996.
- David J.C. MacKay. Introduction to Gaussian processes. *NATO ASI Series F Computer and Systems Sciences*, 168:133–166, 1998.
- Pattie Maes. Agents that reduce work and information overload. *Communications of the ACM*, 37(7):30–40, 1994.
- Ann Majchrzak and Lynne M. Markus. *Methods for Policy Research: Taking Socially Responsible Action*. Sage Publications, 2013.
- Salvatore T. March and Gerald F. Smith. Design and natural science research on information technology. *Decision Support Systems*, 15(4):251–266, 1995.
- M. Lynne Markus, Ann Majchrzak, and Les Gasser. A design theory for systems that support emergent knowledge processes. *Management Information Systems Quarterly*, pages 179–212, 2002.
- S. Massoud Amin and Bruce F. Wollenberg. Toward a smart grid: Power delivery for the 21st century. *IEEE Power and Energy Magazine*, 3(5):34 – 41, September/October 2005.
- Kevin A. McCabe, Stephen J. Rassenti, and Vernon L. Smith. Smart computer-assisted markets. *Science*, 254(5031):534–538, 1991.

- Nigel P. Melville. Information systems innovation for environmental sustainability. *Management Information Systems Quarterly*, 34(1):1–21, 2010.
- Bertrand Meyer. *Object-oriented software construction*. Prentice Hall New York, second edition, 1998.
- Tom M. Mitchell. *Machine Learning*. McGraw-Hill, 1997.
- Robert E. Mittelstaedt. Benchmarking: How to learn from best-in-class practices. *National Productivity Review*, 11(3):301–315, 1992.
- Alan L. Montgomery and Michael D. Smith. Prospects for personalization on the internet. *Journal of Interactive Marketing*, 23(2):130–137, 2009.
- Richard H. Moss, Jae A. Edmonds, Kathy A. Hibbard, Martin R. Manning, Steven K. Rose, Detlef P. van Vuuren, Timothy R. Carter, Seita Emori, Mikiko Kainuma, Tom Kram, et al. The next generation of scenarios for climate change research and assessment. *Nature*, 463(7282):747–756, 2010.
- Kevin P. Murphy. *Machine Learning: A probabilistic perspective*. The MIT Press, 2012.
- BPS Murthi and Sumit Sarkar. The role of the management sciences in research on personalization. *Management Science*, 49(10):1344–1362, 2003.
- Vishnu Nanduri and Tapas K. Das. A reinforcement learning model to assess market power under auction-based energy pricing. *IEEE Transactions on Power Systems*, 22(1):85–95, 2007.
- Reza Nanoha. Confidence assessment of an agent-based simulation – operational validation of the Power TAC simulation platform. Master’s thesis, Rotterdam School of Management, 2013.
- Oded Netzer, Olivier Toubia, Eric T. Bradlow, Ely Dahan, Theodoros Evgeniou, Fred M. Feinberg, Eleanor M. Feit, Sam K. Hui, Joseph Johnson, John C. Liechty, et al. Beyond conjoint analysis: Advances in preference measurement. *Marketing Letters*, 19(3):337–354, 2008.
- James Nicolaisen, Valentin Petrov, and Leigh Tesfatsion. Market power and efficiency in a computational electricity market with discriminatory double-auction pricing. *IEEE Transactions on Evolutionary Computation*, 5(5):504–523, 2001.

- Michael J. North and Charles M. Macal. *Managing Business Complexity: Discovering Strategic Solutions with Agent-Based Modeling and Simulation*. Oxford University Press, Inc., New York, NY, USA, 2007.
- Jay F. Nunamaker, Jr. and Robert O. Briggs. Toward a broader vision for information systems. *ACM Transactions on Management Information Systems*, 2(4):20:1–20:12, January 2012.
- Christopher Painter-Wakefield and Ronald Parr. L1 regularized linear temporal difference learning. Technical Report TR-2012-01, Duke University, Computer Science, 2012.
- Peter Palensky and Dietmar Dietrich. Demand side management: Demand response, intelligent energy systems, and smart loads. *IEEE Transactions on Industrial Informatics*, 7(3):381–388, 2011.
- Gabriele Paolacci, Jesse Chandler, and Panagiotis Ipeirotis. Running experiments on Amazon Mechanical Turk. *Judgment and Decision Making*, 5(5):411–419, 2010.
- David Pardoe, Peter Stone, Maytal Saar-Tsechansky, Tayfun Keskin, and Kerem Tomak. Adaptive auction mechanism design and the incorporation of prior knowledge. *INFORMS Journal on Computing*, 22(3):353–370, 2010.
- David C. Parkes. *Algorithmic Game Theory*, chapter Online Mechanisms, pages 411–439. Cambridge University Press, 2007.
- Ronald Parr, Lihong Li, Gavin Taylor, Christopher Painter-Wakefield, and Michael L. Littman. An analysis of linear models, linear value-function approximation, and feature selection for Reinforcement Learning. In *Proceedings of the 25th International Conference on Machine Learning*, pages 752–759. ACM, 2008. ISBN 978-1-60558-205-4.
- Markus Peters. Autonomous data-driven decision-making in smart electricity markets. Master’s thesis, Erasmus University Rotterdam, August 2012.
- Markus Peters and Wolfgang Ketter. Modeling sparse heterogeneous user preferences. In *Workshop on Information Systems and Technology*, Orlando, Florida, USA, December 2012.
- Markus Peters and Wolfgang Ketter. A probabilistic multi-user preference model. In *Third International Conference on Algorithmic Decision Theory (ADT)*, Brussels, Belgium, November 2013a.

Markus Peters and Wolfgang Ketter. A probabilistic preference model with applications to electricity tariff choice. Conference of the Institute for Operations Research and the Management Sciences (INFORMS), Minneapolis, MN, October 2013b.

Markus Peters and Wolfgang Ketter. A probabilistic multi-user preference model. In *Workshop on Information Technology and Systems (WITS)*, Milan, Italy, December 2013c.

Markus Peters and Wolfgang Ketter. A probabilistic model for learning multi-user preferences with applications to electricity tariff choice. In *Multikonferenz Wirtschaftsinformatik (MKWI)*, Paderborn, Germany, February 2014.

Markus Peters, Wolfgang Ketter, Maytal Saar-Tsechansky, and John Collins. Autonomous data-driven decision-making in smart electricity markets. In Peter A. Flach, Tijl Bie, and Nello Cristianini, editors, *Machine Learning and Knowledge Discovery in Databases*, volume 7524 of *Lecture Notes in Computer Science*, pages 132–147. Springer Berlin Heidelberg, 2012a.

Markus Peters, Wolfgang Ketter, Maytal Saar-Tsechansky, and John E. Collins. Autonomous data-driven decision-making in smart electricity markets. In *Conference on Information Systems and Technologies (CIST)*, Phoenix, AZ, October 2012b.

Markus Peters, Wolfgang Ketter, Maytal Saar-Tsechansky, and John E. Collins. Autonomous data-driven decision-making in smart electricity markets. Conference of the Institute for Operations Research and the Management Sciences (INFORMS), Phoenix, AZ, October 2012c.

Markus Peters, Wolfgang Ketter, Maytal Saar-Tsechansky, and John E. Collins. Autonomous data-driven decision-making in smart electricity markets. In *Poster at the Workshop on Statistical Challenges in eCommerce Research (SCECR)*, Montreal, QC, June 2012d.

Markus Peters, Wolfgang Ketter, and John E. Collins. Design by competitive benchmarking: Tackling the smart grid challenge with innovative artifacts. In *Conference on Information Systems and Technologies (CIST)*, Minneapolis, MN, October 2013a.

Markus Peters, Wolfgang Ketter, and John E. Collins. Autonomous retail electricity trading: Strategies and economic benefits. Conference of the Institute for Operations Research and the Management Sciences (INFORMS), Minneapolis, MN, October 2013b.

Markus Peters, Wolfgang Ketter, Maytal Saar-Tsechansky, and John E. Collins. A reinforcement learning approach to autonomous decision-making in smart electricity markets. *Machine Learning*, 92:5–39, 2013c.

Kaare Brandt Petersen and Michael Syskind Pedersen. The matrix cookbook. Technical report, Technical University of Denmark, 2008.

Marek Petrik, Gavin Taylor, Ronald Parr, and Shlomo Zilberstein. Feature selection using regularization in approximate linear programs for markov decision processes. In *International Conference on Machine Learning (ICML)*, 2010.

Michael E. Porter and Mark R. Kramer. The link between competitive advantage and corporate social responsibility. *Harvard Business Review*, 11, 2006.

William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical recipes: The art of scientific computing*. Cambridge University Press, third edition, 2007.

Jan Pries-Heje and Richard Baskerville. The design theory nexus. *Management Information Systems Quarterly*, 32(4):731–755, 2008.

Danny Pudjianto, Charlotte Ramsay, and Goran Strbac. Microgrids and Virtual Power Plants: Concepts to support the integration of distributed energy resources. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, 222(7):731–741, 2008.

Larry D. Pyeatt and Adele E. Howe. Decision tree function approximation in Reinforcement Learning. In *Proceedings of the Third International Symposium on Adaptive Systems: Evolutionary Computation and Probabilistic Graphical Models*, volume 2, pages 70–77, 2001.

Andreas Pyka and Giorgio Fagiolo. Agent-based modelling: a methodology for neo-schumpeterian economics. *Elgar Companion to Neo-Schumpeterian Economics*, 467, 2007.

Joaquin Quiñonero-Candela and Carl Edward Rasmussen. A unifying view of sparse approximate Gaussian process regression. *The Journal of Machine Learning Research*, 6:1939–1959, 2005.

Joaquin Quinonero-Candela, Carl Edward Rasmussen, and Christopher K.I. Williams. Approximation methods for Gaussian process regression. *Large-scale kernel machines*, pages 203–224, 2007.

- Morteza Rahimiyan and Habib Rajabi Mashhadi. An adaptive Q-learning algorithm developed for agent-based computational modeling of electricity market. *IEEE Transactions on Systems, Man, and Cybernetics*, 40(5):547–556, 2010.
- Vivek Ramavajjala and Charles Elkan. Policy iteration based on a learned transition model. *Machine Learning and Knowledge Discovery in Databases*, pages 211–226, 2012.
- Carl E. Rasmussen and Christopher Williams. *Gaussian processes for Machine Learning*. MIT Press, 2006.
- Prashant Reddy and Manuela Veloso. Learned Behaviors of Multiple Autonomous Agents in Smart Grid Markets. In *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, 2011a.
- Prashant Reddy and Manuela Veloso. Strategy Learning for Autonomous Agents in Smart Grid Markets. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1446–1451, 2011b.
- Prashant Reddy and Manuela Veloso. Factored Models for Multiscale Decision Making in Smart Grid Customers. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, 2012.
- Prashant P. Reddy. *Semi-Cooperative Learning in Smart Grid Agents*. PhD thesis, Carnegie Mellon University, 2013.
- Horst W.J. Rittel and Melvin M. Webber. Dilemmas in a general theory of planning. *Policy Sciences*, 4(2):155–169, 1973.
- Joshua J. Romero. Blackouts illuminate India’s power problems. *IEEE Spectrum*, 49(10):11–12, 2012.
- Michael Rosemann and Iris Vessey. Toward improving the relevance of information systems research to practice: the role of applicability checks. *Management Information Systems Quarterly*, 32(1):1–22, 2008.
- Asim Roy, Patrick Mackin, Jyrki Wallenius, James Corner, Mark Keith, Gregory Schymik, and Hina Arora. An interactive search method based on user preferences. *Decision Analysis*, 5(4):203–229, 2008.
- Gavin A. Rummery and Mahesan Niranjana. On-line Q-learning using connectionist systems. Technical Report CUED/F-INFENG/TR 166, University of Cambridge, 1994.

- Maytal Saar-Tsechansky and Foster Provost. Active sampling for class probability estimation and ranking. *Machine Learning*, 54(2):153–178, 2004.
- Yunus Saatci. *Scalable Inference for Structured Gaussian Process Models*. PhD thesis, University of Cambridge, 2011.
- YK Shetty. Aiming high: competitive benchmarking for superior performance. *Long Range Planning*, 26(1):39–44, 1993.
- Anna Sidorova, Nicholas Evangelopoulos, Joseph S. Valacich, and Thiagarajan Ramakrishnan. Uncovering the intellectual core of the information systems discipline. *Management Information Systems Quarterly*, 32(3):467–482, 2006.
- David Silver, Leonard Newnham, David Barker, Suzanne Weller, and Jason McFall. Concurrent reinforcement learning from customer interactions. In *Proceedings of the 30th International Conference on Machine Learning*, pages 924–932, 2013.
- Herbert A. Simon. *The sciences of the artificial*. MIT press, 1996.
- Ramteen Sioshansi, Paul Denholm, Thomas Jenkin, and Jurgen Weiss. Estimating the value of electricity storage in PJM: Arbitrage and some welfare effects. *Energy Economics*, 31(2):269–277, 2009.
- Edward Snelson, Carl Edward Rasmussen, and Zoubin Ghahramani. Warped Gaussian processes. *Advances in Neural Information Processing Systems*, 16:337–344, 2004.
- Cyrill Stachniss, Christian Plagemann, and Achim J. Lilienthal. Learning gas distribution models using sparse Gaussian process mixtures. *Autonomous Robots*, 26(2-3): 187–202, 2009.
- Steven Stoft. *Power System Economics*. IEEE, 2002.
- Detmar Straub and Soon Ang. Editor’s comments: rigor and relevance in IS research: redefining the debate and a call for future research. *Management Information Systems Quarterly*, 35(1):iii–xi, 2011.
- Goran Strbac. Demand side management: Benefits and challenges. *Energy Policy*, 36(12):4419–4426, Dezember 2008. ISSN 03014215.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An introduction*. Cambridge University Press, 1998.

- Csaba Szepesvári. Algorithms for Reinforcement Learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 4(1):1–103, 2010.
- Leigh Tesfatsion. Agent-based Computational Economics: A constructive approach to economic theory. *Handbook of Computational Economics*, 2:831–880, 2006.
- Louis L. Thurstone. A law of comparative judgment. *Psychological Review*, 34(4): 273, 1927.
- Kenneth Train. *Discrete choice methods with simulation*. Cambridge University Press, 2003.
- Amos Tversky and Itamar Simonson. Context-dependent preferences. *Management Science*, 39(10):1179–1189, 1993.
- Daniel Urieli and Peter Stone. TacTex13: A champion adaptive Power Trading Agent. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS14)*, Paris, France, May 2014.
- US Department of Energy. Grid 2030: A national vision for electricity’s second 100 years, 2003. URL http://energy.gov/sites/prod/files/oeprod/DocumentsandMedia/Electric_Vision_Document.pdf.
- US Energy Information Administration. Electric Power Annual 2011. <http://www.eia.gov/electricity/annual/>, January 2013a.
- US Energy Information Administration. Monthly Energy Review February 2013. <http://www.eia.gov/mer>, February 2013b.
- Konstantina Valogianni, Wolf Ketter, and John Collins. Learning to schedule electric vehicle charging given individual customer preferences. In *Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS14)*, pages 1591–1592, Paris, France, May 2014a.
- Konstantina Valogianni, Wolfgang Ketter, John Collins, and Dmitry Zhdanov. Effective management of electric vehicle storage using smart charging. In *Proceedings of 28th AAAI Conference on Artificial Intelligence*, 2014b.
- Gerrit H. van Bruggen, Ale Smidts, and Berend Wierenga. Improving decision making by means of a marketing decision support system. *Management Science*, 44(5): 645–658, 1998.

Charles F. van Loan. The ubiquitous Kronecker product. *Journal of Computational and Applied Mathematics*, 123(1):85–100, 2000.

Luigi Vanfretti, Dirk Van Hertem, L Nordstrom, and Jan Ove Gjerde. A smart transmission grid for europe: Research challenges in developing grid enabling technologies. In *Power and Energy Society General Meeting*, pages 1–8. IEEE, 2011.

Jarno Vanhatalo, Ville Pietiläinen, and Aki Vehtari. Approximate inference for disease mapping with sparse Gaussian processes. *Statistics in Medicine*, 29(15):1580–1607, 2010.

Jarno Vanhatalo, Jaakko Riihimäki, Jouni Hartikainen, Pasi Jylänki, Ville Tolvanen, and Aki Vehtari. GPstuff: Bayesian modeling with Gaussian processes. *The Journal of Machine Learning Research*, 14(1):1175–1179, 2013.

John Venable and Richard Baskerville. Eating our own cooking: Toward a more rigorous design science of research methods. *Electronic Journal of Business Research Methods*, 10:141–153, 2012.

John Venable, Jan Pries-Heje, and Richard Baskerville. A comprehensive framework for evaluation in design science research. In *Design Science Research in Information Systems. Advances in Theory and Practice*, pages 423–438. Springer, 2012.

Ganesh Kumar Venayagamoorthy. Potentials and promises of computational intelligence for Smart Grids. In *Power & Energy Society General Meeting*, pages 1–6. IEEE, 2009.

Friedrich August von Hayek. The pretence of knowledge. *The American Economic Review*, 79(6):3–7, 1989.

John von Neumann, Oskar Morgenstern, Ariel Rubinstein, and Harold W. Kuhn. *Theory of games and economic behavior (Commemorative Edition)*. Princeton University Press, 2007.

Kiri L. Wagstaff. Machine learning that matters. In *Proceedings of the 29th International Conference on Machine Learning*, pages 529–536, 2012.

Joseph G. Walls, George R. Widmeyer, and Omar A. El Sawy. Building an information system design theory for vigilant EIS. *Information Systems Research*, 3(1): 36–59, 1992.

- Wei-quan Wang and Izak Benbasat. Trust in and adoption of online recommendation agents. *Journal of the Association for Information Systems*, 6(3):72–101, 2005.
- Richard T. Watson, Marie-Claude Boudreau, and Adela J. Chen. Information systems and environmentally sustainable development: Energy Informatics and new directions for the IS community. *Management Information Systems Quarterly*, 34(1):23, 2010.
- Richard T. Watson, J. Corbett, M.C. Boudreau, and J. Webster. An information strategy for environmental sustainability. *Communications of the ACM*, 55(7):28–30, 2012.
- Andrew R. Webb. *Statistical pattern recognition*. Wiley, 2003.
- Michael P. Wellman. *Trading Agents*. Morgan & Claypool, 2011.
- Paul J. Werbos. Putting more brain-like intelligence into the electric power grid: What we need and how to do it. In *International Joint Conference on Neural Networks*, pages 3356–3359. IEEE, 2009.
- Shimon Whiteson, Peter Stone, Kenneth O. Stanley, Risto Miikkulainen, and Nate Kohl. Automatic feature selection in neuroevolution. In *Proceedings of the 2005 conference on Genetic and evolutionary computation*, pages 1225–1232. ACM, 2005.
- Shimon Whiteson, Brian Tanner, Matthew E. Taylor, and Peter Stone. Protecting against evaluation overfitting in empirical reinforcement learning. In *IEEE Symposium on Adaptive Dynamic Programming and Reinforcement Learning (ADPRL)*, April 2011.
- Steven E. Widergren, Joseph M. Roop, Ross T. Guttromson, and Zhenyu Huang. Simulating the dynamic coupling of market and physical system operations. In *IEEE Power Engineering Society General Meeting*, pages 748–753. IEEE, 2004.
- Chris M. Wilson and Catherine W. Price. Do consumers switch to the best supplier? *Oxford Economic Papers*, 62(4):647–668, 2010.
- Worldbank. CO2 emissions (kt) catalog sources world development indicators, July 2013. URL <http://data.worldbank.org/indicator/EN.ATM.CO2E.KT/countries>.
- Hai Yang. Multiple equilibrium behaviors and advanced traveler information systems with endogenous market penetration. *Transportation Research Part B: Methodological*, 32(3):205–218, 1998.

Appendix A

A Reinforcement Learning Approach to Autonomous Decision-Making in Smart Electricity Markets

A.1 SEMS Example Run

We referred to the tendency of feature selection to overfit SELF instances to their target environment in Section 2.5.2. In Figure A.1 we give a concrete example of this behavior. The Figure depicts the development of tariff rates (panel (b)) and cash account balances (panel (c)) for various strategies in one example run of SEMS against a stylized step-function wholesale market (panel (a)). After the jump in wholesale prices at timeslot 25, the TableRL strategy first fails to adjust its rate upwards; and while SELF first increases its rate based on its target margin over the prevailing wholesale price, it quickly reverts to its initially successful strategy of offering the lowest margin in the market (recall from Section 2.4.1 that SELF learns target *margins*, not target *rates*). In the process, it undersells the unprofitable TableRL strategy and falls behind Fixed in terms of cash balance.

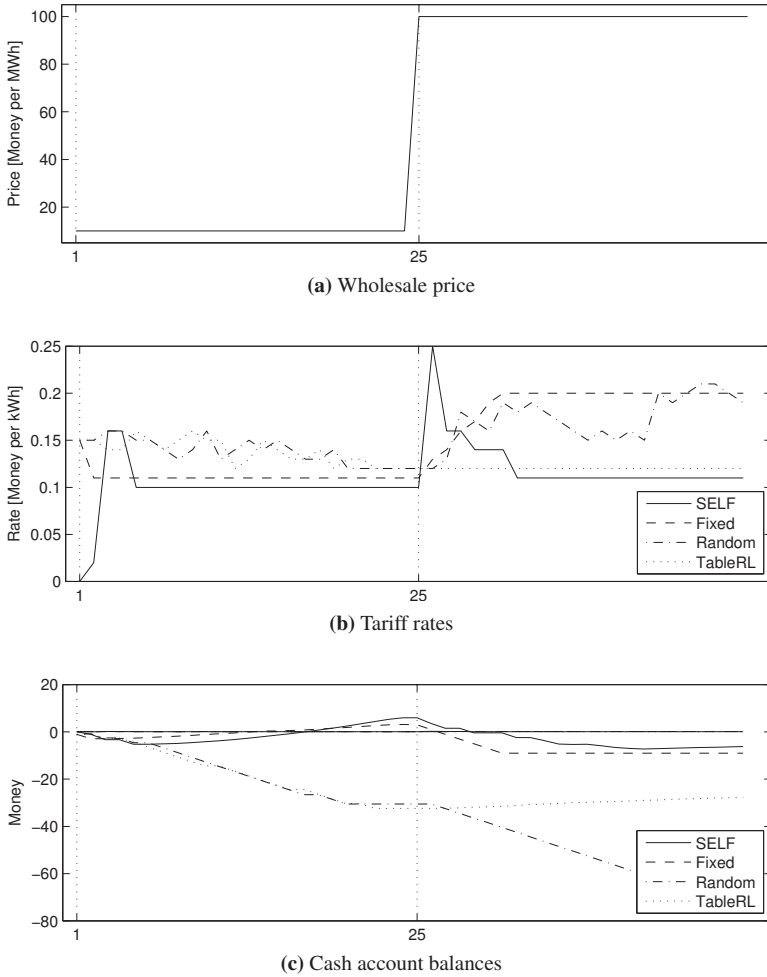


Figure A.1. Simulation using an overfitted strategy against a stylized step-function wholesale market

A.2 Robustness and Self-Play

In the experiments above we have considered SELF instances in competition against the benchmark strategies available in the literature to date. As another test of robustness we let two instances of SELF compete against each other, as well as against the benchmarks used earlier to explore whether our broker's strategy remains stable under self-play. The results of this evaluation are presented in Figure A.2.

As shown, both SELF brokers perform mostly better than their benchmark strategies

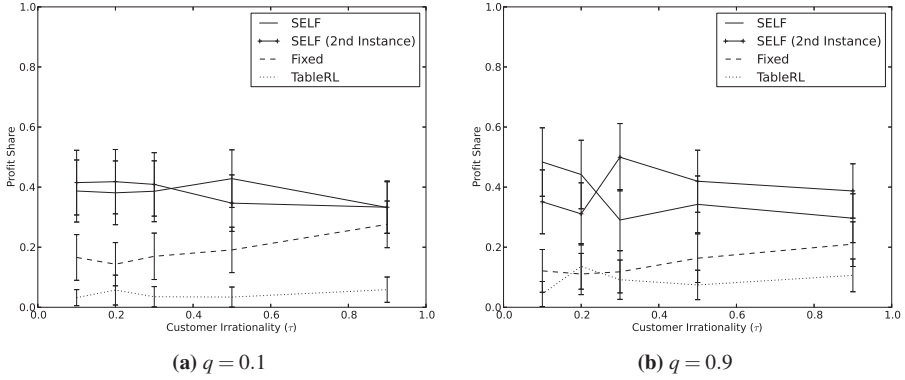


Figure A.2. Performance of two SELF instances in self-play

under self-play, and overall exhibit consistent performance in comparison to the situation when only a single SELF broker is present. The two brokers' performances are similar as well: When customers exhibit low switching rate (q), the SELF brokers performances are statistically indistinguishable; when q is high their performances are barely distinguishable, but with no clear advantage of one or the other instance. Rather, the SELF brokers' respective profit shares are lower than in previous settings, as they simply share the profit available in the market among themselves. Overall, we interpret the consistent, good performance shown by SELF in this more challenging setting as further evidence for the robustness of our broker.

A.3 Miscellaneous Experimental Results

In Section 2.5.3 we explored the role that bootstrapping and noise injection can play in counteracting overfitting tendencies in the feature and parameter selection process. The detailed configurations of the SELF configuration obtained under bootstrapping and noise injection are given in Tables A.1 and A.2, respectively.

An interesting property of bootstrapping and noise injection is that they work in a complementary fashion. Figure A.3 shows the performance of a SELF instance obtained through greedy feature selection using both add-on techniques. The results indicate that the strategy benefited in terms of both, performance in the target environment and generalizability.

| Feature | Plain | RBF | RBF(T) | Bin | Parameter | Value |
|------------------------|-------|-----|--------|-----|---------------------|-------|
| Bias | | | | | α_{max} | 0.37 |
| ActionIndex | | | | | α' | 1.41 |
| ActionOneInK | | | | | ε_{max} | 0.34 |
| BetterConsumptionRates | | | | | ε' | 1.09 |
| CashGradient | | | | | | |
| CustomerGradient | | | | | γ | 0.76 |
| MarketBreadth | | | | | | |
| MarketShare | | | | | | |
| MarkupLeader | | | | | | |
| NumberCustomers | | | | | | |
| RateChangeIndicator | | | | | | |
| TargetMargin | | | | | | |
| WholesalePrice | | | | | | |
| WorseConsumptionRates | | | | | | |

Table A.1. SELF instance obtained through greedy feature selection with bootstrapping

| Feature | Plain | RBF | RBF(T) | Bin | Parameter | Value |
|------------------------|-------|-----|--------|-----|---------------------|-------|
| Bias | | | | | α_{max} | 0.37 |
| ActionIndex | | | | | α' | 1.41 |
| ActionOneInK | | | | | ε_{max} | 0.34 |
| BetterConsumptionRates | | | | | ε' | 1.09 |
| CashGradient | | | | | | |
| CustomerGradient | | | | | γ | 0.76 |
| MarketBreadth | | | | | | |
| MarketShare | | | | | | |
| MarkupLeader | | | | | | |
| NumberCustomers | | | | | | |
| RateChangeIndicator | | | | | | |
| TargetMargin | | | | | | |
| WholesalePrice | | | | | | |
| WorseConsumptionRates | | | | | | |

Table A.2. SELF instance obtained through greedy feature selection with added noise

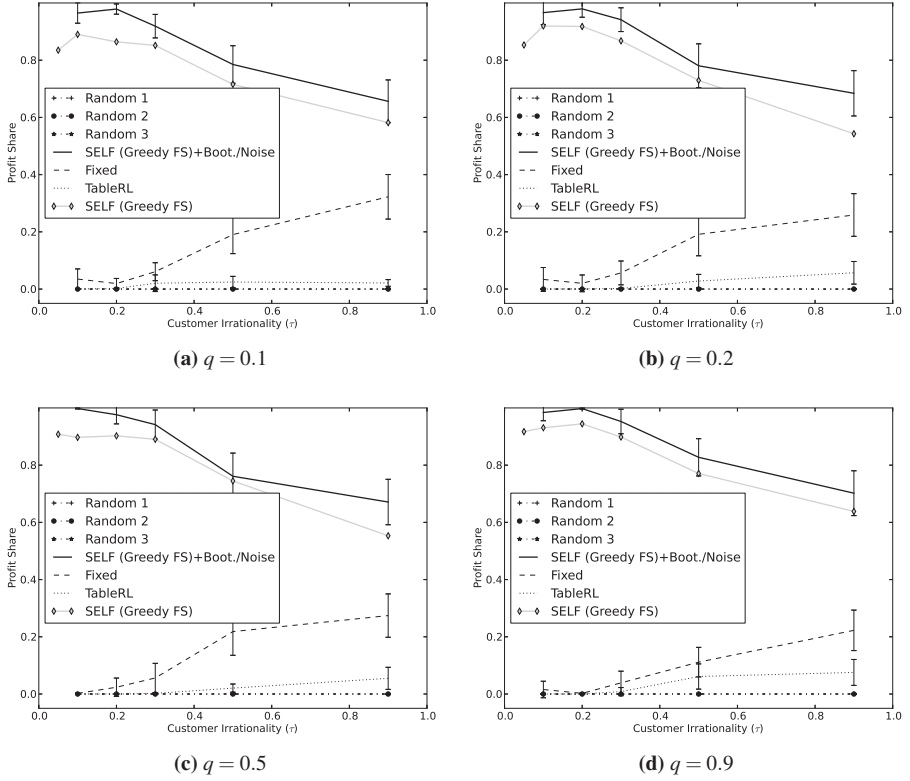


Figure A.3. Performance of a SELF instance obtained through greedy feature selection with noise injection and bootstrapping. Note, that all three random strategies are superimposed at the zero-line.

Appendix B

A Scalable Preference Model for Autonomous Decision-Making Involving Consumer Choices

B.1 Derivations used in Fast Inference

B.1.1 Probit Likelihood

The Laplace mode finding procedure (Algorithm 1) requires computation of the log likelihood $\log p(C|f^c)$, and of its first two derivatives with respect to the values f_t^c of the characteristic evaluations f^c at all trade-offs t (i.e., the Jacobian $\nabla \log p(\cdot)$, and the Hessian $\nabla \nabla \log p(\cdot)$). The Probit likelihood of a single observation is given by Equation (3.3) as:

$$\begin{aligned}\log p(y|f_t^c) &= \log \Phi(y \cdot f_{u,t}) \\ &= \log \Phi \left(y \cdot \left[\sum_{c=1}^{n_c} \gamma_u^c \cdot f_t^c \right] \right)\end{aligned}$$

where Φ denotes the cumulative distribution function (CDF) of the standard normal distribution. The argument to Φ is sometimes scaled by a precision factor σ^{-2} . But because our interpretation of the f^c is invariant under scaling, we can set $\sigma^{-2} = 1$ without loss of

generality. The Jacobian and the Hessian of the log likelihood are given by:

$$\nabla \log p(C|f^c) = \frac{\partial \log p(y|f_t^c)}{\partial f_{t_i}^{c_1}} = \frac{y_i \gamma_u^{c_1} N(f_{t_i})}{\Phi(y f_{t_i})} \quad (\text{B.1})$$

$$\nabla \nabla \log p(C|f^c) = \frac{\partial^2 \log p(y|f_t^c)}{\partial f_{t_i}^{c_1} \partial f_{t_i}^{c_2}} = \frac{-y f_{t_i} \gamma_u^{c_1} \gamma_u^{c_2} N(f_{t_i})}{\Phi(y f_{t_i})} - \frac{\gamma_u^{c_1} \gamma_u^{c_2} N^2(f_{t_i})}{\Phi^2(y f_{t_i})} \quad (\text{B.2})$$

where the derivatives are with respect to the evaluations of characteristics c_1 and c_2 for trade-off t_i , and u denotes the user making choice y .

B.1.2 Laplace Mode Finding

In Laplace mode finding, we approximate the posterior $p(f|C)$ using a single Gaussian

$$q(f|C) = N(f|\hat{f}, A^{-1})$$

centered on the true mode $\hat{f} = \arg \max_f p(f|C)$, and with a precision of $A = -\nabla \nabla \log p(C|f)|_{f=\hat{f}}$ obtained through a second-order Taylor expansion (Rasmussen and Williams, 2006). This mode is unique for the Probit because the Hessian of the log-likelihood is negative definite, and we can find it by setting the first derivative $\nabla \Psi$ of the unnormalized log posterior $\Psi = \log p(C|f) + \log p(f)$ to zero:

$$\nabla \Psi = \nabla \log p(C|f) - K^{-1} f \stackrel{!}{=} 0 \quad (\text{B.3})$$

The second term in Equation (B.3) results from differentiating the GP prior $p(f)$. The mode can then be found using the Newton-Raphson algorithm (Press et al., 2007) with the update step:

$$\begin{aligned} f^{new} &= f - (\nabla^2 \Psi)^{-1} \nabla \Psi \\ &= (K^{-1} + W)^{-1} \underbrace{(Wf + \nabla \log p(y|f))}_b \\ &= K \underbrace{(b - L(I + L^T K L)^{-1} L^T K b)}_a \end{aligned}$$

The last step uses the matrix inversion lemma (Petersen and Pedersen, 2008), and is valid for any symmetric decomposition $W = LL^T$.

Imagine having to choose between the following two tariffs for the household that you currently spend most of your time in. Which one would you prefer?

1. A Fixed tariff with 100 percent renewable energy content. Your monthly cost of electricity will be
 - 57.00\$ if you consume 500 kWh,
 - 106.00\$ if you consume 1000 kWh, and
 - 204.00\$ if you consume 2000 kWh

under this tariff. You pay your monthly electricity bill at the end of each month. A 12 months notice period applies before you can cancel this tariff.

2. A Variable tariff with 0 percent renewable energy content. The cost of electricity in the first month will be
 - 54.50\$ if you consume 500 kWh,
 - 101.00\$ if you consume 1000 kWh, and
 - 194.00\$ if you consume 2000 kWh

under this tariff. After the first month, the price of electricity may go up or down in accordance with the tariff's terms (and within legal bounds). You will have to pre-pay your monthly electricity bill at the beginning of the month. You can cancel your tariff anytime.

Table B.1. Example Choice Situation Used for Collecting the Tariffs Dataset.

B.2 Tariffs Dataset Collection

For this study, we collected a dedicated set of pairwise choice data on Amazon Mechanical Turk (MTurk, <http://www.mturk.com>), a commercial crowdsourcing platform. Several scholars have studied the demographics of MTurk workers, and have proposed guidelines for assuring the quality of data collected through MTurk tasks (Paolacci et al., 2010). These studies give reason to believe that (1) MTurk data can be of equal or better quality than data selected through channels such as student surveys, (2) MTurk workers are highly diverse (increasing external validity), and (3) the unsupervised nature of MTurk tasks may reduce the risk of experimenter bias (increasing internal validity), all if proper precautions are taken against distractions and random responses.

Eighty adult American participants were invited to fill in an academic survey about their electricity tariff preferences in exchange for a payment of \$0.30. All American MTurk workers could theoretically preview our survey through the MTurk platform, and 80 workers ultimately self-selected to participate. The survey consisted of three parts:

1. First, we reviewed basic electricity tariff concepts: fixed, variable, indexed tariffs, and those guaranteeing that a certain percentage of delivered electricity is produced from renewable sources.
2. Next, participants were asked to make ten choices between pairs of tariffs (see Ta-

ble B.1 for an example). Each pair was randomly generated from a total of 261 tariffs offered in Austin, Texas in February 2013. Texas has one of the most advanced retail electricity markets in the United States and provides daily information on available tariffs, see <http://www.powertochoose.org>.

3. Finally, we asked participants to answer ten questions on their demographics and electricity consumption behavior, some of which were attention checkers for which the correct answer had to be consistent with an answer given to another question.

Participants had a maximum of thirty minutes to fill out all questions, but could submit their results before that time. Participants could also withdraw, allowing another MTurk worker to fill out the survey instead. Next to the given answers, we recorded the time between self-selecting for participation and the submission of results. In pretests among colleagues, we had established that it took a quick reader at least three minutes to process all provided information. We therefore discarded surveys submitted before that time. As a further safeguard against random answers, we asked two pairs of attention check questions in the demographics section where the answers to one question depended on the answer of the other. We also discarded surveys where at least one of the attention check pairs was answered inconsistently, leaving us with a total of 61 surveys that met our quality standards.

Appendix C

Competitive Benchmarking of Electricity Brokers

C.1 Power TAC and IS Research Agendas

Several scholars have recently outlined research agendas for addressing societal challenges through IS research (Melville, 2010, Watson et al., 2010), and many of the abstract research questions they propose can be instantiated for the case of sustainable energy provisioning using Power TAC. In particular,

Energy Informatics is concerned with “analyzing, designing, and implementing systems to increase the efficiency of energy demand and supply systems” (Watson et al., 2010), an agenda that resonates with Power TAC’s goal of improving power systems through IS research. The framework consists of a generalized supplier–consumer relationship, mediated through Information Systems, and moderated through policies, regulations, the laws of economics, and corporate and social norms, all of which are possible subjects for study in Power TAC.

IS for Environmental Sustainability is defined as “IS-enabled organizational practices and processes that improve environmental and economic performance” (Melville, 2010). An explicit account of the complexity arising from social interactions at the intersection of IS and sustainability phenomena is at the heart of the proposed Belief – Action – Outcome (BAO) framework, and we contend that Power TAC’s agent-based model is well-suited to study these complex interactions at scale.

Smart Markets refer to an emerging class of fast-paced, information-rich markets in which participants routinely use IS to support their trading and decision-making (Bichler et al., 2010). Interestingly, the notion of **real-time intelligence** proposed by the authors resembles the centralized, optimization-based view taken in Energy Informatics, whereas the notion of **collective intelligence** is more closely related to the decentralized, complexity-oriented perspective of Melville's work. This suggests that the Smart Markets idea can serve to connect the two frameworks, and to identify blind spots at their seams.

Figure C.1 shows a model of an electric distribution system similar to that in Figure 4.8 where we indicated relationships to research questions from all three frameworks. We now consider the significance of each of these research questions in turn.

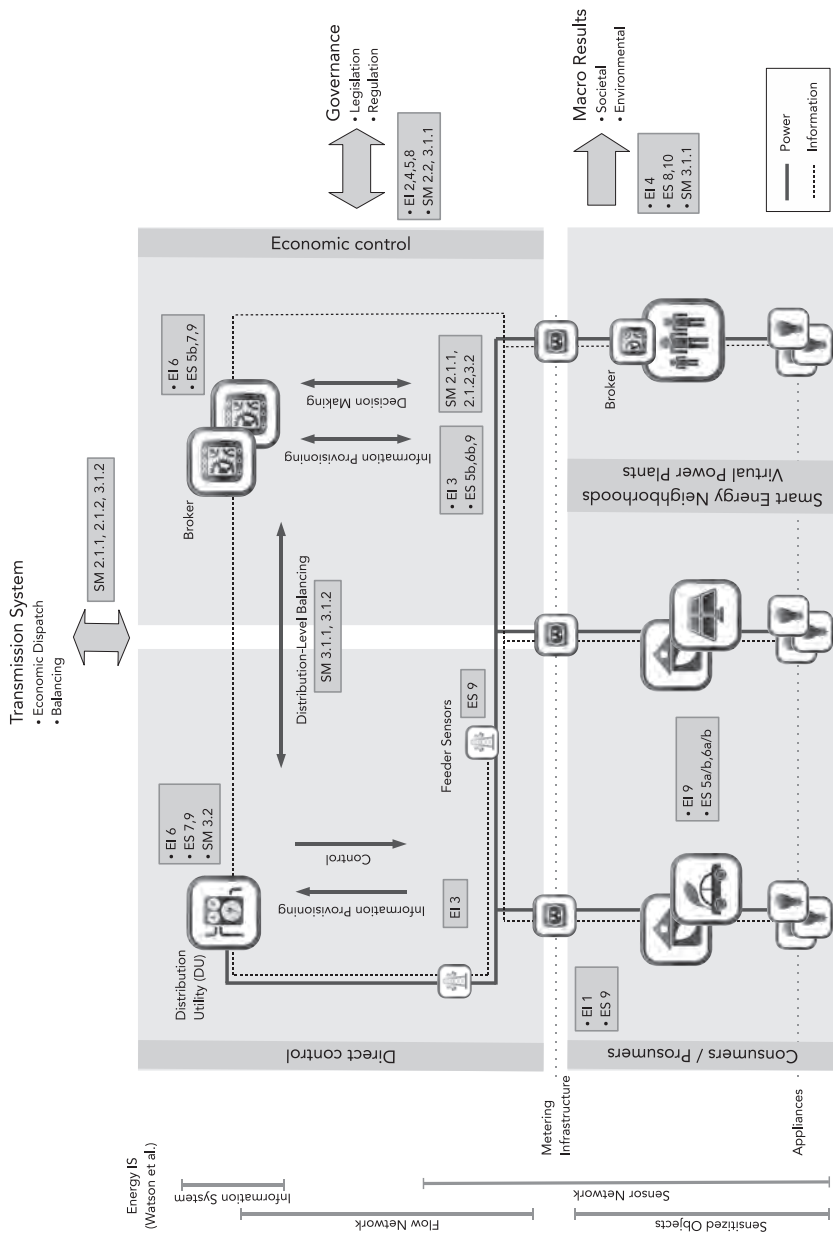


Figure C.1. Instantiation of the research agendas for Energy Informatics (EI; Watson et al. 2010), IS for Environmental Sustainability (ES; Melville 2010), and Smart Markets (SM; Bichler et al. 2010) for electric distribution networks.

Consumers / Prosumers Smart appliances, bidirectional customer-to-grid connectivity, and the transformation of consumers into active “prosumers” with local production and storage capabilities will have a significant impact on the role of customers in future power systems. The optimal granularity of information collection on customer premises is currently contentious (EI1), as is the business case for providing customers with smart metering capabilities (ES9) in exchange for greater visibility and control. Perhaps more importantly, research is needed to establish (a) how customers form their beliefs about electricity consumption, (b) how they derive actions from these beliefs, and (c) how information systems can be designed to effectively influence such actions and beliefs (ES5a/b, ES6a/b, EPRI - Electric Power Research Institute 2012).

Distribution Utility and Distribution Network Future distribution systems will possess improved sensory and control capabilities at the network level. Regulators and DUs need reliable insights into the effects of deploying these capabilities at the firm and economy level (EI6, ES7) so they can decide on optimal deployment levels (EI3, ES9). An important factor in this equation is the degree to which they can anticipate customer preferences and make control decisions that increase grid efficiency while respecting customer convenience constraints (SM3.2).

Brokers Retail electricity markets create opportunities for intermediaries that offer personalized, sustainable, and affordable electricity services to digitally connected customers. While this intermediation can, to some extent, be performed by traditional retailers, we propose that the data-intensity and complexity of future retail electricity markets will breed a sophisticated class of IS artifacts that we call Brokers, and that will fulfill the role of the intermediary, first in support of a human decision-maker but increasingly also in an autonomous fashion (Davenport and Harris, 2005). The development of a design theory for Brokers is one of the key challenges in creating a new decentralized control paradigm for electric grids. Using work on electronic marketplaces (EI6, SM2.1.1, SM2.1.2) and the design of IS artifacts for human use and their affordances (ES5b, ES6b, SM3.2), IS scholars can contribute to creating this new class of artifacts (ES5b, ES6b), and to understanding its implications for customers and the electric grid (ES7, ES9).

Smart Energy Neighborhoods Closely related to Brokers are questions about optimal control structures. Given the availability of decentralized small-scale generation, storage, and decentralized control strategies, the concept of Smart Energy Neighborhoods (also called Virtual Power Plants or Microgrids, Pudjianto et al. 2008) comes within reach. These partially autonomous self-organizing communities possess local high-efficiency

power generation capabilities and only draw on the upstream grid when unable to meet their requirements locally. The research questions surrounding Brokers can be extended to artifacts for the organization and control of such smaller communities as well.

Governance and Macro Results The confluence of retail electricity markets, Brokers, and new forms of consumer choice entails questions about how these entities and their interactions should be governed to attain socially desired outcomes. Mechanism designers have studied such questions (Parkes, 2007) but they are particularly difficult to answer given the real-time nature of the electric grid and its complicated control structure. It is unclear what degree and granularity of disclosure will allow regulators to determine and enforce effective policies (EI2, EI8); how policies should be constructed and validated (EI4, EI5, SM2.2, SM3.1.1); and what the relationship is between policies, IS artifacts, and macro-level results (ES8, ES10).

Distribution-Level Balancing / Transmission System The most challenging cases of mechanism design arise where multiple interdependent markets operate concurrently. New mechanisms will have to be explored (SM2.1.1, SM2.1.2); and new tools are needed to understand the complex interplay between control at the distribution level and in upstream markets (SM3.1.1, SM3.2.2).

ERIM PhD 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://hdl.handle.net/1765/1>. 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).

E. Abbink. Crew Management in Passenger Rail Transport. Promotor(s): Prof.dr. L.G. Kroon & Prof.dr. A.P.M. Wagelmans. <http://hdl.handle.net/1765/1>, EPS-2014-325-LIS.

O.A. Acar. Crowdsourcing for innovation: Unpacking Motivational, Knowledge and Relational Mechanisms of Innovative Behavior in Crowdsourcing Platforms. Promotor: Prof.dr. J.C.M. van den Ende. <http://hdl.handle.net/1765/1>, EPS-2014-321-LIS.

M. Acciaro. Bundling Strategies in Global Supply Chains. Promotor(s): Prof.dr. H.E. Haralambides. <http://hdl.handle.net/1765/19742>, EPS-2010-197-LIS.

E. Akpinar. Consumer Information Sharing; Understanding Psychological Drivers of Social Transmission. Promotor(s): Prof.dr.ir. A. Smidts. <http://hdl.handle.net/1765/50140>, EPS-2013-297-MKT.

A. Alexiev. Exploratory Innovation: The Role of Organizational and Top Management Team Social Capital. Promotor(s): Prof.dr. F.A.J. van den Bosch & Prof.dr. H.W. Volberda. <http://hdl.handle.net/1765/20632>, EPS-2010-208-STR.

M. Akin Ates. Purchasing and Supply Management at the Purchase Category Level: Strategy, Structure, and Performance. Promotor: Prof.dr. J.Y.F. Wynstra. <http://hdl.handle.net/1765/1>, EPS-2014-300-LIS.

R.J. de Almeida. Conditional Density Models Integrating Fuzzy and Probabilistic Representations of Uncertainty. Promotor Prof.dr.ir. Uzay Kaymak. <http://hdl.handle.net/1765/1>, EPS-2014-310-LIS.

K. Bannouh. Measuring and Forecasting Financial Market Volatility using High-Frequency Data. Promotor: Prof.dr.D.J.C. van Dijk. <http://hdl.handle.net/1765/38240>, EPS-2013-273-F&A.

T.M. Benning. A Consumer Perspective on Flexibility in Health Care: Priority Access Pricing and Customized Care. Promotor: Prof.dr.ir. B.G.C. Dellaert. <http://hdl.handle.net/1765/23670>, EPS-2011-241-MKT.

S.M. Ben-Menahem. Strategic Timing and Proactiveness of Organizations. Promotor(s): Prof.dr. H.W. Volberda & Prof.dr.ing. F.A.J. van den Bosch. <http://hdl.handle.net/1765/39128>, EPS-2013-278-S&E.

W.E. van den Berg. Understanding Salesforce Behavior Using Genetic Association Studies. Promotor: Prof.dr. W.J.M.I. Verbeke. <http://hdl.handle.net/1765/51440>, EPS-2014-311-MKT.

N.E. Betancourt. Typical Atypicality: Formal and Informal Institutional Conformity, Deviance, and Dynamics. Promotor: Prof.dr. B. Krug. <http://hdl.handle.net/1765/32345>, EPS-2012-262-ORG.

J.L.G. Binken. System Markets: Indirect Network Effects in Action or Inaction. Promotor: Prof.dr. S. Stremersch. <http://hdl.handle.net/1765/21186>, EPS-2010-213-MKT.

D.C. Blitz. Benchmarking Benchmarks. Promotor(s): Prof.dr. A.G.Z. Kemna & Prof.dr. W.F.C. Verschoor. <http://hdl.handle.net/1765/226244>, EPS-2011-225-F&A.

M. Boons. 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: Prof.dr. H.G. Barkema. <http://hdl.net/1765/50711>, EPS-2014-306-S&E.

W.A.M. Borst. Understanding Crowdsourcing: Effects of Motivation and Rewards on Participation and Performance in Voluntary Online Activities. Promotor(s): Prof.dr.ir. J.C.M.

van den Ende & Prof.dr.ir. H.W.G.M. van Heck. <http://hdl.handle.net/1765/21914>, EPS-2010-221-LIS.

D.P. Budiono. The Analysis of Mutual Fund Performance: Evidence from U.S. Equity Mutual Funds. Promotor: Prof.dr. M.J.C.M. Verbeek. <http://hdl.handle.net/1765/18126>, EPS-2010-185-F&A.

M.J. Burger. Structure and Cooptition in Urban Networks. Promotor(s): Prof.dr. G.A. van der Knaap & Prof.dr. H.R. Commandeur. <http://hdl.handle.net/1765/26178>, EPS-2011-243-ORG.

E. Byington. Exploring Coworker Relationships: Antecedents and Dimensions of Interpersonal Fit, Coworker Satisfaction, and Relational Models. Promotor: Prof.dr. D.L. van Knippenberg. <http://hdl.handle.net/1765/41508>, EPS-2013-292-ORG.

N.M. Camacho. Health and Marketing; Essays on Physician and Patient Decision-making. Promotor: Prof.dr. S. Stremersch. <http://hdl.handle.net/1765/23604>, EPS-2011-237-MKT.

P. Cankurtaran. Essays On Accelerated Product Development. Promotor: Prof.dr.ir. G.H. van Bruggen. <http://hdl.handle.net/1765/1>, EPS-2014-317-MKT.

E.A.M. Caron. Explanation of Exceptional Values in Multi-dimensional Business Databases. Promotor(s): Prof.dr.ir. H.A.M. Daniels & Prof.dr. G.W.J. Hendrikse. <http://hdl.handle.net/1765/50005>, EPS-2013-296-LIS.

L. Carvalho. Knowledge Locations in Cities; Emergence and Development Dynamics. Promotor: Prof.dr. L. van den Berg. <http://hdl.handle.net/1765/38449>, EPS-2013-274-S&E.

L. Carvalho de Mesquita Ferreira. Attention Mosaics: Studies of Organizational Attention. Promotor(s): Prof.dr. P.M.A.R. Heugens & Prof.dr. J. van Oosterhout. <http://hdl.handle.net/1765/19882>, EPS-2010-205-ORG.

R.H.G.M. Cox. To Own, To Finance, and to Insure; Residential Real Estate Revealed. Promotor: Prof.dr. D. Brounen. <http://hdl.handle.net/1765/40964>, EPS-2013-290-F&A.

E.F. Defilippi Angeldonis. Access Regulation for Naturally Monopolistic Port Terminals: Lessons from Regulated Network Industries. Promotor: Prof.dr. H.E. Haralambides. <http://hdl.handle.net/1765/19881>, EPS-2010-204-LIS.

D. Deichmann. Idea Management: Perspectives from Leadership, Learning, and Network Theory. Promotor: Prof.dr.ir. J.C.M. van den Ende. <http://hdl.handle.net/1765/31174>, EPS-2012-255-ORG.

P.T.M. Desmet. In Money we Trust? Trust Repair and the Psychology of Financial Compensations. Promotor: Prof.dr. D. De Cremer & Prof.dr. E. van Dijk. <http://hdl.handle.net/1765/23268>, EPS-2011-232-ORG.

R.C. Dietvorst. Neural Mechanisms Underlying Social Intelligence and Their Relationship with the Performance of Sales Managers. Promotor: Prof.dr. W.J.M.I. Verbeke. <http://hdl.handle.net/1765/21188>, EPS-2010-215-MKT.

T.A.B. Dollevoet. Delay Management and Dispatching in Railways. Promotor: Prof.dr. A.P.M. Wagelmans. <http://hdl.handle.net/1765/38241>, EPS-2013-272-LIS.

S. van Doorn. 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. <http://hdl.handle.net/1765/32166>, EPS-2012-258-STR.

M.G. Douwens-Zonneveld. Animal Spirits and Extreme Confidence: No Guts, No Glory. Promotor: Prof.dr. W.F.C. Verschoor. <http://hdl.handle.net/1765/31914>, EPS-2012-257-F&A.

E. Duca. The Impact of Investor Demand on Security Offerings. Promotor: Prof.dr. A. de Jong. <http://hdl.handle.net/1765/26041>, EPS-2011-240-F&A.

H. Duursema. Strategic Leadership; Moving Beyond the Leader-follower Dyad. Promotor: Prof.dr. R.J.M. van Tulder. <http://hdl.handle.net/1765/39129>, EPS-2013-279-ORG.

N.J. van Eck. Methodological Advances in Bibliometric Mapping of Science. Promotor: Prof.dr.ir. R. Dekker. <http://hdl.handle.net/1765/26509>, EPS-2011-247-LIS.

M. van Essen. An Institution-Based View of Ownership. Promotor(s): Prof.dr. J. van Oosterhout & Prof.dr. G.M.H. Mertens. <http://hdl.handle.net/1765/22643>, EPS-2011-226-ORG.

L. Feng. Motivation, Coordination and Cognition in Cooperatives. Promotor: Prof.dr. G.W.J. Hendrikse. <http://hdl.handle.net/1765/21680>, EPS-2010-220-ORG.

S.P.L. Fourné. 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 & dr. T.J.M. Mom. <http://hdl.handle.net/1765/21680>, EPS-2014-318-S&E.

A.H. Gharehgozli. Developing New Methods for Efficient Container Stacking Operations. Promotor: Prof.dr.ir. M.B.M. de Koster. <http://hdl.handle.net/1765/37779>, EPS-2012-269-LIS.

S. van Gils. Morality in Interactions: On the Display of Moral Behavior by Leaders and Employees. Promotor: Prof.dr. D.L. van Knippenberg. <http://hdl.handle.net/1765/38028>, EPS-2012-270-ORG.

M.N.G. van Ginkel-Bieshaar. The Impact of Abstract versus Concrete Product Communications on Consumer Decision-making Processes. Promotor: Prof.dr.ir. B.G.C. Dellaert. <http://hdl.handle.net/1765/31913>, EPS-2012-256-MKT.

X. Gkougkousi. Empirical Studies in Financial Accounting. Promotor(s): Prof.dr. G.M.H. Mertens & Prof.dr. E. Peek. <http://hdl.handle.net/1765/37170>, EPS-2012-264-F&A.

K.M. Glorie. Clearing Barter Exchange Markets: Kidney Exchange and Beyond. Promotor(s): Prof.dr. A.P.M. Wagelmans & Prof.dr. J.J. van de Klundert. <http://hdl.handle.net/1765/1>, EPS-2014-329-LIS.

N.A Hakimi. Leader Empowering Behaviour: The Leader's Perspective: Understanding the Motivation behind Leader Empowering Behaviour. Promotor: Prof.dr. D.L. van Knippenberg. <http://hdl.handle.net/1765/17701>, EPS-2010-184-ORG.

M. Hensmans. A Republican Settlement Theory of the Firm: Applied to Retail Banks in England and the Netherlands (1830-2007). Promotor(s): Prof.dr. A. Jolink & Prof.dr. S.J. Magala. <http://hdl.handle.net/1765/19494>, EPS-2010-193-ORG.

C. Hernandez Mireles. Marketing Modeling for New Products. Promotor: Prof.dr. P.H. Franses. <http://hdl.handle.net/1765/19878>, EPS-2010-202-MKT.

D. von der Heyde Fernandes. The Functions and Dysfunctions of Reminders. Promotor: Prof.dr. S.M.J. van Osselaer. <http://hdl.handle.net/1765/41514>, EPS-2013-295-MKT.

M.L.M. Heyden. Essays on Upper Echelons & Strategic Renewal: A Multilevel Contingency Approach. Promotor(s): Prof.dr. F.A.J. van den Bosch & Prof.dr. H.W. Volberda. <http://hdl.handle.net/1765/32167>, EPS-2012-259-STR.

I.J. Hoever. Diversity and Creativity: In Search of Synergy. Promotor(s): Prof.dr. D.L. van Knippenberg. <http://hdl.handle.net/1765/37392>, EPS-2012-267-ORG.

B. Hoogendoorn. Social Entrepreneurship in the Modern Economy: Warm Glow, Cold Feet. Promotor(s): Prof.dr. H.P.G. Pennings & Prof.dr. A.R. Thurik. <http://hdl.handle.net/1765/26447>, EPS-2011-246-STR.

N. Hoogervorst. On The Psychology of Displaying Ethical Leadership: A Behavioral Ethics Approach. Promotor(s): Prof.dr. D. De Cremer & Dr. M. van Dijke. <http://hdl.handle.net/1765/26228>, EPS-2011-244-ORG.

X. Huang. An Analysis of Occupational Pension Provision: From Evaluation to Redesign. Promotor(s): Prof.dr. M.J.C.M. Verbeek & Prof.dr. R.J. Mahieu. <http://hdl.handle.net/1765/19674>, EPS-2010-196-F&A.

K.A. Hytönen. Context Effects in Valuation, Judgment and Choice. Promotor(s): Prof.dr.ir. A. Smidts. <http://hdl.handle.net/1765/30668>, EPS-2011-252-MKT.

P. den Iseger. Fourier and Laplace Transform Inversion with Application in Finance. Promotor: Prof.dr.ir. R.Dekker. <http://hdl.handle.net/1765/1>, EPS-2014-322-LIS.

W.L. van Jaarsveld. Maintenance Centered Service Parts Inventory Control. Promotor(s): Prof.dr.ir. R. Dekker. <http://hdl.handle.net/1765/39933>, EPS-2013-288-LIS.

M.N. Jalil. Customer Information Driven After Sales Service Management: Lessons from Spare Parts Logistics. Promotor(s): Prof.dr. L.G. Kroon. <http://hdl.handle.net/1765/22156>, EPS-2011-222-LIS.

M. Kagie. Advances in Online Shopping Interfaces: Product Catalog Maps and Recommender Systems. Promotor(s): Prof.dr. P.J.F. Groenen. <http://hdl.handle.net/1765/19532>, EPS-2010-195-MKT.

E.R. Kappe. The Effectiveness of Pharmaceutical Marketing. Promotor(s): Prof.dr. S. Stremersch. <http://hdl.handle.net/1765/23610>, EPS-2011-239-MKT.

B. Karreman. Financial Services and Emerging Markets. Promotor(s): Prof.dr. G.A. van der Knaap & Prof.dr. H.P.G. Pennings. <http://hdl.handle.net/1765/22280>, EPS-2011-223-ORG.

J.C.M. Kil. Acquisitions Through a Behavioral and Real Options Lens. Promotor(s): Prof.dr. H.T.J. Smit. <http://hdl.handle.net/1765/50142>, EPS-2013-298-F&A.

E. van't Klooster. Travel to Learn: The Influence of Cultural Distance on Competence Development in Educational Travel. Promotors: Prof.dr. F.M. Go & Prof.dr. P.J. van Baalen. <http://hdl.handle.net/1765/151460>, EPS-2014-312-MKT.

S.R. Koendjibiharie. The Information-Based View on Business Network Performance Revealing the Performance of Interorganizational Networks. Promotors: Prof.dr.ir. H.W.G.M. van Heck & Prof.mr.dr. P.H.M. Vervest. <http://hdl.handle.net/1765/1>, EPS-2014-315-LIS.

M. Koning. The Financial Reporting Environment: taking into account the media, international relations and auditors. Promotor(s): Prof.dr. P.G.J.Roosenboom & Prof.dr. G.M.H. Mertens. <http://hdl.handle.net/1765/1>, EPS-2014-330-F&A.

D.J. Konter. Crossing borders with HRM: An inquiry of the influence of contextual differences in the adaption and effectiveness of HRM. Promotor: Prof.dr. J. Paauwe. <http://hdl.handle.net/1765/1>, EPS-2014-305-ORG.

E. Korkmaz. Understanding Heterogeneity in Hidden Drivers of Customer Purchase Behavior. Promotors: Prof.dr. S.L. van de Velde & dr. R.Kuik. <http://hdl.handle.net/1765/1>, EPS-2014-316-LIS.

J.J. Kroezen. The Renewal of Mature Industries: An Examination of the Revival of the Dutch Beer Brewing Industry. Promotor: Prof. P.P.M.A.R. Heugens. <http://hdl.handle.net/1765/1>, EPS-2014-333-S&E.

K.Y. Lam. Reliability and Rankings. Promotor(s): Prof.dr. P.H.B.F. Franses. <http://hdl.handle.net/1765/22977>, EPS-2011-230-MKT.

M.W. Lander. Profits or Professionalism? On Designing Professional Service Firms. Promotor(s): Prof.dr. J. van Oosterhout & Prof.dr. P.P.M.A.R. Heugens. <http://hdl.handle.net/1765/30682>, EPS-2012-253-ORG.

B. de Langhe. Contingencies: Learning Numerical and Emotional Associations in an Uncertain World. Promotor(s): Prof.dr.ir. B. Wierenga & Prof.dr. S.M.J. van Osselaer. <http://hdl.handle.net/1765/23504>, EPS-2011-236-MKT.

J.A. Larco Martinelli. Incorporating Worker-Specific Factors in Operations Management Models. Promotor(s): Prof.dr.ir. J. Dul & Prof.dr. M.B.M. de Koster. <http://hdl.handle.net/1765/21527>, EPS-2010-217-LIS.

J.M. Leunissen. All Apologies: On the Willingness of Perpetrators to Apologize. Promotor: Prof.dr. D. De Cremer. <http://hdl.handle.net/1765/1>, EPS-2014-301-ORG.

Q. Liang. Governance, CEO Identity, and Quality Provision of Farmer Cooperatives. Promotor: Prof.dr. G.W.J. Hendrikse. <http://hdl.handle.net/1765/1>, EPS-2013-281-ORG.

K.C. Liket. Why Doing Good is not Good Enough: Essays on Social Impact Measurement. Promotor: Prof.dr. H.R. Commandeur. <http://hdl.handle.net/1765/51130>, EPS-2014-307-S&E.

M.J.H.M. van der Loos. 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. <http://hdl.handle.net/1765/40081>, EPS-2013-287-S&E.

M. Lovric. Behavioral Finance and Agent-Based Artificial Markets. Promotor(s): Prof.dr. J. Spronk & Prof.dr.ir. U. Kaymak. <http://hdl.handle.net/1765/22814>, EPS-2011-229-F&A.

Y. Lu. Data-Driven Decision Making in Auction Markets. Promotors: Prof.dr.ir.H.W.G.M. van Heck & Prof.dr.W.Ketter. <http://hdl.handle.net/1765/1>, EPS-2014-314-LIS.

T.D. Markwat. Extreme Dependence in Asset Markets Around the Globe. Promotor: Prof.dr. D.J.C. van Dijk. <http://hdl.handle.net/1765/22744>, EPS-2011-227-F&A.

H. Mees. Changing Fortunes: How China's Boom Caused the Financial Crisis. Promotor: Prof.dr. Ph.H.B.F. Franses. <http://hdl.handle.net/1765/34930>, EPS-2012-266-MKT.

J. Meuer. Configurations of Inter-Firm Relations in Management Innovation: A Study in China's Biopharmaceutical Industry. Promotor: Prof.dr. B. Krug. <http://hdl.handle.net/1765/22745>, EPS-2011-228-ORG.

O.R. Mihalache. 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. <http://hdl.handle.net/1765/32343>, EPS-2012-260-S&E.

V. Milea. New Analytics for Financial Decision Support. Promotor: Prof.dr.ir. U. Kaymak. <http://hdl.handle.net/1765/38673>, EPS-2013-275-LIS.

I. Naumovska. Socially Situated Financial Markets: a Neo-Behavioral Perspective on Firms, Investors and Practices. Promoter(s) Prof.dr. P.P.M.A.R. Heugens & Prof.dr. A.de Jong. <http://hdl.handle.net/1765/1>, EPS-2014-319-S&E.

L.K. Nielsen. Rolling Stock Rescheduling in Passenger Railways: Applications in Short-term Planning and in Disruption Management. Promotor: Prof.dr. L.G. Kroon. <http://hdl.handle.net/1765/22444>, EPS-2011-224-LIS.

M.H. Nijdam. Leader Firms: The Value of Companies for the Competitiveness of the Rotterdam Seaport Cluster. Promotor(s): Prof.dr. R.J.M. van Tulder. <http://hdl.handle.net/1765/21405>, EPS-2010-216-ORG.

L.H.J. Noordegraaf-Eelens. Contested Communication: A Critical Analysis of Central Bank Speech. Promotor: Prof.dr. Ph.H.B.F. Franses. <http://hdl.handle.net/1765/21061>, EPS-2010-209-MKT.

A.L.P. Nuijten. Deaf Effect for Risk Warnings: A Causal Examination applied to Information Systems Projects. Promotor: Prof.dr. G. van der Pijl & Prof.dr. H. Commandeur & Prof.dr. M. Keil. <http://hdl.handle.net/1765/34928>, EPS-2012-263-S&E.

M. van Oosterhout. Business Agility and Information Technology in Service Organizations. Promotor: Prof.dr.ir. H.W.G.M. van Heck. <http://hdl.handle.net/1765/19805>, EPS-2010-198-LIS.

S.E. Osadchiy. The Dynamics of Formal Organization: Essays on Bureaucracy and Formal Rules. Promotor: Prof.dr. P.P.M.A.R. Heugens. <http://hdl.handle.net/1765/23250>, EPS-2011-231-ORG.

A.H.J. Otgaar. Industrial Tourism: Where the Public Meets the Private. Promotor: Prof.dr. L. van den Berg. <http://hdl.handle.net/1765/21585>, EPS-2010-219-ORG.

M.N. Ozdemir. Project-level Governance, Monetary Incentives and Performance in Strategic R&D Alliances. Promotor: Prof.dr.ir. J.C.M. van den Ende. <http://hdl.handle.net/1765/23550>, EPS-2011-235-LIS.

Y. Peers. Econometric Advances in Diffusion Models. Promotor: Prof.dr. Ph.H.B.F. Franses. <http://hdl.handle.net/1765/30586>, EPS-2011-251-MKT.

C. Pince. Advances in Inventory Management: Dynamic Models. Promotor: Prof.dr.ir. R. Dekker. <http://hdl.handle.net/1765/19867>, EPS-2010-199-LIS.

J.P. Porck. No Team is an Island. Promotor: Prof.dr. P.J.F. Groenen & Prof.dr. D.L. van Knippenberg. <http://hdl.handle.net/1765/50141>, EPS-2013-299-ORG.

M. Porras Prado. The Long and Short Side of Real Estate, Real Estate Stocks, and Equity. Promotor: Prof.dr. M.J.C.M. Verbeek. <http://hdl.handle.net/1765/30848>, EPS-2012-254-F&A.

D. Potthoff. Railway Crew Rescheduling: Novel Approaches and Extensions. Promotor(s): Prof.dr. A.P.M. Wagelmans & Prof.dr. L.G. Kroon. <http://hdl.handle.net/1765/21084>, EPS-2010-210-LIS.

P.V. Poruthiyil. Steering Through: How Organizations Negotiate Permanent Uncertainty and Unresolvable Choices. Promotor(s): Prof.dr. P.P.M.A.R. Heugens & Prof.dr. S. Magala. <http://hdl.handle.net/1765/26392>, EPS-2011-245-ORG.

M. Pourakbar. End-of-Life Inventory Decisions of Service Parts. Promotor: Prof.dr.ir. R. Dekker. <http://hdl.handle.net/1765/30584>, EPS-2011-249-LIS.

E.S. Pronker. Innovation Paradox in Vaccine Target Selection. Promotor(s): Prof.dr. H.R. Commandeur & Prof.dr. H.J.H.M. Claassen. <http://hdl.handle.net/1765/39654>, EPS-2013-282-S&E.

M.J. Retel Helmrich. Green Lot-Sizing. Promotor: Prof.dr. A.P.M. Wagelmans. <http://hdl.handle.net/1765/41330>, EPS-2013-291-LIS.

C.A. Rietveld. Essays on the Intersection of Economics and Biology. Promotor(s): Prof.dr. P.J.F. Groenen, Prof.dr. A. Hofman, Prof.dr. A.R. Thurik & Prof.dr. P.D. Koellinger. <http://hdl.handle.net/1765/1>, EPS-2014-320-S&E.

J.A. Rijsenbilt. CEO Narcissism; Measurement and Impact. Promotor: Prof.dr. A.G.Z. Kemna & Prof.dr. H.R. Commandeur. <http://hdl.handle.net/1765/23554>, EPS-2011-238-STR.

E.M. Roelofsen. The Role of Analyst Conference Calls in Capital Markets. Promotor(s): Prof.dr. G.M.H. Mertens & Prof.dr. L.G. van der Tas RA. <http://hdl.handle.net/1765/18013>, EPS-2010-190-F&A.

M.W. Roza. 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. <http://hdl.handle.net/1765/22155>, EPS-2011-214-STR.

G. Rubbiany. Investment Behavior of Institutional Investors. Promotor: Prof.dr. W.F.C. Verschoor. <http://hdl.handle.net/1765/40068>, EPS-2013-284-F&A.

G.A.C. Schellekens. Language Abstraction in Word of Mouth. Promotor: Prof.dr.ir. A. Smidts. <http://hdl.handle.net/1765/21580>, EPS-2010-218-MKT.

K. Shahzad. Credit Rating Agencies, Financial Regulations and the Capital Markets. Promotor: Prof.dr. G.M.H. Mertens. <http://hdl.handle.net/1765/39655>, EPS-2013-283-F&A.

F. Sotgiu. Not All Promotions are Made Equal: From the Effects of a Price War to Cross-chain Cannibalization. Promotor(s): Prof.dr. M.G. Dekimpe & Prof.dr.ir. B. Wierenga. <http://hdl.handle.net/1765/19714>, EPS-2010-203-MKT.

M. Sousa. Servant Leadership to the Test: New Perspectives and Insight. Promoters: Prof.dr. D. van Knippenberg & Dr. D. van Dierendonck. <http://hdl.handle.net/1765/1>, EPS-2014-313-ORG.

R. Spliet. Vehicle Routing with Uncertain Demand. Promotor: Prof.dr.ir. R. Dekker. <http://hdl.handle.net/1765/41513>, EPS-2013-293-LIS.

F.J. Srour. Dissecting Drayage: An Examination of Structure, Information, and Control in Drayage Operations. Promotor: Prof.dr. S.L. van de Velde. <http://hdl.handle.net/1765/18231>, EPS-2010-186-LIS.

J.L. Staadt. Leading Public Housing Organisation in a Problematic Situation: A Critical Soft Systems Methodology Approach. Promotor: Prof.dr. S.J. Magala. <http://hdl.handle.net/1765/50712>, EPS-2014-308-ORG.

M. Stallen. Social Context Effects on Decision-Making; A Neurobiological Approach. Promotor: Prof.dr.ir. A. Smidts. <http://hdl.handle.net/1765/39931>, EPS-2013-285-MKT.

M. Tarakci. Behavioral Strategy; Strategic Consensus, Power and Networks. Promotor(s): Prof.dr. P.J.F. Groenen & Prof.dr. D.L. van Knippenberg. <http://hdl.handle.net/1765/39130>, EPS-2013-280-ORG.

M. Teixeira de Vasconcelos. Agency Costs, Firm Value, and Corporate Investment. Promotor: Prof.dr. P.G.J. Roosenboom. <http://hdl.handle.net/1765/37265>, EPS-2012-265-F&A.

M.P. Tempelaar. Organizing for Ambidexterity: Studies on the Pursuit of Exploration and Exploitation through Differentiation, Integration, Contextual and Individual Attributes. Promotor(s): Prof.dr.ing. F.A.J. van den Bosch & Prof.dr. H.W. Volberda. <http://hdl.handle.net/1765/18457>, EPS-2010-191-STR.

V. Tiwari. Transition Process and Performance in IT Outsourcing: Evidence from a Field Study and Laboratory Experiments. Promotor(s): Prof.dr.ir. H.W.G.M. van Heck & Prof.dr. P.H.M. Vervest. <http://hdl.handle.net/1765/19868>, EPS-2010-201-LIS.

C. Tröster. Nationality Heterogeneity and Interpersonal Relationships at Work. Promotor: Prof.dr. D.L. van Knippenberg. <http://hdl.handle.net/1765/23298>, EPS-2011-233-ORG.

D. Tsekouras. No Pain No Gain: The Beneficial Role of Consumer Effort in Decision Making. Promotor: Prof.dr.ir. B.G.C. Dellaert. <http://hdl.handle.net/1765/37542>, EPS-2012-268-MKT.

I.A. Tuncdogan. Decision Making and Behavioral Strategy: The role of regulatory focus in corporate innovation processes. Promotor(s) Prof. F.A.J. van den Bosch, Prof. H.W. Volberda, Prof. T.J.M. Mom. <http://hdl.handle.net/1765/76978>, EPS-2014-334-S&E.

S. Tzioti. Let Me Give You a Piece of Advice: Empirical Papers about Advice Taking in Marketing. Promotor(s): Prof.dr. S.M.J. van Osselaer & Prof.dr.ir. B. Wierenga. <http://hdl.handle.net/1765/21149>, EPS-2010-211-MKT.

I.G. Vaccaro. Management Innovation: Studies on the Role of Internal Change Agents. Promotor(s): Prof.dr. F.A.J. van den Bosch & Prof.dr. H.W. Volberda. <http://hdl.handle.net/1765/21150>, EPS-2010-212-STR.

D. Vagias. Liquidity, Investors and International Capital Markets. Promotor: Prof.dr. M.A. van Dijk. <http://hdl.handle.net/1765/41511>, EPS-2013-294-F&A.

L.P. Veelenturf. Disruption Management in Passenger Railways: Models for Timetable, Rolling Stock and Crew Rescheduling. Promotor: Prof.dr. L.G. Kroon. <http://hdl.handle.net/1765/1>, EPS-2014-327-LIS.

H.J.J. Verheijen. Vendor-Buyer Coordination in Supply Chains. Promotor: Prof.dr.ir. J.A.E.E. van Nunen. <http://hdl.handle.net/1765/19594>, EPS-2010-194-LIS.

M. Venus. Demystifying Visionary Leadership; In Search of the Essence of Effective Vision Communication. Promotor: Prof.dr. D.L. van Knippenberg. <http://hdl.handle.net/1765/40079>, EPS-2013-289-ORG.

V. Visser. Leader Affect and Leader Effectiveness; How Leader Affective Displays Influence Follower Outcomes. Promotor: Prof.dr. D. van Knippenberg. <http://hdl.handle.net/1765/40076>, EPS-2013-286-ORG.

A.J. Vlam. Customer First? The Relationship between Advisors and Consumers of Financial Products. Promotor: Prof.dr. Ph.H.B.F. Franses. <http://hdl.handle.net/1765/30585>, EPS-2011-250-MKT.

E.J. de Waard. Engaging Environmental Turbulence: Organizational Determinants for Repetitive Quick and Adequate Responses. Promotor(s): Prof.dr. H.W. Volberda & Prof.dr. J. Soeters. <http://hdl.handle.net/1765/18012>, EPS-2010-189-STR.

L. Waltman. Computational and Game-Theoretic Approaches for Modeling Bounded Rationality. Promotor(s): Prof.dr.ir. R. Dekker & Prof.dr.ir. U. Kaymak. <http://hdl.handle.net/1765/26564>, EPS-2011-248-LIS.

Y. Wang. Information Content of Mutual Fund Portfolio Disclosure. Promotor: Prof.dr. M.J.C.M. Verbeek. <http://hdl.handle.net/1765/26066>, EPS-2011-242-F&A.

Y. Wang. Corporate Reputation Management; Reaching Out to Find Stakeholders. Promotor: Prof.dr. C.B.M. van Riel. <http://hdl.handle.net/1765/38675>, EPS-2013-271-ORG.

T.C. Weenen. On the Origin and Development of the Medical Nutrition Industry. Promotors: Prof.dr. H.R. Commandeur & Prof.dr. H.J.H.M. Claassen. <http://hdl.handle.net/1765/51134>, EPS-2014-309-S&E.

M. Wolfswinkel. Corporate Governance, Firm Risk and Shareholder Value of Dutch Firms. Promotor: Prof.dr. A. de Jong. <http://hdl.handle.net/1765/39127>, EPS-2013-277-F&A.

Y. Xu. Empirical Essays on the Stock Returns, Risk Management, and Liquidity Creation of Banks. Promotor: Prof.dr. M.J.C.M. Verbeek. <http://hdl.handle.net/1765/18125>, EPS-2010-188-F&A.

S.Y. Yang. Information aggregation efficiency of prediction market. Promotor: Prof.dr.ir. H.W.G.M. van Heck. <http://hdl.handle.net/1765/1>, EPS-2014-323-LIS.

N. Zaerpour. Efficient Management of Compact Storage Systems. Promotor: Prof.dr. M.B.M. de Koster. <http://hdl.handle.net/1765/38766>, EPS-2013-276-LIS.

D. Zhang. Essays in Executive Compensation. Promotor: Prof.dr. I. Dittmann. <http://hdl.handle.net/1765/32344>, EPS-2012-261-F&A.

X. Zhang. Scheduling with Time Lags. Promotor: Prof.dr. S.L. van de Velde. <http://hdl.handle.net/1765/19928>, EPS-2010-206-LIS.

About the author

Markus Peters was born on July 12, 1978 in Neuss, Germany. From 1999 to 2005 he studied Computer Science (Data Mining) and Business Administration (Operations Research) at the University of Aachen. Supported by a Fulbright fellowship, he spent the 2003/2004 academic year at Rensselaer Polytechnic Institute in Troy, NY where he obtained a Master's degree in Information Technology.



After his graduation, Markus worked as IT consultant, first for Deloitte's Business Intelligence (BI) service line and later independently. He advised clients in manufacturing, logistics, marketing, and software development on the design and implementation of enterprise BI systems, and on software engineering topics more generally.

Markus entered the graduate program at the Erasmus Research Institute of Management (ERIM) in 2011 to work on Machine Learning algorithms for future retail electricity markets. For this work, he was awarded with the ERIM Master's degree in Business Research in 2012.

Markus' work has appeared in *Data and Knowledge Engineering* and in the *Machine Learning Journal*, and it has been presented at various conferences such as the Conference of the Association for the Advancement of Artificial Intelligence (AAAI), the Conference on Information Systems and Technology (CIST), the European Conference on Machine Learning (ECML), and the Workshop on Information Technology and Systems (WITS).

MACHINE LEARNING ALGORITHMS FOR SMART ELECTRICITY MARKETS ESSAYS ON AUTONOMOUS ELECTRICITY BROKER DESIGN, PROBABILISTIC PREFERENCE MODELING, AND COMPETITIVE BENCHMARKING

The shift towards sustainable power systems is one of the grand challenges of the twenty-first century. Decentralized production from renewable sources, electric mobility, and related advances are at odds with traditional power systems where central large-scale generation of electricity follows inelastic consumer demand. Smart Markets and intelligent Information Systems (IS) could alleviate these issues by providing new forms of coordination that leverage real-time consumption information and prices to incentivize consumer behaviors that remain within the grid's operational bounds. But the best design for these systems, and the societal implications of different design choices is largely unclear. This dissertation makes three contributions to the debate. First, we propose and evaluate a design theory for Brokers, a novel class of IS-based intermediaries in retail electricity markets that provide participants with additional information and fine-grained economic incentives. Second, we study in greater detail the Broker's core problem of learning from past customer choices as basis for future decisions. We propose a probabilistic preference model that addresses important features of electricity markets, and we demonstrate the performance of this model on electricity tariff choice tasks. And third, we propose and study Competitive Benchmarking, a novel research method for effective IS artifact design in complex environments like power systems, where the social cost of failure is prohibitive. Our results provide guidance on IS design choices for sustainable electricity systems, and they contribute to the foundations for new Smart Grid business models.

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