# Agent Interactions & Mechanisms in Markets with Uncertainties

electricity markets in renewable energy systems

# Agent Interactions & Mechanisms in Markets with Uncertainties: Electricity Markets in Renewable Energy Systems

## Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus, Prof.dr.ir. T.H.J.J. van der Hagen,
chair of the Board for Doctorates
to be defended publicly on
Wednesday 15th July 2020 at 10:00 o'clock

by

Georgios METHENITIS

Master of Science in Artificial Intelligence,
University of Amsterdam, The Netherlands
born in Amarousio, Greece

This dissertation has been approved by the promotors.

Composition of the doctoral committee:

| | |
|---|---|
| Rector Magnificus | chairperson |
| Prof.dr.ir. J.A. La Poutré | Delft University of Technology and Centrum Wiskunde & Informatica, promotor |
| Dr. Michael Kaisers | Centrum Wiskunde & Informatica, copromotor |

Independent members:

| | |
|---|---|
| Prof.dr. S.N. Singh | Indian Institute of Technology Kanpur |
| Prof.dr. J.L. Hurink | University of Twente |
| Prof.dr. P. Palensky | Delft University of Technology |
| Prof.dr. A. Rogers | University of Oxford, UK |
| Prof.dr. C. Witteveen | Delft University of Technology |

The research in this thesis was performed at the Intelligent & Autonomous Systems (IAS) group of Centrum Wiskunde & Informatica (CWI), the national research institute for mathematics and computer science in the Netherlands.

SIKS Dissertation Series No. 2020-18 - The research reported in this thesis has been carried out under the auspices of SIKS, the Dutch Research School for Information and Knowledge Systems.

This work was part of the research programme Uncertainty Reduction in Smart Energy Systems (URSES) with project number 408-13-012, which was partly financed by the Netherlands Organisation for Scientific Research (NWO).

ISBN 978-946-40236-33

An electronic version of this dissertation is available at http://repository.tudelft.nl/.

# Summary

Electricity consumption is highly correlated with the level of human development, which alongside electrification is expected to significantly increase global demand for electricity in the coming decades. In current electricity systems, most of the electricity is generated by large fossil-fuel power plants on-demand and it is distributed by centrally-managed electricity grids. The increasing demand for electricity, however, should not go hand in hand with the simultaneous intensification of fossil-fuel mine and use, which is a driving cause of rising average temperatures on Earth's surface. Natural sources such as the sun and wind are expected to replace conventional sources of electricity, such as coal and gas power plants, in the near future, providing a key measure to address climate change and abate the effects of global warming. However, the intermittent and distributed nature of renewable electricity sources requires a redesign of conventional electricity grids that were originally designed following a top-down approach.

The smart grid is an electricity grid innovation that uses digital communication, measuring, and distributed control to facilitate efficient electricity usage and utilization of renewable electricity sources. Next to adopting more and more renewable electricity sources, users have an active role in the smart grid, both of which pose new and significant challenges. One key challenge is the design of economic mechanisms that encourage active participation of users, and at the same time can deal with the increasing uncertainty of both demand and supply. It is also crucial to analyze the behavior of future electricity systems since the collective efficiency of such systems may be influenced by the decision-making of self-interested users (agents).

In this thesis we focus on fundamental multi-agent systems that are motivated by the transition in electricity systems and relevant settings of the smart grid. In such systems we study strategic interactions and derive optimal strategies for agents in the presence of uncertainty; we further design economic mechanisms for resource allocation that yield efficient outcomes for all participating agents; and last, using tools from game theory, we analyze the behavior of these systems in both micro- and macro-levels. The contributions of this thesis advance state-of-the-art methods that: incentivize uncertainty reduction in the demand of customers (Chapter 2), generalize demand response mechanisms under uncertainty (Chapter 3), facilitate electricity trading under uncertainty in the supply (Chapter 4), and last, analyze the behavior of retail markets under different economic decision-making behavior of buyers (Chapter 5).

In both current and future electricity systems there is a need for continuously balancing supply and demand. To this end, in Chapter 2 we study the design of an innovative electricity tariff to incentivize customers to reduce the uncertainty of their demand, the *risk-sharing tariff*: a two-step parameterized payment scheme that provides the customer the choice to assume a fraction of the retailer's costs

associated with balancing supply and demand. More specifically, this chapter studies a multi-agent system in which a customer wants to purchase a continuously divisible good from a retailer and has a direct influence on the balancing costs of the retailer. Within a game-theoretical analysis, we show that the risk-sharing tariff provides the customer incentives to assume a fraction of the balancing costs. We also show that our proposed tariff is acceptable for both the retailer and the customer, i.e., both have economic incentives to adopt such tariff scheme. In this chapter we further study the influence of the customer's stochastic decision-making on the design of the risk-sharing tariff, since the latter provides the customer with the choice of how much risk to assume from the retailer. Overall, we show that novel tariff schemes, such as the risk-sharing tariff, can enable indirect control of customers' demand and tackle demand uncertainty for retailers in future smart grid systems.

In similar settings, Chapter 3 studies mechanisms to incentivize small-scale users to resolve last-minute imbalances between the available supply and the realization of the demand. More specifically, in this chapter we consider small-scale flexible assets that can alter their demand or generation behavior, e.g., electric vehicles, if they prepare ahead of the realization of the demand. Such flexible assets can be used by retailers on-demand to minimize over-generation and demand peaks that often cause excessive balancing costs to retailers. Building upon previous work, this chapter advances state-of-the-art economic mechanisms to incentivize a number of flexible users to prepare ahead and respond (last-minute) if requested by the retailer. The proposed mechanisms guarantee that both demand response agents and the retailer benefit in expectation, which alongside their simplicity and low computational complexity provide a promising avenue for using the available flexibility of small-scale users and complement existing demand response programs.

In contrast to previous technical chapters that study mechanisms to deal with the problem of balancing supply and demand, Chapter 4 presents a contracting framework to facilitate electricity trading in settings where supply depends on volatile sources, and thus delivery cannot be guaranteed. More specifically, we propose the adoption of service-level agreements (SLAs) that comprise the following features: quantity, reliability, and price. In this chapter, first, we define a family of utility functions for customers with regards to the probability of satisfying their demand, thus extending the concept of the value of lost load (VoLL) with the extra costs associated to the risk of failed delivery. Next, we study the design of economic mechanisms in order to specify and allocate these contracts (SLAs) to different customers, each of which has a different utility function. We demonstrate that the proposed mechanisms dominate alternative allocations that use only the VoLL, and vastly improve the efficiency of the studied system. Overall, the proposed mechanisms can facilitate distributed electricity trading under uncertainty in the supply, adding an essential component to future smart grid systems.

In the last technical chapter of this thesis, Chapter 5, we consider retail markets that enable automated software agents to participate instead of human buyers. The discrepancy between the non-perfect decision-making of human buyers due to information or time limitations and software agents that act optimally with regards to individual interactions may have adverse effects in such settings. In this chapter we investigate the effects of different economic decision-making of buyers on retail markets with regards to the resulting market dynamics and prices. By modeling buyers' different levels of rationality and the competition between sellers, we derive analytically best response strategies for the sellers and we analyze the evolutionary behavior of retail markets under different degrees of buyers' rationality. The theoretical and empirical results of this chapter suggest that perfect rationality have undesirable effects on market competition, which raises the need to revisit design objectives for software agents in future retail markets.

Overall, in this thesis we study agent-based interactions and propose novel economic mechanisms within fundamental models that comprise strategic situations and are motivated by the transition towards the smart grid. Our findings can be used as innovative components of future smart grid systems, which are characterized by the increasing uncertainty on both demand and supply and actively participating users. In addition, our technical contributions provide insights that transfer to the design and analysis of multi-agent systems with similar characteristics of uncertainty in resource allocation.

# Samenvatting

Het verbruik van elektriciteit hangt sterk samen met de mate van menselijke ontwikkeling. Naar verwachting zal dit, samen met toenemende elektrificatie, in de komende decennia een aanzienlijke toename in de wereldwijde vraag naar elektriciteit teweeg brengen. In de elektriciteitssytemen van vandaag wordt de meeste elektriciteit geproduceerd door 'grijze' energiecentrales die op fossiele brandstoffen draaien. Deze grijze centrales volgen de elektriciteitsvraag en de opgewekte elektriciteit wordt gedistribueerd door centraal geregelde elektriciteitsnetwerken. De groeiende vraag naar elektriciteit zou echter niet hand in hand moeten gaan met een toename in winning en verbruik van fossiele brandstoffen; deze zijn immers een belangrijke oorzaak voor wereldwijd stijgende temperatuurgemiddelden. Naar verwachting zullen hernieuwbare elektriciteitsbronnen, zoals zon- en windcentrales, in de nabije toekomst conventionele elektriciteitsbronnen, zoals kolen- en gascentrales, vervangen. Deze transformatie zal een sleutelrol spelen in de inspanningen om klimaatverandering het hoofd te bieden en de effecten van de opwarming van de aarde te verminderen. De onregelmatige en gedistribueerde aard van hernieuwbare elektriciteitsbronnen vraagt echter om een hernieuwd ontwerp van conventionele elektriciteitsnetwerken, die oorspronkelijk een *top-down* ontwerp volgen.

Het *smart grid* is een innovatie van het elektriciteitsnet. Het *smart grid* elektriciteitsnet faciliteert het efficiënte gebruik van elektriciteit en hernieuwbare elektriciteitsbronnen, door gebruik te maken van digitale communicatie, meting en gedistribueerde controle. Naast het in gebruik nemen van steeds meer hernieuwbare bronnen, krijgen gebruikers een actieve rol in het *smart grid*. Allebei deze ontwikkelingen stellen ons voor nieuwe en betekenisvolle uitdagingen. Één belangrijke uitdaging is het ontwerp van economische mechanismen die de actieve deelname van gebruikers aanmoedigen en tegelijkertijd om kunnen gaan met de toenemende onzekerheid van zowel vraag als aanbod. Daarnaast is het cruciaal om het gedrag van toekomstige elektriciteitssystemen te analyseren, omdat de effectiviteit van het gehele system beïnvloed zou kunnen worden door de beslissingen van individuele gebruikers (agenten), die uit eigenbelang handelen.

In dit proefschrift richten we ons op fundamentele multi-agentsystemen, gemotiveerd door de transitie in elektriciteitssystemen naar het *smart grid*. In zulke systemen bestuderen we strategische interacties en leiden hieruit optimale strategieën af voor agenten, in de aanwezigheid van onzekerheid. Bovendien ontwerpen we economische mechanismen voor de toewijzing van middelen die efficiënte resultaten geven voor alle deelnemers. Ten slotte analyseren we het gedrag van deze systemen op zowel micro- als macroniveau, waarbij we gebruikmaken van resultaten uit de speltheorie. De bijdragen van dit proefschrift breiden de stand van de nieuwste methoden uit op vier manieren: het stimuleren van het verminderen van de onzekerheid in de vraag van afnemers (Hoofdstuk 2), het generaliseren van vraagresponsmechanismen

onder onzekerheid (Hoofdstuk 3), het faciliteren van elektriciteitshandel onder onzekerheid in het aanbod (Hoofdstuk 4) en ten slotte, het analyseren van het gedrag van retail-markten onder verschillende economische besluitvormingsprocessen van kopers (Hoofdstuk 5).

In zowel huidige als toekomstige elektriciteitssystemen is het noodzakelijk om continu vraag en aanbod te balanceren. Daarom bestuderen we in Hoofdstuk 2 het ontwerp van een innovatief elektriciteitstarief, dat afnemers stimuleert om de onzekerheid in hun vraag te verkleinen. Dit tarief noemen we het *risico-delingstarief*: een twee-staps geparameteriseerde betalingsregeling die de afnemers de keuze biedt om een deel van de kosten, die de leverancier maakt voor het balanceren van vraag en aanbod, op zich te nemen. In het bijzonder behandelt dit hoofdstuk multi-agentsystemen, waarin een afnemer een continu deelbaar goed van een leverancier wil kopen en zelf een directe invloed heeft op de balanceerkosten van de leverancier. In een speltheoretische analyse laten we zien dat het risico-delingstarief de afnemer stimuleert om een deel van de balanceerkosten op zich te nemen. Daarnaast laten we zien dat het door ons voorgestelde tarief acceptabel is voor zowel de leverancier als de afnemer. Dat wil zeggen, beiden hebben economishe reden om een dergelijke tarievenregeling in te voeren. In dit hoofdstuk gaan we verder in op de invloed van de stochastische besluitvorming van de afnemer op het ontwerp van het risico-delingstarief. Het risico-delingstarief biedt de afnemer namelijk de keuze hoeveel risico deze over wil nemen van de leverancier. Uiteindelijk laten we zien dat nieuwe tariefregelingen, zoals het risico-delingstarief, indirecte controle over de vraag van afnemers kan uitoefenen en de onzekerheid in de vraag aan leveranciers kan aanpakken in de *smart grid*-systemen van de toekomst.

In vergelijkbare omstandigheden bestudeert Hoofdstuk 3 mechanismen die kleinschalige gebruikers stimuleren om onbalans tussen het beschikbare aanbod en de realisaties van de vraag, die op het laatste moment ontstaat, op te lossen. In het bijzonder behandelen we in dit hoofdstuk kleinschalige flexibele deelnemers die hun vraag of productie kunnen aanpassen, zoals bijvoorbeeld elektrische voertuigen, wanneer die zich voorbereiden op de realisatie van de vraag. Zulke flexibele deelnemers kunnen door leveranciers op aanvraag gebruikt worden om overproductie en pieken in de vraag, die vaak hoge balanceringskosten voor leveranciers tot gevolg hebben, te minimaliseren. Voortbouwend op voorgaand werk, vordert dit hoofdstuk de nieuwste economische mechanismen door een aantal flexibele gebruikers te stimuleren zich voor te bereiden en (op het laatste moment) te reageren, indien gevraagd door de leverancier. De voorgestelde mechanismen garanderen dat zowel vraagrespons agenten als de leverancier hier in verwachting baat bij hebben. Samen met hun simpliciteit en lage computationele complexiteit biedt dit een veelbelovende weg naar het gebruik van de beschikbare flexibiliteit van kleinschalige gebruikers en het aanvullen van bestaande vraagresponsprogramma's.

In contrast met voorgaande technische hoofdstukken, die mechanismen voor het balanceren van vraag en aanbod bestudeerden, behandelt Hoofdstuk 4 een contractenkader voor het faciliteren van energiehandel in omstandigheden waar de vraag afhangt van onzekere bronnen en waar levering dus niet gegarandeerd kan worden. In het bijzonder stellen we het gebruik van *service-level agreements* (SLAs) voor, die de volgende kenmerken hebben: kwantiteit, betrouwbaarheid en prijs. In dit hoofdstuk stellen we eerst een familie van nutsfuncties voor afnemers voor. Deze nutsfuncties hebben betrekking tot de kans dat de vraag van de afnemer vervuld wordt. Hiermee wordt het concept van de waarde van verloren lading (*value of lost load* of VoLL) uitgebreid met de extra kosten die het risico van mislukte levering met zich meebrengt. Vervolgens bestuderen we het ontwerp van economische mechanismen, met als doel het specificeren en toewijzen van deze contracten (SLAs) aan verschillende afnemers, waarvan elk een andere nutsfunctie heeft. We laten zien dat de voorgestelde mechanismen alternatieve toewijzingen, die alleen de VoLL gebruiken, domineren en dat ze de efficiëntie van het bestudeerde systeem ruim verbeteren. Zo zien we dat de voorgestelde mechanismen de gedistribueerde energiehandel kunnen faciliteren onder onzekerheid in het aanbod en zo een essentiele component toevoegen aan *smart grid*-systemen van de toekomst.

In het laatste technische hoofdstuk van dit proefschrift, Hoofdstuk 5, bekijken we toekomstige retail markten die geautomatiseerde software-agenten als deelnemers hebben, in plaats van menselijke kopers. In dergelijke situaties kan de discrepantie, met betrekking tot individuele acties, tussen imperfecte beslissingen van menselijke kopers ten gevolge van beperkingen in informatie of tijd enerzijds en software-agenten die optimaal acteren anderzijds, tegengestelde effecten hebben. In dit hoofdstuk onderzoeken we de effecten van verschillende economische besluitvormingsprocessen van kopers in retail markten, met betrekking tot de resulterende marktdynamieken en -prijzen. Door het modelleren van verschillende niveaus van rationaliteit van kopers en de competitie tussen verkopers, leiden we analytisch strategieën met beste antwoorden voor de verkopers af. Daarnaast analyseren we het evolutionaire gedrag van retail markten onder verschillende niveaus van rationaliteit van kopers. De theoretische en empirische resultaten van dit hoofdstuk suggereren dat perfecte rationaliteit ongewilde effecten heeft op marktcompetitie. Dit betekent dat ontwerpdoelen moeten worden herzien voor software-agenten in retail markten van de toekomst.

Al met al bestuderen we in dit proefschrift agent-gebaseerde interacties en stellen we nieuwe economische mechanismen voor. Dit doen we in de context van fundamentele modellen die strategische situaties omvatten en gemotiveerd zijn door de transitie naar het *smart grid*. Onze bevindingen kunnen gebruikt worden als innovatieve componenten voor *smart grid*-systemen van de toekomst, die gekarakteriseerd worden door de toenemende onzekerheid in zowel vraag en aanbod als actief deelnemende gebruikers. Daarnaast bieden onze technische bijdragen inzichten die toepasbaar zijn op het ontwerp en de analyse van multi-agentsystemen, wanneer die gelijksoortige karakteristieken van onzekerheid op het gebied van toewijzing van middelen hebben.

# Acknowledgements

This thesis concludes a big chapter of my life; during the last five years, I grew both as a scientific researcher and as an individual by keeping an open mind and trying to embrace all the tough, challenging, rewarding, and relaxing moments that a PhD can offer. I enjoyed every single moment of this journey, each one making me realize yet another thing about myself and life in general.

For all the above and many more, first, I would like to thank my promotor Han La Poutré and my copromotor Michael Kaisers. Dear Han, thank you for all the advice, support, attention to detail, and challenging questions you were asking me every time I was presenting you a new idea or yet another mathematical proof. I acknowledge your passion for scientific integrity and excellent scientific research. Dear Michael, all this would have been impossible without your feedback which was of great importance to me especially in the early days of my PhD. At that time, I was lost in the plethora of possibilities and research directions. I would like to thank you both for the trust in my work and guidance.

I would also like to acknowledge the guidance I have received earlier in my academic life and was of great importance to my decision to pursue a doctoral degree. I would like to thank Michail G. Lagoudakis for inspiring and encouraging me to continue my studies in artificial intelligence in the Netherlands when I was a student at the Technical University of Crete. I am also grateful to Arnoud Visser from the University of Amsterdam for supervising my master thesis and introducing me to the activities of the Intelligent Robotics Lab and the Dutch Nao Team. Last, I would like to thank Daniel Hennes, Dario Izzo and all the members of the Advanced Concepts Team at the European Space Agency where I spend the summer of 2014 working on my master thesis, this was one of the biggest learning experiences of my life and something that I will never forget.

I also wish to thank all members of the STASCADE project: Frances Brazier, Martijn Warnier, Nina Voulis for their excellent collaboration, and the members of the user committee: Gabriel Bloemhof, George Trienekens, and Mark van Assem for their valuable input during user meetings. Dear Nina, thank you for all the discussions we had during this joint project.

Next, I would like to thank my friend Nicos Voskarides. Dear Nicos, your help was invaluable during this journey. This was not limited to proof-reading my manuscripts and trying to understand the complexities of a totally different domain than yours; your "perfectly rational" perspective was and is still clearing things up for me every time I am asking for your help and advice. Talking over all the research challenges we were facing in our academic lives was a weight off my shoulders.

In addition, I would like to thank Mihail Bazhba. Dear Mihail, thank you for your help in proving Lemma 5.2 and for all interesting discussions we had during our breaks from work.

Furthermore, I would like to express my gratitude towards all former and current members of the Autonomous and Intelligent Systems group at CWI and the rest of the staff at CWI, everyone contributing to a relaxed and supportive work environment. Big thanks to Eric Pauwels, Jasper Hoogland, Felix Claessen, Hoang Luong, Pablo Hernandez-Leal, Daan Bloembergen, Tim Baarslag, Satish Sharma, Dimitrios Tsimpoukis, Iliana Pappi, and Jieting Luo. Special thanks to my office-mates, Aliene van der Veen, André Vaz, Brinn Hekkelman, and Roland Saur not just for sharing an office but for all the lengthy discussions on new research ideas, challenging whiteboard problems, and for keeping it fun. Dear Brinn, I am so grateful to you for translating this thesis' summary in Dutch.

Last and foremost, I would like to thank my friends and family. I am grateful for all the amazing people I met in Amsterdam (Fee, Ίρις, Ιωάννα, Σοφία, Νίκος, Μάριος, Πάρις, Γιάννης, Αχιλλέας, Στάθης, Μιχαήλ, Tommaso, Amogh), Chania (Βασίλης, Βαγγέλης, Δημοσθένης, Δημήτρης Λ., Δημήτρης Μ., Κωνσταντίνος, Στάθης, Χρήστος), and my hometown (Σταμάτης, Δημήτρης Μ., Δημήτρης Κ., Βαγγέλης, Κωστής). You have all supported me in many different ways during this journey, thank you for your love.

Ευχαριστώ Ελένη.

This thesis is dedicated to my family. Ευχαριστώ πατέρα και μάνα για την αγάπη και στήριξη σας. Ευχαριστώ Μελίνα. Αυτή η διπλωματική είναι αφιερωμένη σε εσάς.

# Contents

**Summary** vii

**Samenvatting** xi

**Acknowledgements** xv

**1 Introduction** 1

1.1 Multi-Agent Systems . . . 3
1.1.1 Intelligent agents . . . 4
1.1.2 Multi-agent interactions . . . 5
1.1.3 Game theory . . . 6
1.1.4 Mechanism design . . . 7
1.1.5 Markets . . . 12
1.2 Transition in Electricity Systems . . . 14
1.2.1 Electricity markets & balancing requirements . . . 14
1.2.2 Renewable electricity sources . . . 16
1.2.3 Smart grid . . . 17
1.2.4 Smart grid challenges & solutions . . . 18
1.3 Research Questions . . . 23
1.3.1 Demand & supply uncertainty . . . 23
1.3.2 Agent-based decision making . . . 25
1.3.3 Summary of research questions . . . 26
1.4 Thesis Outline . . . 26
1.5 List of Publications . . . 28

**2 Reducing Demand Uncertainty through Risk-Sharing** 31

2.1 Introduction . . . 32
2.2 The Risk-Sharing Game . . . 34
2.2.1 Optimal quantity of procurement . . . 36
2.2.2 Determining the price for risk-sharing . . . 37
2.2.3 Optimal strategies for flexible customers . . . 39
2.2.4 Comparison of the utilities . . . 44
2.3 Nash Equilibrium Strategies . . . 44
2.3.1 Bounded rational customer . . . 46
2.4 Conclusions . . . 47

2.4.1 Future work . . . . . . . . 48

3 Forecast-Based Mechanisms for Demand Response 51
3.1 Introduction . . . . . . . . 52
3.2 Related Work . . . . . . . . 54
3.3 Problem Formulation . . . . . . . . 54
3.3.1 Retailer's balancing responsibility . . . . . . . . 54
3.3.2 Demand response agents . . . . . . . . 56
3.3.3 Model of demand response . . . . . . . . 57
3.4 Demand Response Mechanism $M$ . . . . . . . . 57
3.4.1 Request probability & interdependencies . . . . . . . . 59
3.4.2 Sequential-task mechanism . . . . . . . . 61
3.4.3 Independent-task mechanism . . . . . . . . 62
3.4.4 General flexibility mechanisms . . . . . . . . 64
3.5 Experimental Evaluation . . . . . . . . 65
3.6 Conclusions . . . . . . . . 69
3.6.1 Future work . . . . . . . . 69

4 SLA Allocation for Renewable Electricity Trading 73
4.1 Introduction . . . . . . . . 74
4.2 Related Work . . . . . . . . 76
4.3 SLA Contracting Framework . . . . . . . . 77
4.3.1 Problem formulation . . . . . . . . 77
4.3.2 Service-level agreements . . . . . . . . 77
4.3.3 Critical & tolerant buyers . . . . . . . . 79
4.4 Auction-Based SLA Allocation . . . . . . . . 81
4.4.1 Sequential second-price auction . . . . . . . . 82
4.4.2 Vickrey-Clarke-Groves . . . . . . . . 83
4.5 Evaluation & Discussion . . . . . . . . 86
4.5.1 Diversity in the criticality . . . . . . . . 86
4.5.2 Demand over supply ratio . . . . . . . . 88
4.5.3 Social value & social welfare . . . . . . . . 89
4.6 Conclusions . . . . . . . . 89
4.6.1 Future work . . . . . . . . 90

5 Bounded-Rationality & Retail Markets 93
5.1 Introduction . . . . . . . . 94
5.2 Related Work . . . . . . . . 96
5.3 Market Model . . . . . . . . 98
5.3.1 Degree of buyers' rationality . . . . . . . . 98
5.3.2 k-level reasoning & competition . . . . . . . . 99
5.4 k-Level Best Response Strategies . . . . . . . . 100
5.4.1 Best response strategies in duopolies . . . . . . . . 100
5.4.2 Analytical best response strategies & rationality . . . . . . . . 100
5.4.3 Duopoly markets . . . . . . . . 103
5.5 Evolutionary Dynamics . . . . . . . . 105

5.5.1 Dynamic belief of competition . . . . . . . . . . . . . . . . . . . . . . 106
5.5.2 Optimal pricing & generalized replicator dynamics . . . . . . . . . . . . 106
5.6 Conclusions . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 112
5.6.1 Future work . . . . . . . . . . . . . . . . . . . . . . . . . . . 113

**6 Conclusions 115**
6.1 Main Contributions . . . . . . . . . . . . . . . . . . . . . . . . . 116
6.1.1 Overview . . . . . . . . . . . . . . . . . . . . . . . . . . . . 116
6.1.2 Risk-sharing tariff & demand uncertainty reduction . . . . . . . . . . . . 117
6.1.3 Demand response mechanisms for unreliable agents . . . . . . . . . . . 118
6.1.4 SLAs for renewable electricity trading . . . . . . . . . . . . . . . . . . 119
6.1.5 Decision-making of buyers in retail markets . . . . . . . . . . . . . . . 120
6.2 Concluding Remarks . . . . . . . . . . . . . . . . . . . . . . . . 121
6.2.1 Future research directions . . . . . . . . . . . . . . . . . . . . . 122

**Bibliography 125**

**SIKS Dissertation Series 139**

**Curriculum Vitae 153**

# 1 Introduction

Electricity consumption is highly correlated with the level of human development, and it is therefore evident that electricity demand will dramatically increase in the following decades (Niu et al., 2013). More specifically, both the development of our societies and the increasing electrification are expected to triple demand for electricity by 2050 (Farhangi, 2010). Currently, most of electricity supply comes from traditional fossil fuel power sources, such as coal, oil and natural gas; these are mined and mainly used as burning fuels for electricity generation, heating, and transportation. However, fossil fuels are responsible for most of the $CO_2$ emissions that are related to human-activity (Raupach et al., 2007). The increasing levels of greenhouse gases, such as $CO_2$, in the atmosphere is the major reason for the significant increase of the average temperature on Earth's surface and the effects of *global warming* (Lashof and Ahuja, 1990; Meinshausen et al., 2009).

The vast increase of electricity demand should not go hand in hand with the simultaneous increase in fossil fuel mining and use. Natural sources such as the sun and wind are expected to replace conventional fossil fuel sources in the future, which alongside other technological advances in electrification of transportation (e.g., railways, electric vehicles), and heating (Moraga-González and Mulder, 2018) have the potential to reduce $CO_2$ emissions and thus abate the effects of global warming (Jenkinson et al., 1991; Mora et al., 2017).

Current electricity grids have been originally designed following a top-down approach: electricity supply is provided by few centrally located large fossil-fuel power plants on-demand (Ramchurn et al., 2011), the supply of which is pooled and traded in electricity markets. The increasing penetration of renewable electricity generation, on the contrary to conventional centralized power plants, is distributed and it can further be adopted by consumers on a local level, i.e., making them prosumers.[1] In addition, consumers are expected to take an active role in future electricity systems by being able to control their own net demand with the use of energy storage technologies. The above reasons induce higher uncertainty on both supply and demand-sides, and therefore pose many challenges with regards to the balancing requirements (i.e., supply and demand should be equal) of future electricity systems. However, they also introduce an opportunity for the transition towards fully sustainable electricity systems without greenhouse emissions, in which active demand will follow the available supply of renewable electricity sources.

[1] Prosumers are entities that not only consume but actively participate in the production of goods (Toffler, 1990).

The increasing uncertainty in future electricity systems raises the need for the redesign of current electricity grids and the corresponding markets that facilitate electricity trading between producers and consumers. To this end, the *smart grid* is an electricity grid innovation away from the traditional paradigm of passive distribution and consumption (see Section 1.2.3). More specifically, in the smart grid, generation is decentralized and thus can be closely located to the demand load it serves. In addition, the presence of sensors and two-way communication between aggregators and electricity consumers enable active participation of consumption or generation entities (e.g., households, electric vehicles, solar PVs), direct or indirect control of loads, and high-resolution pricing schemes for electricity customers based on real-time consumption data. Real-time communication within the smart grid can further enable monitoring electrical characteristics of the network, and therefore not only optimize its function, but also enable fully autonomous operation of small parts of the grid (micro-grids) to mitigate wider system disturbances if necessary.

The smart grid innovation changes the way conventional electricity grids operate. Communication, coordination and economic mechanisms that go beyond centralized conventional systems and traditional flat electricity tariffs need to ensure that the smart grid is an efficient system design. More specifically, economic mechanisms need to ensure that:

- Demand follows the available renewable electricity supply and thus economic risks that are associated to balancing are alleviated.
- Costs and risks that are connected to balancing and the demand behavior of users are allocated in an acceptable manner and charged directly on those who cause them.
- The efficiency of future electricity markets is not affected by the increasing uncertainty of both supply and demand, or the strategic behavior of participating agents.

To this end, in this thesis we consider fundamental situations that are motivated by the transition in electricity systems and some settings of the smart grid. We then propose economic mechanisms to tackle the uncertainty in the demand or supply, and we further analyze the interactions between multiple decision-makers. For instance, we study the strategic interaction of producers and consumers when there is uncertainty about either the demand of consumers or the supply (renewable supply) of producers (see Chapters 2, 3 and 4). To analytically study such situations, each individual decision-maker (e.g., retailer, customer) is modeled as a self-interested agent that chooses its actions in order to maximize its utility. In addition, we study the effects of different economic decision-making of consumers on the design of electricity retail tariffs (see Chapters 2) and the resulting prices of competitive sellers in electricity retail markets (see Chapter 5). Overall, the contributions of this thesis are primarily connected to the field of computer science and economics with an application domain of current and future electricity markets.

The rest of this chapter is organized as follows: In Section 1.1 we provide a thorough discussion on fundamental concepts of multi-agent systems which we use

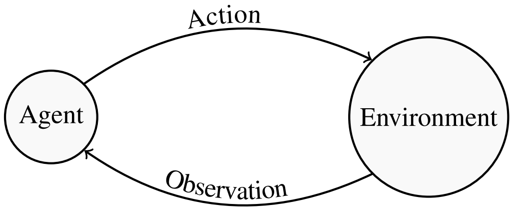


**Figure 1.1** An agent acts upon incoming observations within an environment in an ongoing process.

to formalize different situations in envisioned electricity markets.[2] In Section 1.2 we elaborate on our motivations, which are aligned to the transition of electricity systems towards generation portfolios that heavily depend on natural sources such as the sun and wind, alongside the solution concept of the smart grid. In Section 1.3 we formalize our problem statement and we outline the main research questions of this thesis. We conclude this chapter by providing an outline of the research topics and the overall structure of this thesis in Section 1.4, and present the research output of this thesis in Section 1.5.

## 1.1 Multi-Agent Systems

Throughout this thesis, we model situations that arise in the context of the smart grid, in which multiple decision-makers (e.g., retailer, customers) seek to maximize their utility and collectively influence the efficiency of such systems. *Multi-agent systems* are well-suited to study these complex settings of envisioned smart grid systems; not only do they model the decisions of independent and self-interested agents, but also provide the solution framework for problems that may be beyond the capabilities of single agents (Coelho et al., 2017; Kantamneni et al., 2015). In addition, the application of theoretical solution concepts of game theory, such as the *Nash equilibrium*, provide the means to study strategic situations that arise between multiple self-interested decision-makers (Fadlullah et al., 2011; Saad et al., 2012).

In this section we provide a thorough discussion on fundamental concepts of agents and multi-agent systems which we use throughout this thesis to analyze different scenarios of envisioned electricity systems.

[2] Readers that are familiar with concepts of multi-agent systems, such as game theory, mechanism design, and auctions, may skip Section 1.1.

### 1.1.1 Intelligent agents

Any social interaction environment, such as the smart grid, comprise autonomous decision-makers that observe the dynamics of the environment and act in order to have the best possible outcomes. For instance, buyers participate in electricity markets in order to purchase electricity at the minimum possible price, and sellers to maximize their profits. In artificial intelligence research, an autonomous decision-maker is described with the term *intelligent agent* (Russell and Norvig, 2009).

**Definition 1.1** (Intelligent agent)**.** *An intelligent agent is an entity that senses the environment through sensors and acts upon the environment using actuators.*

The above definition is not limited to computer software (e.g., automated trading software) or hardware that exhibits intelligent behavior; it also applies to human agents, where a central aspect of the definition includes the notion of *agency* (i.e., ability to act).

Figure 1.1 illustrates an abstract model of an intelligent agent that acts upon its environment and is influenced by incoming observations. By explicitly modeling and analyzing the behavior (actions) of individual agents offers a bottom-up (micro-scale) approach to study the emerging behavior of complex systems that comprise multiple agents (Macal and North, 2010).

Throughout this thesis, we consider *utility-based agents* that try to maximize a performance metric, which is described by its *utility function* (Russell and Norvig, 2009).

**Definition 1.2** (Utility)**.** *Utility of agent $i$ is the output of a utility function $u_i$ that measures the desirability of an outcome $x$ from the set of possible outcomes $X$, such that $u_i : X \to \mathbb{R}$.*

Intuitively, utility is a measure of satisfaction and can be used to determine the decision of an agent with regards to multiple available actions it can choose from. In this case, each action yields a utility, or an expected utility in uncertain environments, to the agent.

#### On the rationality of agents

One fundamental notion of intelligent agents is *rationality*. Based on the ethical theory of *utilitarianism* (Mill, 2014), we have the following definition:

**Definition 1.3** (Rational agent)**.** *Given a set of possible outcomes, a rational agent chooses the outcome that maximizes its utility, or, when there is uncertainty, its expected utility.*

Given two possible outcomes $x, y \in X$, a rational agent would always choose outcome $x$ if $u_i(x) > u_i(y)$. Fundamental models that study interactions between agents (see Section 1.1.3) usually assume the presence of rational agents (Nisan et al., 2007). The notion of rational agents is also related to economic agents that participate in free markets (Blume and Easley, 2016): a rational agent would always choose the item at the cheapest price given a set of heterogeneously priced identical items.

### Bounded rational agents

However, rationality comes with the precondition of *perfect information* and *unlimited computational capacity*. Motivated by the fact that perfect information and unlimited computational capacity of agents are not realistic assumptions in practice, Simon (1972) introduced the concept of *bounded rationality*:

**Definition 1.4** (Bounded rationality)**.** *Bounded rationality models the imperfect decision-making of otherwise rational agents due to: imperfect information, limited computational capacity or decision time constraints.*

The concept of bounded rationality is further supported by several works that study the economic decisions of human buyers in markets (Conlisk, 1996; Rubinstein, 1998). However, the following question emerges: *How can we analytically model the imperfect choice of bounded rational agents?* The answer is based on diverse scientific fields, such as psychology, economics and mathematics (Ortega and Braun, 2011; Ortega et al., 2015; Puranam et al., 2015). The first mathematical model proposed to express the stochastic decision-making of an agent over a finite set of choices was the *Luce's axiom* (Luce, 1959). Consider $x$ as a vector of available choices; $x_i$ is the $i$-th choice that an agent can choose and $u(x_i)$ is the utility of choice $i$. Following the Luce's axiom, the probability of choosing $i$ is proportional to the utility it brings to the agent.

$$\mathbb{P}_i(x) = \frac{u(x_i)}{\sum_j u(x_j)} \tag{1.1}$$

The vast majority of models proposed after Luce's axiom (Mattsson and Weibull, 2002; McFadden, 1973; Meginniss, 1976), are logit choice models based on the Boltzmann distribution. In Chapters 2 and 5 we use the *Softmax* rule (Sutton and Barto, 1998), which is also based on the Boltzmann distribution, to model the imperfect decision making of agents in electricity markets.[3]

$$\mathbb{P}_i(x) = \frac{e^{u(x_i)/\tau}}{\sum_j e^{u(x_j)/\tau}}, \ \forall \tau \in (0, \infty), \tag{1.2}$$

where $\tau$ is called the irrationality parameter. For $\tau \to 0$, Softmax approximates the decision-making of a rational agent that chooses the best option with probability one; for $\tau \to \infty$, an agent that chooses each of its options with equal probability (random).

## 1.1.2 Multi-agent interactions

So far, we have discussed the notion of a single agent that interacts within an environment and makes decisions with regards to possible outcomes rationally or under limited information (bounded rationality). However, there are only limited instances of real-world situations that involve a single intelligent agent that can alone influence all outcomes. In most scenarios, there are multiple agents that

[3] Softmax is primarily used in *Reinforcement Learning* to determine the probability of actions given their expected reward.

interact with each other and influence each other. Such systems are called *multi-agent systems* (Weiss, 2013; Wooldridge, 2001).

**Definition 1.5** (Multi-agent systems)**.** *Multi-agent systems are composed of multiple intelligent agents that interact in order to coordinate, solve complex problems, or determine the division of a common-pool resource.*

Some of the main research topics in the field of multi-agent systems include but are not limited to: learning (Tuyls and Weiss, 2012), communication (Foerster et al., 2016), cooperation (Olfati-Saber et al., 2007), and negotiation (Baarslag et al., 2017). The following two sections, however, discuss fundamental fields of research that are used throughout this thesis. Section 1.1.3 introduces *game theory* which is used to study interactions between self-interested agents, and Section 1.1.4 provides a brief overview of *mechanism design* which is used to allocate resources in multi-agent systems.

### 1.1.3 Game theory

Game theory is the study of mathematical models of strategic interaction between rational decision-makers (Myerson, 2013).[4] In game theory a strategic interaction is formalized as a *game*, in which each *player* (agent) can choose from a set of actions. A *payoff* (utility) matrix determines the utility of each player for any possible outcome. In general, there exist many types of games, e.g., cooperative/non-cooperative, zero-sum/general-sum, symmetric/asymmetric, see Myerson (2013) for more details. In this thesis we consider non-cooperative games (see Chapters 2 and 5), in which there is no cooperation while agents want to maximize their own utility. Non-cooperative games can be represented as *extensive* or *normal* form games, where the time sequencing of players' actions in the former distinguish it from the latter, in which players choose their actions simultaneously. In Chapter 2 we consider an extensive-form game between a retailer and a customer, while in Chapter 5 we study a Bertrand market model which comprises a normal-form game.

Consider the following two-player normal form game which is known as the *prisoners' dilemma* (Rapoport et al., 1965):

**Example 1.1** (Prisoners' dilemma)**.** *Two suspects, player A and player B, are accused of a crime. Both suspects are placed in confinement. However, they are placed in separate rooms and they cannot communicate with each other. Suspects have two choices: either to cooperate and confess their crime, or defect and betray the other suspect. Due to lack of strong evidence with regards to the investigated crime, the prosecution is willing to convict them with a minor infraction (*$1$ *year) if both cooperate and confess. If no suspect cooperates they will face a jail time of* $2$ *years. However, if one player cooperates while the other defects, the defector walks away free while the cooperator will face the maximum conviction of* $3$ *years.*

Table 1.1 shows the payoff matrix of the Prisoners' dilemma game. Players $A$ and $B$ have two possible actions: $Cooperate$ and $Defect$. For each combination of actions

[4] Foundational work in game theory was developed in the 1950s (Nash, 1950, 1951; von Neumann et al., 1944).

| | | Player $B$ | |
|---|---|---|---|
| | | $Cooperate$ | $Defect$ |
| Player $A$ | $Cooperate$ | $-1, -1$ | $-3, \ 0$ |
| | $Defect$ | $0, -3$ | $-2, -2$ |

**Table 1.1** Payoff matrix for the prisoners' dilemma game: player A chooses the row and player B chooses the column while each entry of the matrix indicates the utility for players A and B respectively.

chosen by the two players, the utility matrix gives the utility pairs for both players (note that the first entry of each utility pair corresponds to the utility of player A).

**Nash equilibrium** The most important solution concept when studying games between two or more players is the *Nash equlibrium* (NE). Given an $N$-player game and a payoff matrix (e.g., see Table 1.1):

**Definition 1.6** (Nash equilibrium)**.** *A strategy profile (i.e., a set of strategies that are used by the players) is a Nash equilibrium if no player can gain by unilaterally deviating from its strategy (Nash, 1950).*

Revisiting the Prisoners' dilemma (see Example 1.1) and its corresponding payoff matrix in Table 1.1. First, note that in case both players choose to cooperate both get utility $-1$, which is the best possible outcome with regards to the social welfare, i.e., sum of players' utilities. Consider player A, if player A cooperates, the best move for player B is to defect. If player A defects, the best move for B is again to defect. With this rationale, defection is the optimal strategy for both players. It is also the Nash equilibrium since no player can yield a better outcome deviating from this strategy given that the opponent chooses defection.

The solution concept of the Nash equilibrium is used in Chapter 2 of this thesis to determine stable strategy pairs in a two-player extensive-form game between a retailer and a customer in an electricity market setting.

### 1.1.4 Mechanism design

In the previous section we discussed basic concepts of game theory and the notion of the Nash equilibrium. In this section we provide a discussion on the basic theory of *mechanism design* (Myerson, 2013). On the contrary to game theory that attempts to analyze given games, mechanism design is considered as the reverse of game theory; its goal is to design games (e.g., on how to cooperate, divide a common-resource pool, reach mutually beneficial agreements) that have some desired properties in competitive settings where there exist no common goals.

**What is a mechanism?** Many real-world situations involve self-interested agents that wish to divide a common pool of items. However, agents usually have conflicting interests that cannot be resolved without some common-ground rules that can be provided by some protocol. For instance, consider the following scenario: a common-resource pool (e.g., the electricity generation of a wind-power turbine for the next hour) should be divided among $n$ agents. Each agent has a *type* $\theta_i \in \Theta_i$ that is *private information* (i.e., only agent $i$ knows its type), where $\Theta_i$ is the space of types of agent $i$. $\theta_i$ is also called the *true type* of agent $i$. Let $\theta = \{\theta_0, \theta_1, \dots, \theta_{n-1}\}$ denote the vector of agents' types, and $\Theta = \Theta_0 \times \Theta_1 \times \dots \times \Theta_{n-1}$ the space of all possible type vectors. We define $x$ as an *allocation* out of the set of all possible allocations $X$, where each entry $x_i$ is the allocation of agent $i$ (e.g., $x_i$ is the quantity of electricity that agent $i$ gets under allocation $x$).

**Definition 1.7** (Mechanism). *A mechanism is a function of the agents' reported types, $y(\hat{\theta})$, that maps the space of reports $\Theta$ to an outcome space $X$, i.e., $y(\hat{\theta}) : \Theta \to X$ (Nisan et al., 2007).*[5]

Intuitively, a mechanism takes as input the types of agents $\hat{\theta}$ and outputs an allocation $x$, where $\hat{\theta} = \{\hat{\theta}_0, \hat{\theta}_1, \dots, \hat{\theta}_{n-1}\}$ is the vector of the reported types of agents (i.e., agents communicate their types to the mechanism). Note that $\hat{\theta} = \theta$ hold only if agents report their true types.

The reported type of agent $i$ can also be a reported valuation function that maps the allocation $x_i$ to a real value, i.e., $v_i : (\hat{\theta}_i, x_i) \to \mathbb{R}$, and depends on the type $\hat{\theta}_i$. The reported valuation function states the desirability of the allocation $x_i$ to agent $i$ and it can also be written as $\hat{v}_i(x_i)$. Now consider that the allocation $x_i$ also includes the price $p_i \in \mathbb{R}$ that agent $i$ has to reimburse the mechanism. The utility of agent $i$ can be written as $u_i = v_i(\theta_i, x_i) - p_i$ and depends on the true type $\theta_i$. Since the mechanism has limited information with regards to the true types of agents, agent $i$ can misreport its type to the mechanism in order to maximize its utility.

A mechanism is called *incentive compatible* (IC) if agents achieve the best outcome for themselves (with regards to the allocation of the mechanism) if they report their true types, thus when $\hat{\theta} = \theta$. We proceed to illustrate the strongest incentive compatibility property that a mechanism can satisfy.

**Definition 1.8** (Dominant-Strategy Incentive Compatible). *A mechanism is called Dominant-Strategy Incentive Compatible (DSIC) if no agent can gain a better allocation outcome by misreporting its type to the mechanism regardless the reports of other agents. Such mechanisms are also called Truthful or Strategy proof (Nisan et al., 2007).*[6]

The aforementioned truthful implementation property (DSIC) is the most fundamental in mechanism design. However, as we see later in this section, there exist

[5] $y$ is also known as the *social choice* function, a theoretical framework for analyzing the combination of individual opinions and preferences to decide collective outcomes (Arrow, 1951).

[6] A weaker property for incentive compatibility is the *Bayes-Nash incentive compatibility* (BNIC): a mechanism satisfies BNIC if no agent can gain a better allocation outcome given that all other agents report truthfully (Nisan et al., 2007).

non-truthful implementations of mechanisms that are used in practice. Another fundamental property in mechanism design is *Individual Rationality*.

**Definition 1.9** (Individually Rational)**.** *A mechanism is called Individually Rational (IR) if a truthful agent gets non-negative utility in expectation for participating in the mechanism.*

Intuitively, a rational agent would choose to participate only if the mechanism satisfies the property of individual rationality.

In the following two sections we outline the fundamental, in mechanism design, *Groves* family of mechanisms, and the *Vickrey-Clarke-Groves* (VCG) mechanism which is an instance of Groves mechanisms; VCG satisfies both DSIC and IR properties and is used in Chapters 3 and 4 of this thesis.

**Groves mechanisms** Following (Weiss, 2013) we proceed to the following definition:

**Definition 1.10** (Groves mechanisms)**.** *Groves mechanisms are direct mechanisms for which*

$$x^{opt}(\hat{\theta}) = \operatorname*{argmax}_{x \in X} \sum_i \hat{v}_i(x_i), \tag{1.3}$$

$$p_i(\hat{\theta}) = h(\hat{v}_{-i}) - \sum_{j \neq i} \hat{v}_j(x^{opt}(\hat{\theta})), \tag{1.4}$$

where $x^{opt}(\hat{\theta})$ is the allocation that maximizes the summation of agents' reported valuations, and $h(\hat{v}_{-i})$ is an arbitrary function that depends only on the reported valuations of agents other than $i$. The price that agent $i$ pays to the mechanism, $p_i(\hat{\theta})$, is determined by the difference between the quantity $h(\hat{v}_{-i})$ and the sum of all other agents' reported valuations: $\sum_{j \neq i} \hat{v}_j(x(\hat{\theta}))$. Note that the price $p_i$ is independent of agent's $i$ own report, and therefore Groves mechanisms satisfy DSIC: the dominant strategy for agents is to report their true valuation function.[7]

**VCG mechanism** Every choice of the function $h(\hat{v}_{-i})$ yields a different mechanism in the Groves family. The *Clarke pivot rule*,

$$h(\hat{v}_{-i}) = \sum_{j \neq i} \hat{v}_j(x^{opt}(\hat{\theta}_{-i})), \tag{1.5}$$

yields the *Vickrey-Clarke-Groves* (VCG) mechanism (Clarke, 1971; Groves, 1973; Vickrey, 1961):

**Definition 1.11** (VCG mechanism)**.** *VCG is a direct mechanism for which*

$$x^{opt}(\hat{\theta}) = \operatorname*{argmax}_{x \in X} \sum_i \hat{v}_i(x_i), \tag{1.6}$$

$$p_i(\hat{\theta}) = \sum_{j \neq i} \hat{v}_j(x^{opt}(\hat{\theta}_{-i})) - \sum_{j \neq i} \hat{v}_j(x^{opt}(\hat{\theta})), \tag{1.7}$$

[7] For more details, see proof of Theorem 7.3 in (Weiss, 2013).

where the payment of agent $i$ to the mechanism, $p_i(\hat{\theta})$, depends on the Clarke pivot rule (see equation 1.5), which computes the valuation of agents other than $i$ under the optimal allocation without agent $i$ present, $x^{opt}(\hat{\theta}_{-i})$. Intuitively, the payment of agent $i$ to the mechanism is equal to the loss that is incurred to the rest of the society by its presence, which is formally called the *externality* of agent $i$. As an instance of Groves mechanisms' family, VCG is DSIC: agents maximize their utilities by reporting truthfully to the mechanism. Furthermore, if the following mild conditions apply: (i) *no negative externalities*, i.e., agents have non-negative utility for any outcome of the mechanism in which they are not included in the allocation, and (ii) the set of possible outcomes $X$ never increases by removing an agent (*choice-set monotonicity*), VCG further satisfies IR (Weiss, 2013).

### Auctions

*Auctions* are an important part of mechanism design since they define protocols for the allocation of resources among self-interested agents (McAfee and McMillan, 1987; Parsons et al., 2011). Agents participating in auctions can indicate their interest through *bids* for the available resources, bids are then used by the *auctioneer* to determine both the allocation and the payments. Auctions are commonly used in many recourse-allocation problems, e.g., bandwidth allocation (Zhang et al., 2013), public assets (Janssen and Janssen, 2004), and competitive electricity markets that are discussed throughout this thesis (Contreras et al., 2001).

In this section we provide a brief introduction in auction theory and some fundamental types of auctions, some of which are used in later chapters of this thesis. We proceed to provide a classification of auction types as these are described in (Parsons et al., 2011). Auctions can be:

- *Single or double-sided*. In single-sided auctions, one seller receives bids from $n$ buyers (demand auction), or one buyers receives bids (*asks* in this case) from $n$ sellers (supply auction). In double-sided auctions there are $n$ sellers and $m$ buyers both bidding for supply or demand.
- *Single or multi-dimensional*. In single-dimensional auctions the bids are only determined by the price, while in multi-dimensional auctions, bids can include several characteristics (e.g., price and quality).
- *Open or sealed-bid*. In open-bid auctions, bidders place their bids openly to other bidders and can participate further in the auction process. In sealed-bid auctions, bidders place sealed bids such that no other bidder knows their bid.
- *First price or $k$-th price*. In first price auctions, the winner (i.e., the one with the highest bid) pays its own bid. In $k$-th price auctions, the winner pays the price of the $k$-th highest bid.
- *Single-unit or multi-unit*. In single-unit auctions, bidders can place their bids for a single unit at a time, and for multiple units of the same type at the same time in multi-unit auctions.

- *Single-item or multi-item.* In single-item auctions, homogeneous items are auctioned off. In multi-item auctions, items can differ and bidders can have valuations for a *bundle* of items.

**Open-bid auctions** The most well-known family of auctions is the *English* auction. In an English auction, the auctioneer auctions off an item (this can also be a bundle of items that are auctioned as one) by announcing a starting price (also known as the *reserve* price) to the buyers. Then, the auctioneer accepts increasing bids from the buyers usually in pre-specified minimum increments. In any given moment of the auction the last bidder is considered the winning bidder of the auction. The auction ends when no bidder is willing to bid higher than the last placed bid. The winner pays a price that is equal to the highest bid. An English auction is therefore an open, single-sided, single-item, single-dimensional, first-price auction. Other open-bid auctions include the *Japanese* and the *Dutch* auctions (Parsons et al., 2011).

**Sealed-bid & Vickrey auctions** In open-bid auctions bidders have some knowledge of the competition since they can observe the behavior of other bidders. In this section we discuss sealed-bid auctions.

The most used type of sealed-bid auction is the first-price sealed-bid auction (McAfee and McMillan, 1987). The auctioneer collects sealed bids from the participating bidders, where each bid represent the price the bidder is willing to pay to acquire the item that is auctioned off. The winner is determined as the bidder with the highest valuation for the item, and the price that is paid is equal to that highest bid. An auction is called $k$-th price sealed-bid auction, when the price that the winner pays is determined by the $k$-th highest bid.

A second-price ($k = 2$) sealed-bid auction is also called Vickrey. It was first discussed and proposed by the 1996 Nobel Memorial Prize in Economic Sciences winner William Vickrey (Vickrey, 1961). A Vickrey auction is incentive compatible: a bidder maximizes its utility by bidding its true valuation. In a Vickrey auction, the winner cannot increase its utility by increasing its bid since the price is determined by the second highest bid, while bidding lower can result in losing the auction. Any other bidder can increase the probability of winning the auction by increasing its bid, in this case however, the bidder bids higher than its valuation and thus gets negative utility in case of winning.

**Combinatorial auctions** In combinatorial auctions, multiple items are auctioned off by the auctioneer at the same time while bidders are allowed to specify the price they are willing to pay for combinations (bundles) of items (de Vries and Vohra, 2003). It is easy to understand that combinatorial auctions are hard for both bidders and the auctioneer. For bidders because it is difficult to place a valuation over all possible bundles of items, where the number of bundles grows exponentially in the number of items. For the auctioneer because computing an optimal allocation (with regards to its potential revenue) has often intractable computational complexity. This is known as the *Winner Determination Problem* (WDP) and it lies in the complexity space of NP-hard problems (Lehmann et al., 2006). A VCG mechanism (see Definition 1.11) can be used in combinatorial auctions and it holds both DSIC and IR properties. However, unless specific instances of the WDP problem that can be solved in polynomial time are considered, the VCG auction faces the computational complexity barrier of the WDP.

**Sequential auctions** So far we have discussed auctions where items are all auctioned at the same time (*simultaneous auctions*). In *sequential* auctions, multiple items are auctioned off one after the other to the same group of bidders (Boutilier et al., 1999; Leme et al., 2012). For example, in a sequential first-price auction a first-price auction is held for each item one after the other. In practice, sequential auctions are more adopted than combinatorial auctions since they are easier to implement, e.g., internet advertising, wireless spectrum (Bae et al., 2008). In sequential auctions, strategic considerations may arise for the bidders given that: (i) a bidder can choose to wait and therefore choose the other bidders with which it competes (Parkes, 2007), (ii) *externalities* (i.e., propagated information) induced by previous auction outcomes to future auctions, e.g., a bidder can have different expected utilities for future auctions depending on who wins the current auction (Jehiel et al., 1999; Leme et al., 2012). In Chapters 3 and 4 we study practical settings of future smart grid systems in which sequential auctions can be used without strategic implications between consecutive rounds.

### 1.1.5 Markets

*Markets* are substantial components of human societies as they facilitate the exchange of goods, e.g,. food, electricity, water, information and services, between different parties. Markets comprise *buyers* and *sellers*; both participate in markets to obtain information and exchange goods under pre-specified set of rules that are determined both by the nature of the product to be exchanged and the market. In this thesis, we study markets with *commodities*.

**Definition 1.12** (Commodity)**.** *A commodity is an economic good or service of which each instance of a particular quantity holds the same value with no regards to who produced it (Geman, 2005; Smith, 1817).*

In electricity markets, a unit of electricity is an example of a commodity. Commodities are exchanged in commodity markets, which are responsible to transfer commodities from producers to consumers. However, commodities are usually not exchanged in a single market, instead they flow within a *market chain* and different types of markets on their way from the production site to the end consumer (Gereffi and Korzeniewicz, 1994). For most commodities, there exist multiple levels of markets; The most important types of commodity markets are the *retail* and *wholesale* markets. In retail markets, *retailers* (sellers), buy and stock bulk quantities of goods from producers or other intermediate sellers (*mediators*) to satisfy the demand of consumers. On the other hand, wholesale markets facilitate the distribution of goods from producers to retailers.

Markets can be modeled as multi-agent systems since they comprise interactions of multiple self-interested economic entities. For instance, buyers and sellers try to maximize their profits or minimize the price respectively. Most strategic interactions in markets regard the *price determination*. The price that commodities are exchanged for in markets is determined by the two most fundamental concepts in *economics*,

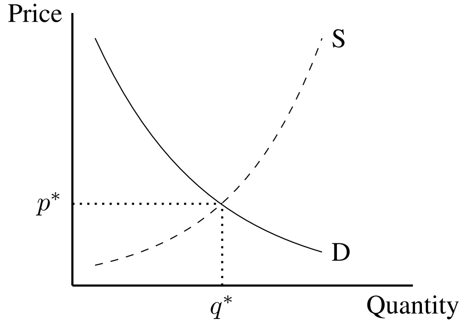


**Figure 1.2** Demand (D) and supply (S) curves that determine the quantity $q^*$ to be exchanged at price $p^*$.

namely, the *supply* and the *demand*. Supply refers to the quantity and the price that a market can offer, e.g., the aggregated inventory of all suppliers in a market. Demand is the quantity that is desired by buyers at certain price levels.

Figure 1.2 presents supply and demand curves: the continuous line presents the quantity of the demand at a respective price, the dashed line shows the offered supply at different price points. Demand and supply curves follow opposite gradients according to the laws of demand (i.e., less buyers are willing to buy items at a high price) and supply (i.e., supply increases if the price that buyers are willing to pay increases) (Henderson, 1922; Landsburg, 2013; Nicholson and Snyder, 2011). We have a *market equilibrium* where supply and demand curves meet, at which point the price and the quantity to be exchanged are determined by the market (Marshall, 2005), e.g., see price $p^*$ and quantity $q^*$ in Figure 1.2. In practice, supply and demand curves can be constructed out of sellers' and buyers' bids in double-sided auctions (that are commonly used in electricity wholesale markets) (Parsons et al., 2011).

In Chapter 5 of this thesis we consider retail markets where sellers compete by offering prices for identical items to buyers (e.g., electricity retail markets). In retail markets, each participating seller has a private cost for the items (e.g., production cost, procurement cost) and an infinite inventory. Sellers decide only on the price of each unit of the items, while buyers choose the seller with the lowest price from whom they will buy their demand. This is known as the *Bertrand* competition (Bertrand, 1988). A similar model in which sellers decide on the quantity of items to produce is the *Cournot* competition (Allaz and Vila, 1993). Both market models are commonly used in the economics literature to study interactions between sellers in retail markets (Singh and Vives, 1984).

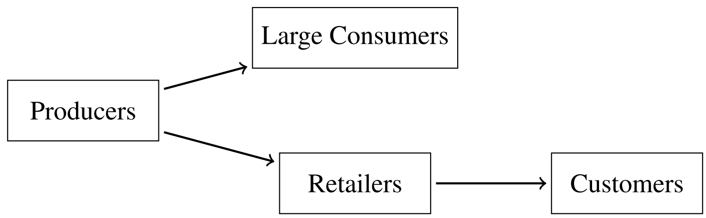


**Figure 1.3** Model of electricity commerce between producers, large consumers and retailers that pool customers into larger portfolios.

# 1.2 Transition in Electricity Systems

Electricity systems are in transition towards sustainable and distributed generation portfolios that primarily depend on natural sources such as the sun and wind. At the same time, in envisioned electricity systems the demand-side is expected to take a more active role: consumers will be able to control their own net demand with the use of energy storage technologies alongside the adoption of privately owned renewable electricity generation (i.e., prosumers). Both the increasing uncertainty on the supply-side and active demand-side management, which decreases the overall predictability of consumers' demand behavior, pose many challenges with regards to balancing requirements of future systems. The transition in electricity systems thus requires the redesign of current electricity systems and markets that electricity is traded in, and novel market mechanisms that need to facilitate the integration and utilization of intermittent electricity sources. Throughout this thesis we study fundamental problems that are motivated by challenges that come with the transition in electricity systems.

In the remainder of this section we provide an overview on basic concepts of current electricity markets that comprise the main application domain of this thesis. We further discuss the role and characteristics of renewable electricity generation in future electricity systems, we outline the solution concept of the smart grid, and last, we present an extensive literature review where we discuss recent methodologies that have been proposed in order to tackle challenges of envisioned electricity systems.

## 1.2.1 Electricity markets & balancing requirements

In current electricity grids, centrally located large fossil power plants provide electricity supply on-demand; this supply is pooled and traded in electricity wholesale markets (Ramchurn et al., 2011). Due to high capacity requirements of these markets, only large consumers (e.g., industrial sites) can participate there to satisfy their demand; large-capacity consumers can further rely on bilateral agreements with producers. On the other hand, smaller-scale consumers (customers), such as households or other service-sector demand entities, purchase their electricity demand in electricity retail markets. In retail markets, customers usually subscribe to long-term flat-rate tariff schemes with *retailers*.[8] Retailers act therefore as aggregators pooling customers into larger portfolios to satisfy their demand with electricity purchased in electricity wholesale markets. Figure 1.3 illustrates the relation between producers, large consumers, retailers and customers.

*Day-ahead* and *balancing* markets (also known as reserve capacity markets) are the main markets to facilitate commerce of electricity between retailers, large consumers and producers of electricity. Current electricity markets in Europe also include several adjustment markets (Conejo et al., 2010). For generality, however, electricity market models we consider throughout this thesis include the day-ahead and balancing markets. In addition, we focus on electricity retailers since the demand of their customers is generally not as predictable as the demand of large consumers that usually are big industrial sites with very specific demand loads. Electricity retail markets thus serve as fundamental market setting in later chapters of this thesis since on the contrary to heavily regulated wholesale markets, liberalization of retail markets can enable innovative economic mechanisms that facilitate the propagation of incentives to the demand-side (Stagnaro, 2017).

Retailers not only pool customers into larger portfolios to satisfy their demand, but also act as *balancing responsible parties* (BRPs): retailers procure electricity in day-ahead markets based on demand forecasts of their portfolios of customers (typically based on weather patterns, historical demand data, etc.), and balance supply with demand in balancing markets, i.e., the difference between the procured quantity in day-ahead market and the actual demand of customers. In practice, imbalances are moderated either by the use of high-cost storage or fast-ramping conventional supply, e.g., gas-turbines, the balancing power of which is traded in balancing markets. However, the increasing peak and general volatility of the demand result in increasing balancing costs for retailers (Kirschen and Strbac, 2018; Palensky and Dietrich, 2011), and are further associated with increasing $CO_2$ emissions (Hintermann, 2016).

In current electricity systems, the main strategy for retailers to control balancing costs is to reduce deviations from electricity demand predictions, and thus improve demand forecasting techniques. However, the presence of prosumers (i.e., consumers with privately owned generation that can also feed excess electricity generation in the grid) and more unpredictable sources of demand (e.g., EVs) in future electricity systems make demand forecasting challenging. In addition, the increasing dependency of electricity supply on natural sources may have a significant impact on wholesale market prices (of day-ahead and balancing markets), especially during shortages of renewable supply (Ketterer, 2014). All above reasons can potentially magnify balancing costs for retailers in future electricity systems.

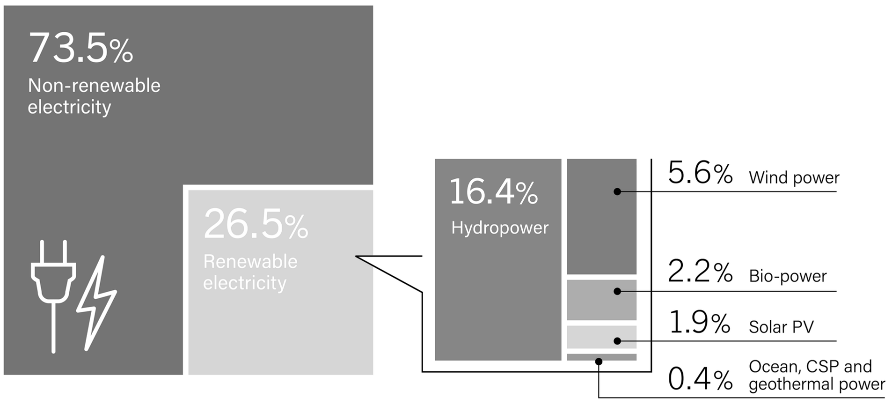


**Figure 1.4** Estimated share of renewable electricity production, end-2017. Image source: REN21's Renewables 2018 Global Status Report.

### 1.2.2 Renewable electricity sources

Currently, renewable electricity sources take up approximately a quarter of the total global electricity supply, where approximately 8% are from natural and volatile sources, such as the sun and wind. An estimated share of the current renewable electricity production globally is shown in Figure 1.4. According to the *REN21's Renewables 2018 Global Status Report (GSR)*, renewable electricity generation accounted for 70% of net additions to global electricity capacity in 2017, which is the largest increase in renewable electricity capacity in recent years.[9] In addition, 2018's *Bloomberg New Energy Outlook* expects that renewable electricity generation specifically from natural sources (solar and wind) will take up 50% of the total electricity generation by 2050.[10] Some of the most important characteristics of renewable electricity from natural sources are the following:

- Renewable electricity generation is not dispatchable (it cannot be used on-demand) or can be deferred (it cannot be shifted in time). Also, its output cannot be regulated to meet the demand.
- Renewable supply is not fully predictable since it is subject to stochastic conditions (weather conditions).

[8] In Europe, most retail contracts have a set price per consumption unit and typically have one year duration.

[9] http://www.ren21.net/gsr-2018/

[10] https://about.bnef.com/new-energy-outlook/

- Renewable electricity supply cannot be centralized, the placement of renewable generators (e.g., wind turbines, solar photovoltaic systems) is subject to spatial constraints.

The easier availability of electricity from natural renewable sources further enables small-scale consumers to produce their own electricity (prosumers), which consequently implies higher risks of shortages or overproduction since generation is volatile and locally highly-correlated. At the same time, this adds extra forecasting complexity with regards to the net electricity demand of prosumers that may actively control their demand loads based on their privately owned renewable generation (Pepermans et al., 2005). Therefore, in future electricity systems, demand needs to follow the available supply.

Overall, the introduction of more sustainable generation portfolios involves the redevelopment of conventional electricity grids, not only with regards to the infrastructure and control, but also the markets that facilitate electricity trading (Phuangpornpitak and Tia, 2013). The characteristics of renewable sources of electricity require more *flexibility*[11] on the demand-side, and at the same time decentralized approaches, since renewable generation is distributed and therefore cannot be centrally monitored and managed (Alanne and Saari, 2006; Goldthau, 2014).

### 1.2.3 Smart grid

The *smart grid* is an electricity grid innovation away from the traditional paradigm of passive distribution and consumption (Amin, 2015; Ellabban et al., 2014; Fang et al., 2012; Farhangi, 2010):

**Definition 1.13** (Smart grid)**.** *Smart grid is an interconnected electricity network that uses digital communication, measuring, and control to facilitate efficient electricity usage and renewable electricity generation.*

On the contrary to conventional electricity networks, the smart grid enables active participation of generation and consumption entities, e.g., households, distributed small-scale generation, electric vehicles (EVs), since these are interconnected and can share local information to system operators. Active participation and control of such consumption entities can therefore facilitate the utilization of supply from renewable electricity sources. Some of the most important ambitions of the future smart grid, as these are outlined in (Farhangi, 2010), are:

- *Sensors and two-way communication.* Sensors (e.g., smart-meters) as well as intelligent home devices enable remote sensing, control, and two-way communication between aggregators of electricity and customers in the smart grid (Depuru et al., 2011).
- *Distributed generation and control.* In contrast to conventional power plants, smart grid assumes decentralized generation that is located close to the demand load it serves; this enables more local control but requires efficient management and coordination.

[11] Demand or supply flexibility is a term to denote the ability of electrical loads to be deferred or curtailed.

- *Demand-side management and customer choices.* Two-way communication and active participation of demand and consumption entities (e.g., households, EVs) enables direct or indirect control of local loads. In the smart grid, electricity customers are also able to choose different plans of electricity billing based on their consumption patterns.
- *Self-monitoring and self-healing.* Real-time communication within the smart grid enables monitoring electrical characteristics of the network, and thus enhances the ability to discover small changes that can trigger bigger systems disturbances and isolate them, and furthermore optimize the overall system operation.
- *Adaptivity and islanding.* To help mitigate wider system disturbances the smart grid has the ability to disconnect parts of it from the centralized grid (i.e., islanding mode) and operate autonomously depending solely on local generation, management, and control (self-sufficient).

Overall, the smart grid innovation changes the way electricity grid operates; it assumes the technological framework to utilize the increasing adoption of renewable electricity generation, and thus aims to make electricity cheaper, more accessible (addressing *electricity poverty*), and further facilitate self-sufficient and autonomous electricity grids that depend their electricity supply only on local generation and control. Two-way communication and sensors in the smart grid also enable active participation of end-users opening up the opportunity for novel economic mechanisms that go beyond current flat-tariffs towards meeting balancing requirements of future electricity networks.

### 1.2.4 Smart grid challenges & solutions

Towards the implementation of the smart grid there exist many challenges with regards to regulations, privacy, and security issues (McDaniel and McLaughlin, 2009; Yan et al., 2013). However, some of the most important challenges of future smart grid systems are associated with:

- The seamless integration of renewable electricity generation, which comes at the expense of higher uncertainty, alongside the increasing peak and volatility of active demand (e.g., prosumers, EVs).
- The design of mechanisms to incentivize demand-side management with regards to balancing requirements and the efficient operation of smart grid systems.
- The analysis of smart grid systems' complex dynamics in which strategic decision-making of multiple stakeholders may decrease the collective efficiency.

All the above challenges comprise situations that require multiple stakeholders making decisions under the presence of different sources of uncertainty. To this end, multi-agent systems and its sub-fields (see Section 1.1) provide the main theoretical tools to study the design and implementation of future electricity systems (Jun et al., 2011; Pipattanasomporn et al., 2009). Within multi-agent systems, agent-based modeling enables the analysis of complex dynamics in electricity markets under the presence of uncertainty (Karnouskos and De Holanda, 2009; Weidlich and Veit, 2008), market implementations that facilitate active participation of end-users (Ilic et al., 2012), agent-based markets for supply and demand matching (Kok et al., 2005), and last, the design of simulation platforms to evaluate new regulations and different market behaviors in electricity markets (Praca et al., 2003).

In addition, mechanism design (see Section 1.1.4) and the application of theoretical solution concepts of game theory (see Section 1.1.3) provide the means for the implementation of distributed economic protocols (Alibhai et al., 2004) and the analysis of strategic interactions between self-interested agents in future smart grid systems (Fadlullah et al., 2011; Pettersen et al., 2005; Saad et al., 2012). For instance, non-cooperative games are commonly used in the literature to model competitive instances between individual agents and analyze the resulting efficiency (Perrault and Boutilier, 2014), or to provide solution concepts (e.g., Nash equilibria) within systems where privately owned assets, such as storage, are individually controlled by strategic agents (Vytelingum et al., 2010).

In the following sections we present an extended review of research literature that is motivated by challenges of envisioned electricity systems and aligned to the topics of this thesis; this is organized as follows:

**Demand-side management** introduces the concept of demand-side management and presents relevant research studies.

**Tackling demand uncertainty** presents recent methodologies to incentivize uncertainty reduction in the demand-side.

**Utilization of renewable sources** outlines related work that studies the integration of renewable generation into future electricity systems.

**Electric vehicles as flexible assets** discusses the role of electric vehicles in future electricity systems.

### Demand-side management

In current electricity grids most of the demand is passive: customers are usually not concerned about their electricity usage patterns and the underlying complexities of electricity systems, instead they want on-demand electricity availability and understandable tariffs. Current retail tariffs for electricity provide flat or day-night rates that do not depend (by design) on the demand patterns of customers, except for total demand volume. Such tariffs, however, do not directly represent the actual costs for the supply (e.g., last-minute balancing costs, network loses) and thus customers neither sense peaks in the demand and the resulting prices in balancing markets nor have economic incentives to change their demand behavior. Advancements in communication, sensing (smart meters), and affordable autonomous and intelligent control, within the smart grid innovation, enable novel pricing schemes that can utilize the potential flexibility of electricity users (i.e., making demand active) assisting in balancing requirements and the efficient functioning of envisioned electricity systems.

To this end, *demand-side management*, which is one of the most prominent research topics in the smart grid literature, assumes flexible customers or assets that can alter their demand or generation behavior given economic incentives, and therefore contribute to the balancing problem extensively described earlier in this thesis (see Section 1.2). *Demand response* is yet another term used to denote the ability of customers to reduce their demand. Both demand response and demand-side management describe methods to influence the electricity usage patterns of customers. Han and Piette (2008) distinguish two types of demand response, namely *time-based* and *incentive-based*.

Time-based demand response is a means to encourage favorable changes in demand patterns of customers by dynamically changing the price (dynamic pricing) of electricity (Borenstein et al., 2002; Roozbehani et al., 2010). For instance, retailers can forward wholesale market prices to customers, where peaks in demand directly influence the balancing market prices and consequently customers. Time-of-use, critical peak pricing, and real-time pricing are some types of dynamic pricing schemes that have been used in practice to stimulate favorable customer behavior (Owen and Ward, 2010), e.g., in Ontario and California, and have also been studied in more fundamental settings (Chakraborty and Khargonekar, 2014; Meng and Zeng, 2013; Oruc et al., 2012). For instance, Agarwal and Cui (2012) study equilibrium strategies for autonomous load balancing under real-time pricing schemes as non-cooperative games between consumers. Similarly, Ramchurn et al. (2011) propose agent-based control for decentralized demand-side management in the smart grid under time-of-use and real-time pricing schemes. Real-time pricing can also be based solely on the local grid frequency (i.e., frequency measurements can be used to sense imbalances between supply and demand), which provides all the necessary information to match supply and demand (Schäfer et al., 2015).

However, dynamic pricing approaches are no substitute for planning and can only resolve residual imbalances. Furthermore, they may result in disruptive and chaotic market behavior, power outages, uncertain availability of electricity, and unexpected/unfair high prices for electricity customers (Herter and Wayland, 2010; Roozbehani et al., 2012). All these affect the stability and predictability of electricity systems, and thus planning and ahead pricing are required (Braithwait et al., 2007). In addition, defining mathematical relations between the exact costs of imbalances and the maximum prices that customers would find acceptable is challenging (Roozbehani et al., 2010).

On the contrary to time-based demand response, in Chapters 2 and 3 we propose incentive-based demand response models in which customers do not react to incoming price signals; instead, they decide whether to provide their flexibility based on economic terms that are agreed prior to the time that flexibility is needed (Han and Piette, 2008). Incentive-based demand response can be achieved in various ways, such as direct control of the demand loads by the grid operator in exchange for predefined payments, or voluntary participation in demand response programs based on discounts or financial rewards to curtail (reduce) excess demand, but also penalties for not contributing.

In practice, retailers currently use long-term bilateral contracts to incentivize large-capacity consumers to reduce their demand if necessary for the system (Kim and Shcherbakova, 2011). However, He et al. (2013) show that shorter-term agreements can also enable smaller-scale user to participate in demand response programs, and further discuss the issue from a policy point of view. Haring and Andersson (2014) study the design of such contracts that provide economic incentives to consumers in order for system operators to directly control their demand loads. Similarly, the work by Ma et al. (2016, 2017) is one of the first to study the problem of incentivizing small-scale demand response agents from a mechanism design perspective: it contributes truthful mechanisms that select a set of flexible users and determine the rewards and penalties such that the selected customers reduce a fixed demand reduction target with high reliability. Similarly to (Ma et al., 2016, 2017), Meir et al. (2017) propose the design of demand response contracts tailored towards more practical applications. Last, Muthirayan et al. (2017) studies the problem of a demand response aggregator, which calls on a subset of its recruited demand response users to reduce their electricity consumption during a demand response event, and computes the rewards (for reducing their demand) based on users' reported baseline consumption.

Last, different demand response methods include but are not limited to: autonomous and distributed mechanisms that optimally schedule the consumption patterns of consumers based on incentive-based pricing schemes (Chen et al., 2010; Mohsenian-Rad et al., 2010), pricing schemes that are based on consumers' reported utility functions for electricity consumption and incentivize consumers to reduce or shift their loads to off-peak hours (Samadi et al., 2012), or the formation of consumer groups that can collectively shift their consumption in time (Akasiadis and Chalkiadakis, 2017).

### Tackling demand uncertainty

Demand response can not only be used to steer the demand behavior of customers, but also to reduce the inherent uncertainty of the demand, since demand is locally highly-correlated causing imbalances that result in excessive balancing costs for retailers. For instance, Rose et al. (2012) propose an incentive compatible mechanism using scoring rules to incentivize customers report their true anticipated demand and consequently improve retailers' demand forecast. Similarly to scoring rules, Vinyals et al. (2014) introduce a *prediction-of-use* (POU) tariff that requires consumers to predict their future baseline consumption and then charge them based on their actual consumption. Based on these earlier results, Perrault and Boutilier (2017) extend POU tariffs for customers that may report multiple consumption profiles, while in another follow-up work, Robu et al. (2018) study group formation strategies for consumers under POU tariffs.

Similarly to related work that is discussed in this section, Chapter 2 proposes a novel electricity tariff that uses the demand prediction of a customer to charge a price that is determined by the actual demand outcome.

### Utilization of renewable sources

Apart from uncertainty in the demand, another line of research studies solutions for the integration and utilization of intermittent electricity sources by explicitly modeling their uncertainty. For instance, Maity and Rao (2010) propose pricing mechanisms for microgrids that solely depend on electricity supply from renewable generation. In addition, Dash et al. (2007) study both centralized (based on the VCG mechanism) and decentralized (double auction) market mechanisms to utilize the supply of distributed and limited-capacity renewable generators by incentivizing them reporting their true predicted output. Scoring rules can also be used to incentivize the formation of virtual power plants and provide an alternative to feed-in tariffs for small-scale renewable producers (Chalkiadakis et al., 2011; Robu et al., 2012). Other related works have studied selection strategies for optimal renewable generation portfolios (Gärttner et al., 2018) and online market mechanisms for matching uncertain supply with flexible demand-loads (Ströhle and Flath, 2016).

An interesting aspect, especially when considering systems with no balancing capabilities (e.g., islanding grids where excess demand cannot be satisfied), is electricity trading under uncertain supply of renewable generators, such as wind turbines. The first work to propose an alternative to traditional trading of electricity that assumes guaranteed delivery was by Bitar et al. (2012), where the authors proposed to package random wind power into electricity with different levels of reliability and sell them at different prices. *Service-level agreements* (SLAs) have also been used in related literature as a contracting framework to facilitate electricity trading under intermittent electricity supply; for intance, Hussain et al. (2018) study the viability of using service-level agreements between consumers and prosumers with wind electricity generation in the smart grid. In line with previous works, in Chapter 4 we study the adoption of SLAs as a direct extension of current electricity tariffs to facilitate trading under uncertain supply, we further propose mechanisms to allocate the available electricity supply to consumers with heterogeneous demand requirements.

The utilization of intermittent sources can also be treated as a scheduling problem, in which deferrable demand loads can be shifted in time to match the output of renewable generators. He et al. (2011) investigates both day-ahead and real-time scheduling of deferrable loads under volatile wind generation as a Markov decision process. In addition, Neely et al. (2010) use the Lyapunov optimization method to schedule demand loads of consumers that can tolerate some delay in their service. Last, other works use scenarios over the output of wind turbines to schedule demand loads based on the likelihood of each scenario (Ströhle et al., 2014; Walraven and Spaan, 2015a,b).

### Electric vehicles as flexible assets

Another emerging line of the research within the smart grid literature is dedicated to the role of electric vehicles (EVs) in current and future electricity networks. The mass adoption of electric vehicles, despite its environmental benefits (Stephan and Sullivan, 2008), poses many challenges: EVs are equipped with large capacity batteries that require high load charging, and thus resulting increasing peak-demand given that charging takes place at similar times in a day (i.e., due to correlated human behavior).

The adoption of electric vehicles poses not only challenges to future electricity systems but also provides additional flexibility that can be used to alleviate the increasing uncertainty of renewable electricity sources or the volatility of the demand. When parked and plugged into charging infrastructure, EVs can provide flexibility to smart grid systems, in a concept that is known in the smart grid literature as *vehicle-to-grid* (feeding electricity into the grid) or *grid-to-vehicle* (storing excess generation) (Kempton and Tomić, 2005; Lopes et al., 2011).

Demand response methods can also rely on the use of EVs as flexible demand response agents to reduce imbalances between supply and demand (Vandael et al., 2013). For instance, Hayakawa et al. (2015); Vandael et al. (2011) propose mechanisms and scheduling strategies for EV charging at times when there is enough renewable electricity generation that cannot be deferred. Similarly, Gerding et al. (2016) propose online mechanisms for optimizing the charging schedules and the revenue of an electric vehicle parking lot that can exchange any electricity excess or shortage with the main grid. Last, Saad et al. (2011) analyze the Nash equilibrium of a non-cooperative game between groups of electric vehicles that are willing to sell (feed-in) a maximum amount of electricity to the grid with regards to their associated costs. In Chapter 3, we study demand response mechanisms that instantiate EVs as flexibile assets, which can be incentivized to reduce imbalances between supply and demand.

## 1.3 Research Questions

So far in this chapter, we have outlined essential theoretical concepts of multi-agent systems (see Section 1.1) that are used throughout this thesis, and we have further introduced the main research domain and one of the main motivations of this thesis, which is related to the transition in electricity systems and the solution concept of the smart grid (see Section 1.2). In this section, we delve deeper into specific issues that arise in settings of future electricity systems, and we outline the main research questions of this thesis that are arranged in the following two sections.

### 1.3.1 Demand & supply uncertainty

The design and implementation of future electricity systems, such as the smart grid, are shaped by the increasing introduction of renewable electricity generation. In such settings, the induced uncertainty of the supply can increase costs for retailers since they hold the responsibility to balance demand with the available or procured supply in balancing markets (Meibom et al., 2009). In addition, current flat-rate tariffs for electricity do not represent high balancing prices during peaks in the demand, and thus customers do not have economic incentives to reduce their demand during these peaks.

**Research Question 1.** *Can we design electricity tariffs that explicitly incorporate the balancing responsibility of the retailer and incentivize heterogeneous customers to reduce the uncertainty of their demand?*

This research question is addressed in Chapter 2.

In line with the above research question, previous works have proposed novel tariff schemes that incentivize customers to provide accurate predictions of their future demand improving retailers' demand forecasts, and thus reducing excessive costs related to balancing supply and demand (Rose et al., 2012; Vinyals et al., 2014). Our first research question further considers the actual balancing costs resulting from the demand uncertainty of electricity customers and seeks to incorporate those explicitly in the design of electricity tariffs. Furthermore, we consider heterogeneous customers with regards to their ability to reduce the uncertainty of their demand, and thus such electricity tariffs need to elicit uncertainty reduction regardless of the specific ability of customers.

Similarly to related works by Rose et al. (2012) and Vinyals et al. (2014), Research Question 1 considers incentives towards the demand-side in order to reduce the uncertainty related to the actual demand of customers. Such tariff schemes, however, cannot guarantee to diminish last-minute imbalances between supply and demand, since demand is privately controlled and supply may not depend solely on planned electricity generation and thus be uncertain. In current systems, last-minute imbalances are resolved either in balancing markets with high prices or by large-capacity consumers (e.g., industrial sites) that can reduce their demand when needed (Kim and Shcherbakova, 2011). The presence of smaller-scale flexible assets in the form of EVs or smart home devices in future smart grid systems introduces additional flexibility in the demand-side, and may comprise the only source of demand flexibility in small-scale electricity grids. For instance, EVs can feed electricity into the grid out of their charged batteries when parked at times where there is supply shortage in the grid (Kempton and Tomić, 2005). To this end, the following research question considers demand response programs for small-scale flexible assets, e.g., EVs, the availability of which is not certain.

**Research Question 2.** *Based on the demand forecast of a retailer, can we design economic mechanisms that incentivize small-scale and unreliable demand response agents to prepare and reduce imbalances between supply and demand if necessary?*

This research question is addressed in Chapter 3.

Similarly to the above research question, previous work studies mechanisms to incentivize demand response of uncertain availability in the smart grid (Ma et al., 2016, 2017). However, as outlined in Chapter 3, our research question differentiates substantially by also considering the demand forecast, which may influence the selection of demand response agents and the requirements of agents to alter their demand. In addition, we further consider the resulting balancing cost of the retailer under such demand response mechanisms.

The above two research questions consider either the design of novel electricity tariff schemes that aim to reduce the uncertainty in the demand-side, or the design of mechanisms to incentivize small-scale flexible assets to reduce last-minute imbalances between supply and demand.

We proceed to a different scenario, in which we consider the uncertainty on the supply-side. More specifically, we consider electricity systems that depend their supply on local renewable generation (e.g., islanding microgrids). In such settings, on the contrary to current systems where demand is always satisfied, supply is uncertain, and thus the delivery of electricity cannot be guaranteed. In the research question that follows we consider that these uncertain quantities of electricity can be traded ahead of time while the corresponding demand can only be satisfied if supply is sufficient.

Intuitively, consumers can agree to satisfy their future demand with some probability that depends on the forecasted renewable electricity supply.

**Research Question 3a.** *How can we model the utility function of consumers to include the uncertainty of electricity delivery and consequently the probability of having their demand satisfied?*

**Research Question 3b.** *Can service-level agreements provide the contracting framework for the allocation and trading of uncertain quantities of electricity between renewable generators and consumers?*

These research questions are addressed in Chapter 4.

A relevant scenario has been addressed by previous work that considers a market setting where electricity can be traded at different levels of reliability and prices (Bitar et al., 2012). However, in the above research questions we further consider the viability of using service-level agreements in place of current electricity retail tariffs. We additionally consider the allocation of service-level agreements to consumers as a mechanism design problem, and last, we aim to characterize consumers (through utility functions) with regards to the uncertainty of satisfying their demand in the above setting.

### 1.3.2 Agent-based decision making

Research questions outlined in the previous section are motivated by challenges that arise in future electricity systems, and consider the design of mechanisms to cope with the increasing uncertainty in the demand or supply. Research questions presented in this section focus on the decision-making of agents participating in future electricity markets.

In envisioned smart grid systems, two-way communication between retailers (sellers) and customers (buyers) enables higher-resolution pricing schemes in which customers may need to choose over different tariffs in very short time intervals. In such markets, however, software agents will need to participate in place of human customers and assist or represent them in making multiple decisions within the course of a day. The research questions outlined in the remainder of this section consider the implications of replacing human customers with software agents in future (electricity) markets.

**Research Question 4a.** *What are the effects of representing buyers with autonomous (economic) decision-making agents in retail markets on the competitive dynamics and the resulting prices?*

This research question is addressed in Chapter 5.

Related to the above research question, previous works have investigated the effects of buyers' decision-making on markets focusing more on the sellers' perspective (Ait Omar et al., 2017; Basov and Danilkina, 2015; Zhang et al., 2009). In contrast, we focus on buyers' market behavior and the effect it has on retail markets with regards to the competition between sellers and on the resulting prices buyers face.

We proceed to outline our last research question, which is related to Research Question 1 that examines the viability of designing tariffs to incentivize uncertainty reduction in the demand of customers. For the design of an electricity tariff or any other general pricing mechanism, it is crucial to consider the economic decision-making of customers, since this can influence the efficiency of a pricing mechanism.

**Research Question 4b.** *What are the implications of considering the economic decision-making of customers when designing tariffs?*

This research question is addressed in Chapter 2.

### 1.3.3 Summary of research questions

Within Section 1.3 we have outlined the main research questions of this thesis. These are motivated by future and current challenges of electricity systems and related to: tariffs to tackle the increasing peak and general volatility of demand, demand response mechanisms to incentivize last-minute balancing through flexible assets, SLAs for electricity trading under uncertain supply from renewable electricity generation, and last, the effects of different economic decision-making of customers on the design of retail tariffs and the behavior of retail markets.

## 1.4 Thesis Outline

In this section we provide an outline of the remaining five chapters of this thesis (Chapters $2-6$):

**Chapter 2** studies the design of an innovative electricity tariff to incentivize uncertainty reduction in the demand-side. More specifically, in this chapter we analyze a multi-agent system in which a buyer agent (customer) wants to purchase a continuously divisible good from a seller agent (retailer). We further consider that the customer has a direct or a representative influence on the balancing costs of the retailer. The main contribution of this chapter is the *risk-sharing tariff*, which is a two-step parameterized payment scheme that provides the customer the choice to assume and alleviate a fraction of the balancing risk from the retailer. It consists of a prepayment based on the expected demand of the customer, and a supplementary payment for any observed deviation from the anticipated demand. Last, this chapter investigates the influence of the customer's stochastic decision-making on the design of the risk-sharing tariff.

The contents of this chapter are based on (Methenitis et al., 2016).

**Chapter 3** studies mechanisms to incentivize small-scale flexible users to resolve last-minute imbalances between supply and demand and thus reduce balancing costs for retailers. In this chapter we build upon previous work assuming demand response agents that can respond (alter their demand) with some probability if they prepare prior to the realization of the demand (Ma et al., 2016). We additionally consider the balancing responsibility of a retailer under a given demand forecast and imbalance price: the retailer is responsible to purchase additional reserve capacity at a high imbalance price to cover any excess in the demand. In this chapter we propose and evaluate two market mechanisms to incentivize demand respond agents to prepare and respond if requested by the retailer. Both of the proposed mechanisms guarantee non-negative utility for both demand response agents and the retailer, and can further be used for simultaneous downward and upward flexibility, i.e., when both demand reduction and demand increase may be necessary to reduce balancing costs.

The contents of this chapter are based on (Methenitis et al., 2019b).

**Chapter 4** studies electricity trading in settings of envisioned smart grid systems, where electricity supply may solely depend on renewable electricity sources. In these settings, uncertain electricity quantities may be available, the delivery of which cannot be guaranteed. However, if not traded, the electricity might need to be curtailed, foregoing potential benefits for both supply and demand-sides. To this end, this chapter proposes the adoption of service-level agreements (SLAs) that comprise the following features: quantity, reliability, and price. The proposed SLAs can be used instead of current flat-rate electricity tariffs in settings where electricity delivery cannot be guaranteed. In this chapter, first, we characterize buyers' varying degrees of criticality with regards to the probability of satisfying their demand, and next, we study the design of mechanisms in order to specify and allocate these contracts (SLAs) to buyers.

The contents of this chapter are based on (Methenitis et al., 2017, 2018).

**Chapter 5** investigates the effects of different economic decision-making of buyers on the competitive dynamics and the resulting prices in retail markets. More specifically, this chapter considers retail markets with identical items (e.g., electricity retail markets) using the Bertrand competition model. The collective decision-making of the participating buyers is modeled with a parameterized function such that it can approximate different levels of rationality (i.e., from random to perfectly rational buyers). At the same time, sellers have heterogeneous beliefs with regards to the competition they are facing (prices that are offered by other sellers). In this chapter we first derive interesting analytical results for the optimal pricing strategy of a reasoning seller with regards to the competition and the degree of buyers' rationality. Lastly, using concepts from evolutionary game theory, we show some counter-intuitive results with regards to the resulting market dynamics under perfect buyers' rationality that provide insights for the design of future retail markets.

The contents of this chapter are based on (Methenitis et al., 2019a).

**Chapter 6** serves as an epilogue of this thesis. In this chapter we connect the contributions of each technical chapter to the research questions posed in Section 1.3 and we discuss potential directions of our research.

## 1.5 List of Publications

In this section we present an overview of the research publications comprising this thesis. For detailed references, we refer the reader to the Bibliography chapter.

i. Georgios Methenitis, Michael Kaisers, and Han La Poutré. **Incentivizing intelligent customer behavior in smart-grids: a risk-sharing tariff & optimal strategies**. 2016. International Joint Conference on Artificial Intelligence (IJCAI).[12]

ii. Georgios Methenitis, Michael Kaisers, and Han La Poutré. **SLA-Mechan-isms for Electricity Trading under Volatile Supply and Varying Criticality of Demand (Extended Abstract)**. 2017. International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS).[13]

iii. Georgios Methenitis, Michael Kaisers, and Han La Poutré. 2018. **Renewable Electricity Trading through SLAs**. Springer - Energy Informatics.

iv. Georgios Methenitis, Michael Kaisers, and Han La Poutré. 2019a. **Degrees of Rationality in Agent-Based Retail Markets**. Springer - Computational Economics.

v. Georgios Methenitis, Michael Kaisers, and Han La Poutré. 2019b. **Fore-cast–Based Mechanisms for Demand Response**. International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS).

[12] An extended abstract of this work was published at Belgium-Netherlands Artificial Intelligence Conference (BNAIC) 2016.

[13] This work was presented at the International Workshop on Agent-Mediated Electronic Commerce and Trading Agents Design and Analysis (AMEC/TADA) 2017.

# 2 Reducing Demand Uncertainty through Risk-Sharing

**Preface** In this chapter we propose a novel electricity tariff to incentivize uncertainty reduction on the demand-side, the *risk-sharing tariff*: a two-step parameterized payment scheme, in which the customer can choose a portion of the balancing risk to assume from the retailer. It consists of a prepayment based on the expected demand of a customer, and a supplementary payment for any observed deviation from the anticipated demand. We present a game-theoretical analysis of the risk-sharing tariff, which captures the strategic conflict of interest between the retailer and the customer, and we present optimal strategies for both players. Last, we show analytically that the proposed tariff provides a customer of varying ability to reduce its demand uncertainty with incentives to assume and alleviate a fraction of the balancing risk, contributing in this way to the uncertainty reduction in the envisioned smart grid.

ℹ *This chapter presents work that was published in the proceedings of the International Joint Conference of Artificial Intelligence (IJCAI) 2016 in New York, US (Methenitis et al., 2016). This work was also published as an extended abstract in the proceedings of the Belenux conference on Artificial Intelligence (BNAIC) 2016 in Amsterdam, Netherlands.*

# 2.1 Introduction

In the previous chapter we provided an extensive discussion with regards to the challenges of future electricity systems, which include more sustainable generation portfolios (see Section 1.2). In particular, we argued that many potential and existing challenges that electricity grids are facing are mainly connected to the continuous need for balancing supply and demand, as well as the increasing demand-peaks. Maintaining balance becomes even more challenging in face of generation from natural resources, such as the sun and wind, that are subject to stochastic availability, and the increasing uncertainty of privately controlled demand in the smart grid.

Balancing supply and demand is a major factor with regards to the efficiency of electricity systems and markets. Stochastic fluctuations and deviations from predictions on both demand and supply sides should be matched with reserve electricity generation coming from fast-ramping conventional generators, e.g., gas turbines, which is traded in balancing markets. Imbalances between supply and demand have therefore a direct impact on the balancing prices and the increasing $CO_2$ emissions.

In current electricity systems, retailers pool customers into larger portfolios and are responsible for balancing the demand of their customers with supply in balancing markets. Consequently, imbalances between supply and demand can result in high economic risks for retailers due to the volatility of reserve prices for balancing. One strategy to control costs is to avoid the need to purchase balancing power by actively reducing deviations from estimated demand, and thus reduce the uncertainty of the demand. However, existing flat-rate electricity tariffs by retailers, especially in Europe, cannot be used to provide incentives for uncertainty reduction on the demand-side. This precludes flexible customers from assuming some of the high costs related to the participation in the balancing markets (Oualmakran et al., 2017), and consequently customers may use their flexibility, e.g., from storage, primarily to their own interest rather than the interest of the retailers' balancing requirements (Vytelingum et al., 2010).

One way to encourage favorable changes in demand patterns by the customers is dynamic pricing (Borenstein et al., 2002; Roozbehani et al., 2010). Time-of-use, critical peak, and real-time pricing are some of the pricing schemes used to stimulate favorable customer behavior in practice (Owen and Ward, 2010). However, as it was also outlined in Section 1.2, dynamic pricing approaches may introduce disruptive and unfavorable market behavior (Herter and Wayland, 2010; Roozbehani et al., 2012), and thus planning and ahead pricing are required (Braithwait et al., 2007).

In this chapter, we present the *risk-sharing tariff*, a novel approach to incentivize uncertainty reduction on the demand-side by giving customer the choice of assuming balancing risk by the retailer. We consider a multi-agent system in which a buyer agent wants to purchase an uncertain quantity of a continuously divisible good from a seller agent. We refer to the buyer and the seller as the *retailer* and the *customer* respectively. Based on the forecasted demand of the customer, the retailer procures a fixed quantity of electricity in the day-ahead market at a low rate. After the actual demand of the customer is observed, the retailer holds the responsibility to balance any deviation between its procured quantity and the actual demand of the customer in the balancing market at a higher rate than the day-ahead price. In line with related literature, which studies optimal procurement strategies for the retailer under the presence of either uncertain demand (Nair et al., 2014) or uncertain prices (Hoogland et al., 2015), we consider a two-step market setting, where the prices are fixed but the demand is uncertain. We further assume that the customer has a direct or a representative (see cases i and ii below) influence on the balancing requirements of the retailer, this is the case in:

i. *Service-level agreements* (SLAs) formally define an agreement between a service provider and the service user, specifying the service and its characteristics, e.g., quality, risk. In the context of electricity markets we interpret SLAs as a direct extension of conventional electricity tariffs: While current electricity tariffs ensure delivery (100% quality) and a fixed kWh price (0% risk), SLAs may provide customer further choices, such as assuming parts of the balancing risk, as discussed in this chapter. Such SLAs may further enable decentralized trading of electricity between small-scale producers and individual customers.

ii. *Highly correlated demand* can be the result of similar demand behavior of customers, influenced for instance by weather conditions in specific locations. The higher the correlation, the closer the deviations of one customer to the deviations of other customers, i.e., changes in the demand behavior of one customer predicts the same change in the behavior of other customers. Therefore, the portfolio distribution may closely resemble the demand distribution of an individual customer for any specific location.

iii. *Local balancing*, current market-based balancing strategies do not consider spatial characteristics of customers. However, it is in the retailer's own interest to balance customers locally. This can lower the costs corresponding to energy losses, transportation costs, network load, and congestion.

In this setting, we propose a two-step parameterized payment: the customer precommits and prepays for its expected demand and later pays for any deviation between the observed and the anticipated load. In addition, the customer has the choice to select the portion of risk that it is willing to assume from the retailer. For instance, in case the customer chooses to assume all the risk from the retailer, it is exposed to the full balancing cost that is determined by the difference between its predicted and actual demand.

The main contributions of this chapter can be summarized as follows:

- We formalize the interaction between the retailer and the customer as a two-player game.
- We define a two-step payment scheme, where the customer first pays for its expected demand, and later pays for any imbalances.
- We study optimal strategies for both players.
- We show that the proposed tariff provides variable incentives and elicits intelligent behavior by the customer.
- We further demonstrate the existence of Nash equilibria in this game, considering that the retailer has access to the private costs of the customer.

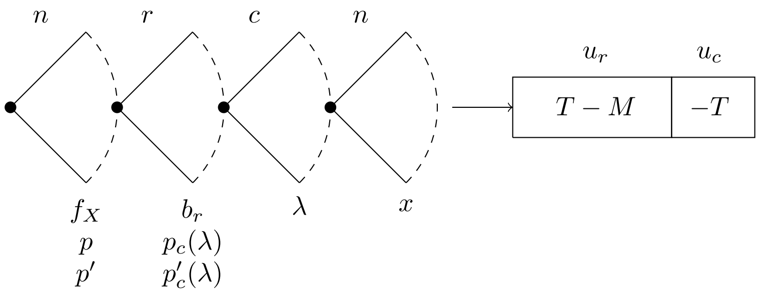


**Figure 2.1** Extensive form representation of the risk-sharing game. The retailer's (r), customer's (c) and nature's (n) moves set the respective decision variables, that together determine the utilities.

- Last, we discuss the concept of bounded rationality and show that the retailer may provide higher incentives to customers that choose over different risk-levels stochastically.

To the best of our knowledge, the work presented in this chapter is the first game-theoretical study that considers incentives for intelligent customer behavior where the customer has the choice of how much risk to assume from the retailer. In the closest state-of-the-art work, Vinyals et al. (2014) propose a *prediction-of-use* (POU) tariff that requires customers to predict their future baseline consumption and then charges them based on their actual consumption. However, the POU tariff does not model the balancing responsibility of the retailer based on the demand forecast and the prices for ahead and balancing markets as we do in this chapter.

The rest of this chapter is organized as follows: First, in Section 2.2 we formalize the game-theoretical model of our problem setting, upon which we study the proposed risk-sharing tariff and the strategies of both players participating, the retailer and the customer. In Section 2.3 we study Nash equilibrium strategy pairs for the risk-sharing game as well as the effects of stochastic tariff selection from the customer. Last, in Section 2.4 we conclude this chapter and outline research directions that future work may investigate further.

## 2.2 The Risk-Sharing Game

To formalize the setting described in the previous section, we capture the strategic interactions between the retailer and the customer in a two-player game. Figure 2.1 illustrates the extensive-form representation of the risk-sharing game, showing the time sequencing of the actions. We consider a two-step market model: the retailer first procures electricity in the ahead market with the unit price $p$ and later pays for any absolute[1] deviation between the observed demand of the customer and the procured quantity with the unit price $p' > p$ in the balancing market. We assume that the prices $p, p'$ are determined by an exogenous process and cannot be influenced by the retailer (i.e., *price-taker*).

Let $X$ denote the random variable of the customer's demand, and $f_X$ the probability density function (PDF) of the specified random variable. We denote with $x \in \mathbb{R}^+$ the realization of the demand. We also consider the distribution $f_X$ as the default behavior by the customer. The distribution $f_X$ is known to both players, since it can be observed in practice, e.g., through smart-meters, and can be approximated given enough observations. The proposed tariff requires the customer to precommit to and prepay its anticipated demand $b_c = \mathbb{E}_{f_X}[x]$, where $\mathbb{E}_{f_X}[x]$ is the expected value of the random variable $X$ under the distribution $f_X$. For ease of reading, in the remainder of this chapter we use the following simplified notation: $\mathbb{E}_f[x] \triangleq \mathbb{E}_{f_X}[x]$.

The retailer, based on the customer's demand distribution $f_X$, procures the quantity $b_r$ in the ahead market. Any absolute deviation between the quantity $b_r$ and the observed demand $x$ of the customer is balanced by the retailer in the balancing market. We consider the expected balancing costs as the *balancing risk* for the retailer (Ferguson, 1967), which is equal to: $\mathbb{E}_f[|b_r - x|]p'$. Recall that we assume a direct influence of the customer's demand to the balancing requirements of the retailer.

In current electricity systems retailers holds all the risk of balancing supply and demand. However, in the risk-sharing tariff, the balancing risk can be shared between the retailer and the customer. Let $\lambda \in [0, 1]$ denote the share of risk that remains with the retailer and $(1 - \lambda)$ the share of risk that is assumed by the customer. The risk-sharing tariff comprises two price functions:

i. The precommitment price $p_c(\lambda)$ for the quantity $b_c$, which we assume is equal to the anticipated load, i.e., $b_c = \mathbb{E}_f[x]$,
ii. and the imbalance price $p'_c(\lambda)$, which is the price that is paid for any absolute deviation between the anticipated and the observed demand $|b_c - x|$.

The retailer decides the price functions $p_c(\lambda)$ and $p'_c(\lambda)$ based on its procurement decision $b_r$, and the probability density function of the demand $f_X$. The customer then chooses the risk share $\lambda$ to be covered by the retailer. The utilities $u_r$ and $u_c$ for the retailer and the customer respectively are determined after the realized demand $x$ is observed.

Let $T$ denote the payment from the customer to the retailer and $M$ the market costs of the retailer. The utilities can be written as: $u_r = T - M$ and $u_c = -T$. Analytically,

$$u_c = -b_c p_c(\lambda) - |b_c - x| p'_c(\lambda) \tag{2.1}$$

[1] In practice, both power excess and shortages can result in the increase of balancing costs for the retailers, since they may be charged for the deployment of upwards or downwards regulation power by the TSO.

is the utility of the customer, including the cost for the precommitted quantity $b_c$ and the cost for any absolute deviation from the anticipated load. Similarly,

$$u_r = b_c p_c(\lambda) + |b_c - x| p'_c(\lambda) - b_r p - |b_r - x| p' \tag{2.2}$$

is the utility of the retailer, which is equal to the payment by the customer deducting the market costs of the retailer.

So far we have described the risk-sharing game between the retailer and the customer, further defining the utilities for both players. We can generalize and say that the risk-sharing tariff approximates the current retail flat tariff situation when no risk is assumed by the customer (i.e., $\lambda = 1$). Let $x$ be the demand of the customer and $N$ the number of payments during one year from the customer to the retailer under the current flat tariff market. Given the law of large numbers we know that for large $N$, $\sum_N x \approx N \times \mathbb{E}_f[x]$ holds. Therefore, the total payment of the customer approximates the payment under the risk-sharing tariff when the retailer holds all the risk, for $\lambda = 1$.

### 2.2.1 Optimal quantity of procurement

After the prices $p, p'$ and the distribution $f_X$ are determined, the retailer procures the quantity $b_r$ in the ahead market. In this section we compute the optimal procurement $b_r^*$ that maximizes the expected utility of the retailer in equation (2.2). Let $U_r^f$ denote the expected utility of the retailer with respect to the random variable of the demand $X$ under the distribution $f_X$.

$$U_r^f = b_c p_c(\lambda) + \mathbb{E}_f[|b_c - x|] p'_c(\lambda) - b_r p - \mathbb{E}_f[|b_r - x|] p' \tag{2.3}$$

The price functions $p_c(\lambda)$ and $p'_c(\lambda)$ are free parameters, since they determine the profit. We treat the price functions as independent of $b_r$ and therefore we minimize the market costs $M$ of the retailer.

**Lemma 2.1.** *The first derivative of the expected utility of the retailer in equation* (2.3) *with respect to $b_r$ is:*

$$\frac{d}{db_r} U_r^f = -p - 2p' F_X(b_r) + p', \tag{2.4}$$

*where $F_X$ is the cumulative distribution function (CDF) of the random variable $X$.*

*Proof.* We compute the derivative of the expected utility of the retailer (with regards to the random variable of the demand $X$) with respect to the procurement quantity $b_r$.

$$\begin{aligned}
\frac{d}{db_r} U_r^f &= -\frac{d}{db_r}(b_r p + \mathbb{E}_f[|b_r - x|] p') \\
&= -b_r - p' \left( \frac{d}{db_r} \left( \int_0^{b_r} (b_r - x) f_X(x) dx \right) \right. \\
&\quad \left. + \frac{d}{db_r} \left( \int_{b_r}^{\infty} (x - b_r) f_X(x) dx \right) \right)
\end{aligned} \tag{2.5}$$

Given the following equalities:

$$\frac{d}{da}\left(\int_0^a (a-x) f_X(x) da\right) = F_X(a)$$

and

$$\frac{d}{da}\left(\int_a^\infty (x-a) f_X(x) da\right) = -(1 - F_X(a)),$$

equation (2.5) becomes:

$$\begin{aligned}\frac{d}{db_r} U_r^f &= -b_r - p'\left(F_X(b_r) - (1 - F_X(b_r))\right)\\ &= p + 2p' F_X(b_r) - p'.\end{aligned}$$

□

We proceed to compute the procurement quantity $b_r$ that maximizes the expected utility of the retailer.

**Theorem 2.1.** *The quantity $b_r^*$ maximizes the expected utility of the retailer.*

$$b_r^* = F_X^{-1}\left(\frac{p' - p}{2p'}\right), \tag{2.6}$$

*where $F_x^{-1}$ is the inverse cumulative distribution (ICDF) function.*

*Proof.* Equation (2.6) follows from $\frac{d}{db_r} U_r^f = 0$. The expected utility of the retailer is a strictly concave function, and thus $b_r^*$ is a unique optimum since the following holds:

$$\frac{d^2}{db_r^2} U_r^f = -2p' f_X(b_r) < 0.$$

□

For any given $p' > p$, the quantity $b_r^*$ is lower than the expected demand due to the absolute imbalance quantity.

### 2.2.2 Determining the price for risk-sharing

In this section we define the requirements and the properties of the risk-sharing tariff and we propose how to choose the price functions. An important requirement for the price functions $p_c(\lambda), p_c'(\lambda)$ is that the expected utility of the retailer for any given $\lambda \in [0, 1)$ should be greater or equal to the expected utility when $\lambda = 1$. More specifically,

$$U_r^f(\lambda) \geq U_r^f(\lambda = 1) \geq b_c \varphi,\ \forall \lambda \in [0, 1),\ \varphi \in \mathbb{R}^+, \tag{2.7}$$

where $\varphi$ denotes an extra profit for the retailer per expected unit of demand. The quantity $\varphi$ approaches business costs in a perfect competition and arbitrarily large values in a monopoly.

Given the requirement in equation (2.7) and using equation (2.3), we derive the following inequality:

$$p_c(\lambda) \geq \frac{1}{b_c}(b_r^* p + \mathbb{E}_f[|b_r^* - x|]p' - \mathbb{E}_f[|b_c - x|]p_c'(\lambda)) + \varphi. \tag{2.8}$$

To find functions $p_c(\lambda)$ and $p_c'(\lambda)$ that satisfy the above inequality, we define the minimum imbalance price function:

$$\varrho_c'(\lambda) \triangleq (1-\lambda)p', \tag{2.9}$$

which is equal to the price the customer would pay by participating in the balancing market for its share $(1-\lambda)$ of balancing risk. Since $p_c(\lambda)$ is a free choice, we propose the minimum ahead price function that satisfies the inequality in equation (2.8) when replacing $p_c'(\lambda)$ with equation (2.9):

$$p_c(\lambda) \triangleq \frac{1}{b_c}(b_r^* p + \mathbb{E}_f[|b_r^* - x| + (\lambda - 1)|b_c - x|]p') + \varphi. \tag{2.10}$$

We proceed to show that this proposed price function guarantees the minimum profit margin $\varphi$ for the retailer.

**Theorem 2.2.** *Any tariff* $(p_c(\lambda), p_c'(\lambda))$*, using* $p_c(\lambda)$ *as defined in equation* (2.10) *and satisfying* $p_c'(\lambda) \geq \varrho_c'(\lambda), \forall \lambda \in [0,1]$*, and* $p_c'(1) = 0$*, satisfies the requirement in equation* (2.7).

*Proof.* For simplicity, let $U_r^{p'}(\lambda)$ denote the expected utility of the retailer with regards to the probability density function $f_X$ and imbalance price $p'(\lambda)$. First note that $p_c(\lambda)$ is defined such that $U_r^{\varrho}(\lambda) = U_r(1)$ when $p_c'(\lambda) = \varrho'(\lambda)$. For any function $p_c'(\lambda)$ that satisfies $p_c'(\lambda) \geq \varrho'(\lambda)\ \forall \lambda \in [0,1]$, $U_r^{p'}(\lambda) \geq U_r^{\varrho'}(\lambda)$ holds, since only the profit from the term $p_c'(\lambda)\mathbb{E}_X[|b_c - x|]$ increases while all other terms are fixed. □

The function $p_c'(\lambda)$ refers to the price per unit for any absolute deviation of the customer's demand given the choice of $\lambda$. We propose $p_c'(\lambda)$ to embrace some additional desired properties with regards to the ability of the customer to reduce its demand uncertainty.

Consider a customer that can alter the probability distribution of the random variable $X$, $f_X \to g_X$, such that $\mathbb{E}_g[|b_c - x|] \leq \mathbb{E}_f[|b_c - x|]$ (i.e., the expected absolute deviation of the random variable $X$ under the distribution $g_X$ is lower than the default demand behavior $f_X$). Note that $\mathbb{E}_f[|b_c - x|]$ denotes the expected value of $|b_c - x|$ with regards to the random variable $X$ under the distribution $f_X$, and $\mathbb{E}_g[|b_c - x|]$ denotes the expected value of $|b_c - x|$ with regards to the random variable $X$ under the distribution $g_X$. We define $g_X$ as the demand response of the customer. We propose a tariff that additionally imposes the constraint $\mathbb{E}_g[x] = \mathbb{E}_f[x]$. Let $\lambda^*(g_X)$ denote the risk that maximizes the utility of the customer with regards to the distribution $g_X$. The following two properties are common sense conditions for demand response tariffs:

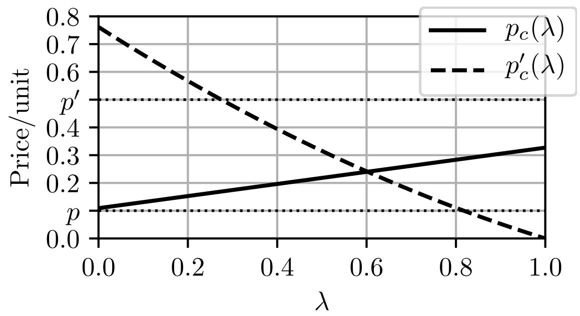


**Figure 2.2** The precommitment $p_c(\lambda)$ and imbalance $p'_c(\lambda)$ price functions for all values of $\lambda$.

i. *No demand response, no risk incentive*: If $\mathbb{E}_g[|b_c - x|] = \mathbb{E}_f[|b_c - x|]$ then $\lambda^*(g_X) = 1$.

ii. *Demand response proportional risk*: Consider the distribution $z_X$, if $\mathbb{E}_g[|b_c - x|] < \mathbb{E}_z[|b_c - x|] < \mathbb{E}_f[|b_c - x|]$ then $0 \le \lambda^*(g_X) < \lambda^*(z_X) < 1$.

We propose the following imbalance price function that satisfies the above properties (see Section 2.2.3) under $\vartheta > 0$.

$$p'_c(\lambda) = (1 - \lambda)(p' + \vartheta\Phi(\lambda)), \tag{2.11}$$

where $\Phi(\lambda)$ denotes the penalty that is equal to the discount in the precommitment price the retailer offers,

$$\Phi(\lambda) = p_c(1) - p_c(\lambda). \tag{2.12}$$

The parameter $\vartheta \in \mathbb{R}^+$ scales the penalty term $\Phi(\lambda)$. Figure 2.2 illustrates the shape of the price functions $p_c(\lambda)$ and $p'_c(\lambda)$ that are computed for $f_X = \mathcal{N}(0.15, 0.1)$, truncated to $x \in [0, 0.79]$, $p = 0.1$, $p' = 0.5$, $\varphi = 0.02$, and $\vartheta = 1$.

The tariff composed of $p_c(\lambda)$ and $p'_c(\lambda)$ guarantees a minimum acceptable utility for the retailer, which is equal to the current flat tariff situation ($\lambda = 1$). The imbalance price function $p'_c(\lambda)$ proposed in equation (2.11) also satisfies desirable properties with respect to the upcoming discussion, associated with the strategy of a customer that can reduce the uncertainty of its demand.

### 2.2.3 Optimal strategies for flexible customers

Demand response in electricity systems refers to the ability of customers to adjust their demand behavior in response to financial incentives provided by electricity providers. In this chapter we interpret demand response as the ability of the customer to reduce the uncertainty of its demand. Let $\Delta$ denote the action of the customer, which affects the distribution of the demand $f_X$, such that the observed demand $x$ is sampled from the new distribution $g_X$. Recall that the expected demand remains the same $\mathbb{E}_g[x] = \mathbb{E}_f[x]$ and the expected absolute deviations may become lower $\mathbb{E}_g[|b_c - x|] \leq \mathbb{E}_f[|b_c - x|]$. Let $C_\Delta(g_X) \triangleq C_\Delta(f_X \to g_X)$ denote the costs associated with reducing the uncertainty, e.g., capturing customer's discomfort or costs of smart devices and batteries.

We show that for any distribution $g_X$ there is a unique $\lambda^* \in [0, 1]$ that maximizes the expected utility of the customer. Let $U_c^g$ denote the expected utility of the customer with demand response $\Delta$ and resulting demand distribution function $g_X$.

$$U_c^g = -b_c p_c(\lambda) - \mathbb{E}_g[|b_c - x|]p'_c(\lambda) - C_\Delta(g_X) \tag{2.13}$$

where the expected absolute imbalance is computed given the distribution $g_X$. Note that the prices $p_c(\lambda)$ and $p'_c(\lambda)$ are computed by the retailer given the distribution $f_X$, since the retailer does not hold information about the demand response action of the customer $\Delta$.

**Lemma 2.2.** *The first derivative of the expected utility of the customer in equation* (2.13) *with respect to* $\lambda$ *is:*

$$\frac{d}{d\lambda} U_c^g = a_g p' - a_f p' - 2(\lambda - 1)\frac{\vartheta}{b_c} a_g a_f p', \tag{2.14}$$

*where* $a_f = \mathbb{E}_f[|b_c - x|]$, *and* $a_g = \mathbb{E}_g[|b_c - x|]$.

*Proof.* We compute the derivative of the quantity in equation (2.13) with respect to $\lambda$.

$$\begin{aligned}\frac{d}{d\lambda} U_c^g &= \frac{d}{d\lambda}(-b_c p_c(\lambda) - \mathbb{E}_g[|b_c - x|]p'_c(\lambda) - C_\Delta(g_X)) \\ &= \frac{d}{d\lambda}(-b_c p_c(\lambda) - \mathbb{E}_g[|b_c - x|]p'_c(\lambda))\end{aligned} \tag{2.15}$$

By using equation (2.15) and $a_g = \mathbb{E}_g[|b_c - x|]$,

$$\frac{d}{d\lambda} U_c^g = b_c \frac{dp_c(\lambda)}{d\lambda} + a_g \frac{dp'_c(\lambda)}{d\lambda}. \tag{2.16}$$

We first consider the first term of equation (2.16).

$$\begin{aligned}b_c \frac{dp_c(\lambda)}{d\lambda} &= b_c \frac{d}{d\lambda}\left(\frac{1}{b_c}(b_r^* p + \mathbb{E}_f[|b_r^* - x| + (\lambda - 1)|b_c - x|]p') + \varphi\right) \\ &= \frac{d}{d\lambda}((\lambda - 1)\mathbb{E}_f[|b_c - x|]p') \\ &= a_f p',\end{aligned} \tag{2.17}$$

where $a_f = \mathbb{E}_f[|b_c - x|]$.

We proceed to the second term of equation (2.16).

$$a_g \frac{dp'_c(\lambda)}{d\lambda} = a_g \frac{d}{d\lambda}((1-\lambda)(p' + \vartheta\Phi(\lambda)))$$
$$= -a_g p' + a_g \frac{d}{d\lambda}((1-\lambda)\vartheta\Phi(\lambda)), \tag{2.18}$$

where $\Phi(\lambda)$ can be written as follows (using equation 2.12):

$$\begin{aligned}\Phi(\lambda) &= p_c(1) - p_c(\lambda)\\ &= \frac{1}{b_c}(b_r^* p + \mathbb{E}_f[|b_c - x|]p') + \varphi\\ &- \frac{1}{b_c}(b_r^* p + \mathbb{E}_f[|b_r^* - x| + (\lambda-1)|b_c - x|]) + \varphi\\ &= \frac{1}{b_c}((1-\lambda)\mathbb{E}_f[|b_c - x|]p')\\ &= \frac{1}{b_c}(1-\lambda)a_f p'.\end{aligned} \tag{2.19}$$

By replacing equation (2.19) in equation (2.18) the second term of equation (2.16) becomes

$$\begin{aligned}a_g \frac{dp'_c(\lambda)}{d\lambda} &= -a_g p' + a_g \frac{d}{d\lambda}\left((1-\lambda)\vartheta\frac{1}{b_c}(1-\lambda)a_f p'\right)\\ &= -a_g p' + \frac{\vartheta}{b_c}a_g \frac{d}{d\lambda}\left((1-\lambda)^2 a_f p'\right)\\ &= -a_g p' + 2(\lambda-1)\frac{\vartheta}{b_c}a_g a_f p'.\end{aligned} \tag{2.20}$$

Last, by replacing equations (2.20) and (2.17) in equation (2.16) we get:

$$\frac{d}{d\lambda}U_c^g = a_g p' - a_f p' - 2(\lambda-1)\frac{\vartheta}{b_c}a_g a_f p'.$$

□

**Theorem 2.3.** *The quantity $\lambda^*$ maximizes the expected utility of the customer for any given $g_X$ and $C_\Delta$.*

$$\lambda^*(g_X) = \left[\frac{a_g - a_f}{2\frac{\vartheta}{b_c}a_g a_f} + 1\right]_0^1, \tag{2.21}$$

where $[x]_l^h = \max(l, \min(h, x))$.

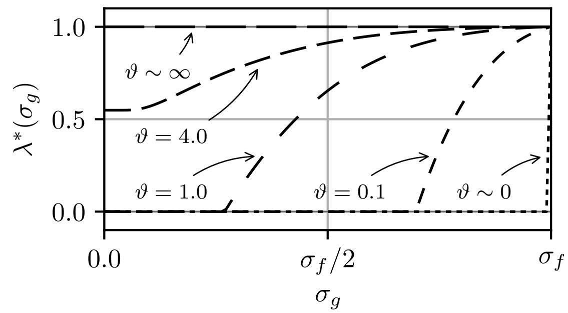


**Figure 2.3** The mapping between $\sigma_g$ and the optimal share $\lambda^*$ that maximizes in expectation the utility of the customer.

*Proof.* Equation (2.21) follows from $\frac{d}{d\lambda}U_c^g = 0$. The utility function of the customer with regards to the risk choice $\lambda$ is strictly concave, since

$$\frac{d^2}{d\lambda^2}U_c^g = -2\frac{\vartheta}{b_c}a_g a_f p' < 0.$$

Therefore, $\lambda^*$ is the unique optimum. □

Under the assumption of a cost-free demand response model, i.e., $C_\Delta(\cdot) = 0$, we proceed to show that a customer with uncertain demand response has incentives to participate in the risk-sharing tariff contributing its maximum available demand response. Consider the distribution $z_X$ such that:

$$\mathbb{E}_g[|b_c - x|] < \mathbb{E}_z[|b_c - x|] < \mathbb{E}_f[|b_c - x|], \tag{2.22}$$

where $z_X$ provides a threshold ability of the customer to reduce the expected absolute deviation of the demand.

**Theorem 2.4.** *For a customer with uncertain demand response $g_X$ that can at least reduce the uncertainty of its demand to the level of $z_X$ such that equation* (2.22) *is satisfied, it holds that:*

$$U_c^g(\lambda^*(z_X)) > U_c^z(\lambda^*(z_X)) > U_c^f(1),$$

*and*

$$U_c^g(\lambda) \geq U_c^f(1),\ \forall\lambda \geq \lambda^*(z_X).$$

*Proof.* The inequality in equation (2.22) indicates that the imbalance payment in equation (2.13) follows the same ranking, as it is a product of unequal expectations

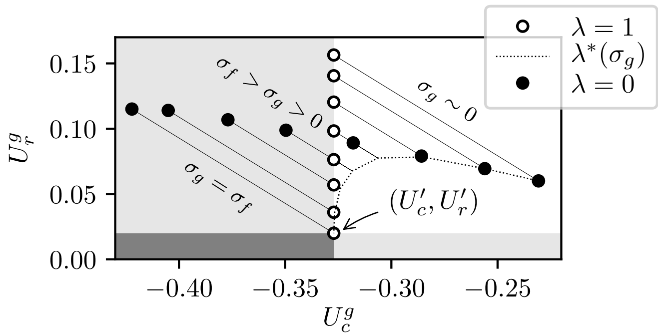


**Figure 2.4** The utilities of the customer $U_c^g$ and the retailer $U_r^g$. Each line segment corresponds to a different $\sigma_g$, starting from $\lambda = 1$ and ending in $\lambda = 0$.

with the identical imbalance price $p_c(\lambda^*(z_X))$. Since all other terms remain the same, this directly induces the inequalities of the resulting utilities stated in the theorem. □

Theorem 2.4 implies that a customer with uncertain demand response $g_X$, which is however bounded by $z_X$, can only benefit by contributing all available demand response. Furthermore, any choice of $\lambda \geq \lambda^*(z_X)$ ensures a lower bound for the utility of the customer.

For the remainder of this chapter we assume that both $g_X$ and $f_X$ are normal distributions with mean $\mu_g = \mu_f$ and standard deviation $\sigma_f$ and $\sigma_g \in (0, \sigma_f]$, respectively. For this restricted case we apply a simplified notation: $\lambda^*(\sigma_g)$ denotes the optimal strategy for the customer, similarly $C_\Delta(\sigma_g)$ denotes the cost. Figure 2.3 presents the function $\lambda^*(\sigma_g)$ for different values of $\vartheta$ in equation (2.11). For $\vartheta \sim 0$, the utility of the customer becomes a linear function that is monotonically increasing in $\lambda$ when $\sigma_g = \sigma_f$. Therefore, the optimal choice of the customer becomes $\lambda^*(\sigma_f) = 1$, and $\lambda^*(\sigma_g < \sigma_f) = 0$. For $\vartheta \sim \infty$, the optimal choice of the customer is to assume no risk ($\lambda^*(\sigma_g) = 1, \ \forall \sigma_g \in (0, \sigma_f]$), since the penalty term $\Phi$ in equation (2.11) is infinitely scaled.

In this section we derived the optimal strategy $\lambda^*(g_x)$ of the customer. We further illustrated how the choice of the parameter $\vartheta$ by the retailer can influence the optimal strategy of the customer. Furthermore, Theorem 2.4 demonstrated that the risk-sharing tariff is attractive to customers with uncertain demand response (eliciting the maximum available uncertainty reduction), who can apply a safe strategy corresponding to their minimum ability.[2]

### 2.2.4 Comparison of the utilities

We compare the expected utilities of both players, again under the assumption of a cost-free demand response model, i.e., $C_\Delta(\sigma_g) = 0, \ \forall \sigma_g \in (0, \sigma_f]$. Figure 2.4 illustrates the expected utilities of both players. Let the tuple $(U'_c, U'_r)$ illustrate the point in the utility space that represents the current flat tariff situation, i.e., $\sigma_g = \sigma_f$ and $\lambda = 1$. Each line segment in the figure represents the utility tuples given a specific demand response $\sigma_g$ and varying $\lambda$. The empty circles represent the utility tuples when the customer chooses to assume no risk ($\lambda = 1$). In this case demand response only yields a benefit to the retailer. On the contrary, filled circles represent the utility tuples when the customer chooses to assume the full share of risk. Increasing the risk assumption (moving across the line segments from $\lambda = 1$ to $\lambda = 0$) requires a certain level of demand response to be profitable for the customer. For high demand response (low $\sigma_g$), it results in the utility increase for the customer. For low demand response (high $\sigma_g$), only the retailer benefits from the decreasing uncertainty of the demand. Reduced uncertainty on the demand side can therefore contribute to the improved social welfare (sum of the players' utilities) through the risk-sharing tariff.

Theorem 2.4 can also be illustrated using Figure 2.4. Note that for normal distributions, the following holds: $\mathbb{E}[|x - \mu|] = \sigma\sqrt{2/\pi}$ (Geary, 1935). Hence, $\sigma_g < \sigma_z < \sigma_f$ implies that the inequalities in equation (2.22) also hold. According to Theorem 2.4, it follows that $U_c^g(\lambda^*(\sigma_z)) > U_c^z(\lambda^*(\sigma_z)) > U_c^f(1)$. Intuitively, the customer can increase its utility by switching from $\sigma_z$ to $\sigma_g$, or more generally by switching from $z_X$ to $g_X$.

## 2.3 Nash Equilibrium Strategies

In this section we study the Nash equilibria (NE) of the risk-sharing game. Where necessary, we make the dependence of utilities on both strategies more explicit by using notation $U_c(\pi_r, \pi_c)$ and $U_r(\pi_r, \pi_c)$, where $\pi_r$ and $\pi_c$ denote the strategies of the retailer and the customer respectively. NE are pairs of strategies $(\pi_r^*, \pi_c^*)$, such that $U_c(\pi_r^*, \pi_c^*) \geq U_c(\pi_r^*, \pi_c), \forall \pi_c$ and $U_r(\pi_r^*, \pi_c^*) \geq U_r(\pi_r, \pi_c^*), \forall \pi_r$. Let $C_\Delta(\sigma_g) \geq 0, \ \forall \sigma_g \in (0, \sigma_f)$ be an arbitrary cost model for demand response and $C_\Delta(\sigma_f) = 0$, i.e., cost without demand response is zero. We assume that the demand response cost model is known by the retailer.

First, we define the strategies of the two players. For the retailer, the only free choice is the scalar $\vartheta$ that parameterizes the proposed tariff in equation (2.11). The strategy of the retailer is denoted by $\pi_r = \vartheta$. For the customer, the strategy $\pi_c = (\lambda, \sigma_g)$ refers to the choice of risk $\lambda$ and demand response $\sigma_g$. Furthermore, the strategy includes the credible threat of returning to the flat tariff without any demand response, $\pi_c^{\text{threat}} = (1, \sigma_f)$, if the utility drops below the reference utility

[2] Safe strategy denotes a risk share $\lambda$ that guarantees a utility that is at least as high as the reference retail tariff without risk.

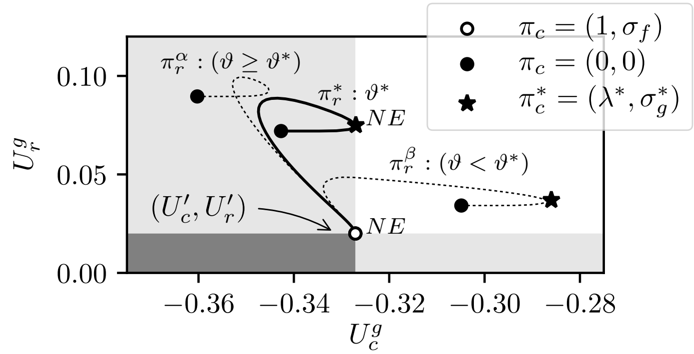


**Figure 2.5** Each curve represents the utility tuples, given the strategies for the retailer: $\pi_r^*$, $\pi_r^\alpha$, $\pi_r^\beta$, and all possible strategies $\pi_c = (\lambda^*(\sigma_g), \sigma_g)$ for the customer.

$U_c'$. The threat is possible due to the action sequence that is indicated in Figure 2.1 and credible since the threat strategy outperforms the protected equilibrium strategy. Given any $C_\Delta(\sigma_g)$ we know from Theorem 2.3 that for any given strategy $\pi_r$ there always exists a strategy $\pi_c^*$ that maximizes the expected utility of the customer: $\pi_c^* = (\lambda^*(\sigma_g^*)\ ,\ \sigma_g^* = \arg\max_{\sigma_g}[U_c^g(\lambda^*(\sigma_g))])$.

Figure 2.5 illustrates the utilities of the two players, which are computed using the quadratic demand response cost model: $C_\Delta(\sigma_g) = w|\sigma_f - \sigma_g|^2$, for $w = 10$. Each curve corresponds to one of the following three retailer strategies: $\pi_r^* = \vartheta^*$, $\pi_r^\alpha = (\vartheta \geq \vartheta^*)$, $\pi_r^\beta = (\vartheta < \vartheta^*)$. The utility tuples along each curve correspond to the customer strategies $\pi_c = (\lambda^*(\sigma_g), \sigma_g)$. The curves start from the utility tuple $(U_c', U_r')$ denoted by the empty circle, where the customer does not assume any risk, and hence performs no demand response, i.e., $\pi_c = (1, \sigma_f)$. The curves end where $\pi_c = (0, 0)$ denoted by the filled circles. Note that the strategy $\pi_r^* = \vartheta^*$, which maximizes the utility of the retailer (solid curve), yields $U_c^g(\lambda^*(\sigma_g^*), \sigma_g^*) = U_c^f(1, \sigma_f)$ for the customer. The utility of the customer using demand response (star on solid curve) becomes equal to the utility without demand response (open circle).

**Theorem 2.5.** *The strategy pair I,* $(\pi_r^*, \pi_c^* = (\lambda^*, \sigma_g^*))$*, and the set of pairs II,* $(\pi_r^\alpha, \pi_c^{threat})$*, are the only two types of NE in the risk-sharing game.*

*Proof.* Any positive change $\varepsilon$ in the strategy of the retailer, such that $\vartheta \leftarrow \vartheta^* + \varepsilon$ (e.g., $\pi_r^\alpha$), will cause the customer to adopt the strategy $\pi_c^{\text{threat}}$ since $U_c^g(\pi_c^*, \vartheta) < U_c'$, leading both players to the utility pair $(U_c', U_r')$. On the other hand, any negative change $\varepsilon$ (e.g., $\pi_r^\beta$) will directly reduce the retailer's expected utility. Hence, the

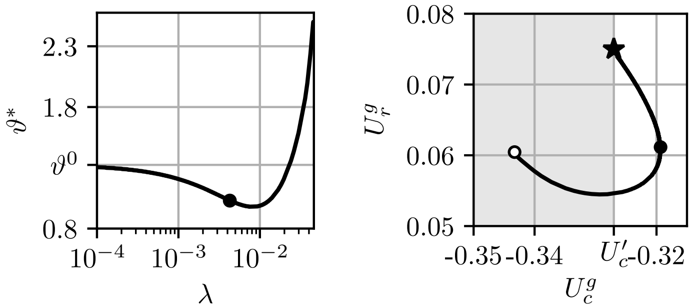


**Figure 2.6** Left: The choice $\vartheta^*$ by the retailer depending on the irrationality parameter $\tau$. The value of $\vartheta^0$ corresponds to NE I. Right: The utility pairs starting from the rational NE pair (star) and ending at the open circle, where the customer acts almost randomly (the curve would continue in a straight zero-sum line to the upper left). The filled circle indicates $\tau^*$, which maximizes the utility of the customer.

retailer has no incentive to deviate from I. In addition, the customer strategy is the best response by definition, and thus the customer has no incentive to deviate either, making I a NE. Now, consider any strategy pair $(\pi_r^\beta, \pi_c)$. The retailer can gain by deviating from this strategy pair by increasing $\vartheta$, making sure that no equilibrium containing $\pi_r^\beta$ exists. Finally, consider any pair $(\pi_r^\alpha, \pi_c^{\text{threat}})$. Since $U_r$ is unaffected by $\vartheta$ given $\lambda^*(\pi_r^\alpha) = 1$, providing no incentive to apply demand response, no player can gain by deviating unilaterally. Hence, each pair of strategies in the set II is a NE. $\square$

We showed the existence of two types of NE within the risk-sharing game, where $(\pi_r^*, \pi_c^* = (\lambda^*, \sigma_g^*))$ Pareto dominates $(\pi_r^\alpha, \pi_c^{\text{threat}})$ and therefore is favorable for both players.

### 2.3.1 Bounded rational customer

The concept of bounded rationality was previously discussed in Section 1.1.1 and it assumes that agents, automated or not, do not behave as perfectly rational decision-makers, bounded by imperfect information and their limited computational capacity (Simon, 1972). The same applies to the economic behavior of customers, which can also be modeled using the bounded rationality paradigm (McFadden, 1975): customers do not always subscribe to the cheapest tariff but the probability of doing so is high.

In this section we use the Softmax function (see Section 1.1.1) to model the decision of the customer with regards to the choice of the risk it is willing to assume in the risk-sharing tariff setting.

$$\mathbb{P}(\pi_c^k) = \frac{\exp(U_c(\pi_c^k)/\tau)}{\sum_{\pi_c^i} \exp(U_c(\pi_c^i)/\tau)}, \tag{2.23}$$

where $\mathbb{P}(\pi_c^k)$ is the probability of the customer to use strategy $\pi_c^k$. We apply this function to a discrete sampling of the continuous parameter $\sigma_g$ to probabilistically mix between strategies $\pi_c = (\lambda^*(\sigma_g), \sigma_g), \sigma_g \in (0, \sigma_f]$. Figure 2.6 illustrates the numerical approximation of the optimal choice of $\vartheta^*$ for the retailer and the utility tuples under various degrees of irrationality $\tau$. For $\tau = 10^{-4}$, $\vartheta^0$ approximates the value of $\vartheta^*$ computed earlier for the rational customer, since for low $\tau$ Softmax approximates the rational choice. Beyond $\tau \approx 10^{-2}$, the retailer starts increasing $\theta^*$ to infinity as the customer becomes random, resulting in the infinite increase in the utility of the retailer at the cost of the customer. Note that $\tau^* = 4.27 \cdot 10^{-3}$ maximizes the utility of the customer. At the same time, it results in a higher utility than the one is obtained by a rational customer in NE since the retailer offers a higher incentive for demand response if believes to be facing a bounded rational customer.

The above result has several implications for implementing automated tariff selection algorithms. In particular, the irrationality parameter $\tau^*$ can be seen as in equilibrium with $\vartheta^*$ suggesting that automated strategies may perform better by adopting probabilistic Softmax selection algorithms instead of greedy approaches.

# 2.4 Conclusions

In this chapter we proposed a tariff where the balancing risk can be shared to incentivize uncertainty reduction on the demand side. First, we defined an extensive-form two-player game between the retailer and the customer in settings where the customer has a direct influence on the balancing requirements and costs of the retailer (see Section 2.2). Within this game-theoretical framework we proposed the risk-sharing tariff, which is a parameterized two-payment tariff scheme: the customer first pays a precommitment price for its anticipated demand, while after its actual demand is observed it pays for any absolute deviation between its anticipated and observed demand. In the risk-sharing tariff the customer can choose the portion of the balancing risk it is willing to assume from the retailer and therefore determine the price for precommitment and imbalances.

We showed that the proposed tariff is acceptable for both the retailer (see Section 2.2.2) and the customer (see Section 2.2.4), since the tariff yields benefits for both players in expectation. In addition, we studied best response strategies that are computable as presented in Sections 2.2.1 and 2.2.3. We further illustrated how social welfare is improved (see Figure 2.4) due to the uncertainty reduction on the demand-side, and we provided arguments (see Theorem 2.4) why the proposed tariff elicits all freely available demand response even in cases where the ability of the customer to demand respond is not certain. In Section 2.3 we showed the existence of NE within the risk-sharing game and illustrated them with computations. Last, we argued that bounded rationality can be a valuable concept for the design of tariffs. More specifically, we showed that the retailer offers higher incentives for demand response (through the risk-sharing tariff) in face of a customer that chooses the amount of risk to assume stochastically (see Section 2.3.1).

Overall, the proposed risk-sharing tariff is a step towards a key motivation of envisioned smart grid systems: this can be delineated by the active participation of customers in more elaborate tariff schemes which can enable indirect demand control, reducing uncertainty on the demand-side.

### 2.4.1 Future work

Throughout this chapter we have made some simplifying assumptions in order to show relevant theoretical properties of the proposed risk-sharing tariff scheme. Relaxing some of these assumptions, the work presented in this chapter further serves as a basis for future research directions.

First, we assume that prices $p$ and $p'$ are known in advance and only determined by nature. The studied setting can be extended to include stochastic prices for the ahead and balancing markets. We further consider that the demand distribution $f_X$ is known a-priori by the retailer. This may not be a realistic assumption in higher time resolutions (e.g., the demand of a quarter hour). Future work can investigate the design of a risk-sharing tariff where the demand distribution is reported by the customer directly.

Another assumption of our setting is that the expected demand after demand response should be equal to the initially anticipated demand (see Section 2.2.3). This may not be applicable in settings where a customer is not able to reduce the uncertainty of its demand without reducing or increasing its overall demand. Future work can study mechanisms that also allow the customer to report its expected demand (after demand response) in addition to the risk share.

Last, we consider settings where one customer has a direct influence on the balancing requirements of the retailer, and therefore our game-theoretical analysis considers the demand uncertainty of a single customer. It is of great interest to consider settings in which multiple customers participate in the risk-sharing tariff, where incentives are shared among different users that may have different private costs for demand response.

# 3 Forecast-Based Mechanisms for Demand Response

**Preface** In this chapter we study mechanisms to incentivize uncertain demand response in future electricity systems. We assume agents (e.g., EVs) that can respond (reduce or increase their demand) with some probability if they prepare prior to the realization of the demand. Previous work studies truthful mechanisms that select a minimal set of agents to prepare and respond such that a fixed demand reduction target is achieved with high probability. In this chapter we additionally consider the balancing responsibility of a retailer under a given demand forecast and imbalance price: the retailer is responsible to purchase additional reserve capacity at a high imbalance price to cover any excess in the demand. In this extended setting we study mechanisms that request only a subset of prepared agents to respond, since the reduction target depends on the realization of the demand and may not require all agents to respond. We propose: (i) a sequential mechanism that in each round embeds a second-price auction and is truthful under some mild assumptions for the setting, and (ii) a truthful combinatorial mechanism that runs in polynomial time and uses VCG payments. We show that both mechanisms guarantee non-negative utility in expectation for both agents and the retailer (mechanism), and can further be used for simultaneous downward and upward flexibility. Last, we verify our theoretical findings in an empirical evaluation over a wide range of mechanism parameters.

ℹ *This chapter presents work that was published in the proceedings of the International Conference of Autonomous Agents and Multiagent Systems (AAMAS) 2019 in Montreal, Canada (Methenitis et al., 2019b).*

## 3.1 Introduction

There are two important types of electricity markets that facilitate commerce of electricity between energy producers and customers: *day-ahead* and *balancing* markets (Conejo et al., 2010). Retailers (aggregators), based on demand forecasts, procure electricity in day-ahead markets to satisfy the demand of their customers. Imbalances between the procured quantities and the actual (intra-day) demand are moderated in balancing markets, in which the reserve power of high-cost storage units and conventional fast-ramping generators, e.g., gas-turbines, is traded. As a result, imbalances result in increasing costs for retailers and they are further associated with excessive $CO_2$ emissions (Hintermann, 2016). Maintaining balance between supply and demand is therefore one of the main factors that determine the efficiency of both existing and envisioned *smart grid* systems (Palensky and Dietrich, 2011).

To this end, *demand-side management* assumes that electricity users can alter their demand or generation behavior given economic incentives (e.g., smart tariffs, dynamic pricing), and therefore assist in reducing imbalances between supply and demand (Palensky and Dietrich, 2011). In practice, retailers currently agree upon long-term contracts to incentivize large-capacity users to reduce their demand if necessary (Kim and Shcherbakova, 2011). The introduction of smaller-scale flexible users, such as intelligent home appliances, electric vehicles (EVs) and home batteries, can further alleviate retailers from costs related to balancing supply and demand (Kempton and Tomić, 2005). However, the uncertain availability of such users requires more flexible contracts, such as short-term bilateral agreements that can be agreed upon and executed if necessary in a day-to-day manner.

In this chapter we consider the following setting: a retailer based on a demand forecast procures electricity in the day-ahead market to satisfy the demand of its customers. Since the demand is not certain, there is no guarantee that the actual demand (at the delivery time) is equal to the procured quantity. As a result, any imbalance between the procured quantity and the demand at the time of delivery should be adjusted in the balancing market with a much higher price than the procurement price. We consider agents that can reduce imbalances, after the realization of the demand (i.e., when demand is known but not finalized) and before the time of delivery, if requested by the retailer. Agents decide whether to prepare with some cost before the realization of the demand; prepared agents are able (with some probability) to respond if requested after the realization of the demand. Agents' responses can be observed and incur extra costs to agents. The following example illustrates an instantiation of the model of agents used in our setting:

**Example 3.1.** *Consider a neighborhood with multiple EVs that are parked and plugged into charging stations. Some of the vehicles may be fully charged, while others may be charging. In case of excess demand fully charged vehicles can be utilized to provide the extra needed electricity out of their battery, while vehicles that undergo charging can pause their charging. Each vehicle has a preparation cost, which is the opportunity/planning cost caused by extending its stay in a charging station. The probability of response refers to the uncertain availability of a vehicle to* *reduce its demand upon request. Last, the response cost is associated to the operating cost of response, such as the cost of battery degradation.*

In this setting we design mechanisms to incentivize demand response for agents that their availability to reduce/increase their demand if requested is not certain (see Example 3.1). We additionally consider the balancing responsibility of a retailer under a given demand forecast and imbalance price. More specifically, given the demand forecast, the imbalance price and the characteristics of agents (i.e., preparation cost, response probability and response cost), we design mechanisms that: (i) elicit truthful information with regards to the characteristics of agents, (ii) select a subset of agents to prepare, (iii) do not require all prepared agents to respond but only upon request (i.e., until imbalance is resolved), (iv) determine the order that prepared agents are asked to respond, (v) compute rewards and penalties for selected agents in order to incentivize them to prepare and respond (if able) if requested, and (vi) reduce the expected balancing cost of the retailer and overall increase social welfare.

The main contributions of this chapter can be summarized as follows:

- We study implications in mechanism design that arise by interdependencies between agents that are requested to respond sequentially and respond stochastically.
- To select a subset of all available agents and incentivize them to prepare and respond if requested, we propose a sequential mechanism that in each round embeds a second-price auction and is truthful under some assumptions for the setting, and a truthful combinatorial mechanism that runs in polynomial time and uses VCG payments.
- We show that both mechanisms guarantee non-negative utility in expectation for both agents and the mechanism, and can further be used for simultaneous downward and upward flexibility.
- We empirically evaluate the proposed mechanisms over a wide range of parameters and find that they achieve up to $16\%$ reduction in the balancing cost of the retailer and up to $14\%$ increase in social welfare compared to the case when no demand response is used. Last, we provide an evaluation of related work (Ma et al., 2016) in our extended setting.

To the best of our knowledge, the work presented in this chapter is the first to propose mechanisms that connect incentives for uncertain demand response with the balancing responsibility of the retailer, given the demand forecast of the customers and the imbalance price for the retailer.

The rest of this chapter is organized as follows: In Section 3.2 we present previous works that study mechanisms to incentivize demand response and we discuss how the work presented in this chapter differentiates. In Section 3.3 we formalize our problem setting. In Section 3.4 we define the general demand response mechanism that considers rewards and penalties for all selected agents and asks agents to respond sequentially, we further propose two truthful mechanisms that are based on the general demand response mechanism. Section 3.5 presents our empirical evaluation of the proposed mechanism over a wide range of mechanism parameters. Last, in

Section 3.6 we conclude this chapter and discuss future research directions of this work.

# 3.2 Related Work

In the closest state-of-the-art work, Ma et al. (2016) propose mechanisms to incentivize reliable demand reduction in electricity grids. The proposed mechanisms, that are based on greedy allocation and critical value payments (Lehmann et al., 1999), achieve a fixed demand reduction target with some reliability while minimizing the number of selected agents (see Section 3.5 for a detailed description). Our work differentiates from the aforementioned work in the following ways: (i) Ma et al. (2016) use a fixed demand reduction target. On the contrary, we propose mechanisms to allocate demand response when the demand reduction target is not known and only the demand forecast (distribution) is used to select agents. (ii) Ma et al. (2016) propose a two-stage setting: in the first stage, a set of agents is selected, and in the second stage, all selected agents are asked to respond. In contrast, we consider mechanisms that request prepared agents to respond until an imbalance between supply and demand is resolved (or cannot be reduced further), and thus the probability of requesting selected agents to respond is less or equal to one. (iii) Ma et al. (2016) do not consider the actual need for demand response (fixed demand reduction target). In this chapter, however, we consider that agents are requested to respond based on the realization of the demand, which can prevent excessive costs for demand response. Last, (iv) Ma et al. (2016) evaluate their proposed mechanisms with regards to the resulting payments to agents. On the contrary, we consider both the expected balancing cost for the retailer and the social cost of demand response, and thus the overall social welfare.

The aforementioned work by Ma et al. (2016) considers unit responses by agents (i.e., each agent can reduce one unit of demand), while a later extension of this work generalizes to multi-unit responses, uncertainties in preparation costs of agents, as well as a multi-effort probability of response (Ma et al., 2017). Other related work studies demand response contracts under uncertainty, where the reserve cost for the retailer (i.e., cost for not reaching the reduction target) is considered (Haring and Andersson, 2014; Meir et al., 2017). However, no prior work has studied mechanisms to incentivize uncertain demand response given the additional information of the demand forecast, as we do in this chapter.

# 3.3 Problem Formulation

In this section we outline our problem setting: Section 3.3.1 formulates the balancing responsibility of the retailer, Section 3.3.2 introduces demand response agents and Section 3.3.3 illustrates how demand response is used by the retailer.

## 3.3.1 Retailer's balancing responsibility

We consider a single retailer of electricity that is the balancing responsible party. The demand of the retailer's portfolio of customers is described by the discrete random

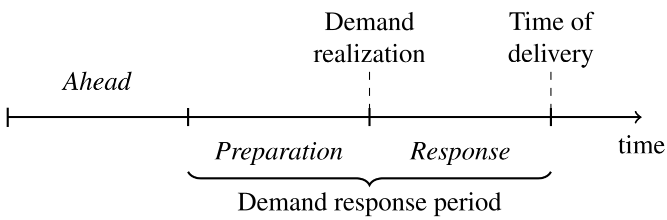


**Figure 3.1** Timeline of the model formulation.

variable $X$. The demand forecast $f_X(x) = \mathbb{P}_X(X = x)$ is the probability mass function (PMF) of $X$. We denote with $x \in \mathbb{Z}_0^D$ the realization of the demand, where $D$ is the upper bound of the support of $X$.

Consider the timeline in Figure 3.1, similarly to Chapter 2 we consider that the retailer procures the quantity $b \in \mathbb{Z}^+$ at unit price $p \in \mathbb{R}^+$ ahead of the realization of the demand, during the *ahead* period, in the day-ahead market. When no demand response is used, the retailer pays any positive imbalance between the demand realization $x$ and the procured quantity $b$ at unit price $p' \in \mathbb{R}^+$ (*imbalance price*) at the time of delivery in the balancing market. We assume that $p' > p$, and that the prices $p$ and $p'$ are determined by an exogenous process and cannot be influenced by the retailer (*price-taker*). We further assume that the procurement quantity $b$ is predetermined (see Assumption 3.1 in Section 3.4) and we focus on the expected balancing cost of the retailer:

$$C_{\neg DR} = p' \, \mathbb{E}_X[x - b | x > b], \tag{3.1}$$

where only positive imbalance from the procured quantity incurs a balancing cost to the retailer.

In practice, both demand excess (positive imbalance) and shortages (negative imbalance) result in balancing costs for retailers. However, to align our model with related work (Ma et al., 2016), we first consider the expected positive imbalance (see equation 3.1). In Section 3.4.4 we generalize to the case where both positive and negative imbalances incur balancing cost to the retailer.

**Remark 3.1.** *Our choice for a discrete model, where both the demand $x$ and the procurement $b$ are discrete variables, is motivated by markets that have trading volumes that are multiples of a unit quantity (as it is usual in day-ahead and balancing electricity markets) or markets with discrete items.*[1]

[1] Our model can be generalized to continuous variables if we neglect the need for decimal reduction, or if there is a minimum price to participate in the balancing market.

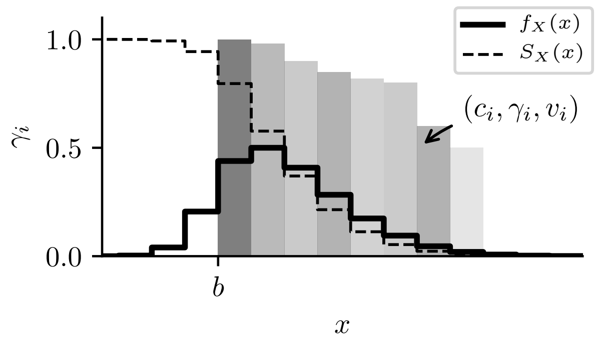


**Figure 3.2** Demand response model: bars represent prepared agents prior to demand realization. For demand realization $x > b$, the retailer requests prepared agents to respond sequentially (from left to right) until the positive imbalance $(x - b)$ is resolved. The ordering of the agents is an example ordering with regards to the response probability.

### 3.3.2 Demand response agents

We consider agents that are flexible and can reduce or increase their demand and consequently any positive imbalance between the procured quantity and the realized demand, after demand realization and before the time of delivery, during the *response* period (see Figure 3.1).[2]

Let $A = \{0, 1, \ldots, n-1\}$ denote the finite set of demand response agents. Let also $d_i \in \{-1, +1\}$, $\forall i \in A$ denote the flexibility of agent $i$, i.e., for $d_i = -1$ agent $i$ can reduce its demand by one unit, while for $d_i = +1$ agent $i$ can increase its demand by one unit. We assume unit downward flexibility, i.e., $d_i = -1, \forall i \in A$. Later in this chapter (see Section 3.4.4) we also consider the case of both downward and upward unit flexibility where $d_i \in \{-1, +1\}$.

We follow previous work to define the timing of demand response and additional characteristics of agents with regards to their ability to respond (Ma et al., 2016). The type of agent $i$, $\theta_i$, is the triplet $(c_i, \gamma_i, v_i)$. Prior to demand realization and during the *preparation* period (see Figure 3.1), agent $i$ decides whether to prepare with preparation cost $c_i \geq 0$. After demand realization and during the response period, if agent $i$ is prepared, it is able with probability $\gamma_i \in (0, 1]$ to respond. If agent $i$ is

[2] Demand response retail market programs take place in short time periods (e.g.,15-min) and are based on time-ahead "realization" of the demand (Wang et al., 2015), i.e., when demand is very close to the real value. Real-time imbalances are not handled by demand response agents since this requires time (notification and response), but instead by automatic generation control or spinning reserves.

able to respond, it can decide either to respond with response cost $v_i \geq 0$, or not without any cost. The decision of agent $i$ to prepare and the ability to respond cannot be observed, while the response can be observed.

### 3.3.3 Model of demand response

In this section we discuss how demand response agents can be used by the retailer. Consider a set of agents that decide to prepare prior to demand realization, and demand realization $x > b$, i.e., positive imbalance. During the response period (see Figure 3.1), the retailer requests prepared agents to respond (reduce their demand) in some order; each agent that responds reduces the imbalance quantity $(x - b)$ by one unit. The response of an agent can be observed before the next agent is asked to respond. If the imbalance is resolved (enough agents have responded), the retailer stops requesting agents to respond. Otherwise, if imbalance is not resolved (there are no more prepared agents to respond), the retailer pays the remaining imbalance quantity with price $p'$ at the time of delivery.

**Example 3.2.** *Figure 3.2 presents the demand forecast $f_X(x)$, the dashed curve illustrates the survival function $S_X(x) = \mathbb{P}_X(X > x)$. Bars of different color intensity (darker means higher total cost for demand response) represent prepared (prior to demand realization) agents starting from the procured quantity of the retailer $b$. The height of the bars show the probability that agent $i$ is able to respond, $\gamma_i$, and bars' width show the quantity that each agent can reduce its demand. For demand realization $x > b$, the retailer requests sequentially (from left to right) prepared agents to respond until imbalance $(x - b)$ is zero or no more agents can be requested.*

In contrast to related work that does not consider the realization of the demand (Ma et al., 2016, 2017), in the following sections we design mechanisms that request agents to respond only if there is a positive imbalance; therefore, the probability that a selected agent is requested to respond is not equal to one but it is influenced by: (i) the demand forecast $f_X(x)$, (ii) the order in which is asked to respond, and (iii) the response probabilities of preceding agents.

## 3.4 Demand Response Mechanism $M$

In this section we define the general mechanism $M$ in which selected agents are asked to respond sequentially to reduce a positive imbalance. In Section 3.4.1 we compute the probability that agents are requested to respond, the expected utility of both agents and the retailer (mechanism), and we analyze dependencies between agents that arise in our setting. Sections 3.4.2 and 3.4.3 outline our proposed mechanisms. Last, Section 3.4.4 generalizes our proposed mechanisms to the case where both positive and negative imbalances incur balancing cost to the retailer.

Recall that $p'$ is the imbalance price, $X$ is the random variable of the demand, $b$ is the procurement quantity, and $\theta_i$ is the type of agent $i$. We define the general mechanism $M(X, b, p', \hat{\theta}) \rightarrow (s_i, o_i, r_i, t_i)$, $\forall i \in A$, in which all available agents

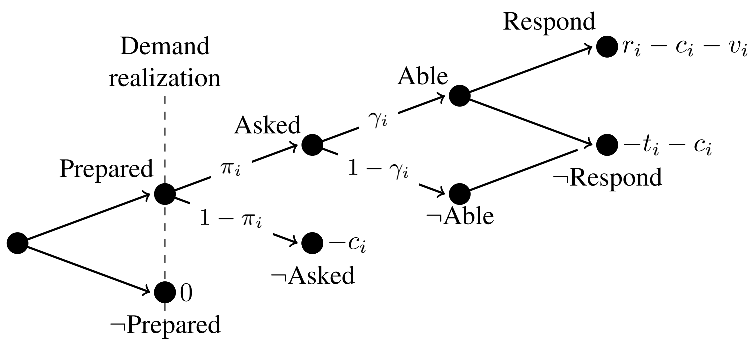


**Figure 3.3** Schematic sequence of agent's $i$ decisions (preparation and response), stochastic transitions (the probability of which is placed midway in arrows), and costs (illustrated on the right side of end nodes). Reward $r_i$ is transferred from the mechanism to agent $i$ in case of response, penalty $t_i$ is paid to the mechanism otherwise.

report their types $\hat{\theta} = \{\hat{\theta}_0, \hat{\theta}_1, \ldots, \hat{\theta}_{n-1}\}$ ($\hat{\theta}_i$ is the reported type of agent $i$) to the mechanism $M$ during the preparation period (see Figure3.1).

**Assumption 3.1.** *The retailer does not have access to the available flexibility (reports $\hat{\theta}$) during the ahead period, and thus the procurement quantity $b$ is already determined before the preparation period.*

The quadruplet $(s_i, o_i, r_i, t_i)$ is the resulting allocation for agent $i$, where $s_i \in \{0, 1\}$ denotes the selection of agent $i$, $o_i \in \mathbb{Z}_0^{n-1}$ the order in which agent $i$ may be requested to respond by the mechanism, $r_i \geq 0$ the reward that is transferred from the mechanism to agent $i$ in case agent $i$ responds after is requested by the mechanism, and $t_i \geq 0$ the penalty that is paid to the mechanism if agent $i$ does not respond after request. Payments from and towards the mechanism take place after the realization of the demand and after observing the response of agents if requested by the mechanism (*contingent payments*).

Consider agent $i$ that is selected by the mechanism ($s_i = 1, o_i \in \mathbb{Z}_0^{n-1}$) with reward $r_i$ and penalty $t_i$. Figure 3.3 illustrates the general mechanism $M$, where $\pi_i$ denotes the probability that agent $i$ is requested to respond by the mechanism (see Section 3.4.1). With knowledge of $r_i$ and $t_i$, agent $i$ decides whether to prepare prior to the realization of the demand (preparation period). During the response period, the mechanism asks agent $i$ to respond with probability $\pi_i$. If agent $i$ is able to respond (with probability $\gamma_i$), it decides whether to respond during the response period.

### 3.4.1 Request probability & interdependencies

In this section we compute the probability $\pi_i$ with which mechanism $M$ requests agent $i$ to respond. We further compute the utilities of both agents and the retailer (mechanism) under mechanism $M$. Last, we study implications in mechanism design that arise by dependencies between agents that are requested to respond sequentially.

#### Probability of response request

To compute the probability $\pi_i$ we consider that agents report their true types $\theta$. Consider agent $i$ with unit flexibility (recall that $d_i = -1, \forall i \in A$) and type $\theta_i = (c_i, \gamma_i, v_i)$. We assume w.l.o.g. that $s_i = 1, o_i = i$, $\forall i \in A$, i.e., all agents are selected and the order that are requested to respond follows the indexing of agents. We further assume that all agents prepare and respond if requested (agent $i$ is able to respond with probability $\gamma_i$). Let $a_i \in \{0, 1\}$ denote the observed action of agent $i$ after request, it is equal to one in case of response and zero otherwise.

$$\pi_i = \begin{cases} S_X(b), & i = 0 \\ S_X(b+i) + \sum_{k=0}^{i-1} f_X\big(b+k+1\big)\mathbb{P}\Big(\sum_{j=0}^{i-1} a_j \leq k\Big), & i > 0 \end{cases} \tag{3.2}$$

is the probability of response request from the mechanism to agent $i$, where $\mathbb{P}(\sum_j a_j \leq k)$ is the probability that less than or equal to $k$ agents respond from the agents preceding $i$, i.e., $\forall j < i$. For $i = 0$, $\pi_i$ is equal to the probability that demand is larger than $b$, i.e., the first agent in the order is always asked to respond if there is positive imbalance. For $i > 0$, equation (3.2) further accounts for failures (inability to respond) of agents preceding $i$ in case $x < (b + i + 1)$.

The quantity $\sum_j a_j$ in equation (3.2) is the sum of independent Bernoulli variables and follows a Poisson binomial distribution (Wang, 1993). The probability $\mathbb{P}(\sum_j a_j \leq k)$ is the cumulative distribution function of a Poisson binomial distribution for $k$ successes.

$$\mathbb{P}\Big(\sum_j a_j \leq k\Big) = \sum_{l=0}^{k} \sum_{L \in F_l} \prod_{q \in L} \gamma_q \prod_{m \in L^c} (1 - \gamma_m), \tag{3.3}$$

where for each number of successes $l \in [0, k]$, $F_l$ contains all sets of size $l$ in the powerset of $A$ and $L^c$ is the complementary set of $L$, i.e., $L \cup L^c = A$. For experiments presented later in this chapter (see Section 3.5), we compute the probability in equation (3.3) using a closed-form expression based on the discrete Fourier transform of the characteristic function of the distribution (Hong, 2013).

#### Utilities under mechanism $M$

Let $C_{DR}$ denote the expected cost for the mechanism $M$ (expected balancing cost plus the cost for demand response) under allocation $(s_i, o_i, r_i, t_i), \forall i \in A$. To derive the cost $C_{DR}$, we assume that agents report their true types $\theta$ and w.l.o.g. that $s_i = 1, o_i = i$, $\forall i \in A$, i.e., all agents are selected and the order that are requested to respond follows the indexing of agents. We further assume that all selected agents prepare and respond if requested.

$$C_{DR} = \sum_{i \in A} \pi_i \gamma_i r_i - \pi_i (1-\gamma_i) t_i$$
$$+ p' \sum_{x=b+1}^{D} f_X(x) \sum_{k=0}^{\min\{x-b-1,n\}} \mathbb{P}\Big(\sum_{i \in A} a_i = k\Big)(x-b-k), \tag{3.4}$$

where the first term is equal to the expected payments towards and from the agents. The expected payment to agent $i$ is equal to: $\pi_i \gamma_i r_i$, and the expected income from agent $i$: $\pi_i(1-\gamma_i)t_i$. The second term of equation (3.4) computes the expected balancing cost for $x \in [b+1, D]$ and $k \in [0, \min\{(x-b-1), n\}]$ ($k$ is the number of responses) and it considers only the cases in which there is a remaining imbalance (thus all agents have been requested to respond). Note that, there is no imbalance for $k \geq (x-b)$, and $k$ cannot be greater than the number of agents ($k \leq n$).

We define the (retailer's) expected utility of the mechanism $U_M$ as the difference between the expected balancing cost when no demand response is used (see equation 3.1) and the expected cost for mechanism $M$ (see equation 3.4).

$$U_M = C_{\neg DR} - C_{DR}, \tag{3.5}$$

where the utility of the mechanism depends on the cost $C_{DR}$ that is influenced by the allocation. Note that $C_{\neg DR}$ depends only on the procurement quantity (see equation 3.1 and Assumption 3.1).

Similarly, given reward $r_i$ and penalty $t_i$ the expected utility of agent $i$ is:

$$u_i = \pi_i \gamma_i (r_i - v_i) - \pi_i (1-\gamma_i) t_i - c_i, \tag{3.6}$$

where $\pi_i$ is the probability of response request in equation (3.2). Agent $i$ pays $c_i$ to prepare. If agent $i$ is asked and able to respond, pays $v_i$ and gets reward $r_i$. Otherwise, if agent $i$ is asked and cannot respond, agent $i$ pays $t_i$ to the mechanism.

**Definition 3.1.** *A mechanism is called* dominant-strategy incentive compatible *(DSIC), or truthful, if no agent can increase its utility by misreporting its type to the mechanism, and* individually rational *(IR) if agents get non-negative utility in expectation (i.e., agents are willing to participate). Furthermore, a mechanism is called individual rational for the center (CR) if the center's (mechanism) expected utility is non-negative (Porter et al., 2008).*[3]

## Interdependent tasks with uncertain executions

As discussed earlier in this chapter, the allocation is determined by both the selection of agent $i$ and the order that the mechanism requests agent $i$ to respond. We showed that the latter influences the probability of response request to agent $i$ (see Section 3.4.1) and therefore tasks (each order in the allocation) are interdependent. The valuation of agent $i$ for a given allocation depends on the probabilities that pre-

[3] We adopt the notion of center's rationality (CR) considering the utility of the mechanism for the retailer; equivalently, we can say that mechanism $M$ satisfies CR if it is IR for the retailer to adopt mechanism $M$.

ceding agents are able to respond. However, there exist no efficient mechanism (in the class of Groves mechanisms) that satisfies DSIC, IR and CR when there are interdependencies between tasks with uncertain executions (Conitzer and Vidali, 2014; Porter et al., 2008; Zhao et al., 2016).

Given the above impossibility result, in Sections 3.4.2 and 3.4.3 we design mechanisms that select agents to perform demand response by removing dependencies between tasks and satisfy all DSIC, IR and CR properties.

## 3.4.2 Sequential-task mechanism

In this section we define the sequential mechanism $Seq_M$, which selects agents for each order in the allocation sequentially.

### Minimum acceptable reward

Given the expected utility of agent $i$ in equation (3.6), we define the *minimum acceptable reward* for agent $i$, $\varrho_i$: the minimum reward for which it is rational for agent $i$ to prepare prior to demand realization (during the preparation period) and respond if it is able with probability $\gamma_i$ during the response period. The minimum acceptable reward $\hat{\varrho}_i$ (based on report $\hat{\theta}_i$) that yields positive expected utility for agent $i$ is:

$$r_i \geq \frac{\pi_i(1-\hat{\gamma}_i)t_i + \hat{c}_i}{\pi_i\hat{\gamma}_i} + \hat{v}_i \triangleq \hat{\varrho}_i, \tag{3.7}$$

where $\hat{\varrho}_i$ is the lower bound of the reward $r_i$. We further set an upper bound for the reward $r_i$, it should not be larger or equal to the imbalance price $p'$ that the retailer pays in case agent $i$ does not respond (and no other agent responds), $r_i < p'$.

### Mechanism $Seq_M$

We define the sequential mechanism $Seq_M$ as follows: $Seq_M(X, b, p', \hat{\theta}, T) \to (s_i, o_i, r_i, t_i = T)$, $\forall i \in A$, where $T \geq 0$ is a fixed penalty. Let $A'$ denote the set of available agents (i.e., agents that are not yet selected) and $\kappa$ the order in which the mechanism requests an agent to respond. Initially, $A' = A$, $\kappa = 0$, and $s_i = 0$, $\forall i \in A$. We detail the steps of the mechanism $Seq_M$ below:

(1) Collect reports from all available agents, $\hat{\theta}_i$, $\forall i \in A'$.

(2) Compute $\hat{\varrho}_i$, $\forall i \in A'$ for $o_i = \kappa$ as in equation (3.7); $\pi_i$ is computed with regards to previously selected agents as in equation (3.2).

(3) Consider in order $\kappa$ agent $w = \arg\min_{i\in A'} \hat{\varrho}_i$ (lowest $\hat{\varrho}_i$), and reward $r_w = \min_{i\neq w\in A'} \hat{\varrho}_i$ (second lowest $\hat{\varrho}_i$).

(4) In case that $r_w < p'$, select agent $w$, i.e., $(s_w = 1, o_w = \kappa, r_w = \min_{i\neq w\in A'} \hat{\varrho}_i, t_w = T)$, and remove agent $w$ from the set of available agents, $A' = A' - w$. Then, go to step (1) and increase the order $\kappa$ by one, $\kappa = \kappa + 1$. For $r_w \geq p'$, stop without selecting agent $w$ ($s_w = 0$).

We consider that $Seq_M$ takes place during the preparation period (see Figure 3.1). At each round, the computed reward and the fixed penalty is communicated to the selected agent that decides whether to prepare before the demand realization and respond if it is requested. After demand realization and in case of positive imbalance, $Seq_M$ requests agents to respond sequentially according to the order that they are selected until imbalance is zero, or there are no more agents to respond. If agent $i$ is asked to respond, $Seq_M$ pays $r_i$ to the agent in case of response, and receives penalty $t_i$ otherwise.

#### Incentives and truthfulness of $Seq_M$

We proceed to discuss agents' incentives to report truthfully and participate in $Seq_M$, and further show that $Seq_M$ can only benefit by selecting agents to perform demand response.

**Assumption 3.2.** *Agents do not have access to: (i) the reports of other agents $\hat{\theta}_{-i}$, the number of participating agents, and the distribution of agents' types (no communication), (ii) the reward that is communicated to the selected agent after each round of the mechanism (no price discovery), and (iii) the demand forecast of the retailer.*

**Theorem 3.1.** *Given Assumption 3.2, $Seq_M$ is DSIC and IR.*

*Proof.* We first show that the mechanism is IR for the agents. For report $\hat{\theta}_i$, the minimum acceptable reward $\hat{\varrho}_i$ as it is computed in equation (3.7) yields zero expected utility for agent $i$, $u_i = 0$. Any reward $r_i \geq \hat{\varrho}_i$ yields positive utility $u_i \geq 0$. Let $\hat{\varrho}_j$ be the second lowest minimum acceptable reward. It holds by definition that $\hat{\varrho}_j \geq \hat{\varrho}_i$ and consequently $u_i \geq 0$. Therefore, it is rational for selected agent $i$ to prepare and respond if requested.

We proceed to show that a selected agent cannot improve its utility by misreporting to the mechanism. Given Assumption 3.2 and the definition of $Seq_M$ each round of the mechanism is an isolated Vickrey (second-price) auction (Vickrey, 1961), since there is no information propagating to the next rounds, i.e., *no externalities*.[4] In each round, agents deterministically choose to report truthfully to the mechanism since any round can be the last round in which they can obtain positive expected utility with reward lower than $p'$. □

Given Theorem 3.1, $\pi_i$ is computed in step (2) of $Seq_M$ (see Section 3.4.2) based on truthful reports.

**Proposition 3.1.** *For any fixed penalty $T \geq 0$, $Seq_M$ is CR.*

*Proof.* Note that rewards are lower than $p'$ and response is requested only if there is an imbalance that otherwise has to be paid with price $p'$. It follows that the utility of $Seq_M$ (see Section 3.4.1) is always non-negative. □

### 3.4.3 Independent-task mechanism

Recall that the probability of response request to agent $i$, $\pi_i$, depends on the response probabilities of preceding agents in the allocation. In this section we design a truthful combinatorial mechanism by removing dependencies between selected agents.

[4] See (Leme et al., 2012) for sequential mechanisms with externalities.

### Mechanism $Ind_M$

We define the mechanism $Ind_M$, which requests agents to respond as follows:

**Definition 3.2.** *For demand realization $x > b$, $Ind_M$ requests all agents up to order $\lambda = x - b - 1$ to respond.*

Intuitively, the mechanism $Ind_M$ asks agents up to a specific order, that corresponds to the imbalance quantity, to respond. It follows that, the probability that agent $i$ is asked to respond in equation (3.2) is independent of preceding agents ($o_j < o_i$) and $\pi_i = S_X(b + o_i)$.

Similarly to $Seq_M$ (see Section 3.4.2), $Ind_M$ fixes a penalty, $t_i = T, \forall i \in A$. In addition, $Ind_M$ fixes a reward $R$, $r_i = R, \forall i \in A$, which is the reward that the mechanism pays agent $i$ after response. Mechanism $Ind_M$ can be written as: $Ind_M(X, b, p', \hat{\theta}, R, T) \to (s_i, o_i, r_i = R, t_i = T, z_i)$, $\forall i \in A$, where $z_i$ is a payment from agent $i$ to the mechanism upon allocation. Each selected agent $i$ gets reward $r_i = R$ if it is requested and responds, and pays penalty $t_i = T$ in case of no response.

Unlike $Seq_M$ that may request all selected agents to respond to resolve some imbalance quantity, failing to respond under $Ind_M$ yields balancing cost $p'$ to the mechanism, since all selected agents have to respond to resolve any positive imbalance. This means that under $Ind_M$, fixing $T = p'$ "transfers" the balancing responsibility to selected agents.

### Optimal allocation & VCG payments

Given reward $R$, penalty $T$, and considering $z_i = 0$, the expected utility of agent $i$ for order $o_i$ and based on the report $\hat{\theta}_i$ is:

$$\hat{u}_i(o_i) = \pi_i \hat{\gamma}_i (R - \hat{v}_i) - \pi_i (1 - \hat{\gamma}_i) T - \hat{c}_i, \tag{3.8}$$

where $\pi_i = S_X(b + o_i)$ (see Definition 3.2). We define the optimal allocation $o^{opt}$ (assignment of each agent to an order) such that:

$$o^{opt} = \arg\max_{o \in O} \sum_{i \in A} \max\left(0, \hat{u}_i(o_i^{opt})\right), \tag{3.9}$$

where $O$ is the set of all permutations. The term $\max(0, \hat{u}_i(o_i^{opt}))$ ensures that no agent is selected with negative expected utility, therefore $s_i = 0$, if $\hat{u}_i(o_i^{opt}) \leq 0$.

**Lemma 3.1.** *The problem of finding the optimal allocation $o^{opt}$ in equation* (3.9) *can be solved in polynomial time $\mathcal{O}(n^3)$.*

*Proof.* When the probability of response request $\pi_i$ is independent of previously allocated agents, the optimal allocation problem can be formulated as the *linear assignment problem* (LAP), where $n$ agents are assigned $n$ tasks. In our setting each task stands for an order $o_i \in \mathbb{Z}_0^{n-1}$. It can be solved optimally in polynomial time, $\mathcal{O}(n^3)$, by the *Hungarian method* (Kuhn, 1955) while there exist implementations that can speed up further the computation of the optimal assignment in LAPs (Date and Nagi, 2016). □

To simplify notation in the remainder of this section, we define $u'_i(o_i^{opt}) \triangleq \max\big(0, \hat{u}_i(o_i^{opt})\big)$. We use the VCG payment rule to compute the payment of agent $i$ to the mechanism upon allocation:

$$z_i = \left( \sum_{j \in A \setminus i} u'_j\big(o_j^{opt_{-i}}\big) - \sum_{j \in A \setminus i} u'_j\big(o_j^{opt}\big) \right), \tag{3.10}$$

where $o^{opt_{-i}}$ is the optimal allocation without agent $i$ present (Nisan et al., 2007). Intuitively, $z_i$ is equal to the difference between the sum of utilities that agents other than $i$ get without agent $i$ present, and the sum of utilities they get under its presence (marginal loss). Note that, agents that are not selected cause zero marginal loss to other agents.

We consider that $Ind_M$ takes place during the preparation period (see Figure 3.1), where each selected agent $i$ pays $z_i$ upon allocation. After demand realization, $Ind_M$ requests selected agents up to the order $\lambda$. Selected agents in the set $\{j : o_j \leq \lambda\}$ are asked to respond, $Seq_M$ pays $R$ to agents that respond, and receives penalty $T$ otherwise.

**Theorem 3.2.** *The mechanism $Ind_M$ is DSIC and IR.*

*Proof.* First, it follows from Lemma 3.1 that $o^{opt}$ maximizes the sum of agents' utilities. By definition of the VCG mechanism (Nisan et al., 2007), $Ind_M$ is DSIC and IR. Agent $i$ maximizes its utility by reporting truthfully, $\hat{\theta}_i = \theta_i$, and it is rational for agent $i$ to prepare and respond if requested under reward $R$ and penalty $T$. □

**Proposition 3.2.** *For any $R \leq p'$ and $T \geq 0$, $Ind_M$ is CR.*

*Proof.* For any $R \leq p'$ and $T \geq 0$, no allocation can yield losses for the mechanism (see proof of Proposition 3.1). □

### 3.4.4 General flexibility mechanisms

In previous sections we have considered a setting where only positive imbalance from the procured quantity $b$ results in balancing cost for the retailer (see Section 3.3.1). In this section we show that mechanisms $Seq_M$ and $Ind_M$ can generalize in settings where both positive and negative imbalances result in balancing costs.

In addition to the set of downward flexibility agents $A$ (also denoted by $A^-$ in the remainder of this chapter), we consider the set of upward flexibility agents $A^+$, with reports $\hat{\theta}^+ = \{(\hat{c}_i, \hat{\gamma}_i, \hat{v}_i) : \forall i \in A^+\}$ and $d_i = +1, \forall i \in A^+$. Both $Seq_M$ and $Ind_M$ can be applied on the set $A^+$, independently from the set $A^-$, under the same imbalance price $p'$ or a different price for negative imbalances, with only a small adjustment of the request probability $\pi_i$. For $Seq_M$, the probability that agent $i$ is asked to respond $\pi_i$ for $i > 0$ in equation (3.2) becomes:

$$\pi_i = F_X(b - i - 1) + \sum_{k=0}^{i-1} f_X\big(b - k - 1\big)\mathbb{P}\Big(\sum_{j=0}^{i-1} a_j \leq k\Big), \tag{3.11}$$

where $F_X(x) = 1 - S_X(x) = \mathbb{P}_X(X \leq x)$. For $i = 0$, equation (3.2) becomes $\pi_i = F_X(b-1)$. Similarly, for $Ind_M$, $\pi_i = F_X(b - o_i - 1)$.

Both $Seq_M$ and $Ind_M$ hold their properties (DSIC, IR, and CR) since they are applied independently on different sets of agents.

# 3.5 Experimental Evaluation

In this section we empirically evaluate the performance of the proposed mechanisms $Seq_M$ and $Ind_M$. We also provide an evaluation of the mechanisms proposed by Ma et al. (2016) (in our extended setting), which we detail below:

**Fixed-Reward/Penalty Ma et al. Mechanisms** Fixed-reward Ma et al. mechanism fixes a reward $R$, a reduction target $Z$ and a reliability target $\tau$. For agent reports $\hat{\theta}$ (reports are the same as in Section 3.4), it computes the maximum penalty that each agent is willing to pay such that the agent retains non-negative utility (similarly to equation (3.7) that computes the minimum acceptable reward). Then, the mechanism sorts agents in a decreasing order with regards to their maximum willingness to pay, and selects the minimum number of agents such that: $\mathbb{P}(\sum a \geq Z) \geq \tau$, i.e., the probability of reaching the reduction target $Z$ is higher than the reliability target $\tau$. The mechanism sets penalty (in case of unsuccessful response) for each selected agent $i$ that is equal to the smallest willingness to pay from the set of agents that would have been selected without agent $i$ present. Intuitively, that is the lowest willingness to pay for agent $i$ to get selected. After allocation, the mechanism asks all selected agents to respond. Similarly, fixed-penalty Ma et al. mechanism fixes penalty $T$ and computes the minimum acceptable reward (as in equation (3.7) for $\pi_i = 1$) (Ma et al., 2017). Both fixed-reward and fixed-penalty mechanisms make a *deep market* assumption, i.e., there are enough agents in the market (economy) to fulfill the requirements with regards to the reduction target and the reliability.

Fixed-reward and fixed-penalty Ma et al. mechanisms do not consider the balancing responsibility of the retailer after the realization of the demand. Hence, the following sections do not serve as a direct comparison of the mechanisms proposed in this chapter and the mechanisms proposed by Ma et al. (2016); instead, they focus on the added value of information considered by our mechanisms (i.e., demand forecast, imbalance price), and the advantage of only requesting agents to respond after the realization of the demand.

**Experimental setup** We consider a market with $n = 200$ agents. Each agent has preparation cost $c_i \sim \mathcal{U}[0, p']$, response probability $\gamma_i \sim \mathcal{U}[0.5, 1]$, and response cost $v_i \sim \mathcal{U}[0, p' - c_i]$, where $\mathcal{U}[\alpha, \beta]$ denotes a uniform distribution from $\alpha$ to $\beta$. Note that $c_i + v_i \leq p'$, i.e., the sum of the costs for preparation and response is lower or equal to the imbalance price $p'$, which is a relevant assumption for the setting. For the demand distribution we use a discretized skew normal distribution $\mathcal{N}(\mu_X, \sigma_X, \alpha_X)$, where $\mu_X = 500, \sigma_X = 100, \alpha_X = 10$ (e.g., see Figure 3.2). The procurement quantity is set to $b = \mathbb{E}_X[x]$, and $p' = 0.6$. Our results are averaged over 200 independent runs where the demand distribution is fixed.

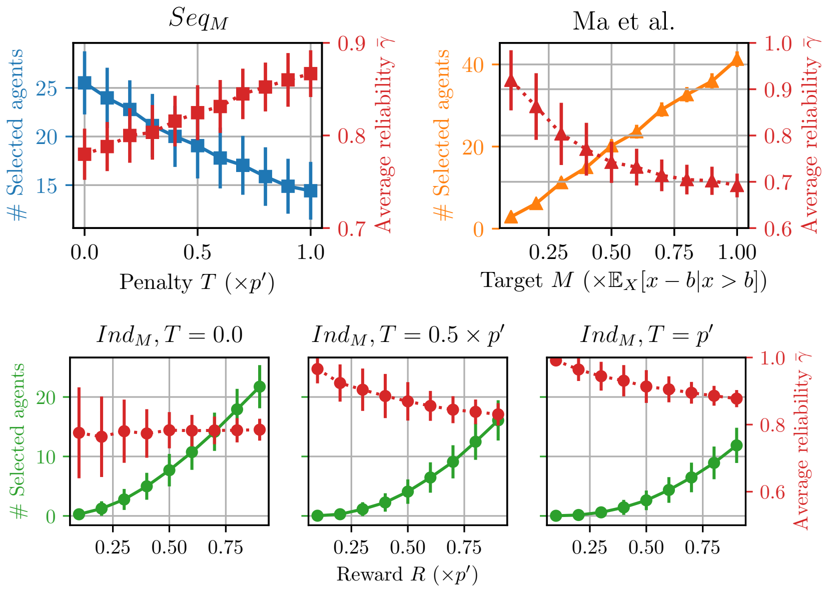


**Figure 3.4** Mean and standard deviation of the number and the average reliability of selected agents. Continuous lines correspond to values on left vertical axes and dotted lines to right vertical axes. $Seq_M$ mechanism (top left), $Ind_M$ mechanism (bottom row), and Ma et al. mechanisms (top right).

**Mechanism parameters** For the mechanism $Seq_M$, we use fixed penalty $T \in \{0.1, 0.2, \ldots, 1.0\} \times p'$. For $Ind_M$, we use $R \in \{0.1, 0.2, \ldots, 0.9\} \times p'$ and $T \in \{0, 0.5, 1\} \times p'$. Furthermore, for fixed-reward Ma et al. mechanism, we use $R \in \{0.4, 0.5, \ldots, 2.0\} \times p'$ (for $R < 0.4 \times p'$, negative penalties are induced to selected agents by the mechanism). For fixed-penalty Ma et al. mechanism, we use $T \in \{0.0, 0.1, \ldots, 2.0\} \times p'$. Last, for both fixed-reward and fixed-penalty Ma et al. mechanisms, we use reduction target $Z \in \{0.1, 0.2, \ldots, 1.0\} \times \mathbb{E}_X[x-b|x>b]$ and reliability $\tau = 0.95$.

**Number of selected agents & average reliability** Figure 3.4 shows the number and the average reliability $\bar{\gamma}$ of selected agents under $Seq_M$, $Ind_M$, and Ma et al. mechanisms. The number of selected agents in both Ma et al. mechanisms is influenced by the reduction target $Z$ and the reliability target $\tau$, and not by the fixed-reward $R$ and fixed-penalty $T$. For $Seq_M$ and $Ind_M$, reward $R$ and penalty $T$ affect the number of selected agents. For $T = 0$, $Seq_M$ selects approximately 25

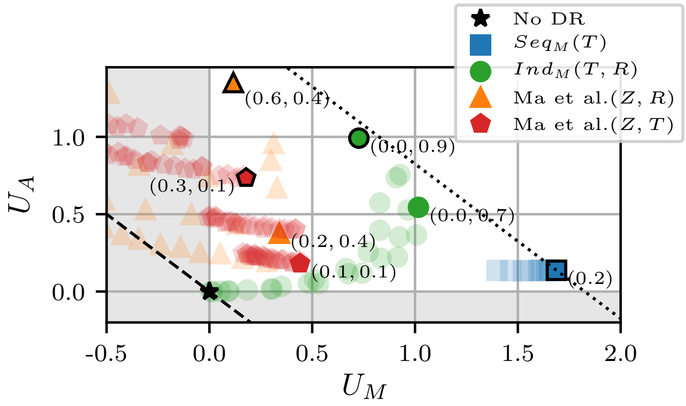


**Figure 3.5** Expected utilities of the mechanism $U_M$ and the agents $U_A$ under all mechanisms for a wide range of parameters. For all mechanisms, we vary reward $R$ and penalty $T$ with regards to the imbalance price $p'$. For Ma et al. mechanisms, the reduction target $Z$ is shown with regards to the expected imbalance quantity $\mathbb{E}_X[x - b|x > b]$.

agents, this corresponds to a reduction target $Z = 0.6 \times \mathbb{E}_X[x - b|x > b]$ for Ma et al. mechanisms. For $T = p'$, $Seq_M$ selects on average 15 agents. The number of selected agents for $Ind_M$ is lower than for $Seq_M$ since $Ind_M$ asks agents up to the quantity of the imbalance to respond (see Definition 3.2); $Ind_M$ does not count for possible failures of agents that otherwise would increase the probabilities of response requests and consequently select more agents. As anticipated, the average reliability $\bar{\gamma}$ is influenced by the reward $R$ and penalty $T$ parameters. For higher penalty $T$, fewer agents with higher reliability are selected by our proposed mechanisms.

**Social welfare & balancing cost** Figure 3.5 illustrates the utility space of the mechanism (retailer) $U_M$ and the agents $U_A = \sum_i u_i$, on the horizontal and the vertical axis respectively. The star marker shows the case when no demand response is used, and thus the mechanism pays positive imbalances with price $p'$. For every drawn set of agents out of 200 independent runs, we compute the analytical expected utility under each mechanism based on equation (3.4). For all mechanisms, the solid color marker shows the point where the utility of agents ($U_A$) is maximum, and the solid marker with black colored borders shows the point where social welfare ($U_M + U_A$) is maximum. The parameters used for the mechanisms are shown in parentheses, where target $Z$ is multiplied with the expected positive imbalance

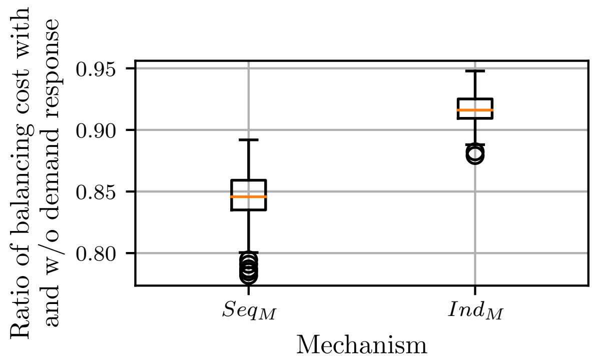


**Figure 3.6** Ratio of expected balancing cost (absolute imbalances case) for the mechanism with and without the use of demand response (lower is better).

$\mathbb{E}_X[x - b|x > b]$ and $R, T$ with the imbalance price $p'$. Transparent markers show points in the utility space for parameters that are not shown in the figure.

The shaded area illustrates the utility space where either the mechanism, the agents, or both have negative utility in expectation (when compared to the case of no demand response). In comparison to Ma et al. mechanisms that only consider incentives for the agents (satisfy IR for participating agents), both $Seq_M$ and $Ind_M$ guarantee non-negative expected utility for both agents and the mechanism (both satisfy IR and CR) since they consider both the demand forecast and the balancing cost of the mechanism.

Next, we evaluate all mechanisms with regards to the utility of the mechanism $U_M$ and the social welfare $(U_M + U_A)$. Parallel lines in Figure 3.5 illustrate points of equal social welfare, the dashed line for the case of no demand respond, and the dotted line for the maximum social welfare under both our proposed mechanisms (almost equal): $Seq_M$ ($T = 0.2$) and $Ind_M$ ($T = 0$, $R = 0.9p'$). When compared to the case of no demand response, the expected social welfare increases by $14\%$ for $Seq_M$ ($T = 0.2p'$), $13\%$ for $Ind_M$ ($T = 0$, $R = 0.9p'$), $11\%$ for fixed-reward Ma et al. ($Z = 0.6 \times \mathbb{E}_X[x - b|x > b], R = 0.4p'$), and last, $6\%$ for Ma et al. with fixed penalty ($Z = 0.3 \times \mathbb{E}_X[x - b|x > b], T = 0.1p'$). The utility of the mechanism increases (i.e., expected balancing cost decreases) by: $13\%$ for $Seq_M$ ($T = 0$), $7\%$ for $Ind_M$ ($T = 0, R = 0.7p'$) and $2 \sim 3\%$ for both fixed-reward and fixed-penalty Ma et al. mechanisms. Compared to Ma et al. mechanisms in this extended setting, $Seq_M$ and $Ind_M$ improve both social welfare and the utility of the mechanism since they request agents to respond only if there is positive imbalance.

**Simultaneous upward & downward flexibility** Last, we show that both $Seq_M$ and $Ind_M$ reduce balancing costs substantially for the retailer in the case where both positive and negative imbalances from the procured quantity incur balancing cost to the retailer.

We consider that any absolute deviation from the procurement quantity ($b = \mathbb{E}_X[x]$) is balanced with price $p'$. As described in Section 3.4.4, $Seq_M$ and $Ind_M$ can be used to allocate both upward and downward flexibility agents. We draw equal number of both types of agents, $|A^-| = |A^+| = 200$. For $Seq_M$, we use $T = 0$. For $Ind_M$, $R = 0.6p'$ and $T = 0$. We keep the distribution of agent types and the demand distribution same as those of earlier experiments. Figure 3.6 presents the ratio of the expected balancing cost for $Seq_M$ and $Ind_M$ with and without demand response ($C_{DR}/C_{\neg DR}$). On average, $Seq_M$ ($T = 0$) mechanism achieves a $16\%$ reduction in the balancing cost of the mechanism, while $Ind_M$ ($R = 0.6p', T = 0$) yields $9\%$ reduction.

# 3.6 Conclusions

In this chapter we studied a highly relevant problem in electricity systems: how to incentivize uncertain demand response under a given demand forecast and imbalance price. We proposed two mechanisms: a sequential mechanism ($Seq_M$) that is truthful under some mild assumptions (see Theorem 3.1), and a truthful combinatorial mechanism ($Ind_M$) that runs in polynomial time and uses VCG payments (see Theorem 3.2). Both mechanisms require only a subset of selected agents to respond, while they guarantee non-negative utility in expectation for both agents and the retailer (mechanism). The proposed mechanisms can further be used in settings where both positive and negative imbalances result in balancing cost for the retailer. Last, we verified the theoretical properties of both mechanisms in an empirical evaluation over different parameters. Our proposed mechanisms achieved up to $16\%$ reduction in the balancing cost of the retailer and $14\%$ increase in social welfare compared to when no demand response is used.

Overall, the proposed mechanisms can be used in practice to utilize the flexibility of small-scale demand response agents, instead of reserve capacity that is traditionally purchased in balancing markets by retailers, and replace or complement demand response emergency programs with large capacity customers. The societal impact of the proposed mechanisms is that less reserve capacity is traded, thus abating the use of fast-ramping generation that is responsible for excessive $CO_2$ emissions. This can further contribute to reducing electricity prices for customers and improving the stability of systems in cases of islanding grids that depend only on local renewable generation.

## 3.6.1 Future work

The presented work provides a basis for future extensions, which can be achieved by relaxing some of the assumptions, or by extending the overall problem setting. For instance, in Section 3.4 we consider that the procurement decision of the retailer is taken without considering the reports (for demand response) of agents. It is of interest to show how reports for demand response affect the procurement decision of the retailer.

Furthermore, in Remark 3.1 we delineated our motivation for considering unit responses. This was based on realistic settings of electricity markets (e.g., ahead and balancing markets) that have trading volumes that are multiples of a unit quantity, while the same applies to markets with indivisible goods. A direct extension of this work is to consider multi-unit demand response. Although our model can be generalized to continuous variables if we neglect the need for decimal reduction or there is a minimum price $p'$ to participate in the balancing market, the derived mechanisms do not generalize to the case of continuous flexibility.

Assumption 3.2 might also be restrictive for the practical implementation of the sequential mechanism $Seq_M$. By dropping Assumption 3.2 we have a sequential setting in which information (externalities) from each round of the mechanism propagates and consequently affects the strategic decisions of agents in later rounds. More specifically, the price of anarchy (PoA) in sequential second-price auctions with externalities can be arbitrarily worse than the optimal (Leme et al., 2012), sequential first-price or options-based auctions are better-suited in such settings since they have some performance guarantees.

Last, in this chapter we have considered mechanisms that request agents to respond only if there is imbalance, and we have studied the implications that arise by dependencies between agents that are requested to respond sequentially. Future research can study the design of mechanisms that request random sets of agents to respond in order to remove dependencies between agents in the allocation.

# 4 SLA Allocation for Renewable Electricity Trading

**Preface** In this chapter we propose to adopt service-level agreements (SLAs) for electricity trading in systems where supply depends on renewable electricity sources and therefore delivery cannot be guaranteed. These service-level agreements (SLAs) comprise the following features: quantity, reliability, and price. First, we define a characterization of the value degradation of tolerant and critical buyers with regards to the uncertainty of electricity delivery, generalizing the widely used value of lost load (VoLL) in our settings. Next, we propose two mechanisms to allocate these SLAs to buyers using either a *sequential second-price auction* or the combinatorial *Vickrey-Clarke-Groves* (VCG) mechanism, and we discuss the settings in which truthfulness can be obtained in the sequential setting. We empirically compare their performance and demonstrate that the proposed mechanisms dominate alternative allocations that use only the VoLL, and vastly improve the efficiency of the studied system. Overall, this chapter contributes an essential component to the future smart grid by facilitating distributed energy trading under uncertainty.

*This chapter presents work that was published as an extended abstract in the proceedings of the International Conference of Autonomous Agents and Multiagent Systems (AAMAS) 2017 in Sao Paulo, Brazil (Methenitis et al., 2017), and presented at the co-located workshop AMEC/TADA 2017. The full version of this work was published in Springer-Open Energy Informatics Journal (Methenitis et al., 2018).*

# 4.1 Introduction

Previous chapters of this thesis propose methods (risk-sharing tariff and mechanisms to incentivize demand response) to tackle the inherent uncertainty on the demand-side and therefore alleviate excessive costs related to imbalances between supply and demand (see Chapters 2 and 3). This chapter focuses instead on the intermittent nature of renewable electricity supply. Natural sources, such as the sun and wind, are subject to stochastic availability and non-dispatchable: their output cannot be regulated to match the demand. In envisioned smart grid systems the increasing reliance on electricity from natural sources implies higher risks of shortages or overproduction since such generation is volatile and locally highly correlated. That is exemplified in microgrids that solely build their electricity supply on local renewable generation (from prosumers and smaller-scale producers) in islanding scenarios of the smart grid.

In contrast to conventional centralized systems in which demand is always satisfied through reserve electricity generation and balancing markets, future smart grid systems assume the utilization of local and renewable electricity supply. In such systems, local exchange (trading) of electricity becomes challenging since delivery cannot be guaranteed. Current electricity tariffs promise certain delivery, and are thus not well-suited to trade these uncertain quantities. However, if not traded the electricity might need to be curtailed, foregoing potential benefits for both supply and demand sides.

To this end, in this chapter we interpret *service-level agreements* (SLAs) as a direct extension of conventional electricity tariffs, which ensure delivery (100% quality) and a fixed kWh price (0% risk). We build on the definition of SLAs as agreements between service providers and service users, specifying the service and its characteristics (Verma, 1999). In contrast to the current straightforward contracts for electricity, SLAs can be extended to include more features (e.g., delivery time, reliability, penalty for no delivery) in order to provide the contracting framework for balancing volatile supply with demand between buyers (e.g., customers participating in retail tariff schemes) and sellers of electricity (e.g., small-scale producers that base their generation portfolio on renewable electricity sources). We specifically study SLAs that comprise the following features:

**Quantity** The quantity of electricity that is subject to be transferred from the service provider to the service user.

**Reliability** The probability of successful delivery of the quantity of electricity that is specified in the SLA.

**Price** The price per unit of the transferred quantity.

The SLAs described above provide a fundamental contracting framework for electricity trading between buyers and sellers in settings where the availability of electricity is not certain, and buyers' ability to cope with uncertainty vary. We proceed to illustrate the concepts discussed so far in the following motivating example:

**Example 4.1.** *A seller holds a prediction of its generation for the next day during 1pm-2pm. The generation is not certain; there is* $90\%$ *probability that the seller* *generates at least* $1$ *unit and* $50\%$ *probability for at least* $2$ *units of electricity. There are two buyers that both have unit demand: the first buyer is a hospital that needs to perform a task, and the second buyer is an electric vehicle (EV) that needs to charge its battery. Assuming that there is no other seller in the system we consider, the two buyers can agree on SLAs of* $90\%$ *or* $50\%$ *reliability with the seller for the unit demand they require. Considering that the importance of the task the hospital needs to complete is higher than the EV's, it is socially optimal to assign the most certain unit of generation to the hospital.*

Under the presence of intermittent supply and SLAs that can prioritize the delivery of electricity to different buyers, it is socially desirable that reliable electricity is allocated to critical demand; consequently, the risk of load-shedding is assigned to less critical buyers that in turn perform this task at lower social cost. The widest adopted concept to measure criticality in the literature as well as in practice is the *value of lost load* (VoLL) (Kariuki and Allan, 1996). The VoLL is defined as the estimated amount that customers receiving electricity through contracts would be willing to pay to avoid a disruption in their electricity service. Revisiting Example 4.1, we can identify that the VoLL for the hospital is higher than for the EV. In this chapter we define a generalized value function that extends the concept of VoLL in our setting. We further study the problem of allocating the outlined SLAs (efficiently) to buyers with different preferences with regards to the uncertainty of electricity delivery, given a seller with uncertain electricity generation. The seller holds a prediction of its generation output (distribution), while the process of specifying and allocating SLAs to buyers participating in such an electricity market can be structured as a mechanism design problem (Nisan and Ronen, 2001).

The main contributions of this chapter can be summarized as follows:

- We define a contracting framework through SLAs that enables electricity trading under uncertain supply.
- We propose a family of exponential functions that characterizes buyers' varying degrees of criticality, thus generalizing the value of lost load with costs associated to the risk of failed delivery.
- We propose two mechanisms to allocate these SLAs to strategic buyers using either a *sequential second-price auction* or the combinatorial *Vickrey-Clarke-Groves* (VCG) mechanism.
- We discuss the settings in which truthfulness can be obtained in the sequential second-price auction.
- Last, we empirically show that the efficiency of the studied system vastly improves in face of buyers with varying abilities to cope with uncertainty.

The remainder of this chapter is organized as follows: in Section 4.2 we present related literature. In Section 4.3 we first formulate the problem of electricity trading under the presence of intermittent supply, we then introduce the structure of the proposed framework through SLAs, and last, we propose a value function that determines different types and preferences of buyers. In Section 4.4 we discuss mechanisms to allocate SLAs to buyers, and we examine incentive compatibility issues arising in the studied setting. In Section 4.5 we evaluate our proposed method through simulations. Last, in Section 4.6 we conclude this chapter and further discuss interesting research directions.

# 4.2 Related Work

Similarly to product (service) *differentiation* in power systems engineering and energy markets (Oren and Smith, 1993; Salah et al., 2017), we propose SLAs that differentiate in the reliability of successful delivery. Product differentiation in electricity trading also includes but is not limited to: (i) deadline-differentiated trading, where customers agree to defer the service of pre-determined loads in exchange for lower prices (Bitar and Xu, 2017), (ii) varying probabilities of electricity service interruptions (Chao et al., 1986), (iii) reliability differentiation for trading spinning reserve capacity (Siddiqi and Baughman, 1995). While all previous works study reliability differentiation in the electricity service with regards to optimal pricing or scheduling policies, our work is more related to the mechanism design approach proposed by Bitar and Xu (2017). In contrast to this work, however, we study the problem of assigning all the available supply to buyers of different levels of criticality and not the scheduling problem of deferrable loads to match the available supply.

Service-level agreements have been used as a tool for monitoring and coordination to ensure trustworthiness between different stakeholders, primarily with regards to business processes (Gustavsson et al., 2011; Hussain et al., 2012), or as negotiation protocols (Amato et al., 2014). In contrast to our work, the discussion remains conceptual, and no quantitative implications on costs or efficiency are given. Service-level agreements have further been used in resource allocation problems in computational grids to ensure the optimal allocation of computational resources and fair satisfaction of the participants through negotiations (Silaghi et al., 2012). In this work the embedding is not in the energy domain while the focus is on strategic negotiation rather than the elicitation of truthful reports.

In the closest to ours work, Bitar et al. (2012) study the viability of selling uncertain quantities of wind generation with variable-reliability. The authors further explore the connection between uncertainty in the generation and the costs for reserve capacity, and real-time markets. In a similar problem setting, Dash et al. (2007) studies task allocation market mechanisms for multiple suppliers of finite or uncertain capacity. Last, a recent work that considers SLAs for trading uncertain quantities of wind-generated electricity studies the optimization of electricity flow between wind electricity generators and smart grid customers using SLAs (Hussain et al., 2018). We follow a similar idea; however, we focus our attention on the characterization of the demand with respect to its criticality, as well as the design of the mechanisms to assign demand through SLAs for electricity trading to strategic buyers with different preferences.

To the best of our knowledge, the work presented in this chapter is the first that adapts SLAs for electricity trading under uncertain supply from a mechanism design perspective, providing a discussion with regards to buyers' incentives to truthfully report their preferences followed by an empirical evaluation in illustrative settings.

# 4.3 SLA Contracting Framework

In this section we propose an SLA contracting framework, in which buyers (electricity customers) can participate by purchasing quantities of electricity through contracts of a specific quantity, reliability, and price from a seller that relies its available supply on renewable electricity generation. This section is structured as follows: in Section 4.3.1 we outline the problem setting. and in Section 4.3.2 we define the basic structure of the proposed SLAs. Last, Section 4.3.3 focuses on the characterization of buyers' preferences with regards to the reliability of the SLAs.

## 4.3.1 Problem formulation

Our fundamental problem setting considers a marketplace with a single seller of electricity and multiple buyers (electricity customers). We further assume that there is no outside option for buyers, and therefore all demand should be satisfied by the seller. Such setting may resemble practical scenarios of envisioned smart grid systems, in which a prosumer may act as a seller that seeks to trade it potential excess generation (that cannot be privately used or stored) to potential buyers locally. Let $B$ denote the set of buyers, there are $n$ buyers in the set $B$, such that $B = \{1, \ldots, n\}$.

We assume that the seller holds a prediction (probability distribution function) ahead of the delivery time in the *ahead* timestep, while its actual generation output is known at the delivery time in the *realization* timestep. Let $Q$ denote the random variable of the available supply, and $q \in \mathbb{R}^+$ the actual realization of the supply. The prediction that is known by the seller in the ahead timestep is denoted by $f_Q(q)$, which denotes the probability density function of the random variable $Q$. We further consider the cumulative density function of the random variable $Q$, $F_Q(q)$, where

$$F_Q(q) = \mathbb{P}(Q \leq q) = \int_0^q f_Q(x)\, dx.$$

**Remark 4.1.** *Our two-step time model serves as a fundamental model of the day-ahead auction process in current electricity markets, and it can also be applied at different horizons. The mechanisms outlined in later sections of this chapter are also applicable for shorter time-horizons, e.g., hour(s) ahead, that may be required to facilitate electricity trading in microgrids with flexible power-to-heat demand.*

Each buyer $i$ from the set $B$ has a demand for electricity $d_i$, which we assume is fixed and known by the buyer ahead of time, considering hardware assets that induce a deterministic demand. The mechanisms proposed later in this chapter, based on the prediction of the available supply in the realization timestep allocate SLAs of some quantity, reliability and price, to participating buyers. The observed realization of the supply $q$ determines how much load can be served, and consequently the set of buyers that are indeed served such that the SLAs are satisfied in expectation with regards to their reliability.

## 4.3.2 Service-level agreements

As outlined in the introduction of this chapter (see Section 4.1), an SLA is a triplet $(d_i, \gamma_i, p_i)$, which comprises the quantity $d_i$, the reliability $\gamma_i$, and the price $p_i$ per transferred unit of electricity for buyer $i \in B$. For the remainder of this chapter

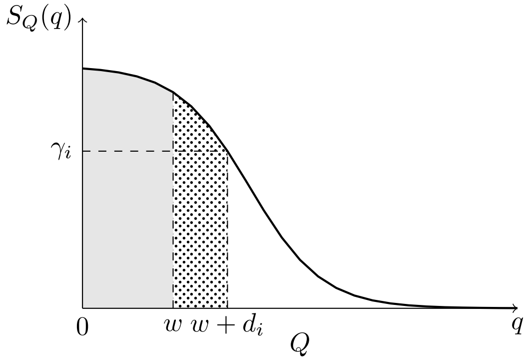


**Figure 4.1** The reliability function of the seller. The thick line illustrates the reliability function $S_Q(q) = \mathbb{P}(Q > q)$ of the random variable $Q$. The dotted area represents the portion of demand $d_i$ of buyer $i$ with reliability $\gamma_i = S_Q(w + d_i)$. The gray shaded area represents already assigned SLAs between the seller and buyers.

we assume buyers with unit-demand, $d_i = 1$, $\forall i \in B$. We further assume that the delivery of the electricity of the assigned SLAs is either successful or not. Let $\hat{d}_i \in \{0, d_i\}$ denote the transferred quantity to buyer $i$. We define $\alpha_i \in \mathbb{R}^+$ as the private value of buyer $i$ per unit of transferred electricity when delivery is assured ($\gamma_i = 1$).

**Remark 4.2.** *The private value $\alpha_i$ resembles the value of lost load (VoLL) since the binary model of electricity delivery can result either in full demand satisfaction (valued at $\alpha_i d_i$) for the buyer or zero.*

Considering the probability of electricity delivery (reliability) the *expected value* of buyer $i$ with regards to $\gamma_i$ is equal to:

$$v_i(\gamma_i) = \alpha_i d_i \gamma_i. \tag{4.1}$$

Let $S_Q(q)$ denote the reliability function (also known as the survival function) of the seller. Figure 4.1 illustrates $S_Q(q)$. Note that $S_Q(q) = 1 - F_Q(q)$, where $F_Q(q)$ is the cumulative density function of the random variable $Q$. The reliability function $S_Q(q)$ determines the probability that the available supply exceeds a certain value $q$. The dotted area represents an SLA (note that no price $p_i$ is determined here) between the seller and buyer $i$. The demand of buyer $i$ is equal to $d_i$ and the reliability of the specific SLA is $\gamma_i = S_Q(w + d_i)$, where $w$ is the demand that is already deducted by previously allocated SLAs.

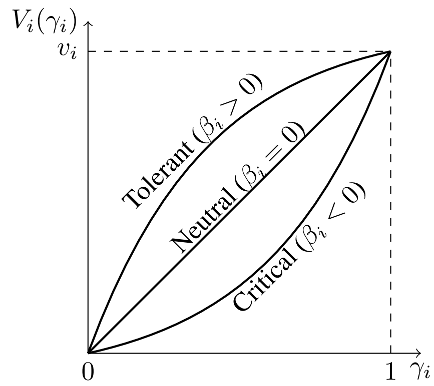


**Figure 4.2** The expected value function $V_i$ of a buyer with regards to the reliability $\gamma_i$ of an SLA. At full reliability ($\gamma_i = 1$) the expected value of the SLA by the buyer is equal to $v_i$ which is determined by the private value $\alpha_i$ and the demand $d_i$. The value of $\beta_i$ distinguishes different attitudes of buyer $i$ towards the reliability $\gamma_i$. The values used for the illustration are $\beta_i = \{-2.5, 0, 2.5\}$.

Given the assumption of unit-demand buyers, we further consider that the demand quantity $d_i$ of buyer $i$ is not comparable to the total demand of the $n$ buyers and the expected generation of the seller, such that the following holds: $S_Q(w) \approx S_Q(w+d_i)$ and $\mathbb{P}(w \leq q \leq w + d_i) \approx 0$. Intuitively, the probability of partial delivery approximates zero in the unit-demand case, justifying the binary model we use for the value that successful delivery brings to the buyer (see equation 4.1).

We proceed to elaborate the setting in which buyers may have different preferences with regards to the probability of being served (reliability), and thus represent demand loads with different criticalities.

### 4.3.3 Critical & tolerant buyers

The expected value of buyer $i$ in equation (4.1) is linearly dependent on the reliability $\gamma_i$ of the SLA. Since the system gives raise to risk, we can distinguish between different attitudes of buyers towards risk, from critical to tolerant, as it is usual in economics and expected utility theory (Ingersoll, 1987). To this end, we define the *generalized expected value function* $V_i(\gamma_i)$, where the reliability $\gamma_i$ induces the risk in the form of uncertain delivery of the specified in the SLA quantity of electricity:

$$V_i(\gamma_i) = \alpha_i d_i \rho_i(\gamma_i), \tag{4.2}$$

where $\rho_i(\gamma_i)$ encompasses the attitude of buyer $i$ with regards to the reliability $\gamma_i$. Note that for $\rho_i(\gamma_i) = \gamma_i$, $V_i(\gamma_i)$ becomes equal to the expected value $v_i(\gamma_i)$ in equation (4.1). The generalized expected value function in equation (4.2) needs to further embrace some common sense properties:

- Buyers have zero value for no reliability, i.e., $V_i(0) = 0, \ \forall i \in B$.
- Buyers have maximum value for no uncertainty, i.e., $V_i(1) = \alpha_i d_i, \ \forall i \in B$.
- Buyers have higher value for more certainty in the delivery, i.e., $V_i(\gamma_i) \geq V_i(\gamma_i - \varepsilon), \forall \varepsilon \in \mathbb{R}^+, \ \forall i \in B$ (monotonicity).
- Buyers have positive value for any positive reliability, i.e., $V_i(\gamma_i) > 0, \forall \gamma_i > 0, \forall i \in B$ (buyers' willingness to participate).

We consider a variation of the exponential utility function proposed by Ingersoll (1987). In line with the aforementioned properties of $V(\gamma_i)$, we define $\rho_i(\gamma_i)$ with regards to $\beta_i \in \mathbb{R}$.

$$\rho_i(\gamma_i) = \begin{cases} \dfrac{1 - e^{-\beta_i \gamma_i}}{1 - e^{-\beta_i}} & , \beta_i \neq 0 \\ \gamma_i & , \beta_i = 0 \end{cases} \tag{4.3}$$

This variation of the exponential utility function in equation (4.3) can be substituted in equation (4.2), yielding the generalized expected value of buyer $i$ with regards to the reliability $\gamma_i$. $\beta_i \in \mathbb{R}$ distinguishes the buyer's type from *critical* ($\beta_i < 0$) to *tolerant* ($\beta_i > 0$). Note that for $\beta_i = 0$, the expected value function $V_i(\gamma_i)$ becomes equal to the expected value in equation (4.1). Figure 4.2 illustrates the function $V_i(\gamma_i)$ for different values of $\beta_i$. We distinguish buyers with regards to their attitudes towards reliability:

**Critical** For $\beta_i < 0$, the generalized expected value function is convex, representing a critical (risk-averse) buyer. There is a stiff degradation of the value with regards to the uncertainty of the delivery for the critical buyer, resulting from opportunity costs that arise in case of failed delivery.

**Tolerant** For $\beta_i > 0$, the generalized expected value function is concave, the buyer is tolerant (risk-seeking). Lower reliability translates to a rather high expected value, resulting from opportunity value that arises in case of failed delivery.

**Neutral** For $\beta_i = 0$, the generalized expected value becomes identical to equation (4.1) that is linearly dependent to the reliability, representing a neutral buyer.

The generalized expected value function outlines a realistic model for capturing buyers' preferences describing the graceful or stiff degradation of the VoLL with regards to the probability of successful delivery. The type of buyer $i$ is characterized by the tuple $(\alpha_i, \beta_i)$. In the context of electricity markets, the same quantity of electricity may have different value for different customers, which is captured in our model by $\alpha_i$. In addition, the incurred value of a lost load with regards to the probability of electricity delivery is determined in our model by $\beta_i$. The function $V_i(\gamma_i)$ further indicates the expected VoLL of buyer $i$ given the reliability of the

SLA as follows: $\mathbb{E}[VoLL_i] = V_i(1) - V_i(\gamma_i)$. Last, the generalized expected value function can be defined in both one-shot and repeated settings, where in the latter buyers could vary their types with regards to the outcome of their earlier assignments.

Let us now proceed with an instance of a realistic scenario where the value of a buyer depends on the probability of being served, and thus the opportunity cost or value that arises in case of failed delivery.

**Example 4.2.** *An EV-taxi is operating its task (transporting people) with a half-full battery capacity with no urgent need to charge its battery. The EV-taxi can represent a tolerant buyer since there is value that is gained in case of charging (future benefits), and opportunity value that is gained in case of not charging (immediate benefits resulting by not pausing its task). When the capacity of its battery is running low, the EV-taxi can better be represented by a critical buyer since the value in case of no delivery (immediate benefits for task continuation) is decreasing relatively to the value that can be gained by charging and pausing its task.*

Given the above example, the generalized value function can distinguish buyers in terms of opportunity cost or value that could be incurred or gained respectively with regards to the probability of successful delivery of electricity.

The price $p_i$ that is specified in the SLA, and hence the expected utility of buyer $i$ is determined by the resulting allocation of the mechanism (see Section 4.4). Let $U_i$ denote the expected utility of buyer $i$,

$$U_i = V_i(\gamma_i) - d_i p_i \gamma_i, \tag{4.4}$$

where the expected payment that is transferred from buyer $i$ to the seller upon delivery is subtracted from its expected value.

In the section that follows we discuss how SLAs can be allocated to buyers with different private values for SLAs of varying reliability.

# 4.4 Auction-Based SLA Allocation

Auctions are commonly used in competitive electricity markets that take place day-ahead (Contreras et al., 2001), and they are known to yield efficient allocations even in cases there is uncertainty about buyers' valuations (Krishna, 2010). In this section we consider auctions as the method to allocate SLAs among buyers of varying types assuming no agency for the seller: the seller serves as the mechanism to allocate SLAs to participating buyers.

Let $A$ denote an allocation from the set of all feasible allocations $\mathcal{A}$, $A \in \mathcal{A}$, as the triplet of vectors $(\mathbf{d}, \boldsymbol{\gamma}, \mathbf{p})$, where each entry $A^{(i)} = (d_i, \gamma_i, p_i)$ is an allocated SLA between the seller and buyer $i$. Considering unit-demand buyers, i.e., $d = d_i = 1, \forall i \in B$, allocation $A$ can be expressed as $A = \{o_1, o_2, \ldots, o_n\}$, where $o_i \in \mathbb{Z}_1^n$ denotes the order of an allocated SLA. Given the order $o_i$ of the SLA $(d_i, \gamma_i, p_i)$ the reliability $\gamma_i$ is given by: $\gamma_i = S(o_i d)$ (see Figure 4.1). Following the definition of the reliability function, $\gamma_i \geq \gamma_j, \forall j \in B$, where $o_i < o_j$ holds. Intuitively, the reliability is monotonically decreasing with the order. The set of feasible allocations $\mathcal{A}$ includes all allocations $A$ for which every element appears only once in the set.

The value of an SLA for buyer $i$ is determined by equation (4.2) with regards to the reliability $\gamma_i$ and consequently the order $o_i$.

We define buyers' expected social value as the sum of the expected values of the set of buyers given the allocated SLAs as follows: $\sum_{i \in B} V_i(\gamma_i)$, where $\gamma_i = S(o_i d)$ is determined according to $A$. The allocation $A$ further determines the order that buyers get served, for each buyer $\forall i \in B$, $served_i = (q \geq o_i d)$, which follows from $\mathbb{P}(served_i) = \gamma_i$.

### 4.4.1 Sequential second-price auction

First, we consider a sequential second-price auction (SSPA) (Leme et al., 2012; Syrgkanis and Tardos, 2012) as the mechanism to allocate SLAs to participating buyers in the ahead timestep for all supply that may become available at the realization timestep. Items, SLAs in this case, are auctioned off one at a time. Given the assumption of unit-demand buyers, the seller auctions off SLAs of quantity $d$. We consider that the seller starts auctioning SLAs of decreasing reliability, such that the first SLA has reliability of $S(d)$, the second $S(2d)$, and so on. In sequential auctions, the order in which items are auctioned off affects the auctioneer's revenue (Elkind and Fatima, 2007). Given the monotonicity property of the generalized value function in equation (4.2), i.e., buyers always value more higher reliability, SSPA of decreasing reliability SLAs maximizes the revenue of the seller. In later sections of this chapter we also evaluate the case where the seller auctions off SLAs of increasing reliability.

A second-price auction (Vickrey, 1961) (also known as Vickrey) is held by the seller in every round $k$ of the auction, in which an SLA of reliability $S(kd)$ is auctioned off. Let $V_i'(S(kd))$ denote the reported value of buyer $i$ with regards to the reliability $S(kd)$ offered in the $k$-th round of the sequential auction. Each buyer $i$ places a sealed bid $z_i$, which is equal to its reported value with regards to the reliability $S(kd)$. The winner of each round $k$ of the sequential auction is buyer $w \in B$ that submits the highest bid, $w = \mathrm{argmax}_i\, z_i$, and pays to the seller the price $p_w$ that is equal to the second highest bid, $p_w = \max_{i \neq w} z_i$. Buyer $w$ is allocated an SLA of unit quantity $d$, reliability $S(kd)$, and price $p_w$, and participates no further in the next rounds of the auction. We assume that there is no *price-discovery*, and thus in each round of the auction only buyer $w$ knows the price of the assigned SLA. We further assume that buyers do not bid higher than their value, i.e., no *over-bidding*, and hence, an allocation cannot result in negative utility for the winning buyer.

We establish the setting of the second-price auction in the following assumption:

**Assumption 4.1.**

*i. Buyers do not communicate their preferences to their competition (other participating buyers).*

*ii. Buyers have no knowledge regarding the number and the distribution of buyers participating in the auction.*

*iii. Buyers do not know the reliability function of the seller; buyers only know that the reliability of the next SLA to be auctioned off is lower or equal to the reliability of the SLA that is being auctioned in the current round.*

*iv. The reliability function $S_Q(q)$ is defined for all demand quantities.*

v. *Not all of the demand is guaranteed to be satisfied within an SSPA, since SLAs may be assigned with zero reliability.*

**Theorem 4.1.** *Given Assumption 4.1, each round of the SSPA is an isolated Vickrey auction and therefore the mechanism is dominant-strategy incentive-compatible (DSIC).*

*Proof.* Each round of SSPA can be the last round or the round before with value arbitrarily close to zero and therefore can be treated as an isolated Vickrey auction (Vickrey, 1961). Buyers' dominant strategy is to report their true value function, i.e., $V_i'(S(kd)) = V_i(S(kd)), \forall i \in B$. □

The proof of Theorem 4.1 exploits the property that no stochastic model can be built by a buyer regarding follow up rounds of SSPA. Given Assumption 4.1, each round in which a buyer can wait without participating (bidding low or zero), does not add any information regarding the distribution of future bids of other buyers. Consequently, there is no stochastic model which can compute an expectation of future utilities in case of waiting the next round to bid truthfully. To prove Theorem 4.1, we assume that buyers are deterministic choosing to participate as this was the last round of the auction to maximize the likelihood of getting assigned an SLA of high reliability.

**Remark 4.3.** *Assumption 4.1 is necessary to show that each round of the sequential mechanism is an isolated Vickrey auction and therefore is truthful. However, points (ii) and (iii) may be strong assumptions in practice (e.g., in repeated settings). We refer the reader to Section 3.6.1 (Externalities) in which we discuss the practical implications of such an assumption.*

The outlined SSPA mechanism auctions off SLAs of decreasing reliability to buyers in a sequential fashion. We showed that given the assumptions of the proposed SSPA mechanism, there is no incentive for a strategic buyer to misreport its value function, and therefore the proposed sequential mechanism elicits truthful reports.

### 4.4.2 Vickrey-Clarke-Groves

Sequential second-price auctions are suitable mechanisms to allocate SLAs to buyers since they provide a simple mechanism framework that can be easily implemented in practical settings of the smart grid. However, the allocation of SLAs may not be optimal, as assignment depends more on the value $\alpha_i$, and less on the criticality of buyer $i$, $\beta_i$. Consider the following example:

**Example 4.3.** *There are two unit-demand buyers, buyer* $1$ *valuates* $90\%$ *of successful delivery* $V_1(90\%) = a$ *and* $V_1(50\%) = 3/4a$ *and buyer* $2, V_2(90\%) = a/2, V_2(50\%) \approx 0,$ *using the sequential auction proposed in Section 4.4.1, buyer* $1$ *is assigned the SLA with* $90\%$ *while buyer* $2$ *is assigned* $50\%$*. Assuming zero payments, the resulting social value of the above assignment is* $a$*. However, the socially optimal allocation would be buyer* $2$ *to be assigned* $90\%$ *and buyer* $1$ *with* $50\%$ *resulting in social value of* $5/4a$*.*

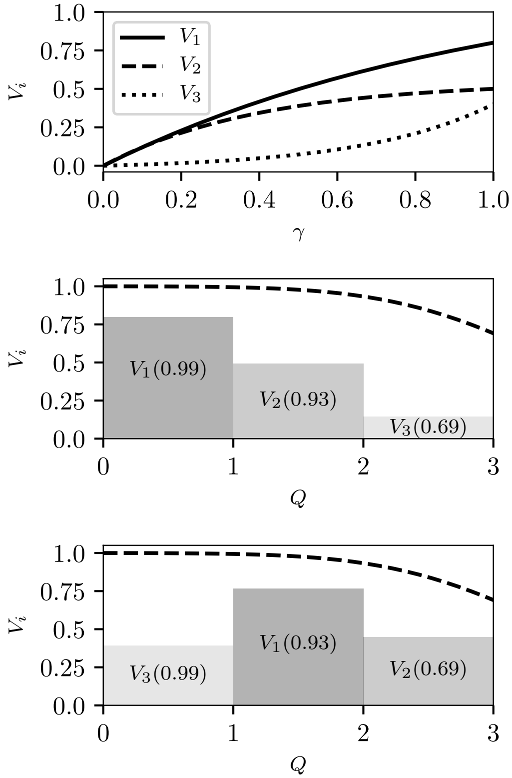


**Figure 4.3** The value functions $V_i$ of buyers with regards to the reliability $\gamma$ (top). The expected values of buyers under the resulting allocation of the sequential second-price auction (middle). Resulting expected values of buyers under the socially optimal allocation (bottom). The dashed line illustrates the reliability over allocated SLAs.

Similarly to the above example, Figure 4.3 (top) illustrates the expected value functions of three buyers of diverse types. In Figure 4.3, middle and bottom bar charts show two different allocations alongside the assigned reliability and the resulting expected value of buyers. The dashed line presents the reliability of the allocated SLAs. In the greedy allocation (middle), each slot is assigned to the buyer who has the highest bid (SSPA). The allocation that yields the optimal social value is illustrated in the bottom figure. SSPA myopically allocates SLAs to the highest bidder in each round of the auction and thus results in suboptimal (with regards to the social value) solutions.

Combinatorial auctions are means to derive socially optimal allocations (Cramton et al., 2007). In a combinatorial auction, the auctioneer computes the allocation after receiving buyers' bids for the whole bundle of items. We consider Vickrey-Clarke-Groves (VCG) which maximizes social welfare (or equivalently social value in our setting) (Nisan et al., 2007). Given the vector $\gamma$ of decreasing reliability of the available SLAs such that $\gamma = \langle S(d), S(2d), \ldots, S(nd)\rangle$, each buyer $i$ submits a vector of bids $\mathbf{z}_i = \langle V_i'(S(d)), V_i'(S(2d)), \ldots, V_i'(S(nd))\rangle$, which comprises the reported value of buyer $i$ for each reliability. Recall that $A \in \mathcal{A}$ is a feasible allocation $A = \{o_1, o_2, \ldots, o_n\}$, where $o_i$ denotes the order over the decreasing reliability SLAs of the allocated buyer. Let $A_{opt}$ denote the optimal allocation such that:

$$A_{opt} = \operatorname*{argmax}_{A\in\mathcal{A}} \sum_{i\in B} V_i'(A).$$

We further define $A^{-i} \in \mathcal{A}^{-i}$ as a feasible allocation for all buyers excluding buyer $i$ while $A^{-i}_{opt}$ denotes the optimal allocation without buyer $i$ present.

Under the VCG mechanism, the price $p_i$ for each buyer is determined by its marginal contribution:

$$p_i = \sum_{j\in B\setminus i} V_j'(A^{-i}_{opt}) - \sum_{j\in B\setminus i} V_j'(A_{opt}). \tag{4.5}$$

The price $p_i$ is determined by the difference between the optimal social value that is achieved when buyer $i$ is excluded from the set $B$ and the sum of the values that buyers other than $i$ achieve under the optimal allocation $A_{opt}$ with buyer $i$ present. Intuitively, each buyer pays the loss incurred to the society by its presence. Under the VCG mechanism it is a dominant strategy for buyers to report their valuations truthfully (Nisan et al., 2007).

In most combinatorial problems finding a socially optimal allocation lies in the class of NP-complete problems. However, we show that finding a socially optimal SLA allocation for the studied setting can be solved optimally in polynomial time.

**Corollary 4.1.** *The problem of SLA allocation for unit-demand buyers can be solved optimally in polynomial time $\mathcal{O}(n^3)$ by the Hungarian method (Kuhn, 1955).*

*Proof.* The unit-demand SLA allocation problem is equivalent to the *linear assignment problem* (LAP), where $n$ agents have to be assigned $n$ tasks while the number of tasks is equal to the number of agents. Each task stands for a slot in allocation $A \in \mathcal{A}$. □

Corollary 4.1 shows that the optimal allocation of SLAs, which maximizes the social value of buyers, can be found in polynomial time (i.e., $\mathcal{O}(n^3)$). Furthermore, there exist implementations that can speed up further the computation of the optimal assignment in LAPs (Date and Nagi, 2016). This makes the VCG mechanism a tractable solution for practical applications of our model in the context of the smart grid.

# 4.5 Evaluation & Discussion

In this section we evaluate the performance of the studied mechanisms to allocate SLAs to buyers of different types. Specifically, we evaluate: the VCG mechanism (see Section 4.4.2), SSPA where the seller auctions off SLAs of decreasing reliability (SPD) and increasing reliability (SPI) as described in Section 4.4.1.[1]

To study the efficiency of the proposed mechanisms against baseline allocations, we consider two mechanisms that use only the value of lost load (VoLL) to compute an allocation. In both baseline mechanisms a simultaneous second-price auction is held by the seller in the ahead timestep, in which each buyer $i$ bids its value for certain delivery (i.e., $V_i(1) = \alpha_i d_i$).[2] In the first baseline mechanism (POB) we consider only neutral buyers (i.e., no added value or cost is generated as a result of the uncertainty) and use the value function in equation (4.1). In the second baseline mechanism (POC) we consider all types of buyers and use the generalized value function in equation (4.2).

We evaluate and compare all mechanisms with regards to the *social value* (SV), and the *social welfare* (SW). Following equation (4.2), the social value is defined as the average value of buyers.

$$\mathrm{SV} = \frac{1}{|B|} \sum_{i \in B} V_i(\gamma_i) \tag{4.6}$$

In addition, social welfare is the average of the expected utilities of buyers (see equation 4.4).

$$SW = \frac{1}{|B|} \sum_{i \in B} U_i \tag{4.7}$$

Last, the expected income of the seller (mechanism's revenue) is equal to the expected payments of buyers to the seller.

$$U_s = \sum_{i \in B} d_i \gamma_i p_i = |B|(SV - SW) \tag{4.8}$$

## 4.5.1 Diversity in the criticality

First, we analyze the influence of the diversity in criticality $\beta$ of buyers on the social value. Recall that $\alpha_i$ determines the value of buyer $i$ for certain delivery, and $\beta_i$ characterizes buyer $i$ in term of the criticality of its demand (with regards to the uncertainty of electricity delivery). In this first evaluation, we consider that buyers

[1] Note that the theoretical property with regards to buyers' truthfulness in SPD (see Theorem 4.1) also holds for the SPI mechanism under Assumption 4.1.

[2] This is similar to a simultaneous second-price auction that is held by the seller after the realization timestep where only the available supply is auctioned off, and thus buyers' bids are equal to their value for certain delivery.

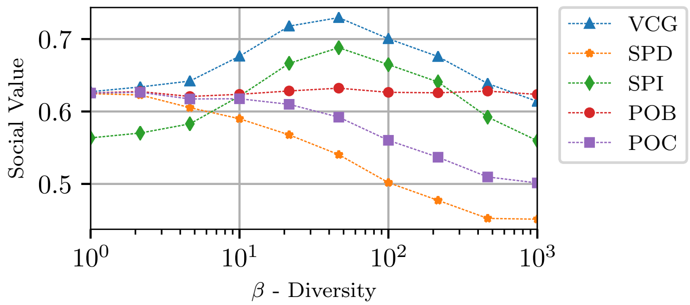


**Figure 4.4** Social value with regards to the diversity of criticality, for very low values of $\beta$-diversity all buyers approximate neutral attitude towards uncertainty in the delivery. For higher values of $\beta$-diversity buyers have increasingly varying criticality.

have comparable private values $\alpha_i \sim \mathcal{U}(0.5, 1.0),\ \forall i \in B$, which captures buyers with similar needs with regards to electricity usage (e.g., households). The random variable of the supply is normally distributed, $Q \sim \mathcal{N}(\mu_Q = 20, \sigma_Q = 5)$, while the total demand exceeds by $20\%$ the expected supply, $\sum_{i \in B} d_i = 1.2\ \mu_Q = |B|$, i.e., there are 24 unit-demand buyers. We consider that $\beta_i \sim \mathcal{U}(-D, D),\ \forall i \in B$, where $\mathcal{U}$ is the uniform distribution; $D \in \mathbb{R}^+$ refers to the diversity of $\beta$ values ($\beta$-Diversity). High values of $D$ result in high probability having extremely critical or tolerant buyers in the set $B$.

Figure 4.4 shows the resulting social value for each mechanism with regards to $D$ ($\beta$-Diversity). For $D \approx 0$ (buyers approximate the neutral behavior, i.e., $\beta \approx 0$), SPI results in lower social value than the rest of the evaluated mechanisms (which yield approximately equal social value) since the likelihood to obtain higher values from buyers assigned SLAs early (at low reliability) decreases. For $D \in [10^1, 10^2]$, there is a clear distinction between the resulting social value under different mechanisms while VCG results in the highest social value that is obtained in this experiment. In the same range, SPI results in a significant increase in social value by prioritizing over tolerant buyers in the allocation.[3] We observe the opposite for SPD which is the result of auctioning off SLAs of decreasing reliability (see Section 4.4.2). In settings where buyers demonstrate extreme behavior with regards to their criticality ($\beta \ll 0$) the efficiency of the system is vastly affected. More specifically, for high diversity $D$ ($D \approx 10^3$) the probability of extremely critical buyers is increased and thus the

[3] In SPI it is more likely that tolerant buyers are assigned low reliability SLAs, since their value for low reliability is higher than other types of buyers.

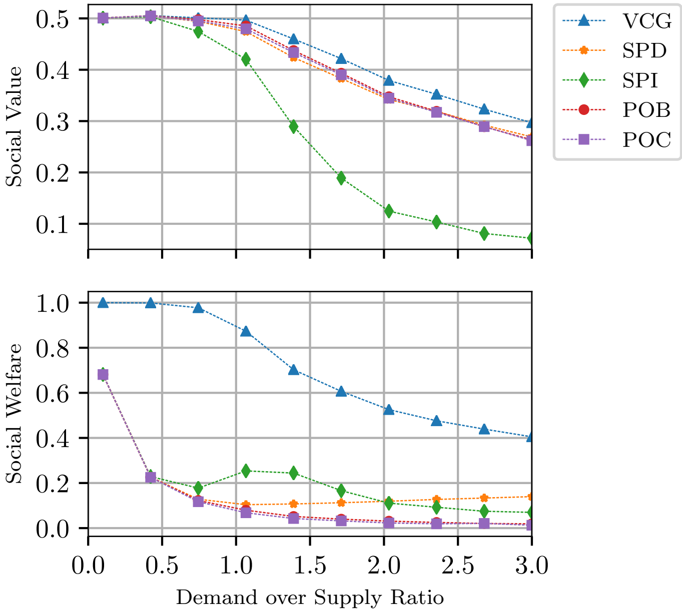


**Figure 4.5** Social value (top) and social welfare (bottom) achieved by different allocation mechanisms, namely, VCG, SSPA with decreasing reliability sequential auctions (SPD), increasing reliability sequential auctions (SPI), and two variants of a baseline method, POB and POC. The social value is computed as the average of all buyers' values with regards to the ratio of demand over the expected supply. Social welfare is normalized with respect to the social value obtained by the optimal VCG allocation.

resulting social value drops below the performance of the baseline mechanism POB under all other mechanisms. POB is not affected by the increasing diversity of $\beta$ since it only considers neutral buyers.

In this section we showed how diversity in criticality $\beta$ affects the resulting social value under all evaluated mechanisms. Overall, VCG outperforms all other mechanisms for almost the whole range of $D$.

### 4.5.2 Demand over supply ratio

We show that even in the case of large variations in the private value $\alpha$, VCG mechanism results in higher social value with regards to the other evaluated mechanisms. We use a diverse set of $\alpha$, where $\alpha_i \sim \mathcal{U}(0.1, 1.0)$, $\forall i \in B$, to captures highly irregular private values of heterogeneous buyers in electricity systems. In addition, we use low $\beta$-diversity, $\beta_i \sim \mathcal{U}(-5, +5)$, for buyers that do not exhibit extremely tolerant or critical risk attitude. Figure 4.5 (top) illustrates the resulting social value with regards to the ratio of the total demand over the expected supply. When the total demand can be satisfied with high probability (ratio $< 0.5$), there is no significant difference between the evaluated mechanisms with regards to the resulting social value. As the ratio increases, social value is decreasing for all evaluated mechanisms. We can observe that the social value obtained under VCG is higher than the other mechanisms. In addition, there is a huge drop in the performance of SPI for ratio higher than one. The social value under all the other evaluated mechanisms does not vary significantly from each other.

### 4.5.3 Social value & social welfare

Last, we show that the social welfare under VCG mechanism approximates the social value when the total demand is lower or approximately equal to the expected supply. Figure 4.5 (bottom) presents the resulting social welfare with regards to the ratio of the total demand over the expected supply under all evaluated mechanisms. The social welfare is normalized with regards to the social value obtained under the optimal allocation (under VCG mechanism), $SW/SV_{VCG}$. Intuitively, this is equal to the ratio of the social value that remains to buyers, while the rest is transferred to the mechanism (seller) through payments. Up to ratio $\approx 1$, the resulting social welfare under VCG mechanism is at least $90\%$ of the social value obtained using the optimal VCG allocation. In addition, the increasing social welfare under SPI mechanism for ratio $> 0.5$ is an intuitive result since more low-reliability SLAs become available when the ratio increases.[1] On the contrary, SPD achieves around $15\%$ of the optimal social value; however, it exhibits a more stable (although lower) performance than SPI. The social welfare under the baseline mechanisms (POB and POC) approximates zero for high values of demand to expected supply ratio, and consequently most of the generated social value is transferred through payments to the mechanism.

In this section we evaluated and illustrated the performance of the studied mechanisms with regards to the resulting social value and social welfare under a wide range of parameters. The VCG mechanism allocates SLAs in a socially optimal manner and consequently the social value and welfare of the system is maximized under this mechanism. Sequential mechanisms, such as the SPD and SPI, result in lower social value and welfare than VCG, however they have a slight advantage over baseline mechanisms that use only the VoLL. In addition, SPD and SPI can be easily implemented in practical settings of future smart grid systems since do not have any computational limitations with regards to scaling issues when compared to VCG.

## 4.6 Conclusions

In this chapter we proposed a contracting framework and mechanisms to allocate SLAs for electricity trading under uncertain supply and varying demand criticality of buyers. More specifically, we adopted SLAs as a direct extension of current conventional tariffs for use in electricity markets under uncertain supply, and we further defined the set of features that SLAs comprise: quantity, reliability, and price (see Section 4.3.2). We further proposed a generalized value function for buyers with regards to the criticality of their demand in face of uncertain delivery that generalizes the concept of the value of lost load (VoLL) with regards to the risk of unsuccessful delivery (see Section 4.3.3). The allocation of SLAs to buyers of heterogeneous types was computed using two Vickrey-based mechanisms: a sequential second-price auction (in which no buyer has an incentive to misreport its value under certain conditions) and the truthful VCG (see Section 4.4). Last, we evaluated both mechanisms in an experimental study: we showed that the VCG dominates all other mechanisms over a wide range of parameters, and vastly improves the efficiency of the proposed system when compared to baseline allocation mechanisms that only consider the VoLL (see Section 4.5).

In view of the attained properties and performance, we believe that using SLAs as we delineated here provides a promising avenue for addressing electricity trading in future smart grid systems. In particular, VCG allocation of SLAs is computable in $O(n^3)$, making it viable to assign the risk of demand curtailment to buyers that perform this task at low social cost. The proposed method can therefore be a tractable solution for peer-to-peer trading to balance local fluctuations in islanding grids or microgrids.

### 4.6.1 Future work

The contributions of this chapter can further serve as a broad basis for future extensions, the most interesting of which are outlined in this section.

Throughout this chapter we have considered no agency for the seller, the seller acts as the mechanism to allocate uncertain quantities of electricity to buyers of different preferences. It is of interest to study mechanisms that do not act on behalf of the seller, but instead they receive the estimated supply of one or multiple sellers before computing the allocation to buyers. In such mechanisms, one can study the possibility of using *scoring rules* to incentivize seller(s) to report their estimated supply truthfully (Robu et al., 2012; Vinyals et al., 2014).

Other research directions may consider the multi-unit demand case where buyers have continuous demand for electricity, and enriched SLA features that include time of delivery, penalty for misreporting the available supply, or unsuccessful delivery. In addition, future work can study the exact characterization of the value function of buyers in realistic settings of the future smart grid, e.g., heat pumps that act as buyers and for whom temperature variations introduce heterogeneous needs for electricity demand (demand of different criticality).

Last, the studied mechanisms of this chapter were not evaluated with regards to the seller's utility but only with regards to the social value and welfare of the buyers. A direct extension of this work is therefore to consider the seller's generation cost, in that case a reserve price auction may be used to ensure a positive expected utility for the seller.

# 5 Bounded-Rationality & Retail Markets

**Preface** In this chapter we investigate the effects of different degrees of rationality in the economic decision-making of buyers on the competition and prices that buyers face in retail markets with identical items. Within a fundamental retail market model that resembles a Bertrand competition, we use the Softmax function to approximate different degrees of buyers' rationality. The competition between sellers is modeled using hierarchical reasoning, in which each seller computes the price to offer to buyers (best response strategy) with regards to its belief for the competition. In the main theoretical contribution of this chapter we derive an analytical best response strategy (price) of a strategic seller given a set of opponent prices and the degree of buyers' rationality, and show that there exists an optimal degree of buyers' rationality that minimizes the price. We further use evolutionary game theory to empirically show the effects of perfect rationality, which results in unstable competition and price dynamics, and thus increasing costs for buyers. In contrast, we show that bounded rationality in the economic decision-making of buyers leads to smoother dynamics and higher benefits for buyers.

ℹ *The chapter presents work that was published in Springer Computational Economics (Methenitis et al., 2019a).*

# 5.1 Introduction

Classic game theoretical models to study strategic interactions between self-interested decision-makers (agents) assume the presence of intelligent and rational agents (Nash, 1950; Nisan et al., 2007; Sutton and Barto, 1998). However, such a rationality assumption misreckons that participants in the interactions do not usually have perfect knowledge of the environment within realistic domains (Russell and Thaler, 1985). Economic markets and consequently economic decisions of human buyers that participate in these markets is one instance where agents do not exhibit rational behavior (Conlisk, 1996; Rubinstein, 1998). As it was first discussed in Section 1.1.1, *bounded rationality* is a fundamental concept that considers the imperfect decision-making of otherwise rational agents due to, e.g., imperfect information (Simon, 1982). Without perfect information, a bounded rational decision-maker may act rationally over a limited set of choices. For the remainder of this chapter we refer to rational agents as *perfectly rational*, while bounded rational agents denotes agents of lower (unspecified) degree of rationality.

Automated agents already operate in agent-mediated e-commerce (Guttman et al., 1999; He et al., 2003; Maes et al., 1999), and it is inevitable that in future economies humans will be replaced by software as a principal agent of economic decision-making (Marwala and Hurwitz, 2017). In addition, recent advancements in e-commerce and fields of Artificial Intelligence such as Deep Learning (Goodfellow et al., 2016) and Automated Negotiation (Baarslag et al., 2017) illustrate the potential to enhance the abilities of agents in the complex settings of economic markets.

The discrepancy between human decision-makers that are bounded rational due to information or time limitations, and potential software agents that are perfectly rational with regards to individual interactions, leaves open the question whether such change is desirable:

> *Should the behavior of automated (software) agents be made perfectly rational?*

From a myopic point of view, increasing the level of rationality seems like a straightforward conclusion; however, this neglects the impact of changing dynamics over repeated interactions. Economic markets can instantiate such repeated settings, in which the behavior of individual decision-makers can change the dynamics and the overall market behavior. It is therefore of great interest to study the effects that perfectly rational decision-makers have on fundamental economic paradigms such as retail markets.

In this chapter we consider retail markets where sellers compete by offering prices for identical items to buyers, e.g., electricity markets. Each seller has a private cost for the items, e.g., procurement or production cost, and an infinite inventory of items. Sellers offer items to buyers at specific prices simultaneously in order to control a high market share and increase their profits. Assuming that buyers are perfectly rational (i.e., they choose the lowest price with probability one), this is known as the *Bertrand competition* (see Section 1.1.5). The Nash equilibrium of the Bertrand competition is the *competitive price* in the case that sellers have identical private costs (Bertrand, 1988). At the competitive price equilibrium, each seller sets a price equal to its private cost and the market is shared equally among the sellers: given that sellers have the same private costs, no seller has an incentive to deviate from the competitive price since a higher price results in zero market share and a lower price in negative utility for the seller.

The resulting competitive price equilibrium is formed under the following assumptions:

i. Sellers have no model of the competition (i.e., opponent sellers), and thus no information regarding the competing prices.
ii. Buyers are perfectly rational, i.e., they select the lowest price with probability one.

However, assumption (i) is not trivial in repeated markets where sellers can observe opponent prices and therefore model their competition, i.e., *opponent modeling* (Albrecht and Stone, 2018). In addition, assumption (ii) does not hold in practice, unless we consider small-scale markets with limited options for buyers and thus perfect knowledge.

Motivated by the above assumptions, this chapter investigates the effects of different degrees of buyers' rationality on the competition and the resulting prices for buyers in retail markets. This, however, requires modeling of both sides of a retail market, namely, the buyers and the sellers. First, we use the *Softmax* function (see equation 1.2) to model varying degrees of buyers' rationality. Furthermore, we use *$k$-level reasoning* to model the competition between sellers (Camerer et al., 2004; Stahl and Wilson, 1995). In $k$-level reasoning, $k$ denotes the depth of strategic reasoning of an agent. A 0-level agent has no model of the opponents and therefore cannot be strategic, i.e., 0-level agent uses a fixed or a random strategy. A $k$-level agent reasons with regards to its belief for the reasoning levels of its opponents. According to the *standard assumption* of $k$-level reasoning, a $k$-level agent believes to be facing $(k-1)$-level agents. In the studied setting, we analyze the best response strategy of a strategic seller (i.e., the price to offers to buyers) with regards to the prices posted by the competition. Last, we use evolutionary game theory to study the evolution of the competition in repeated interactions between sellers for given degrees of buyers' rationality, which has also been used to simulate producers' behavior in electricity markets (Menniti et al., 2008).

The main contributions of this chapter can be summarized as follows:

- First, we derive an analytical best response strategy of a strategic seller given a set of opponent prices and the degree of buyers' rationality.
- Interestingly, we find that buyers maximize their utility by not being perfectly rational in their choices.
- We use evolutionary dynamics to study the evolution of competition between sellers and show an evolutionary advantage of higher-level reasoning sellers when using the standard assumption of $k$-level reasoning.
- We extend the standard assumption of $k$-level reasoning towards a more realistic belief model for the competition (true distribution over lower reasoning levels),

and we observe that perfect rationality contributes to monopolistic behavior of higher-level reasoning sellers and unstable competition dynamics.

- On the contrary to perfect rationality, we show that bounded rationality leads to smoother competition dynamics and higher benefits for buyers.

To the best of our knowledge, this chapter presents the first study that combines bounded rationality in the price selection of buyers and opponent modeling for the sellers ($k$-level reasoning) within the Bertrand competition model to study the effects of different degrees of buyers' rationality on the competition and prices. The main objective of this chapter is not limited to study the consequences of varying degrees of buyers' rationality in retail markets; it also adds fundamental knowledge that can be used for the design of future agent-based automated markets with commodities (e.g., future electricity markets), and general competitive multi-agent settings with heterogeneous agents.

The remainder of this chapter is organized as follows: Section 5.2 provides an overview of the literature that is relevant to the work presented in this chapter. Next, in Section 5.3 we introduce the market model. In Section 5.4 we derive analytical best response strategies for strategic sellers with regards to prices offered by the competition and the degree of buyers' rationality, we also present experiments to verify our theoretical findings. In Section 5.5 we introduce concepts from evolutionary game theory and use them to show the effects of the degree of buyers' rationality in repeated interactions in retail markets. Last, Section 5.6 concludes this chapter providing a discussion on the insights of our results and further proposing promising future directions of this research.

## 5.2 Related Work

Bertrand competition and many of its variants are well-studied market models in the literature (Caragiannis et al., 2017; Dufwenberg and Gneezy, 2000; Spulber, 1995). For instance, Spulber (1995) studies the Nash equilibrium in the Bertrand competition and shows that when rivals' costs are unknown, each seller offers a price above its marginal cost and has positive expected profit. In another work, Caragiannis et al. (2017) study markets with multiple sellers that offer identical items to buyers with different valuations on each seller. The authors model this setting as a two-stage full-information game and show the price of anarchy and the efficiency of computing equilibria in this game. In this chapter we study settings within the Bertrand market model without assuming a full-information setting for sellers: sellers have only a belief about the competition they face.

As discussed in Section 1.1.5, a similar model to Bertrand in which sellers decide on the quantity of items to sell without any knowledge of the competition is the *Cournot* competition (Allaz and Vila, 1993). Singh and Vives (1984) study the connection between the Bertrand and the Cournot competition models by analyzing the duality of prices and quantities in differentiated duopolies. For retail markets we study in this chapter, the Bertrand model is better suited than the Cournot competition, in which sellers can only alter the price for items but not the quantity to sell (Weber, 2006).

The classical price competition model named after Bertrand (1988) prescribes that in equilibrium sellers set prices equal to their private costs. However, this equilibrium outcome is not in line with real-life observations in which buyers are not rational in their choices over prices, and in which sellers model their competition. More specifically, Dufwenberg and Gneezy (2000) show that the resulting prices that sellers offer to buyers further depend on the number of sellers that compete in the market. This is known as the *Bertrand Paradox* (Bruttel, 2009; Dufwenberg and Gneezy, 2000). Aligned with the Bertrand Paradox, we consider buyers that are bounded rational and use a stochastic model of choosing over prices. Similarly to related work (Ait Omar et al., 2017; Basov and Danilkina, 2015), we use the Softmax function to model the stochastic price selection of buyers, while other works make use of the equivalent multinomial logit function or the alternative Luce choice axiom (Anas, 1983; McFadden, 1975).

Previous work has also studied the effects of bounded rationality on Bertrand markets (Ait Omar et al., 2017; Basov and Danilkina, 2015; Zhang et al., 2009). For instance, Zhang et al. (2009) consider a Bertrand model with bounded rational sellers and study convergence properties of the competition. In the closest to ours work, Basov and Danilkina (2015) study price equilibria with regards to the degree of buyers' rationality. They propose a model where sellers can choose to educate or confuse buyers, i.e., increase or decrease their degree of rationality respectively, and present the effects of these choices. Extending previous results (Basov and Danilkina, 2015), Ait Omar et al. (2017) show that within a Bertrand oligopoly, sellers can benefit if buyers have lower degree of rationality. Our model substantially differentiates from the aforementioned work in the following ways. First, we consider automated (software) agents in place of human buyers. In this setting, agents of high computational capacity can reach levels of (almost) perfect rationality, and thus the degree of buyers' rationality cannot be manipulated by sellers.

The effects of bounded rational agents have also been studied with regards to learning agents, as the concept of bounded rationality is associated to the *exploration Vs. exploitation* problem in reinforcement learning (Sutton and Barto, 1998). For instance, Wunder et al. (2010b) study the effects of the exploration rate of players on the resulting players' payoffs in two-player prisoners' dilemma games. The authors show that increasing exploration rate (i.e., lowering the frequency of using a greedy policy) results in higher than in Nash equilibrium payoffs for players.

Last, in this chapter we consider sellers of heterogeneous reasoning levels using hierarchical reasoning to model competition. Hierarchical ($k$-level) reasoning has been also used in other game theoretical domains to model opponents (Hennes et al., 2012; Hu and Wellman, 2001; Lindner and Sutter, 2013; Wunder et al., 2010a). Hu and Wellman (2001) use $k$-level reasoning to learn the strategies of opponent agents (opponent modeling) in double-auctions. The authors conclude that more sophisticated modeling (high hierarchical reasoning level) does not guarantee an improvement in the performance of agents. In contrast to work by Hu and Wellman (2001), we use $k$-level reasoning to compute the best response strategy of a reasoning seller with regards to lower levels of reasoning. Consequently, higher levels of reasoning result in higher performance, since lower levels of reasoning function under limited information with regards to the competition. The work presented in this chapter is further related to literature that uses hierarchical reasoning to model varying information levels. More specifically, Hennes et al. (2012) use $k$-level reasoning to analyze the competitive advantage of high information access in markets. They conclude that random traders achieve in expectation higher gains than traders under partial information, who are in turn exploited by higher information level traders.

In view of the related work, this chapter fills a gap that connects bounded rational buyers and iterative reasoning sellers to model market competition dynamics.

# 5.3 Market Model

In this section we present our basic market setting, we also show how we model different degrees of buyers' rationality and the competition between sellers.

We use the Bertrand model to study retail markets where sellers offer identical items to a finite population of buyers, assuming that sellers have an infinite inventory of items, and equal private costs (Bertrand, 1988). In practice, e.g., in electricity retail markets private costs for electricity do not vary significantly. We define $c_i > 0$ as the private cost of seller $i$, and $p_i$ as the price that seller $i$ offers to buyers ($p_i$ is the decision of seller $i$), $p$ is the vector of prices set by all sellers. Furthermore, $p_{-i}$ denotes the vector of prices set by sellers other that $i$. Both the price $p_i$ and the prices of sellers other than $i$, determine the utility of seller $i$,

$$u_i = (p_i - c_i)s_i(p), \tag{5.1}$$

where $s_i(p)$ is the function that maps the vector of prices $p$ to the market share of seller $i$, i.e., $s_i : p \to [0, 1] \in \mathbb{R}$, such that $\sum_i s_i(p) = 1$. Last, we assume that the price of seller $i$ cannot be lower than its private cost $c_i$, $p_i \geq c_i$, since for any positive market share, $s_i(p) > 0$, $p_i < c_i$ results in negative utility for seller $i$.

## 5.3.1 Degree of buyers' rationality

In the retail market setting we consider, sellers offer identical items at specific prices to buyers. Buyers choose the price and consequently the seller to buy the items from. Assuming that buyers are perfectly rational, they choose the lowest price with probability one. In practice, however, buyers use a stochastic model for choosing over the offered prices, i.e., buyers are bounded rational (Rubinstein, 1998).

We use the Softmax function (see equation 1.2) alongside the Bertrand market model, to study the effects of different degrees of buyers' rationality. The fraction of buyers that choose price $p_i$ (market share of seller $i$) is given by:

$$s_i(p) = \frac{e^{-p_i/\tau}}{\sum_j e^{-p_j/\tau}}, \forall \tau \in (0, \infty), \tag{5.2}$$

where $\tau$ is the coefficient that exaggerates or diminishes the contrast between different prices for the buyers. Note that, in experiments presented later in this chapter we use the $\log(\tau)$ range for ease of illustration.

**Remark 5.1.** *Equation* (5.2) *is identical to the multinomial logit function that is widely used in the economics literature to model buyers' stochastic decision-making when facing different prices (Anas, 1983; McFadden, 1975).*

**Remark 5.2.** *Equation* (5.2) *models the collective degree of buyers' rationality and not the individual degrees of rationality within the population of buyers.*

The quantity $s_i(p)$ can also be interpreted as the probability that an individual buyer out of the buyers' population chooses price $p_i$. For $\tau$ close to zero ($\tau \to 0$), buyers are approximately perfectly rational choosing the lowest price with probability one, while for high values of $\tau$ ($\tau \to \infty$), buyers choose over prices with equal probability (uniformly random). The parameter $\tau$ can be adjusted to model different degrees of buyers' rationality, between (almost) perfect rational buyers and buyers that choose over prices randomly. Last, we compute the cost for the buyers as follows:

$$\sum_i s_i(p) \times p_i, \tag{5.3}$$

where the cost is equal to the sum of sellers' prices weighted by the market share of each seller, i.e., average price for the buyers.

### 5.3.2 k-level reasoning & competition

In the previous section we described the basic market model and outlined the decision of buyers over different prices with regards to their collective degree of rationality $\tau$ (see equation 5.2).

The present and following sections discuss how sellers decide the prices to offer to buyers. The decision of a seller with regards to the price (i.e., strategy) to offer to buyers is not only influenced by the degree of buyers' rationality but also by the prices offered by its competition (other sellers). We consider that sellers model their competition using $k$-level reasoning, where $k$ denotes the reasoning level of a seller (Stahl and Wilson, 1995). This resembles sellers that can have varying information levels or computational resources. For the remainder of the chapter, $Lk$ stands for the $k$-th level of reasoning.

First, we consider $L0$ sellers. A $L0$ seller does not model opponent sellers, and therefore its strategy (price) does not consider opponent prices. For higher levels of reasoning ($k > 0$), standard models of $k$-level reasoning assume the following: A $Lk$ agent believes to be facing $L(k-1)$ agents (Arad and Rubinstein, 2012; Hu and Wellman, 2001). Other models of $k$-level reasoning modify the aforementioned assumption as follows: A $Lk$ agent has a belief with regards to the probability of meeting each of the lower levels (Camerer et al., 2004). Last, in $k$-level reasoning no $Lk$ agent believes that it competes against agents of equal or higher reasoning levels.

In this chapter we use both models described above. For generality, we assume that $Lk$ seller has a belief distribution over lower reasoning levels. Let $x$ denote the vector of the true distribution over levels of reasoning, where each entry $x_k$ denotes the probability (frequency) that $Lk$ appears in the population of sellers. We define $\lambda_k$ as the belief distribution of $Lk$ seller with regards to the true distribution $x$, $\lambda_k$ consists of $k$ entries (the first entry is the frequency of $L0$ in the population), $\lambda_k = \langle \lambda_0, \lambda_1, \ldots, \lambda_{k-1} \rangle$. Each entry $\lambda_k^z$ is the probability of competing against $Lz$ seller, $\sum_{z=0}^{k-1} \lambda_k^z = 1$. Note that, $L0$ does not have a belief for the competition and for $k > 0$, sellers of the same reasoning level have identical beliefs with regards to the competition. Given the belief $\lambda_k$, we proceed to derive the best response strategy of $Lk$ seller, i.e., the price to offer to buyers that maximizes its utility.

# 5.4 k-Level Best Response Strategies

In this section we illustrate the best response strategy (price) of $Lk$ seller $i$ given private cost $c_i$, and belief $\lambda_k$ over its competition. For brevity, we omit $i$ from the notation since $Lk$ is independent of seller $i$.

## 5.4.1 Best response strategies in duopolies

We define $\pi_k^*$ as the best response strategy of $Lk$; $\pi_k^*$ is the function that maps: (i) the private cost $c$, (ii) the belief $\lambda_k$, and (iii) the degree of buyers' rationality $\tau$, to the price $p_k^*$, i.e., $\pi_k^* : (c, \lambda_k, \tau) \rightarrow p_k^*$. To simplify notation, we also use $p_k^*$ rather than $\pi_k^*$ in the remainder of this chapter. Considering a known $L0$ strategy, $p_0$, the strategy of $Lk$ agent is computed by iteratively best respond to lower levels of reasoning. To illustrate this, consider that $Lk$ seller believes that competes against a $L(k-1)$ opponent seller. Then, the best response of $Lk$ is given by:

$$\arg\max_{p_k} (p_k - c) s_k(\langle p_k, p_{k-1}^* \rangle),$$

where $p_{k-1}^*$ is the best response to $p_{k-2}^*$. Next, by taking into account the belief $\lambda_k$,

$$p_k^* = \arg\max_{p_k} \sum_{z=0}^{k-1} \lambda_k^z (p_k - c) s_k(\langle p_k, p_z^* \rangle), \tag{5.4}$$

is the best response of $Lk$ seller with regards to the probability of competing against each of the lower levels $z$. The $Lk$ best response strategy presented here serves as an illustration of the iterated best response model. In the following section we derive an analytical solution for the best response strategy of $Lk$ for any number of opponents with regards to the opponent prices.

## 5.4.2 Analytical best response strategies & rationality

Recall that $p_{-i}$ denotes the vector of prices set by sellers other than $i$. Here, we assume a known $p_{-i}$ since prices of opponent sellers result out of iterated best response strategies in $k$-level reasoning. We make no further assumptions for the private costs of opponent sellers, note that $c_i$ is the private cost of seller $i$.

**Theorem 5.1.** *The price $p_i^*$ maximizes the utility of seller $i$ given the vector of opponent prices $p_{-i}$, the private cost $c_i$, and the degree of buyers' rationality $\tau$:*

$$p_i^*(p_{-i}, c_i, \tau) = \tau W \left( \frac{e^{-\frac{c_i}{\tau} - 1}}{\sum_{j \neq i} e^{-\frac{p_j}{\tau}}} \right) + c_i + \tau, \tag{5.5}$$

*where $W$ is the Lambert function, i.e., $x = f^{-1}(xe^x) = W(xe^x)$* (Corless et al., 1996).

*Proof.* Given seller $i$, and the vector of opponent prices $p_{-i}$, the utility of seller $i$ is equal to:

$$u_i = (p_i - c_i) \frac{e^{-p_i/\tau}}{\sum_j e^{-p_j/\tau}}. \tag{5.6}$$

To derive the price $p_i^*$, we first use the quotient rule to compute the derivative of the utility of seller $i$ in equation (5.6) with respect to the price $p_i$:

$$\frac{\partial u_i}{\partial p_i} = \frac{e^{-\frac{p_i}{\tau}} \sum_{j \neq i} \left( e^{-\frac{p_j}{\tau}} \times \left( \frac{c_i - p_i}{\tau} + 1 \right) + e^{-\frac{p_i}{\tau}} \right)}{\left( \sum_j e^{-\frac{p_j}{\tau}} \right)^2}. \tag{5.7}$$

Equation (5.7) is the derivative of the utility of seller $i$ with respect to the price $p_i$. By solving equation (5.7) to be equal to zero we get equation (5.5), which is the the only solution of $(\partial u_i / \partial p_i) = 0$. Given that the only term that affects the sign of $(\partial u_i / \partial p_i)$ is

$$\sum_{j \neq i} \left( e^{-\frac{p_j}{\tau}} \times \left( \frac{c_i - p_i}{\tau} + 1 \right) + e^{-\frac{p_i}{\tau}} \right),$$

it is easy to see that as $p_i$ decreases further from $p_i^*$, both the quantity

$$e^{-\frac{p_j}{\tau}} \times \left( \frac{c_i - p_i}{\tau} + 1 \right)$$

and $e^{-p_i/\tau}$ increase, and thus $(\partial u_i / \partial p_i) > 0$ for $p_i < p_i^*$. Similarly, as $p_i$ increases further than $p_i^*$, both of the above quantities decrease, and thus $(\partial u_i / \partial p_i) < 0$ for $p_i > p_i^*$. Hence, $p_i^*$ is the price that maximizes the function $u_i$. □

Theorem 5.1 shows the best response strategy of seller $i$ with regards to the opponent prices $p_{-i}$, the private cost $c_i$, and the degree of buyers' rationality $\tau$. The above theorem is relevant for markets where prices are public knowledge, while the degree of buyers' rationality $\tau$ can be approximated.

We proceed to show some interesting theoretical results that follow from Theorem 5.1 under the following assumption:

**Assumption 5.1.** *We consider a Bertrand duopoly with a reasoning seller $i$ with private cost $c_i$ that observes: (i) the price of the opponent $p_{-i}$, which we assume is fixed for all $\tau$, and (ii) the degree of buyers' rationality $\tau$.*

Intuitively, the above assumption considers a duopoly market in which the opponent seller can not observe or estimate the degree of buyers' rationality $\tau$ and uses a fixed price $p_{-i}$. The competitive price $p_{-i}$ can also resemble the price of an outside option for buyers, e.g., their private cost for producing the items on their own, that does not depend on the degree of their rationality $\tau$. In contrast, the reasoning seller can observe the degree of buyers' rationality, motivated by the example of a company with resources for market research.

In the remainder of this section we abbreviate the notation of the best response function in equation (5.5), $p_i^*(p_{-i}, c_i, \tau)$, where possible. First, by using equation (5.5) we get the following lemma:

**Lemma 5.1.** *Given Assumption 5.1,*

$$p_i^* < p_{-i} \Leftrightarrow c_i < p_{-i} - 2\tau. \tag{5.8}$$

*Proof.* We use the property of the Lambert function, $W(f(x)) = g(x) \Leftrightarrow f(x) = g(x)e^{g(x)}$, to solve the following inequality,

$$\tau W\left(\frac{e^{-\frac{c_i}{\tau}-1}}{e^{-\frac{p_{-i}}{\tau}}}\right) + c_i + \tau < p_{-i}, \tag{5.9}$$

which results the inequality in equation (5.8). □

The above lemma shows the upper bound for the private cost $c_i$, such that the best response strategy $p_i^*$ is lower than the opponent price $p_{-i}$, and thus buyers can benefit. A less intuitive bound for the cost $c_i$ than in equation (5.8) can be computed for more than one opponent price.

We proceed to show that buyers benefit if they are not perfectly rational, i.e., $\tau > 0$, under the same setting.

**Lemma 5.2.** *Given Assumption 5.1 and $c_i < p_{-i}$, there exists $\tau^* \in (0, (p_{-i}-c_i)/2)$, such that $p_i^*(\tau^*) \le p_i^*(\tau), \forall \tau \in (0, \infty)$.*

*Proof.* Given that the quantity $(p_{-i} - c_i)$ is fixed for all $\tau$, and $\tau' = (p_{-i} - c_i)/2$, equation (5.8) implies that $p_i^* < p_{-i}$ for $\tau < \tau'$.

Given that equation (5.5) is not defined for $\tau = 0$, we compute the limit as $\tau$ tends to 0,

$$\lim_{\tau\to 0} p_i^* = c_i + \lim_{\tau\to 0}\left[\tau W\left(e^{\frac{p_{-i}-c_i}{\tau}-1}\right) + \tau\right].$$

By the L'Hospital's rule we get that $\lim_{\tau\to 0} p_i^* = c_i + (p_{-i} - c_i) = p_{-i}$. As $\tau \to 0$, $p_i^*$ tends to $p_{-i}$.

Thus, for every $\varepsilon > 0$ sufficiently small, the continuous function $p_i^*$ lies below $p_{-i}$ for every $\tau$ that belongs to $[\varepsilon, \tau' - \varepsilon]$.

Given the extreme value theorem for continuous functions in compact intervals, there is a $\tau^* \in [\varepsilon, \tau' - \varepsilon]$ for which $p_i^*(\tau^*) \le p_i^*(\tau)$, $\forall \tau \in [\varepsilon, \tau' - \varepsilon]$. In addition, we know from equation (5.8) that $\lim_{\varepsilon\to 0} p_i^*(\tau' - \varepsilon) = p_i^*(\tau') \ge p_{-i}$,

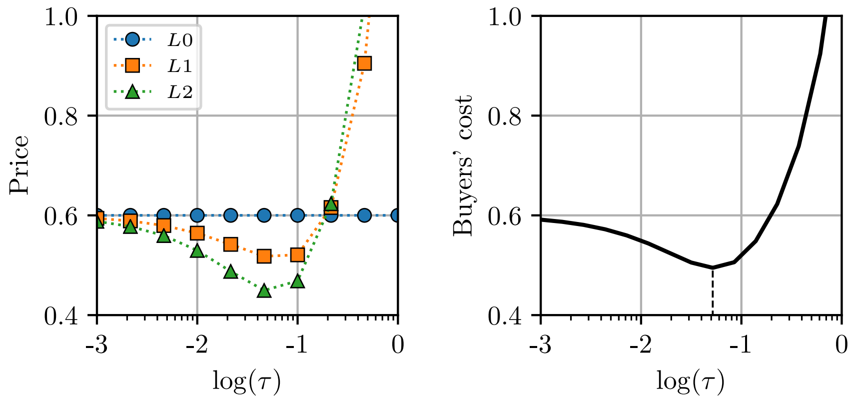


**Figure 5.1** (Left) Best response strategy (price) of reasoning level $Lk$ with regards to $\log(\tau)$. (Right) Buyers' cost with regards to $\log(\tau)$.

and $\lim_{\varepsilon \to 0} p_i^*(\varepsilon) = p_{-i}$. By taking $\varepsilon$ sufficiently small, and by $\inf_{\tau \in [\varepsilon, \tau']} p_i^*(\tau) \leq \inf_{\tau \in [\tau', \infty]} p_i^*(\tau)$, we get that $p_i^*(\tau^*) \leq p_i^*(\tau), \forall \tau \in (0, \infty)$. $\square$

**Theorem 5.2.** *Given Assumption 5.1 and $c_i < p_{-i}$, the optimal price of the reasoning seller $i$, $p_i^*$, is minimum for a degree of buyers' rationality $\tau^*$, with $\tau^* > 0$, and thus not for perfect rational buyers.*

*Proof.* The theorem directly follows from Lemmas 5.1 and 5.2. $\square$

Theorem 5.2 shows that the minimum price of the reasoning seller is obtained for a degree of rationality $\tau > 0$ (not perfect rationality).

In this section we derived analytical results with regards to the best response price of a reasoning seller, and the degree of buyers' rationality that minimize the price of the reasoning seller. We illustrate these results experimentally in the next section.

### 5.4.3 Duopoly markets

In line with our assumptions in the previous section, we consider a duopoly market where both sellers have identical private costs. We further use the standard assumption of $k$-level reasoning, namely, a $Lk$ seller believes to be competing against a $L(k-1)$ opponent seller, and thus $\lambda_k^z = 1$ for $z = (k-1)$ and $\lambda_k^z = 0$ for $z < (k-1)$. To derive the price of each $Lk$ seller we use the iterated best response strategy of $Lk$ similarly to equation (5.4) and the analytical best response price as this was derived in equation (5.5). More specifically, the price of $Lk$ is given by: $p_k^*(p_{k-1}, c_i, \tau)$, where we replace $p_{-i}$ in equation (5.5) with $p_{k-1}$, i.e., the price of the $(k-1)$ reasoning level. For the remainder of this section, we use 3 levels of reasoning; while our results can be generalized to any number of levels of reasoning, levels 0, 1 and 2 exemplify the cases of no, partial and (almost) full information respectively. Note that the number of possible strategies (levels of reasoning) is distinct from the number of sellers. Furthermore, $L0$ is a naive strategy that sells at an arbitrary fixed profitable price $p_0$, i.e. for $L0$ seller $i$, $p_0$ is larger than the private cost $c_i$.

Figure 5.1 (left) presents the best response strategy (price) of the 3 levels of reasoning with regards to the logarithm of the degree of buyers' rationality $\tau$. All sellers have identical private costs, $c = 0.2$, for $L0$ we use $p_0 = 0.6$. Values on the horizontal axis approximate different degrees of rationality from $\log(\tau) = -3$ (almost perfect rationality) to $\log(\tau) = 0$ (almost random price selection). For $\log(\tau) = -3$, the best response strategy of $Lk$ is marginally lower than the price of $L(k-1)$. Given that for $\log(\tau) = -3$, buyers are almost perfectly rational, a marginal decrease in the price of $Lk$ with regards to $L(k-1)$ results in $Lk$ to attain almost the full market share. As $\tau$ increases, the difference between prices becomes larger to counterbalance the stochastic selection of buyers over different prices. Intuitively, sellers choose a lower profit margin in order to achieve a higher market share.

For each reasoning level $k$ for $k > 0$, there exists $\tau_k^*$ for which the price $p_k^*$ becomes minimum. For instance, for $k = 1, 2$ the degree of buyers' rationality that minimizes the price $p_k^*$ is when $\log(\tau_k^*) \approx -1.3$. For higher values of $\tau$, buyers assign more equal probabilities for selecting among different prices. Hence, sellers of varying levels of reasoning achieve almost equal divisions of the market share that are only slightly influenced by the prices, and thus prices inflate in face of maximizing profits.

### Utility of sellers & buyers

We proceed to show the influence of the degree of buyers' rationality $\tau$ on the cost for buyers which we compute as in equation (5.3). Here, we use a uniform distribution for $x$, i.e., $x_0 = x_1 = x_2$ (recall that $x$ is the true distribution over levels of reasoning). Figure 5.1 (right) presents the cost for buyers with regards to logarithm of their collective degree of rationality $\log(\tau)$. For $\log(\tau) = -3$, the cost is marginally lower than the price $p_0$, however, it decreases further as $\tau$ becomes larger. For $\log(\tau) \approx -1.3$, the cost for the buyers is minimum. As $\tau$ increases further buyers choose randomly over prices and thus the cost is increasing since prices inflate.

The results presented throughout this section verify our theoretical findings for the existence of a degree of rationality (not perfect rationality) for which prices of reasoning sellers become minimum (see Theorem 5.2). To compute the cost for buyers we have considered a uniform distribution over levels of reasoning $x$. In the following section we show that the distribution $x$ can be influenced by the success rate of each reasoning level $Lk$ in a repeated setting.

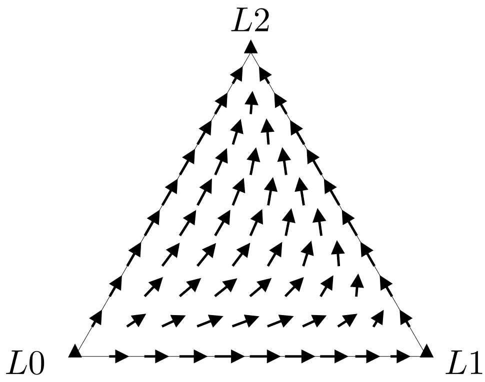


**Figure 5.2** Replicator dynamics over levels of reasoning, for almost perfect buyers' rationality, $\log(\tau) = -3$. Arrows (direction and magnitude) show the derivative of $x$, $\dot{x}$.

# 5.5 Evolutionary Dynamics

Considering repeated interactions that take place in markets, the frequency with which each strategy (i.e., reasoning level) appears in the population is influenced by its success rate (i.e., fitness). In this section we use evolutionary game theory to study the evolutionary dynamics of reasoning levels in the population of sellers.

Evolutionary game theory is a population-based application of game theory in repeated settings into which Darwinian competition can be modeled (Newton, 2018; Smith and Price, 1973; Weibull, 1997). Unlike game theory, it studies the dynamics of strategy change of a population over time, where strategies are influenced by their success rate and individuals from the population cannot select a strategy, instead they are given one.

Recall that $x$ denotes the distribution over reasoning levels (strategies), where each entry $x_k$ denotes the frequency strategy $Lk$ appears in the population. We further denote with $f_k$ the fitness function of $Lk$ that depends on the distribution over reasoning levels $x$, $f_k : x \to \mathbb{R}$. The strategy change $\dot{x}$ is computed by the *replicator equation* (Hofbauer, 1985) as follows:

$$\dot{x}_k = x_k \left[ f_k(x) - \varphi(x) \right], \tag{5.10}$$

where $\varphi(x)$ is the average fitness of the population.

$$\varphi(x) = \sum_z x_z f_z(x) \tag{5.11}$$

Equation (5.10) computes the change in frequency $Lk$ appears in the population at time $t+1$ after interaction at time $t$. The replicator equation comprises only the selection process, and therefore the most successful strategies increase their frequency in the population.

We revisit the duopoly scenario of the previous section (see Section 5.4.3) to apply the replicator equation. We compute the fitness $f_k$ for every possible duopoly as follows:

$$f_k(x) = \sum_{z=0}^{K} x_z (p_k^* - c) s_k(< p_k^*, p_z^* >), \tag{5.12}$$

where $K$ is the highest reasoning level (here, $K = 2$). Figure 5.2 presents the replicator dynamics for the duopoly model of Section 5.4.3. Arrows at each point of the simplex show the derivative $\dot{x}$ (direction and magnitude). We observe that evolution favors the highest reasoning level $L2$, we can similarly say that $L2$ has a competitive advantage.

In this section we used the replicator equation to study the evolution over reasoning levels in the duopoly scenario of Section 5.4.3, assuming that a $Lk$ seller believes to be facing a $L(k-1)$ opponent seller (standard assumption of $k$-level reasoning). We showed that the highest reasoning level has always an evolutionary advantage since the belief of each type is not influenced by changes in the distribution $x$. This result generalizes for any number of reasoning levels.

### 5.5.1 Dynamic belief of competition

In this section we alter the standard assumption of $k$-level reasoning to a dynamic belief model that is influenced by the distribution $x$.

We extend our setting to an oligopoly market with $n$ sellers and identical private costs for sellers. We consider that the belief of a $Lk$ seller with regards to opponent levels of reasoning sellers is the real distribution $x$ for all levels lower than $k$, such that $\lambda_k = \langle x_0, x_1, \ldots, x_{k-1} \rangle$. Note that $\sum_{z=0}^{k-1} \lambda_k^z < 1$, since $x_k > 0$, i.e., only lower than $k$ levels of reasoning are included in the belief distribution of $Lk$. In addition, for $x_k$ close to one (i.e., $Lk$ dominates the population), $\sum_{z=0}^{k-1} \lambda_k^z$ is close to zero. We define $x_{out} = 1 - \sum_{z=0}^{k-1} \lambda_k^z$ as the probability of facing equal or higher levels of reasoning opponents. The probability $x_{out}$ can only be computed for $k > 0$, since $L0$ does not have a belief distribution. Hence, the belief of $Lk$ becomes $\lambda_k = \langle x_0, x_1, \ldots, x_{k-1}, x_{out} \rangle$. We interpret the probability $x_{out}$ as the probability of competing with an unknown opponent, e.g., the outside option for buyers. The opponent price associated with the probability $x_{out}$ is denoted with $p_{out}$. The price $p_{out}$ can be set equal to the maximum price buyers are willing to pay to alleviate the risk of extreme prices set by dominant strategies.

### 5.5.2 Optimal pricing & generalized replicator dynamics

We use equation (5.5) to approximate the price of each reasoning level $p_k^*$. $Lk$ seller draws samples (opponent price vectors $p_{-i}$ of length $n-1$) with regards to its belief $\lambda_k$. In our experiments, the $Lk$ best response (optimal price for $k$-level of reasoning)

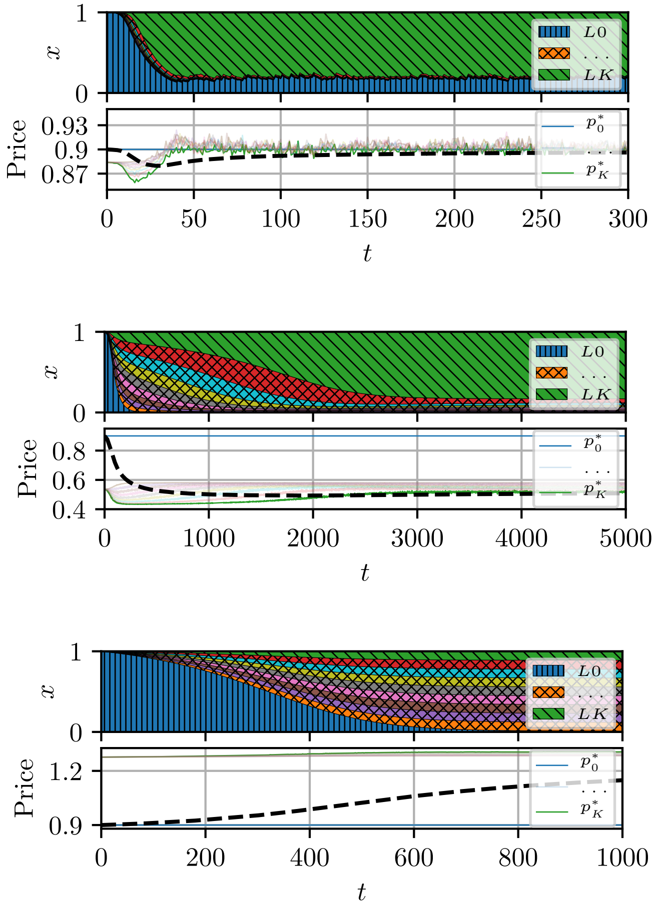


**Figure 5.3** Evolution of levels of reasoning and price for almost perfect rationality (top, $\log(\tau) = -2.7$), bounded rationality (middle, $\log(\tau) = -0.7$), and random behavior (bottom, $\log(\tau) = 0$). Stack plots at the top show the evolution of distribution $x$, and plots at the bottom illustrate the prices set by different levels of reasoning, the dashed line shows the development of the cost for the buyers.

is averaged over 100 sampled opponent price vectors. More samples do not change the behavior of the simulation in experiments presented later in this chapter.

Furthermore, to include *innovation* of strategies in the population of sellers, i.e., new sellers that enter competition or sellers that increase/decrease their level of reasoning, we use the *generalized replicator equation* (Hofbauer and Sigmund, 1998):

$$\dot{x}_k = \sum_z [x_z f_z(x) Q_{z\to k}] - \varphi(x) x_k, \tag{5.13}$$

where $Q_{z\to k}$ is the transition probability of an individual from the population from $Lz$ to $Lk$ (i.e., mutation probability), and $\varphi(x)$ the average fitness of the population (see Equation 5.11). The fitness of $Lk$, $f_k(x)$, is computed as follows:

$$f_k(x) = \frac{1}{M}\sum_{\mu=1}^{M}(p_k^* - c)s_k(\langle p_k^*, p^*_{z_\mu^{(1)}\sim x}, \dots, p^*_{z_\mu^{(n-1)}\sim x}\rangle), \tag{5.14}$$

where each $z_\mu^{(j)} \sim x$ are independent samples (i.e., $n-1$ opponent prices) from the true distribution over reasoning levels $x$, and the fitness is averaged out of $M$ sampled opponent price vectors. Considering that the population of sellers is finite, $\dot{x}$ is not deterministic for a given $x$, therefore computing the average fitness improves the approximation (Kemenade et al., 1998). We use $M = 100$ for experiments presented in the remainder of this chapter.

### Evolution of reasoning levels

Figure 5.3 illustrates the evolution over levels of reasoning and price with regards to time $t$ for $c = 0.2$, $p_0 = 0.9$, $p_{out} = 1$, and 10 levels of reasoning (from the lowest $L0$ to the highest $L9$, here $K = 9$). The initial distribution $x_0$ is set to $\langle 1, 0, \dots, 0\rangle$, only $L0$ is present at time $t = 0$. The mutation probability is set to $0.01$, where transition probabilities are uniformly distributed over all different levels, i.e., $\sum_{z\neq k} Q_{k\to z} = 0.01/(\text{number of levels} - 1)$, and $Q_{k\to k} = 0.99$. Stack plots placed at the top show the evolution of the distribution $x$ over levels of reasoning, and plots at the bottom show the price evolution for $\log(\tau) \in \{-2.7, -0.7, 0\}$. The bold dashed line shows the average cost for the buyers.

First, we discuss the case of almost perfect rationality, $\log(\tau) = -2.7$ (see Figure 5.3, top). Given the positive mutation probability in equation (5.13), higher levels ($L1 - K$) of reasoning "invade" the population of $L0$. $LK$ best responds to all lower levels of reasoning, thus it increases its share in $x$. For $t > 50$, $LK$ becomes dominant in the population, at the same time the frequency of reasoning levels between $L0$ and $LK$ diminish in the distribution $x$. In addition, prices as well as the distribution $x$ are not stable, resulting in price spikes that lead prices higher than the price $p_0$ ($p_0 = 0.9$). Both price spikes and the instability in the evolution of the distribution $x$ are caused due to: (i) the low probability for $LK$ to compete with lower level of reasoning opponents ($\sum_{j=0}^{K-1} x_j \approx 0.2$), and (ii) the high probability $x_{out}$ to face the outside option price $p_{out}$. The level of price spikes is subject to the outside price $p_{out}$, higher values for $p_{out}$ result in higher spikes further away from the price $p_0$. During price spikes, $L0$ benefits due to the high prices of $(L1 - K)$ and

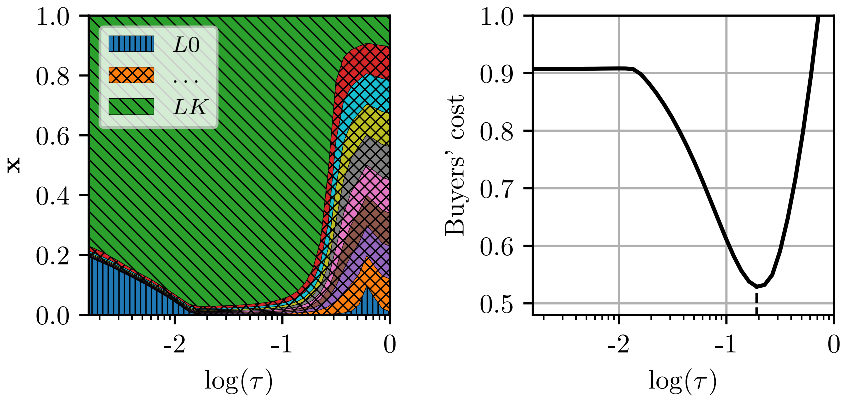


**Figure 5.4** (Left) Distribution of reasoning levels $x$, (right) buyers' cost. Results are computed for 10k steps of evolution and 20 independent evolution runs.

increases its share in $x$. Thereafter, higher levels of reasoning $(L1 - K)$ decrease their price in face of the increasing share of $L0$ in $x$ until $L0$ share decreases again. This results in chaotic evolutionary dynamics while similar behavior is observed for $\log(\tau) < -1.7$.

We observe smoother evolutionary dynamics and lower average price for buyers for lower degrees of buyers' rationality, more specifically, for $\log(\tau) > -1.7$. For instance, for $\log(\tau) = -0.7$ (see Figure 5.3, middle), evolution reaches an equilibrium state at $t > 3$k, where the distribution $x$ and the prices become stable. On the contrary to the case of almost perfect rationality (see Figure 5.3, top), the prices set by higher levels of reasoning $(L1 - K)$ are lower than $p_0$ $(p_0 = 0.9)$, and thus the average cost for the buyers decrease. Note that, the frequency of reasoning levels between $L0$ and $LK$ is not diminished as in the case of almost perfect rationality. The lower average price for buyers is a result of sustaining competition between different levels of reasoning sellers and the smoother dynamics of the evolution.

Last, we show the evolution of the distribution $x$ and the prices when the buyers' price selection is almost random (see Figure 5.3, bottom). For $\log(\tau) = 0$, reasoning levels $(L1 - K)$ share the distribution $x$ equally, where all reasoning sellers offer prices that exceed the price of $L0$, $p_0$, and the price $p_{out}$, and therefore increase the cost for buyers.

Overall, higher degrees of buyers' rationality yield higher average cost for buyers than lower degrees of rationality, e.g., $\log(\tau) = -0.7$. Furthermore, unstable evolutionary dynamics under almost perfect rationality increase prices further due to price spikes. In our experiments, we additionally used gradual updates to the prices in order to study the possibility more stable states can be reached in the evolution even in the case of perfect rationality. When gradual updates were used, results were consistent to the results presented here, however, the evolution of the distribution $x$ was slower.

### Competitive advantage & price

We proceed to show how the degree of buyers' rationality affects the competition in terms of the evolutionary advantage of higher reasoning levels, the resulting prices for buyers, and the stability of the competition.

Figure 5.4 (left) illustrates the distribution $x$ over levels of reasoning after 10k steps (mean of the last 100 steps) of the evolution averaged over 20 independent runs. $LK$ is the dominant in $x$ for almost all values of $\tau$, i.e., $\log(\tau) < -0.25$. For $\log(\tau) \approx -0.25$, all levels $L0$ to $LK$ have approximately equal shares in $x$. This is due to the almost equal prices reasoning levels set (similarly to the duopoly setting examined in earlier sections, see Figure 5.1, left). For $\log(\tau) > -0.25$, the market is shared among levels $L1$ and $LK$, since all levels of reasoning but $L0$ offer very high prices to (almost) random buyers.

We further show the effect of varying degrees of rationality $\tau$ on buyers' cost (see Figure 5.4, right). The cost is averaged over the last 100 out of 10k steps of evolution and over 20 independent evolution runs. For low $\tau$, the average cost for buyers is marginally higher than the cost without the presence of higher than $L0$ reasoning levels, $p_0 = 0.9$. This is the result of unstable competition dynamics that cause price spikes, during which prices become higher than the price of $L0$ strategy, $p_0$. Recall, that $p_{out} = 1$ alleviates the possibility of extreme prices, and thus the cost for buyers would increase further for higher $p_{out}$ due to the increasing level of price spikes. In contrast, from $\log(\tau) = -1.7$ to $\log(\tau) = -0.2$ buyers' cost drops below the price $p_0 = 0.9$, this is mainly caused by the smoother behavior of evolution that converges to stable distributions and alleviate price spikes. In line with our theoretical findings in Section 5.4, we observe that there is a degree of rationality $\log(\tau^*) \approx -0.7$ that minimizes the average cost for buyers (shown in the figure by the dashed vertical line).

In the presented experiments, we demonstrated that lower degrees than perfect buyers' rationality decrease the prices sellers offer to buyers during the evolution of the competition. For almost perfect buyers' rationality, the highest reasoning level sellers exploit instances of monopoly situations and increase their prices, while under bounded buyers' rationality competition is sustained decreasing prices for buyers. In the section that follows, we evaluate the stability of the competition with regards to the degree of buyers' rationality.

### Evolutionary stability

If the dynamics were known in explicit closed form, one could apply analytical notions of stability (e.g., evolutionary stable strategies, asymptotically stable) to analyze equilibrium strategies (Smith, 1972). However, given our implicit dynamics arising from system simulation (see Section 5.5.2), we need to draw on empirical means for characterizing the asymptotic behavior of the evolution. In the remainder

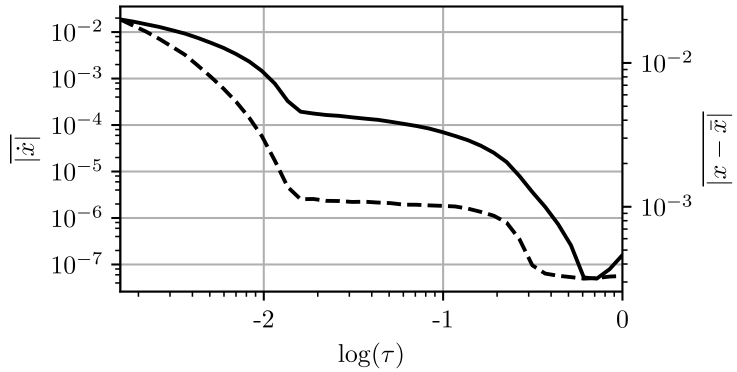


**Figure 5.5** Average magnitude of $\dot{x}$ (solid line, left verical axis) and average Euclidean distance of the distribution $x$ from the average distribution $\bar{x}$ (dashed line, right vertical axis). Results are computed for 10k steps of evolution and 20 independent evolution runs.

of this section we analyze both the first-order derivative and the distribution trajectory $x$, and examine how the degree of buyers' rationality influences the stability of the evolution.

First, we use the average magnitude (Euclidean norm) of $\dot{x}$, $\overline{|\dot{x}|}$, that is shown by the solid line in Figure 5.5 (left vertical axis). We compute this over the last 100 out of 10k steps of the evolution while results are averaged over 20 independent runs. The quantity $\overline{|\dot{x}|}$ is maximum for almost perfect buyers' rationality, specifically, $\overline{|\dot{x}|} > 10^{-3}$, $\forall \log(\tau) < -2$. This is in line with our observations in Figure 5.3 (top), where we showed chaotic behavior in the evolution of $x$ for a low $\tau$ value. As $\tau$ increases, steps in the evolution become smaller and consequently $\overline{|\dot{x}|}$ decreases. For $\log(\tau^*) \approx -0.7$, which minimizes the average cost for buyers in Figure 5.4 (right), $\overline{|\dot{x}|}$ is very low ($10^{-5}$).

Next, we use the Euclidean distance between $x$ and the average distribution $\bar{x}$, $\overline{|x - \bar{x}|}$, which is shown by the dashed line in Figure 5.5 (right vertical axis). The quantities $\bar{x}$ and $\overline{|x - \bar{x}|}$ are computed over the last 100 out of 10k steps of evolution, and averaged over 20 independent runs. Similarly to $\overline{|\dot{x}|}$, $\overline{|x - \bar{x}|}$ decreases as $\tau$ increases, and hence the distribution $x$ stays closer to the average distribution $\bar{x}$ for bounded rational buyers.

Our results suggest that imperfect rationality contributes to smoother competition dynamics, corroborating our observations in Section 5.5.2.

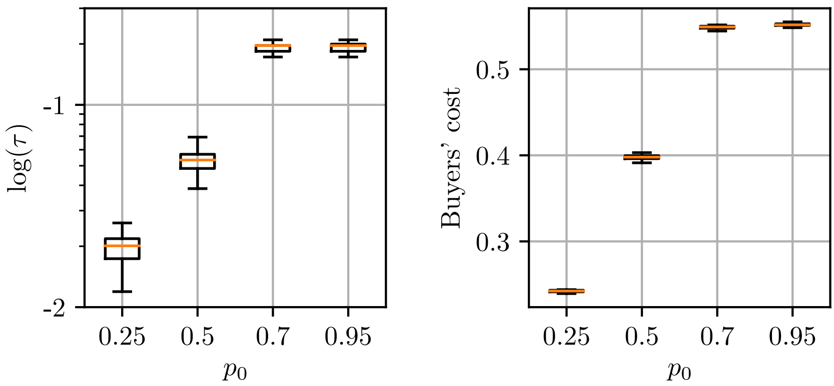


**Figure 5.6** (Left) Optimal degrees of buyers' rationality $\tau^*$ with regards to the price of $L0$, $p_0$. (Right) Buyers' cost when $\tau = \tau^*$ with regards to the price of $L0$, $p_0$. Results are computed for 10k steps of evolution and 20 independent evolution runs.

### Price of zero reasoning level

So far we have shown the effects of different degrees of buyers' rationality on the behavior of retail markets with regards to: the evolution of competition, the resulting prices for buyers, and the stability of evolutionary dynamics. Here, we show that the properties shown in previous sections generalize for different prices of $L0$ strategy, $p_0$. Figure 5.6 illustrates both the degree of rationality $\log(\tau^*)$ that minimizes the cost for buyers (left), and the corresponding cost for the values of $\log(\tau^*)$ (right). The cost for buyers is minimum if buyers are not perfectly rational for all values of $p_0$, however as the difference $(p_0 - c)$ becomes larger, $\log(\tau^*)$ increases (lower degree of rationality). At the same time, buyers' cost is relatively lower than $p_0$ as $p_0$ increases. Intuitively, the margin between the resulting average cost for buyers (computed for the optimal degree of buyers' rationality) and the price $p_0$ increase as the difference $(p_0 - c)$ increase.

## 5.6 Conclusions

In this chapter we illustrated the effects of varying the degree of buyers' rationality in retail markets. In the presented experiments we showed that almost perfect rationality caused spikes in price due to the unstable evolutionary dynamics, and thus increased the cost for buyers. On the contrary, lower degrees of rationality resulted in lower cost for buyers, by both sustaining competition between sellers of varying levels of reasoning and by increasing the stability of evolutionary dynamics. In line with related work (Wunder et al., 2010b), we can also conclude that using a stochastic choice model for decision-making in our setting leads in higher payoffs for the buyers.

Arriving at this non-trivial conclusion, we have made some simplifying assumptions with regards to the market setting and the model of competition between sellers. On the contrary, real-world retail markets involve highly perplexing dynamics and demonstrate extremely complex behavior, which can not be fully delineated in fundamental market models. Our results are thus not conclusive but instead seek to provide insights and add fundamental knowledge that can be used for the design of future retail markets with commodities that enable market participation by software agents, and general competitive multi-agent settings with heterogeneous agents.

Overall, in this chapter we studied the effects of varying degrees of buyers' rationality and sellers' opponent modeling (using $k$-level reasoning), in the Bertrand competition. In Theorem 5.1 we mathematically derived the best response strategy (price) given a set of opponent prices and the degree of buyers' rationality. We further used evolutionary dynamics to show the evolution of competition and prices in both duopoly and oligopoly scenarios. By replacing the standard assumption of $k$-level reasoning with a dynamic belief that depends on the distribution over reasoning levels, we showed that perfect rationality results in monopolistic behavior of higher reasoning level sellers, spikes in price, and unstable competition dynamics. The existence of an optimal degree of rationality stated in Theorem 5.2 and the improved evolutionary dynamics illustrated in our experiments thus provide a rationale for agents' bounded rationality in retail markets, raising the need to revisit design objectives for software agents in retail markets in light of their wider systematic impact.

### 5.6.1 Future work

The work presented in this chapter also serves as a basis for a number of extensions, some of which we detail in this section.

First, in this chapter we have considered the collective degree of buyers' rationality (see Remark 5.2). However, when considering the economic decision-making of an individual buyer given a set of prices, perfect rationality always yields the optimal result. Future extensions of this work may consider finite populations of buyers that each has a degree of rationality while it is of interest to show the possibility that individual buyers could converge to the optimal degree of rationality with regard to the resulting prices.

Throughout this chapter we assume that there is no cost associated with the reasoning level of sellers. A straightforward extension of this work may consider an arbitrary cost model for each reasoning level, or compute bounds up to which it is beneficial for sellers of higher levels of reasoning to enter the retail market competition.

Future extensions of this work may also consider more elaborate market models to resemble realistic settings of retail markets and finite population replicator dynamics (Taylor et al., 2004). Finite population models could result in different insights where the number of sellers can influence the resulting equilibria.

# 6 Conclusions

The rapid development of our societies alongside many technological advancements in electrification of transportation and heating are expected to increase demand for electricity significantly in the near future. In order to satisfy this increasing demand for electricity, natural sources such as the sun and wind are expected to complement and slowly replace conventional sources of electricity, such as fossil fuel and nuclear power plants. Future electricity systems will therefore be characterized especially by:

- The spatial distribution of renewable electricity generation.
- The inherent uncertain and intermittent nature of renewable electricity generation.
- The increasing active role of consumption and generation (e.g., prosumers, electric vehicles, smart home appliances).
- Interconnected electricity networks that use sensors and two-way communication to enable utilization of renewable electricity generation and coordination.

The outlined characteristics of future electricity systems are the main motivations of the innovative solution paradigm of the smart grid, which enables the transition from conventional electricity grids to distributed interconnected systems (see Section 1.2.3). In the smart grid, centralized control and generation, and passive consumption are replaced by distributed renewable electricity generation, local control and the actively participating consumers and small-scale producers. For instance, in the smart grid consumers may adapt their demand behavior based on external signals (e.g., price), or purchase their electricity directly from generators on their local network.

However, the seamless integration of renewable electricity generation and the decentralized nature of the smart grid pose new and significant challenges towards its practical implementation. Most of these challenges focus on the design of coordination protocols (mechanisms) to enable active participation of demand and generation entities (agents). The collective efficiency and the emergent behavior of such systems is the result of individual actions by autonomous and self-interested agents while each agent has its own set of preferences and is interested in maximizing its own utility rather than optimizing a global metric (e.g., social welfare). As a result, the ability of agents to actively control their demand or generation loads further raises strategic considerations and conflict of interest between agents, which can be analyzed and solved through the analysis of the resulting multi-agent systems.

In this thesis we focused on two fundamental challenges that are motivated by the transition in electricity systems:

- The increasing uncertainty on both the demand and supply sides.
- The economic decision-making of agents participating in future electricity markets

More specifically, uncertainty in the demand and supply raises the need for innovative pricing schemes and economic mechanisms that incentivize favorable changes in the demand behavior of customers to alleviate risks related to imbalances. Moreover, the potential participation of automated agents in future electricity markets may have adverse effects in electricity retail markets.

Throughout this thesis, artificial intelligence and its fields, such as game theory and mechanism design, provide the theoretical tools for the analysis and the design of mechanisms that not only take into consideration individuals' preferences, but also satisfy theoretical guarantees with regards to the efficiency of multi-agent systems. Similarly to a large body of scientific literature that is motivated by the transition in electricity systems and the smart grid (see Section 1.2.4), this thesis focused on satisfying a key requirement for the efficient operation of the envisioned and current electricity systems, which is the need for continuous balancing demand and supply.

# 6.1 Main Contributions

In this section we outline the main contributions of this thesis: Section 6.1.1 present a brief overview of our contributions, while Sections 6.1.2-6.1.5 provide a more in depth discussion with regards to the contributions of each technical chapter revisiting the corresponding research questions.

## 6.1.1 Overview

The main contributions of this thesis are in the design of economic mechanisms and the study of agent interactions in fundamental settings within the domain of the smart grid. These fundamental settings resemble instances of multi-agent systems in current or future electricity markets, in which different stakeholders exchange electricity or demand response services. More specifically, in Chapters 2 and 3 we considered the balancing responsibility of an electricity retailer and the uncertainty in the demand of its customers. In this setting, Chapter 2 proposed the risk-sharing tariff, which enables the retailer sharing the balancing responsibility with its customers incentivizing in such a way uncertainty reduction on the demand-side. Chapter 2 further studied the effects of the economic decision-making of customers to the design phase of such a tariff. In a similar setting, Chapter 3 proposed economic mechanisms to incentivize small-scale and unreliable agents to provide their demand response services to the retailer and alleviate excessive costs of the retailer with regards to balancing supply and demand.

Next, Chapter 4 studied a different scenario of the envisioned smart grid. More specifically, we considered a setting in which a seller that depends its generation on a renewable electricity source wishes to sell its potential supply to buyers; however, due to the intermittent nature of its supply the seller cannot guarantee electricity delivery to all the buyers. In such a setting, we proposed a contracting framework with service-level agreements that enables electricity trading when delivery is not guaranteed but is subject to the actual output of a renewable electricity source.

Last, the contributions of Chapter 5 are focused on the effects of different economic decision-making of buyers in retail markets. In this chapter we showed that the potential participation of software agents in these markets can have adverse effects on the competition between sellers, while our results raised some issues concerning the design of future electricity retail markets where software agents could participate in high-resolution pricing schemes.

### 6.1.2 Risk-sharing tariff & demand uncertainty reduction

In Chapter 2 we studied the design of a novel electricity tariff to incentivize uncertainty reduction on the demand-side addressing the following research question:

**Research Question 1.** *Can we design electricity tariffs that explicitly incorporate the balancing responsibility of the retailer and incentivize heterogeneous customers to reduce the uncertainty of their demand?*

To address this question we analyzed an extensive-form two-player game between a retailer and a customer in a fundamental model of an electricity market. In this setting, the retailer has the balancing responsibility and should therefore balance supply with the actual demand of the customer. We further considered that the customer has a direct influence on the balancing requirements of the retailer, and thus any deviation for the predicted demand causes a balancing cost to the retailer. In this game-theoretical framework, we proposed the risk-sharing tariff, which is a parameterized two-payment tariff scheme where the customer first pays a precommitment price for the anticipated demand, and after the actual demand is observed pays for any deviation between the precommitment quantity and the observed demand. The risk-sharing tariff allows the customer to choose the portion of the balancing risk is willing to assume from the retailer, which also affects the prices that the customer faces.

In Chapter 2 we showed that the risk-sharing tariff can contribute to reducing excessive costs of the retailer connected to balancing supply and demand. The customer has an economic incentive to reduce the uncertainty of the demand, since a lower deviation from the anticipated demand also decreases the electricity costs for the customer. Moreover, we showed that the proposed tariff is acceptable for both the retailer and the customer, since the risk-sharing tariff yields benefits for both players. We also illustrated how the overall social welfare can be improved by the reduced uncertainty in the demand of the customer. Furthermore, we provided arguments why the proposed tariff elicits all freely available demand response even in cases where the ability of the customer to reduce the uncertainty of the demand may not be certain, and we showed the existence of Nash equilibrium strategies in the studied two-player game.

With regards to the previous research question, in this chapter we showed that by sharing the risk associated with balancing supply and demand retailers (demand aggregators) can forward financial risks that come with the inherent uncertainty of the demand to customers. Customers can therefore not only reduce the price they pay for electricity by managing (or planning) better their demand loads, but they can also contribute to fewer imbalances between supply and demand. Fewer imbalances between supply and demand can further abate the excessive use of fast-ramping generators that produce $CO_2$ emissions.

**Economic-decision making on tariff design** Chapter 2 further addressed the following research question:

**Research Question 4b.** *What are the implications of considering the economic decision-making of customers when designing tariffs?*

The risk-sharing tariff proposed in Chapter 2 is a pricing scheme for electricity in which the customer chooses the share of balancing risk is willing to assume by the retailer. This choice determines the resulting price for precommitment and imbalances for the customer.

In Chapter 2 we demonstrated that the design of such a tariff can be influenced if we consider the imperfect economic decision-making of the customer. For instance, we showed that in case the customer does not choose the cheapest option with probability one, the retailer is enticed to offer higher economic incentives in order to elicit better equilibrium outcomes. With regards to the above research question, our results provide a valuable insight for the design of (electricity) tariffs in face of customers that select to subscribe into different tariffs stochastically.

Overall, the results of Chapter 2 are not only relevant for the design of electricity tariffs in future electricity markets, but can further be used for pricing models within variants of the classical *newsvendor problem* (Petruzzi and Dada, 1999), in which suppliers have to procure items to satisfy the uncertain demand of their customers.

### 6.1.3 Demand response mechanisms for unreliable agents

Similarly to Chapter 2, by explicitly modeling the balancing responsibility of a retailer in an electricity market, in Chapter 3 we studied economic mechanisms to incentivize uncertain demand response under a given demand forecast and imbalance price for the retailer. Chapter 3 addressed the following research question:

**Research Question 2.** *Based on the demand forecast of a retailer, can we design economic mechanisms that incentivize small-scale and unreliable demand response agents to prepare and reduce imbalances between supply and demand if necessary?*

In this chapter we considered the following problem: A retailer of electricity based on the forecasted demand of its consumers procures a fixed supply in the day-ahead market to satisfy the demand. Since consumers' demand is not certain but it can only be estimated, there is no guarantee that the actual demand is equal to the procured quantity. Therefore, any imbalance between the procured quantity and the actual demand of consumers should be adjusted in the balancing market with a much higher price than the day-ahead price. Following the work by Ma et al. (2016), we considered the presence of flexible agents that can reduce imbalances (by altering their demand or generation behavior) when demand is known but not finalized and before the balancing phase, if requested by the retailer. Agents decide whether to prepare with some preparation cost before the actual demand is known while the availability of prepared agents is not certain when requested to reduce or increase their demand.

With regards to the technical contributions of this chapter, we proposed two demand response mechanisms: a sequential mechanism that is truthful under some mild assumptions, and a truthful combinatorial mechanism that runs in polynomial time and uses Vickrey-Clarke-Groves (VCG) payments. Both mechanisms do not require all agents to respond but only a subset of them. We showed that both mechanisms yield positive expected utility for both the agents and the retailer (mechanism), and can be used in settings where both positive and negative imbalances result in balancing cost for the retailer. Last, we verified the theoretical properties of both mechanisms in an empirical evaluation over a wide range of parameters, where the proposed mechanisms achieved up to $16\%$ reduction in the balancing cost of the retailer and $14\%$ increase in social welfare compared to when no demand response is used.

To the best of our knowledge, the research study presented in Chapter 3 is the first to study the design of mechanisms that select a number of agents to prepare prior when only the forecast of the demand is known, and then request demand response agents to alter their demand or generation only if there is an imbalance between supply and demand (see Section 3.2). The technical contributions of Chapter 3 advance the state-of-the-art by proposing demand response mechanisms that do not require all prepared demand response agents to respond as previous works do (Ma et al., 2016, 2017). Doing so in more practical settings that also include the balancing responsibility of the retailer in electricity markets.

Although our results may be focusing on the advancement and the application of demand response programs for future electricity markets, they can also be used in other domains that require many agents to execute interdependent sub-tasks while agents can fail in each of these sub-tasks. For instance, ride-sharing applications and distributed communication protocols.

### 6.1.4 SLAs for renewable electricity trading

Next, in Chapter 4 we investigated the adoption of a contracting framework, through service-level agreements, to facilitate the allocation of renewable electricity supply to electricity customers that have demand loads of different criticality, addressing the following research question:

**Research Question 3b.** *Can service-level agreements provide the contracting framework for the allocation and trading of uncertain quantities of electricity between renewable generators and consumers?*

On the contrary to Chapters 2 and 3 that considered the balancing responsibility of a retailer that is directly influenced by the inherent uncertainty of the demand, Chapter 4 considered a different scenario, in which the uncertainty is on the side of the supply. This is the case in electricity systems that depend their supply on local renewable generation. In such settings, on the contrary to current electricity systems that demand is always satisfied, supply is uncertain and thus the delivery of electricity cannot be guaranteed. More specifically, we considered the stochastic output of a renewable electricity generator (seller). In practice, this can be a wind farm that depends its future electricity output on the availability of wind and therefore cannot ensure a specific supply. However, the probability of satisfying some demand quantity can be inferred by the wind forecast.

Conventional fixed-rate electricity tariffs cannot be used in such settings since the seller would need to guarantee the delivery of electricity. The main contribution of Chapter 4 is the adoption of SLAs as a direct extension of conventional tariffs for use in electricity markets that electricity delivery cannot be guaranteed. The proposed SLAs comprise the following features: quantity, reliability, and price. In addition, to allocate these SLAs to buyers we defined a generalized value function for the buyers with regards to the criticality of their demand in the face of uncertain delivery, further addressing the following research question:

**Research Question 3a.** *How can we model the utility function of consumers to include the uncertainty of electricity delivery and consequently the probability of having their demand satisfied?*

The proposed value function generalizes the concept of the value of lost load (VoLL) with regards to the risk of unsuccessful delivery. To determine the allocation of SLAs to varying types of buyers, we proposed two mechanisms based on Vickrey mechanisms: a sequential second-price auction and the VCG, where we showed that both mechanisms ensure that no buyer has an incentive to misreport its value under certain conditions. Last, we evaluated the two mechanisms in an experimental study showing that VCG performance dominates all other allocations over a wide range of settings, and vastly improves the efficiency of the proposed system when compared to baseline allocation mechanisms considering only the VoLL.

Overall, the contributions of Chapter 4 illustrate that the adoption of SLAs can facilitate electricity trading in islanding micro-grids that only depend their electricity supply on renewables. Our approach can also be seen as a form of demand-side management, since under this setting some demand loads can only be satisfied and therefore be active if there is enough supply. The mechanisms and the SLAs presented in this chapter can also be used in other domains, in which items or services can be sold while their availability is not certain.

### 6.1.5 Decision-making of buyers in retail markets

Previous chapters have in common the uncertainty of future electricity systems either in the demand of consumers, or in the supply of renewable electricity sources, while their contributions regard economic mechanisms for future electricity marketplaces. In this section we discuss the contributions of the last technical chapter of this thesis, Chapter 5, in which we studied the effects of the economic decision-making of buyers in fundamental retail market settings, addressing the following research question:

**Research Question 4a.** *What are the effects of representing buyers with autonomous (economic) decision-making agents in retail markets on the competitive dynamics and the resulting prices?*

The motivations for the above research question stem from future retail markets within the smart grid, which can enable high-resolution pricing schemes. Whereas in current electricity markets customers subscribe to electricity tariffs for long time-periods (e.g., 1 year), customers participating in future electricity retail markets may need to change over different tariffs in very short time intervals (e.g., 1-day, 1-hour). In order to make such short-term decisions, efficient software agents can participate in place of human customers. To investigate the above research question, Chapter 5 used the Bertrand competition model as a fundamental market setting. In this market model, we modeled the collective degree of buyers' rationality using the multinomial logit function to represent the different economic behavior of buyers. In addition, we considered the competition setting between sellers that participate in this retail market using $k$-level reasoning: each seller computes the price to offer to buyers with regards to its belief over the competition.

One of the main contributions of this chapter is the mathematical analysis of the best response strategy of a strategic seller with regards to the competition, and the degree of buyers' rationality. In Chapter 5 we further used evolutionary game theory to analyze repeated interactions that take place in markets between the competing sellers, and we illustrated the evolution of competition and prices in both duopoly and oligopoly scenarios for different degrees of buyers' rationality. In our main finding, we showed that perfect rationality results in monopolistic behavior of higher reasoning level sellers, spikes in price, and unstable competition dynamics. On the contrary, we observed stable evolutionary dynamics in the competition that resulted in lower prices for buyers, when buyers' choose over prices using a stochastic choice-model (bounded rational buyers). Overall, the contributions of Chapter 5 provide a rationale for agents' bounded rationality in retail markets, since the stability of current real-world retail markets may also be the result of the stochastic decision-making of human buyers. Our results and insights further raise the need to revisit design objectives for software agents in retail markets in light of their wider systematic impacts.

## 6.2 Concluding Remarks

The efficiency of envisioned smart grid systems will be determined by the collective behavior of many autonomous and self-interested stakeholders. In such systems, economic mechanisms will enable electricity trading, demand response programs, and other services in order to facilitate the introduction and the utilization of renewable electricity sources. These economic mechanisms should satisfy theoretical and efficiency guarantees since strategic behavior of users can potentially jeopardize the efficiency of such systems or lead to power outages. Game theory and mechanism design provide the theoretical tools to analyze such situations that may arise in the envisioned systems, and provide solutions that do not depend on deterministic or centralized solutions that may not consider the preferences of individual users.

This thesis proposed economic mechanisms and analyzed agent-based interactions within fundamental models that comprise strategic situations motivated by the transition in electricity systems. The contributions of this thesis comprise state-of-the-art mechanisms that:

- Incentivize uncertainty reduction in the demand-side via electricity tariffs that can share the balancing responsibility of retailers to the customers.

- Facilitate demand response programs for smaller-scale flexible users of future electricity systems.
- Enable electricity trading between consumers and producers in settings where supply is uncertain and therefore producers cannot guarantee delivery to consumers.

An additional contribution was the analysis of fundamental retail market behavior with regards to the economic decision-making of buyers that yielded crucial insights on the design of future (electricity) retail markets.

The methods proposed in this thesis, or future extensions that can be based on the broad-basis of our contributions, may comprise significant components of future smart grid systems since they can be used in practical scenarios of future smart grid networks. However, the results and contributions of this thesis are not only related to the domain of the smart grid and, more general, future electricity systems. Our insights may further be used in other application domains, for instance, e-commerce, transportation, industrial informatics, since the studied models of this thesis also comprise fundamental problems that are relevant to the fields of artificial intelligence, mechanism design, game theory, and economics.

### 6.2.1 Future research directions

In this section we conclude this thesis discussing some general research directions in the domain of the smart grid. In addition, we refer the reader to the end of each technical chapter for a more elaborate discussion with regards to interesting future work that can be based on the broad basis of our technical contributions.

**Distributed mechanism design** Throughout this thesis we proposed economic mechanisms (see Chapters 3 and 4) and pricing schemes (see Chapter 2) under the presence of a central authority, such as an electricity retailer or a seller with renewable electricity generation. Our contributions aim to assist future electricity systems, that although highly distributed in nature, require some form of centralized mechanisms to solve allocation problems on a local-level. Especially in small-scale markets such centralized protocols may suffer from trust issues, or require legal frameworks in order to ensure that the objectives of these mechanisms are not in conflict with the ones of the users. In addition, users may not be willing to share their preferences and characteristics to centralized authorities. To this end, *distributed mechanism design* provides a promising avenue for the design of distributed protocols that could be adopted on a local-level and ensure trustworthiness between the users participating in future electricity systems (Feigenbaum and Shenker, 2002). Moreover, *blockchain* is an emerging technology that can facilitate (peer-to-peer) trading of electricity and demand response services without the need for central mediators where users' reports propagate to (Mengelkamp et al., 2018).

**Cooperation & fairness** In this thesis we studied fundamental problems of future smart grid systems that comprise non-cooperative settings in which each agent (player) tries to maximize its own utility. However, future electricity systems may not always involve strategic considerations for the users; for instance, in electricity cooperatives (i.e., these can be micro-grid systems that connect to the main grid via a single point of common coupling) different users may share common goals with regards to the collective efficiency of the systems they co-exist in. To this end, we believe that the design of cooperative mechanisms is of equal importance towards the practical implementations of the smart grid, whereas non-cooperative mechanisms may heavily influence the efficiency of such systems in face of ensuring truthful participation. Promising research directions also include the dynamic formation of agent coalitions that can facilitate simpler coordination mechanisms to satisfy balancing requirements (Loni and Parand, 2017). Another important aspect that is often overlooked in the design of economic mechanisms for electricity markets is *fairness*. Although hard to be defined, fairness should be an important evaluation metric of economic mechanisms for the exchange of services in the smart grid and markets with commodities (Hekkelman and La Poutré, 2019; Vuppala et al., 2011).

**Automated consumer participation & preference elicitation** Last, future research may also focus on the interconnection between systems' operators and users in the smart grid. For instance, economic mechanisms proposed in this thesis require users reporting their private information with regards to their future demand behavior. This requirement may not be practical in future electricity markets that may function in higher time-resolutions than current markets, where long-term contracts for electricity supply are typically in place. It is of interest to study methods to elicit users' long-term preferences with regards to their electricity usage or other private information without the need users have to frequently report to mechanisms (Baarslag and Gerding, 2015).

# Bibliography


T. Agarwal and S. Cui. Noncooperative Games for Autonomous Consumer Load Balancing over Smart Grid. In *Game Theory for Networks*, pages 163–175, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.

D. Ait Omar, M. Outanoute, M. Baslam, M. Fakir, and B. Bouikhalne. *Joint Price and QoS Competition with Bounded Rational Customers*, pages 457–471. Springer International Publishing, Cham, 2017.

C. Akasiadis and G. Chalkiadakis. Mechanism Design for Demand-Side Management. *IEEE Intelligent Systems*, 32(1):24–31, 2017.

K. Alanne and A. Saari. Distributed Energy Generation and Sustainable Development. *Renewable and Sustainable Energy Reviews*, 10(6):539 – 558, 2006.

S. V. Albrecht and P. Stone. Autonomous Agents Modelling Other Agents: A Comprehensive Survey and Open Problems. *Artificial Intelligence*, 258:66–95, 2018.

Z. Alibhai, W. A. Gruver, D. B. Kotak, and D. Sabaz. Distributed Coordination of Micro-grids using Bilateral Contracts. In *2004 IEEE International Conference on Systems, Man and Cybernetics*, volume 2, pages 1990–1995, 2004.

B. Allaz and J.-L. Vila. Cournot Competition, Forward Markets and Efficiency. *Journal of Economic Theory*, 59(1):1 – 16, 1993.

A. Amato, B. D. Martino, M. Scialdone, and S. Venticinque. Towards a SLA for Collaborating Smart Solar-Powered Micro-Grids. In *2014 International Conference on Intelligent Networking and Collaborative Systems*, pages 592–597, 2014.

M. Amin. Smart Grid. *Public Utilitiies Fortnightly*, 2015.

A. Anas. Discrete Choice Theory, Information Theory and the Multinomial Logit and Gravity Models. *Transportation Research Part B: Methodological*, 17(1):13 – 23, 1983.

A. Arad and A. Rubinstein. The 11-20 Money Request Game: A Level-k Reasoning Study. *The American Economic Review*, 102(7):3561–3573, 2012.

K. J. Arrow. Social Choice and Individual Values. 1951.

T. Baarslag and E. H. Gerding. Optimal Incremental Preference Elicitation during Negotiation. In *Proceedings of the 24th International Conference on Artificial Intelligence*, pages 3–9. AAAI Press, 2015.

T. Baarslag, M. Kaisers, E. H. Gerding, C. M. Jonker, and J. Gratch. When Will Negotiation Agents Be Able to Represent Us? The Challenges and Opportunities for Autonomous Negotiators. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 4684–4690, 2017.

J. Bae, E. Beigman, R. A. Berry, M. L. Honig, and R. Vohra. Sequential Bandwidth and Power Auctions for Distributed Spectrum Sharing. *IEEE Journal on Selected Areas in Communications*, 26(7):1193–1203, 2008.

S. Basov and S. Danilkina. Bertrand Oligopoly with Boundedly Rational Consumers. *The BE Journal of Theoretical Economics*, 15(1):107–123, 2015.

J. Bertrand. *Review of Walrass Théorie Mathématique de la Richesse Sociale and Cournots Recherches sur les Principes Mathematiques de la Theorie des Richesses in Cournot Oligopoly: Characterization and Applications (edited by AF Daughety)*. Cambridge University Press, 1988.

E. Bitar and Y. Xu. Deadline Differentiated Pricing of Deferrable Electric Loads. *IEEE Transactions on Smart Grid*, 8(1):13–25, 2017.

E. Bitar, K. Poolla, P. Khargonekar, R. Rajagopal, P. Varaiya, and F. Wu. Selling Random Wind. In *2012 45th Hawaii International Conference on System Sciences*, pages 1931–1937, 2012.

L. E. Blume and D. Easley. *Rationality*, pages 1–13. Palgrave Macmillan UK, London, 2016.

S. Borenstein, M. Jaske, and A. Rosenfeld. Dynamic Pricing, Advanced Metering, and Demand Response in Electricity Markets. *UC Berkeley: Center for the Study of Energy Markets*, 2002.

C. Boutilier, M. Goldszmidt, and B. Sabata. Sequential Auctions for the Allocation of Resources with Complementarities. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, IJCAI '99, pages 527–523, San Francisco, CA, USA, 1999. Morgan Kaufmann Publishers Inc.

S. Braithwait, D. Hansen, and M. O'Sheasy. Retail Electricity Pricing and Rate Design in Evolving Markets. *Edison Electric Institute*, pages 1–57, 2007.

L. V. Bruttel. Group Dynamics in Experimental Studies - The Bertrand Paradox Revisited. *Journal of Economic Behavior & Organization*, 69(1):51 – 63, 2009.

C. F. Camerer, T.-H. Ho, and J.-K. Chong. A Cognitive Hierarchy Model of Games. *The Quarterly Journal of Economics*, 119(3):861, 2004.

I. Caragiannis, X. Chatzigeorgiou, P. Kanellopoulos, G. A. Krimpas, N. Protopapas, and A. A. Voudouris. Efficiency and Complexity of Price Competition among Single-product Vendors. *Artificial Intelligence*, 248:9–25, 2017.

P. Chakraborty and P. P. Khargonekar. A Demand Response Game and its Robust Price of Anarchy. In *2014 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pages 644–649, 2014.

G. Chalkiadakis, V. Robu, R. Kota, A. Rogers, and N. R. Jennings. Cooperatives of Distributed Energy Resources for Efficient Virtual Power Plants. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS '11, pages 787–794, Richland, SC, 2011. International Foundation for Autonomous Agents and Multiagent Systems.

H. Chao, S. S. Oren, S. A. Smith, and R. B. Wilson. Multilevel Demand Subscription Pricing for Electric Power. *Energy Economics*, 8(4):199 – 217, 1986.

L. Chen, N. Li, S. H. Low, and J. C. Doyle. Two Market Models for Demand Response in Power Networks. In *2010 First IEEE International Conference on Smart Grid Communications*, pages 397–402, 2010.

E. H. Clarke. Multipart Pricing of Public Goods. *Public choice*, 11(1):17–33, 1971.

V. N. Coelho, M. W. Cohen, I. M. Coelho, N. Liu, and F. G. Guimarães. Multi-agent Systems applied for Energy Systems Integration: State-of-the-art Applications and Trends in Microgrids. *Applied Energy*, 187:820 – 832, 2017.

A. J. Conejo, M. Carrión, J. M. Morales, et al. *Decision Making under Uncertainty in Electricity Markets*, volume 1. Springer, 2010.

V. Conitzer and A. Vidali. Mechanism Design for Scheduling with Uncertain Execution Time. In *AAAI*, pages 623–629, 2014.

J. Conlisk. Why Bounded Rationality? *Journal of Economic Literature*, 34(2):669–700, 1996.

J. Contreras, O. Candiles, J. I. de la Fuente, and T. Gomez. Auction Design in Day-ahead Electricity Markets. *IEEE Transactions on Power Systems*, 16(3):409–417, 2001.

R. M. Corless, G. H. Gonnet, D. E. G. Hare, D. J. Jeffrey, and D. E. Knuth. On the Lambert W Function. *Advances in Computational Mathematics*, 5(1):329–359, 1996.

P. Cramton, Y. Shoham, and R. Steinberg. An Overview of Combinatorial Auctions. *SIGecom Exch.*, 7(1):3–14, 2007.

R. K. Dash, P. Vytelingum, A. Rogers, E. David, and N. R. Jennings. Market-Based Task Allocation Mechanisms for Limited-Capacity Suppliers. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 37(3):391–405, 2007.

K. Date and R. Nagi. GPU-accelerated Hungarian Algorithms for the Linear Assignment Problem. *Parallel Computing*, 57(C):52–72, 2016.

S. de Vries and R. V. Vohra. Combinatorial Auctions: A Survey. *INFORMS Journal on Computing*, 15(3):284–309, 2003.

S. S. S. R. Depuru, L. Wang, and V. Devabhaktuni. Smart Meters for Power Grid: Challenges, Issues, Advantages and Status. *Renewable and Sustainable Energy Reviews*, 15(6):2736 – 2742, 2011.

M. Dufwenberg and U. Gneezy. Price Competition and Market Concentration: An Experimental Study. *International Journal of Industrial Organization*, 18(1):7 – 22, 2000.

E. Elkind and S. Fatima. Maximizing Revenue in Sequential Auctions. In *International Workshop on Web and Internet Economics*, pages 491–502. Springer, 2007.

O. Ellabban, H. Abu-Rub, and F. Blaabjerg. Renewable Energy Resources: Current Status, Future Prospects and their Enabling Technology. *Renewable and Sustainable Energy Reviews*, 39:748–764, 2014.

Z. Fadlullah, Y. Nozaki, A. Takeuchi, and N. Kato. A Survey of Game Theoretic Approaches in Smart Grid. In *Wireless Communications and Signal Processing (WCSP), 2011 International Conference on*, pages 1–4, 2011.

X. Fang, S. Misra, G. Xue, and D. Yang. Smart Grid: The New and Improved Power Grid: A Survey. *IEEE Communications Surveys Tutorials*, 14(4):944–980, 2012.

H. Farhangi. The Path of the Smart Grid. *IEEE Power and Energy Magazine*, 8(1):18–28, 2010.

J. Feigenbaum and S. Shenker. Distributed Algorithmic Mechanism Design: Recent Results and Future Directions. In *Proceedings of the 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, DIALM '02, pages 1–13, New York, NY, USA, 2002. Association for Computing Machinery.

T. S. Ferguson. *Mathematical Statistics: A Decision Theoretic Approach*, volume 1. Academic press, 1967.

J. Foerster, I. A. Assael, N. de Freitas, and S. Whiteson. Learning to Communicate with Deep Multi-Agent Reinforcement Learning. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 2137–2145. Curran Associates, Inc., 2016.

J. Gärttner, C. M. Flath, and C. Weinhardt. Portfolio and Contract Design for Demand Response Resources. *European Journal of Operational Research*, 266(1):340 – 353, 2018.

R. C. Geary. The Ratio of the Mean Deviation to the Standard Deviation as a Test of Normality. *Biometrika*, 27(3/4):310–332, 1935.

H. Geman. Commodities and Commodity Derivatives. *Modeling and Pricing for Agriculturals, Metals and Energy, Chichester (Grande-Bretagne): Wiley Finance*, 2005.

E. H. Gerding, S. Stein, S. Ceppi, and V. Robu. Online Mechanism Design for Vehicle-to-grid Car Parks. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, IJCAI'16, pages 286–293, Palo Alto, CA, 2016. AAAI Press.

G. Gereffi and M. Korzeniewicz. *Commodity Chains and Global Capitalism*. ABC-CLIO, 1994.

A. Goldthau. Rethinking the Governance of Energy Infrastructure: Scale, Decentralization and Polycentrism. *Energy Research & Social Science*, 1:134 – 140, 2014.

I. Goodfellow, Y. Bengio, and A. Courville. *Deep Learning*. MIT Press, 2016.

T. Groves. Incentives in Teams. *Econometrica*, 41(4):617–631, 1973.

R. Gustavsson, S. Hussain, and L. Nordström. Engineering of Trustworthy Smart Grids Implementing Service Level Agreements. In *Intelligent System Application to Power Systems (ISAP), 2011 16th International Conference on*, pages 1–6, 2011.

R. Guttman, A. Moukas, and P. Maes. *Agents as Mediators in Electronic Commerce*, pages 131–152. Springer Berlin Heidelberg, Berlin, Heidelberg, 1999.

J. Han and M. Piette. Solutions for Summer Electric Power Shortages: Demand Response and its Applications in Air Conditioning and Refrigerating Systems. *Refrigeration, Air Conditioning, & Electric Power Machinery*, 29(1):1–4, 2008.

T. Haring and G. Andersson. Contract Design for Demand Response. In *IEEE PES Innovative Smart Grid Technologies, Europe*, pages 1–6, 2014.

K. Hayakawa, E. H. Gerding, S. Stein, and T. Shiga. Online Mechanisms for Charging Electric Vehicles in Settings with Varying Marginal Electricity Costs. In *Proceedings of the 24th International Conference on Artificial Intelligence*, IJCAI'15, pages 2610–2616. AAAI Press, 2015.

M. He, N. R. Jennings, and H.-F. Leung. On Agent-mediated Electronic Commerce. *IEEE Transactions on Knowledge and Data Engineering*, 15(4):985–1003, 2003.

M. He, S. Murugesan, and J. Zhang. Multiple Timescale Dispatch and Scheduling for Stochastic Reliability in Smart Grids with Wind Generation Integration. In *2011 Proceedings IEEE INFOCOM*, pages 461–465, 2011.

X. He, N. Keyaerts, I. Azevedo, L. Meeus, L. Hancher, and J.-M. Glachant. How to Engage Consumers in Demand Response: A Contract Perspective. *Utilities Policy*, 27:108 – 122, 2013.

B. Hekkelman and H. La Poutré. Fairness in Smart Grid Congestion Management. In *2019 IEEE PES Innovative Smart Grid Technologies Europe (ISGT-Europe)*, pages 1–5, 2019.

H. D. Henderson. *Supply and Demand*, volume 1. London: Nisbet & Company Limited, 1922.

D. Hennes, D. Bloembergen, M. Kaisers, K. Tuyls, and S. Parsons. Evolutionary Advantage of Foresight in Markets. In *Proceedings of the 14th Annual Conference on Genetic and Evolutionary Computation*, GECCO '12, pages 943–950, New York, NY, USA, 2012. ACM.

K. Herter and S. Wayland. Residential Response to Critical-peak Pricing of Electricity: California Evidence. *Energy*, 35(4):1561 – 1567, 2010.

B. Hintermann. Pass-Through of CO2 Emission Costs to Hourly Electricity Prices in Germany. *Journal of the Association of Environmental and Resource Economists*, 3(4):857–891, 2016.

J. Hofbauer. The Selection Mutation Equation. *Journal of Mathematical Biology*, 23(1): 41–53, 1985.

J. Hofbauer and K. Sigmund. *Evolutionary Games and Population Dynamics*. Cambridge university press, 1998.

Y. Hong. On Computing the Distribution Function for the Poisson Binomial Distribution. *Computational Statistics & Data Analysis*, 59:41 – 51, 2013.

J. Hoogland, M. De Weerdt, and H. La Poutré. Now, Later, or Both: A Closed-Form Optimal Decision for a Risk-Averse Buyer. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '15, pages 1769–1770, Richland, SC, 2015. International Foundation for Autonomous Agents and Multiagent Systems.

J. Hu and M. P. Wellman. Learning about Other Agents in a Dynamic Multiagent System. *Cognitive Systems Research*, 2(1):67 – 79, 2001.

I. Hussain, S. Ali, B. Khan, Z. Ullah, C. Mehmood, M. Jawad, U. Farid, and A. Haider. Stochastic Wind Energy Management Model within Smart Grid Framework: A Joint Bi-directional Service Level Agreement (SLA) between Smart Grid and Wind Energy District Prosumers. *Renewable Energy*, 2018.

S. Hussain, R. Gustavsson, A. Saleem, and L. Nordström. Service Level Agreement: Coordination and Monitoring of Actors in Smart Grid. In *Innovative Smart Grid Technologies (ISGT Europe), 2012 3rd IEEE PES International Conference and Exhibition on*, pages 1–7, 2012.

D. Ilic, P. G. D. Silva, S. Karnouskos, and M. Griesemer. An Energy Market for Trading Electricity in Smart Grid Neighbourhoods. In *2012 6th IEEE International Conference on Digital Ecosystems and Technologies (DEST)*, pages 1–6, 2012.

J. E. Ingersoll. *Theory of Financial Decision Making*, volume 3. Rowman & Littlefield, Lanham, MD, 1987.

M. Janssen and M. C. W. Janssen. *Auctioning Public Assets: Analysis and Alternatives*. Cambridge University Press, 2004.

P. Jehiel, B. Moldovanu, and E. Stacchetti. Multidimensional Mechanism Design for Auctions with Externalities. *Journal of Economic Theory*, 85(2):258 – 293, 1999.

D. S. Jenkinson, D. E. Adams, and A. Wild. Model Estimates of CO2 Emissions from Soil in Response to Global Warming. *Nature*, 351:304, 1991.

Z. Jun, L. Junfeng, W. Jie, and H. Ngan. A Multi-agent Solution to Energy Management in Hybrid Renewable Energy Generation System. *Renewable Energy*, 36(5):1352 – 1363, 2011.

A. Kantamneni, L. E. Brown, G. Parker, and W. W. Weaver. Survey of Multi-agent Systems for Microgrid Control. *Engineering Applications of Artificial Intelligence*, 45:192 – 203, 2015.

K. K. Kariuki and R. N. Allan. Evaluation of Reliability Worth and Value of Lost Load. *IEE Proceedings - Generation, Transmission and Distribution*, 143(2):171–180, 1996.

S. Karnouskos and T. N. De Holanda. Simulation of a Smart Grid City with Software Agents. *EMS*, 9:424–429, 2009.

C. H. Kemenade, J. N. Kok, J. A. L. Poutré, and D. Thierens. Transmission Function Models of Finite Population Genetic Algorithms. Technical report, CWI (Centre for Mathematics and Computer Science), Amsterdam, The Netherlands, 1998.

W. Kempton and J. Tomić. Vehicle-to-grid Power Implementation: From Stabilizing the Grid to supporting Large-scale Renewable Energy. *Journal of Power Sources*, 144(1):280 – 294, 2005.

J. C. Ketterer. The Impact of Wind Power Generation on the Electricity Price in Germany. *Energy Economics*, 44:270 – 280, 2014.

J. H. Kim and A. Shcherbakova. Common Failures of Demand Response. *Energy*, 36(2):873 – 880, 2011.

D. S. Kirschen and G. Strbac. *Fundamentals of Power System Economics*. John Wiley & Sons, 2018.

J. K. Kok, C. J. Warmer, and I. G. Kamphuis. PowerMatcher: Multiagent Control in the Electricity Infrastructure. In *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, AAMAS '05, pages 75–82, New York, NY, USA, 2005. ACM.

V. Krishna. Chapter Two - Private Value Auctions: A First Look. In V. Krishna, editor, *Auction Theory (Second Edition)*, pages 11 – 26. Academic Press, San Diego, second edition edition, 2010.

H. W. Kuhn. The Hungarian Method for the Assignment Problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.

S. Landsburg. *Price Theory and Applications*. Cengage Learning, 2013.

D. A. Lashof and D. R. Ahuja. Relative Contributions of Greenhouse Gas Emissions to Global Warming. *Nature*, 344(6266):529, 1990.

D. Lehmann, L. I. O'Callaghan, and Y. Shoham. Truth Revelation in Approximately Efficient Combinatorial Auctions. In *Proceedings of the 1st ACM Conference on Electronic Commerce*, EC '99, pages 96–102, New York, NY, USA, 1999. ACM.

D. Lehmann, R. Müller, and T. Sandholm. The Winner Determination Problem. *Combinatorial auctions*, pages 297–318, 2006.

R. P. Leme, V. Syrgkanis, and E. Tardos. Sequential Auctions and Externalities. In *Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '12, pages 869–886, Philadelphia, PA, USA, 2012. Society for Industrial and Applied Mathematics.

F. Lindner and M. Sutter. Level-k Reasoning and Time Pressure in the 11-20 Money Request Game. *Economics Letters*, 120(3):542 – 545, 2013.

A. Loni and F. Parand. A Survey of Game Theory Approach in Smart Grid with Emphasis on Cooperative Games. In *2017 IEEE International Conference on Smart Grid and Smart Cities (ICSGSC)*, pages 237–242, 2017.

J. A. P. Lopes, F. J. Soares, and P. M. R. Almeida. Integration of Electric Vehicles in the Electric Power System. *Proceedings of the IEEE*, 99(1):168–183, 2011.

R. D. Luce. *Individual Choice Behavior: A Theoretical analysis*. Wiley, New York, NY, USA, 1959.

H. Ma, V. Robu, N. Li, and D. C. Parkes. Incentivizing Reliability in Demand-side Response. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, IJCAI'16, pages 352–358. AAAI Press, 2016.

H. Ma, D. C. Parkes, and V. Robu. Generalizing Demand Response Through Reward Bidding. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '17, pages 60–68, Richland, SC, 2017. International Foundation for Autonomous Agents and Multiagent Systems.

C. M. Macal and M. J. North. Tutorial on Agent-based Modelling and Simulation. *Journal of Simulation*, 4(3):151–162, 2010.

P. Maes, R. H. Guttman, and A. G. Moukas. Agents That Buy and Sell. *Commun. ACM*, 42 (3):81–ff., 1999.

I. Maity and S. Rao. Simulation and Pricing Mechanism Analysis of a Solar-Powered Electrical Microgrid. *IEEE Systems Journal*, 4(3):275–284, 2010.

A. Marshall. From Principles of Economics. In *Readings In The Economics Of The Division Of Labor: The Classical Tradition*, pages 195–215. World Scientific, 2005.

T. Marwala and E. Hurwitz. *Artificial Intelligence and Economic Theory: Skynet in the Market*. Springer, 2017.

L. G. Mattsson and J. W. Weibull. Probabilistic Choice and Procedurally Bounded Rationality. *Games and Economic Behavior*, 41(1):61 – 78, 2002.

R. P. McAfee and J. McMillan. Auctions and Bidding. *Journal of Economic Literature*, 25 (2):699–738, 1987.

P. McDaniel and S. McLaughlin. Security and Privacy Challenges in the Smart Grid. *IEEE Security Privacy*, 7(3):75–77, 2009.

D. McFadden. Conditional Logit Analysis of Qualitative Choice Behaviour. In P. Zarembka, editor, *Frontiers in Econometrics*, pages 105–142. Academic Press New York, New York, NY, USA, 1973.

D. McFadden. The Revealed Preferences of a Government Bureaucracy: Theory. *The Bell Journal of Economics*, 6(2):401–416, 1975.

J. R. Meginniss. *A New Class of Symmetric Utility Rules for Gambles, Subjective Marginal Probability Functions, and a Generalized Bayes' Rule*. Columbia University, Graduate School of Business, 1976.

P. Meibom, C. Weber, R. Barth, and H. Brand. Operational Costs Induced by Fluctuating Wind Power Production in Germany and Scandinavia. *IET Renewable Power Generation*, 3(1):75–83, 2009.

M. Meinshausen, N. Meinshausen, W. Hare, S. C. Raper, K. Frieler, R. Knutti, D. J. Frame, and M. R. Allen. Greenhouse-gas Emission Targets for Limiting Global Warming to 2°C. *Nature*, 458(7242):1158, 2009.

R. Meir, H. Ma, and V. Robu. Contract Design for Energy Demand Response. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 1202–1208, 2017.

F. L. Meng and X. J. Zeng. A Stackelberg Game-theoretic Approach to Optimal Real-time Pricing for the Smart Grid. *Soft Computing*, 17(12):2365–2380, 2013.

E. Mengelkamp, B. Notheisen, C. Beer, D. Dauer, and C. Weinhardt. A Blockchain-based Smart Grid: Towards Sustainable Local Energy Markets. *Computer Science - Research and Development*, 33(1):207–214, 2018.

D. Menniti, A. Pinnarelli, and N. Sorrentino. Simulation of Producers Behaviour in the Electricity Market by Evolutionary Games. *Electric Power Systems Research*, 78(3):475 – 483, 2008.

G. Methenitis, M. Kaisers, and H. La Poutré. Incentivizing Intelligent Customer Behavior in Smart-grids: A Risk-sharing Tariff & Optimal Strategies. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, IJCAI'16, pages 380–386. AAAI Press, 2016.

G. Methenitis, M. Kaisers, and H. La Poutré. SLA-Mechanisms for Electricity Trading Under Volatile Supply and Varying Criticality of Demand. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '17, pages 1640–1642, Richland, SC, 2017. International Foundation for Autonomous Agents and Multiagent Systems.

G. Methenitis, M. Kaisers, and H. La Poutré. Renewable Electricity Trading through SLAs. *Energy Informatics*, 1(1):57, 2018.

G. Methenitis, M. Kaisers, and H. La Poutré. Degrees of Rationality in Agent-Based Retail Markets. *Computational Economics*, 2019a.

G. Methenitis, M. Kaisers, and H. La Poutré. Forecast-Based Mechanisms for Demand Response. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, AAMAS '19, pages 1600–1608, Richland, SC, 2019b. International Foundation for Autonomous Agents and Multiagent Systems.

J. S. Mill. *Utilitarianism*. Cambridge Library Collection - Philosophy. Cambridge University Press, 2014.

A. Mohsenian-Rad, V. W. S. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia. Autonomous Demand-Side Management Based on Game-Theoretic Energy Consumption Scheduling for the Future Smart Grid. *IEEE Transactions on Smart Grid*, 1(3):320–331, 2010.

C. Mora, B. Dousset, I. R. Caldwell, F. E. Powell, R. C. Geronimo, C. R. Bielecki, C. W. W. Counsell, B. S. Dietrich, E. T. Johnston, L. V. Louis, M. P. Lucas, M. M. McKenzie, A. G. Shea, H. Tseng, T. W. Giambelluca, L. R. Leon, E. Hawkins, and C. Trauernicht. Global Risk of Deadly Heat. *Nature Climate Change*, 7:501, 2017.

J. L. Moraga-González and M. Mulder. Electrification of Heating and Transport: A Scenario Analysis for the Netherlands Up to 2050. 2018. Centre for Energy Economics Research, CEER Policy Papers 2, University of Groningen.

D. Muthirayan, D. Kalathil, K. Poolla, and P. Varaiya. Mechanism Design for Demand Response Programs. *arXiv preprint arXiv:1712.07742*, 2017.

R. B. Myerson. *Game Theory*. Harvard university press, 2013.

J. Nair, S. Adlakha, and A. Wierman. Energy Procurement Strategies in the Presence of Intermittent Sources. *SIGMETRICS Perform. Eval. Rev.*, 42(1):85–97, 2014.

J. F. Nash. Equilibrium Points in N-person Games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950.

J. F. Nash. Non-Cooperative Games. *Annals of Mathematics*, 54(2):286–295, 1951.

M. J. Neely, A. S. Tehrani, and A. G. Dimakis. Efficient Algorithms for Renewable Energy Allocation to Delay Tolerant Consumers. In *2010 First IEEE International Conference on Smart Grid Communications*, pages 549–554, 2010.

J. Newton. Evolutionary Game Theory: A Renaissance. *Games*, 9(2), 2018.

W. Nicholson and C. Snyder. *Microeconomic Theory: Basic Principles and Extensions*. Nelson Education, 2011.

N. Nisan and A. Ronen. Algorithmic Mechanism Design. *Games and Economic Behavior*, 35 (1):166 – 196, 2001.

N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.

S. Niu, Y. Jia, W. Wang, R. He, L. Hu, and Y. Liu. Electricity Consumption and Human Development Level: A Comparative Analysis Based on Panel Data for 50 Countries. *International Journal of Electrical Power and Energy Systems*, 53:338 – 347, 2013.

R. Olfati-Saber, J. A. Fax, and R. M. Murray. Consensus and Cooperation in Networked Multi-Agent Systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.

S. S. Oren and S. A. Smith. *Service Opportunities for Electric Utilities: Creating Differentiated Products*, volume 13. Springer US, New York, US, 1993.

D. A. Ortega and P. A. Braun. *Artificial General Intelligence: 4th International Conference, AGI 2011, Mountain View, CA, USA, August 3-6, 2011. Proceedings*, chapter Information, Utility and Bounded Rationality, pages 269–274. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.

P. A. Ortega, D. A. Braun, J. Dyer, K.-E. Kim, and N. Tishby. Information-Theoretic Bounded Rationality. *arXiv preprint arXiv:1512.06789*, 2015.

S. Oruc, A. Pandharipande, and S. Cunningham. An Electricity Market Incentive Game based on Time-of-use Tariff. In *Fourth Workshop on Game Theory in Energy, Resources and Environment, Montreal, Canada*, 2012.

Y. Oualmakran, J. M. Espeche, M. Sisinni, T. Messervey, and Z. Lennard. Residential Electricity Tariffs in Europe: Current Situation, Evolution and Impact on Residential Flexibility Markets. *Proceedings*, 1(7), 2017.

G. Owen and J. Ward. Smart tariffs and household demand response for Great Britain. *Sustainability First, London*, page 2010, 2010.

P. Palensky and D. Dietrich. Demand Side Management: Demand Response, Intelligent Energy Systems, and Smart Loads. *IEEE Transactions on Industrial Informatics*, 7(3): 381–388, 2011.

D. C. Parkes. Online Mechanisms. *In Algorithmic Game Theory, ed. N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani*, pages 411–439, 2007. Cambridge: Cambridge University Press.

S. Parsons, J. A. Rodriguez-Aguilar, and M. Klein. Auctions and Bidding: A Guide for Computer Scientists. *ACM Comput. Surv.*, 43(2), 2011.

G. Pepermans, J. Driesen, D. Haeseldonckx, R. Belmans, and W. D'haeseleer. Distributed Generation: Definition, Benefits and Issues. *Energy Policy*, 33(6):787 – 798, 2005.

A. Perrault and C. Boutilier. Efficient Coordinated Power Distribution on Private Infrastructure. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '14, pages 805–812, Richland, SC, 2014. International Foundation for Autonomous Agents and Multiagent Systems.

A. Perrault and C. Boutilier. Multiple-Profile Prediction-of-Use Games. In G. Sukthankar and J. A. Rodriguez-Aguilar, editors, *Autonomous Agents and Multiagent Systems*, pages 275–295, Cham, 2017. Springer International Publishing.

N. C. Petruzzi and M. Dada. Pricing and the Newsvendor Problem: A Review with Extensions. *Operations Research*, 47(2):183–194, 1999.

E. Pettersen, A. B. Philpott, and S. W. Wallace. An Electricity Market Game Between Consumers, Retailers and Network Operators. *Decis. Support Syst.*, 40(3-4):427–438, 2005.

N. Phuangpornpitak and S. Tia. Opportunities and Challenges of Integrating Renewable Energy in Smart Grid System. *Energy Procedia*, 34:282 – 290, 2013. 10th Eco-Energy and Materials Science and Engineering Symposium.

M. Pipattanasomporn, H. Feroze, and S. Rahman. Multi-agent Systems in a Distributed Smart Grid: Design and Implementation. In *Power Systems Conference and Exposition, 2009. PSCE '09. IEEE/PES*, pages 1–8, 2009.

R. Porter, A. Ronen, Y. Shoham, and M. Tennenholtz. Fault Tolerant Mechanism Design. *Artificial Intelligence*, 172(15):1783 – 1799, 2008.

I. Praca, C. Ramos, Z. Vale, and M. Cordeiro. MASCEM: A Multiagent System that Simulates Competitive Electricity Markets. *IEEE Intelligent Systems*, 18(6):54–60, 2003.

P. Puranam, N. Stieglitz, M. Osman, and M. M. Pillutla. Modelling Bounded Rationality in Organizations: Progress and Prospects. *The Academy of Management Annals*, 9(1):337–392, 2015.

S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. Jennings. Agent-based Control for Decentralised Demand Side Management in the Smart Grid. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS '11, pages 5–12, Richland, SC, 2011. International Foundation for Autonomous Agents and Multiagent Systems.

A. Rapoport, A. M. Chammah, and C. J. Orwant. *Prisoner's Dilemma: A Study in Conflict and Cooperation*, volume 165. University of Michigan press, 1965.

M. R. Raupach, G. Marland, P. Ciais, C. Le Quéré, J. G. Canadell, G. Klepper, and C. B. Field. Global and Regional Drivers of Accelerating CO2 Emissions. *Proceedings of the National Academy of Sciences*, 104(24):10288–10293, 2007.

V. Robu, R. Kota, G. Chalkiadakis, A. Rogers, and N. R. Jennings. Cooperative Virtual Power Plant Formation Using Scoring Rules. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 3*, AAMAS '12, pages 1165–1166, Richland, SC, 2012. International Foundation for Autonomous Agents and Multiagent Systems.

V. Robu, M. Vinyals, A. Rogers, and N. R. Jennings. Efficient Buyer Groups With Prediction-of-Use Electricity Tariffs. *IEEE Transactions on Smart Grid*, 9(5):4468–4479, 2018.

M. Roozbehani, M. Dahleh, and S. Mitter. Dynamic Pricing and Stabilization of Supply and Demand in Modern Electric Power Grids. In *2010 First IEEE International Conference on Smart Grid Communications*, pages 543–548, 2010.

M. Roozbehani, M. A. Dahleh, and S. K. Mitter. Volatility of Power Grids Under Real-Time Pricing. *IEEE Transactions on Power Systems*, 27(4):1926–1940, 2012.

H. Rose, A. Rogers, and E. H. Gerding. A Scoring Rule-based Mechanism for Aggregate Demand Prediction in the Smart Grid. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS '12, pages 661–668, Richland, SC, 2012. International Foundation for Autonomous Agents and Multiagent Systems.

A. Rubinstein. *Modeling Bounded Rationality*. MIT press, 1998.

S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall Press, Upper Saddle River, NJ, USA, 3rd edition, 2009.

T. Russell and R. Thaler. The Relevance of Quasi Rationality in Competitive Markets. *The American Economic Review*, 75(5):1071–1082, 1985.

W. Saad, Z. Han, H. V. Poor, and T. Basar. A Noncooperative Game for Double Auction-based Energy Trading between PHEVs and Distribution Grids. In *2011 IEEE International Conference on Smart Grid Communications (SmartGridComm)*, pages 267–272, 2011.

W. Saad, Z. Han, H. V. Poor, and T. Basar. Game-Theoretic Methods for the Smart Grid: An Overview of Microgrid Systems, Demand-Side Management, and Smart Grid Communications. *IEEE Signal Processing Magazine*, 29(5):86–105, 2012.

F. Salah, C. M. Flath, A. Schuller, C. Will, and C. Weinhardt. Morphological Analysis of Energy Services: Paving the way to Quality Differentiation in the Power Sector. *Energy Policy*, 106:614 – 624, 2017.

P. Samadi, H. M. Rad, R. Schober, and V. W. S. Wong. Advanced Demand Side Management for the Future Smart Grid Using Mechanism Design. *IEEE Transactions on Smart Grid*, 3 (3):1170–1180, 2012.

B. Schäfer, M. Matthiae, M. Timme, and D. Witthaut. Decentral Smart Grid Control. *New Journal of Physics*, 17(1):015002, 2015.

S. N. Siddiqi and M. L. Baughman. Reliability Differentiated Pricing of Spinning Reserve. *IEEE Transactions on Power Systems*, 10(3):1211–1218, 1995.

G. C. Silaghi, L. D. Serban, and C. M. Litan. A Time-constrained SLA Negotiation Strategy in Competitive Computational Grids. *Future Generation Computer Systems*, 28(8):1303 – 1315, 2012.

H. A. Simon. Theories of Bounded Rationality. *Decision and organization*, 1(1):161–176, 1972.

H. A. Simon. *Models of Bounded Rationality: Empirically Grounded Economic Reason*, volume 3. MIT press, 1982.

N. Singh and X. Vives. Price and Quantity Competition in a Differentiated Duopoly. *The RAND Journal of Economics*, 15(4):546–554, 1984.

A. Smith. *An Inquiry into the Nature and Causes of the Wealth of Nations*, volume 2. W. Strahan and T. Cadell, London, 1817.

J. M. Smith. Game Theory and the Evolution of Fighting. *On evolution*, pages 8–28, 1972.

J. M. Smith and G. R. Price. The Logic of Animal Conflict. *Nature*, 246(5427):15–18, 1973.

D. F. Spulber. Bertrand Competition when Rivals' Costs are Unknown. *The Journal of Industrial Economics*, 43(1):1–11, 1995.

C. Stagnaro. Competition and Innovation in Retail Electricity Markets: Evidence from Italy. *Economic Affairs*, 37(1):85–101, 2017.

D. O. Stahl and P. W. Wilson. On Players' Models of other Players: Theory and Experimental Evidence. *Games and Economic Behavior*, 10(1):218–254, 1995.

C. H. Stephan and J. Sullivan. Environmental and Energy Implications of Plug-in Hybrid-electric Vehicles. *Environmental science & technology*, 42(4):1185–1190, 2008.

P. Ströhle and C. M. Flath. Local Matching of Flexible Load in Smart Grids. *European Journal of Operational Research*, 253(3):811 – 824, 2016.

P. Ströhle, E. H. Gerding, M. M. de Weerdt, S. Stein, and V. Robu. Online Mechanism Design for Scheduling Non-preemptive Jobs Under Uncertain Supply and Demand. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '14, pages 437–444, Richland, SC, 2014. International Foundation for Autonomous Agents and Multiagent Systems.

R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*, volume 1. MIT press Cambridge, 1998.

V. Syrgkanis and E. Tardos. Bayesian Sequential Auctions. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 929–944, New York, NY, USA, 2012. ACM.

C. Taylor, D. Fudenberg, A. Sasaki, and M. A. Nowak. Evolutionary Game Dynamics in Finite Populations. *Bulletin of Mathematical Biology*, 66(6):1621–1644, 2004.

A. Toffler. *Future Shock*. Bantam, 1990.

K. Tuyls and G. Weiss. Multiagent learning: Basics, Challenges, and Prospects. *AI Magazine*, 33(3):41–41, 2012.

S. Vandael, N. Boucké, T. Holvoet, K. De Craemer, and G. Deconinck. Decentralized Coordination of Plug-in Hybrid Vehicles for Imbalance Reduction in a Smart Grid. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS '11, pages 803–810, Richland, SC, 2011. International Foundation for Autonomous Agents and Multiagent Systems.

S. Vandael, B. Claessens, M. Hommelberg, T. Holvoet, and G. Deconinck. A Scalable Three-Step Approach for Demand Side Management of Plug-in Hybrid Vehicles. *IEEE Transactions on Smart Grid*, 4(2):720–728, 2013.

D. Verma. *Supporting Service Level Agreements on IP Networks*. Macmillan Technical Publishing, London, UK, 1999.

W. Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *The Journal of finance*, 16(1):8–37, 1961.

M. Vinyals, V. Robu, A. Rogers, and N. R. Jennings. Prediction-of-use Games: A Cooperative Game Theory approach to Sustainable Energy Tariffs. In *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '14, pages 829–836, Richland, SC, 2014. International Foundation for Autonomous Agents and Multiagent Systems.

J. von Neumann, O. Morgenstern, H. W. Kuhn, and A. Rubinstein. *Theory of Games and Economic Behavior (60th Anniversary Commemorative Edition)*. Princeton University Press, 1944.

S. K. Vuppala, K. Padmanabh, , and S. Paul. Incorporating Fairness within Demand Response Programs in Smart Grid. In *ISGT 2011*, pages 1–9, 2011.

P. Vytelingum, T. D. Voice, S. D. Ramchurn, A. Rogers, and N. R. Jennings. Agent-based Micro-storage Management for the Smart Grid. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1 - Volume 1*, AAMAS '10, pages 39–46, Richland, SC, 2010. International Foundation for Autonomous Agents and Multiagent Systems.

E. Walraven and M. T. Spaan. A Scenario State Representation for Scheduling Deferrable Loads Under Wind Uncertainty. In *MSDM 2015: AAMAS Workshop on Multiagent Sequential Decision Making Under Uncertainty, Istanbul, Turkey, May 2015*, 2015a.

E. Walraven and M. T. J. Spaan. Planning Under Uncertainty with Weighted State Scenarios. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, UAI'15, pages 912–921, Arlington, Virginia, United States, 2015b. AUAI Press.

Q. Wang, C. Zhang, Y. Ding, G. Xydis, J. Wang, and J. Østergaard. Review of Real-time Electricity Markets for Integrating Distributed Energy Resources and Demand Response. *Applied Energy*, 138:695 – 706, 2015.

Y. H. Wang. On the Number of Successes in Independent Trials. *Statistica Sinica*, 3(2): 295–312, 1993.

C. Weber. *Uncertainty in the Electric Power Industry: Methods and Models for Decision Support*, volume 77. Springer Science & Business Media, 2006.

J. W. Weibull. *Evolutionary Game Theory*. MIT press, 1997.

A. Weidlich and D. Veit. A Critical Survey of Agent-based Wholesale Electricity Market Models. *Energy Economics*, 30(4):1728 – 1759, 2008.

G. Weiss. *Multiagent Systems*. The MIT Press, 2013.

M. J. Wooldridge. *Introduction to Multiagent Systems*. John Wiley & Sons, Inc., New York, NY, USA, 2001.

M. Wunder, M. Kaisers, M. Littman, and J. R. Yaros. A Cognitive Hierarchy Model Applied to the Lemonade Game. In *Proceedings of the 3rd AAAI Conference on Interactive Decision Theory and Game Theory*, AAAIWS'10-03, pages 66–73. AAAI Press, 2010a.

M. Wunder, M. Littman, and M. Babes. Classes of Multiagent Q-learning Dynamics with e-greedy Exploration. In *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ICML'10, pages 1167–1174, USA, 2010b. Omnipress.

Y. Yan, Y. Qian, H. Sharif, and D. Tipper. A Survey on Smart Grid Communication Infrastructures: Motivations, Requirements and Challenges. *IEEE Communications Surveys Tutorials*, 15(1):5–20, 2013.

J. Zhang, Q. Da, and Y. Wang. The Dynamics of Bertrand Model with Bounded Rationality. *Chaos, Solitons & Fractals*, 39(5):2048 – 2055, 2009.

Y. Zhang, C. Lee, D. Niyato, and P. Wang. Auction Approaches for Resource Allocation in Wireless Systems: A Survey. *IEEE Communications Surveys Tutorials*, 15(3):1020–1041, 2013.

D. Zhao, S. D. Ramchurn, and N. R. Jennings. Fault Tolerant Mechanism Design for General Task Allocation. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, AAMAS '16, pages 323–331, Richland, SC, 2016. International Foundation for Autonomous Agents and Multiagent Systems.

# SIKS Dissertation Series

2011 01 Botond Cseke (RUN), Variational Algorithms for Bayesian Inference in Latent Gaussian Models
02 Nick Tinnemeier (UU), Organizing Agent Organizations. Syntax and Operational Semantics of an Organization-Oriented Programming Language
03 Jan Martijn van der Werf (TUE), Compositional Design and Verification of Component-Based Information Systems
04 Hado van Hasselt (UU), Insights in Reinforcement Learning; Formal analysis and empirical evaluation of temporal-difference
05 Bas van der Raadt (VU), Enterprise Architecture Coming of Age - Increasing the Performance of an Emerging Discipline.
06 Yiwen Wang (TUE), Semantically-Enhanced Recommendations in Cultural Heritage
07 Yujia Cao (UT), Multimodal Information Presentation for High Load Human Computer Interaction
08 Nieske Vergunst (UU), BDI-based Generation of Robust Task-Oriented Dialogues
09 Tim de Jong (OU), Contextualised Mobile Media for Learning
10 Bart Bogaert (UvT), Cloud Content Contention
11 Dhaval Vyas (UT), Designing for Awareness: An Experience-focused HCI Perspective
12 Carmen Bratosin (TUE), Grid Architecture for Distributed Process Mining
13 Xiaoyu Mao (UvT), Airport under Control. Multiagent Scheduling for Airport Ground Handling
14 Milan Lovric (EUR), Behavioral Finance and Agent-Based Artificial Markets
15 Marijn Koolen (UvA), The Meaning of Structure: the Value of Link Evidence for Information Retrieval
16 Maarten Schadd (UM), Selective Search in Games of Different Complexity
17 Jiyin He (UVA), Exploring Topic Structure: Coherence, Diversity and Relatedness
18 Mark Ponsen (UM), Strategic Decision-Making in complex games
19 Ellen Rusman (OU), The Mind's Eye on Personal Profiles
20 Qing Gu (VU), Guiding service-oriented software engineering - A view-based approach
21 Linda Terlouw (TUD), Modularization and Specification of Service-Oriented Systems
22 Junte Zhang (UVA), System Evaluation of Archival Description and Access
23 Wouter Weerkamp (UVA), Finding People and their Utterances in Social Media
24 Herwin van Welbergen (UT), Behavior Generation for Interpersonal Coordination with Virtual Humans On Specifying, Scheduling and Realizing Multimodal Virtual Human Behavior
25 Syed Waqar ul Qounain Jaffry (VU), Analysis and Validation of Models for Trust Dynamics
26 Matthijs Aart Pontier (VU), Virtual Agents for Human Communication - Emotion Regulation and Involvement-Distance Trade-Offs in Embodied Conversational Agents and Robots

27 Aniel Bhulai (VU), Dynamic website optimization through autonomous management of design patterns
28 Rianne Kaptein (UVA), Effective Focused Retrieval by Exploiting Query Context and Document Structure
29 Faisal Kamiran (TUE), Discrimination-aware Classification
30 Egon van den Broek (UT), Affective Signal Processing (ASP): Unraveling the mystery of emotions
31 Ludo Waltman (EUR), Computational and Game-Theoretic Approaches for Modeling Bounded Rationality
32 Nees-Jan van Eck (EUR), Methodological Advances in Bibliometric Mapping of Science
33 Tom van der Weide (UU), Arguing to Motivate Decisions
34 Paolo Turrini (UU), Strategic Reasoning in Interdependence: Logical and Game-theoretical Investigations
35 Maaike Harbers (UU), Explaining Agent Behavior in Virtual Training
36 Erik van der Spek (UU), Experiments in serious game design: a cognitive approach
37 Adriana Burlutiu (RUN), Machine Learning for Pairwise Data, Applications for Preference Learning and Supervised Network Inference
38 Nyree Lemmens (UM), Bee-inspired Distributed Optimization
39 Joost Westra (UU), Organizing Adaptation using Agents in Serious Games
40 Viktor Clerc (VU), Architectural Knowledge Management in Global Software Development
41 Luan Ibraimi (UT), Cryptographically Enforced Distributed Data Access Control
42 Michal Sindlar (UU), Explaining Behavior through Mental State Attribution
43 Henk van der Schuur (UU), Process Improvement through Software Operation Knowledge
44 Boris Reuderink (UT), Robust Brain-Computer Interfaces
45 Herman Stehouwer (UvT), Statistical Language Models for Alternative Sequence Selection
46 Beibei Hu (TUD), Towards Contextualized Information Delivery: A Rule-based Architecture for the Domain of Mobile Police Work
47 Azizi Bin Ab Aziz (VU), Exploring Computational Models for Intelligent Support of Persons with Depression
48 Mark Ter Maat (UT), Response Selection and Turn-taking for a Sensitive Artificial Listening Agent
49 Andreea Niculescu (UT), Conversational interfaces for task-oriented spoken dialogues: design aspects influencing interaction quality

---

2012 01 Terry Kakeeto (UvT), Relationship Marketing for SMEs in Uganda
02 Muhammad Umair (VU), Adaptivity, emotion, and Rationality in Human and Ambient Agent Models
03 Adam Vanya (VU), Supporting Architecture Evolution by Mining Software Repositories
04 Jurriaan Souer (UU), Development of Content Management System-based Web Applications
05 Marijn Plomp (UU), Maturing Interorganisational Information Systems
06 Wolfgang Reinhardt (OU), Awareness Support for Knowledge Workers in Research Networks
07 Rianne van Lambalgen (VU), When the Going Gets Tough: Exploring Agent-based Models of Human Performance under Demanding Conditions

08 Gerben de Vries (UVA), Kernel Methods for Vessel Trajectories
09 Ricardo Neisse (UT), Trust and Privacy Management Support for Context-Aware Service Platforms
10 David Smits (TUE), Towards a Generic Distributed Adaptive Hypermedia Environment
11 J.C.B. Rantham Prabhakara (TUE), Process Mining in the Large: Preprocessing, Discovery, and Diagnostics
12 Kees van der Sluijs (TUE), Model Driven Design and Data Integration in Semantic Web Information Systems
13 Suleman Shahid (UvT), Fun and Face: Exploring non-verbal expressions of emotion during playful interactions
14 Evgeny Knutov (TUE), Generic Adaptation Framework for Unifying Adaptive Web-based Systems
15 Natalie van der Wal (VU), Social Agents. Agent-Based Modelling of Integrated Internal and Social Dynamics of Cognitive and Affective Processes.
16 Fiemke Both (VU), Helping people by understanding them - Ambient Agents supporting task execution and depression treatment
17 Amal Elgammal (UvT), Towards a Comprehensive Framework for Business Process Compliance
18 Eltjo Poort (VU), Improving Solution Architecting Practices
19 Helen Schonenberg (TUE), What's Next? Operational Support for Business Process Execution
20 Ali Bahramisharif (RUN), Covert Visual Spatial Attention, a Robust Paradigm for Brain-Computer Interfacing
21 Roberto Cornacchia (TUD), Querying Sparse Matrices for Information Retrieval
22 Thijs Vis (UvT), Intelligence, politie en veiligheidsdienst: verenigbare grootheden?
23 Christian Muehl (UT), Toward Affective Brain-Computer Interfaces: Exploring the Neurophysiology of Affect during Human Media Interaction
24 Laurens van der Werff (UT), Evaluation of Noisy Transcripts for Spoken Document Retrieval
25 Silja Eckartz (UT), Managing the Business Case Development in Inter-Organizational IT Projects: A Methodology and its Application
26 Emile de Maat (UVA), Making Sense of Legal Text
27 Hayrettin Gurkok (UT), Mind the Sheep! User Experience Evaluation & Brain-Computer Interface Games
28 Nancy Pascall (UvT), Engendering Technology Empowering Women
29 Almer Tigelaar (UT), Peer-to-Peer Information Retrieval
30 Alina Pommeranz (TUD), Designing Human-Centered Systems for Reflective Decision Making
31 Emily Bagarukayo (RUN), A Learning by Construction Approach for Higher Order Cognitive Skills Improvement, Building Capacity and Infrastructure
32 Wietske Visser (TUD), Qualitative multi-criteria preference representation and reasoning
33 Rory Sie (OUN), Coalitions in Cooperation Networks (COCOON)
34 Pavol Jancura (RUN), Evolutionary analysis in PPI networks and applications
35 Evert Haasdijk (VU), Never Too Old To Learn – On-line Evolution of Controllers in Swarm- and Modular Robotics
36 Denis Ssebugwawo (RUN), Analysis and Evaluation of Collaborative Modeling Processes
37 Agnes Nakakawa (RUN), A Collaboration Process for Enterprise Architecture Creation

38 Selmar Smit (VU), Parameter Tuning and Scientific Testing in Evolutionary Algorithms
39 Hassan Fatemi (UT), Risk-aware design of value and coordination networks
40 Agus Gunawan (UvT), Information Access for SMEs in Indonesia
41 Sebastian Kelle (OU), Game Design Patterns for Learning
42 Dominique Verpoorten (OU), Reflection Amplifiers in self-regulated Learning
43 Withdrawn
44 Anna Tordai (VU), On Combining Alignment Techniques
45 Benedikt Kratz (UvT), A Model and Language for Business-aware Transactions
46 Simon Carter (UVA), Exploration and Exploitation of Multilingual Data for Statistical Machine Translation
47 Manos Tsagkias (UVA), Mining Social Media: Tracking Content and Predicting Behavior
48 Jorn Bakker (TUE), Handling Abrupt Changes in Evolving Time-series Data
49 Michael Kaisers (UM), Learning against Learning - Evolutionary dynamics of reinforcement learning algorithms in strategic interactions
50 Steven van Kervel (TUD), Ontologogy driven Enterprise Information Systems Engineering
51 Jeroen de Jong (TUD), Heuristics in Dynamic Sceduling; a practical framework with a case study in elevator dispatching

---

2013 01 Viorel Milea (EUR), News Analytics for Financial Decision Support
02 Erietta Liarou (CWI), MonetDB/DataCell: Leveraging the Column-store Database Technology for Efficient and Scalable Stream Processing
03 Szymon Klarman (VU), Reasoning with Contexts in Description Logics
04 Chetan Yadati (TUD), Coordinating autonomous planning and scheduling
05 Dulce Pumareja (UT), Groupware Requirements Evolutions Patterns
06 Romulo Goncalves (CWI), The Data Cyclotron: Juggling Data and Queries for a Data Warehouse Audience
07 Giel van Lankveld (UvT), Quantifying Individual Player Differences
08 Robbert-Jan Merk (VU), Making enemies: cognitive modeling for opponent agents in fighter pilot simulators
09 Fabio Gori (RUN), Metagenomic Data Analysis: Computational Methods and Applications
10 Jeewanie Jayasinghe Arachchige (UvT), A Unified Modeling Framework for Service Design.
11 Evangelos Pournaras (TUD), Multi-level Reconfigurable Self-organization in Overlay Services
12 Marian Razavian (VU), Knowledge-driven Migration to Services
13 Mohammad Safiri (UT), Service Tailoring: User-centric creation of integrated IT-based homecare services to support independent living of elderly
14 Jafar Tanha (UVA), Ensemble Approaches to Semi-Supervised Learning Learning
15 Daniel Hennes (UM), Multiagent Learning - Dynamic Games and Applications
16 Eric Kok (UU), Exploring the practical benefits of argumentation in multi-agent deliberation
17 Koen Kok (VU), The PowerMatcher: Smart Coordination for the Smart Electricity Grid
18 Jeroen Janssens (UvT), Outlier Selection and One-Class Classification
19 Renze Steenhuizen (TUD), Coordinated Multi-Agent Planning and Scheduling
20 Katja Hofmann (UvA), Fast and Reliable Online Learning to Rank for Information Retrieval

21 Sander Wubben (UvT), Text-to-text generation by monolingual machine translation
22 Tom Claassen (RUN), Causal Discovery and Logic
23 Patricio de Alencar Silva (UvT), Value Activity Monitoring
24 Haitham Bou Ammar (UM), Automated Transfer in Reinforcement Learning
25 Agnieszka Anna Latoszek-Berendsen (UM), Intention-based Decision Support. A new way of representing and implementing clinical guidelines in a Decision Support System
26 Alireza Zarghami (UT), Architectural Support for Dynamic Homecare Service Provisioning
27 Mohammad Huq (UT), Inference-based Framework Managing Data Provenance
28 Frans van der Sluis (UT), When Complexity becomes Interesting: An Inquiry into the Information eXperience
29 Iwan de Kok (UT), Listening Heads
30 Joyce Nakatumba (TUE), Resource-Aware Business Process Management: Analysis and Support
31 Dinh Khoa Nguyen (UvT), Blueprint Model and Language for Engineering Cloud Applications
32 Kamakshi Rajagopal (OUN), Networking For Learning; The role of Networking in a Lifelong Learner's Professional Development
33 Qi Gao (TUD), User Modeling and Personalization in the Microblogging Sphere
34 Kien Tjin-Kam-Jet (UT), Distributed Deep Web Search
35 Abdallah El Ali (UvA), Minimal Mobile Human Computer Interaction
36 Than Lam Hoang (TUe), Pattern Mining in Data Streams
37 Dirk Börner (OUN), Ambient Learning Displays
38 Eelco den Heijer (VU), Autonomous Evolutionary Art
39 Joop de Jong (TUD), A Method for Enterprise Ontology based Design of Enterprise Information Systems
40 Pim Nijssen (UM), Monte-Carlo Tree Search for Multi-Player Games
41 Jochem Liem (UVA), Supporting the Conceptual Modelling of Dynamic Systems: A Knowledge Engineering Perspective on Qualitative Reasoning
42 Léon Planken (TUD), Algorithms for Simple Temporal Reasoning
43 Marc Bron (UVA), Exploration and Contextualization through Interaction and Concepts

---

2014 01 Nicola Barile (UU), Studies in Learning Monotone Models from Data
02 Fiona Tuliyano (RUN), Combining System Dynamics with a Domain Modeling Method
03 Sergio Raul Duarte Torres (UT), Information Retrieval for Children: Search Behavior and Solutions
04 Hanna Jochmann-Mannak (UT), Websites for children: search strategies and interface design - Three studies on children's search performance and evaluation
05 Jurriaan van Reijsen (UU), Knowledge Perspectives on Advancing Dynamic Capability
06 Damian Tamburri (VU), Supporting Networked Software Development
07 Arya Adriansyah (TUE), Aligning Observed and Modeled Behavior
08 Samur Araujo (TUD), Data Integration over Distributed and Heterogeneous Data Endpoints
09 Philip Jackson (UvT), Toward Human-Level Artificial Intelligence: Representation and Computation of Meaning in Natural Language

10 Ivan Salvador Razo Zapata (VU), Service Value Networks
11 Janneke van der Zwaan (TUD), An Empathic Virtual Buddy for Social Support
12 Willem van Willigen (VU), Look Ma, No Hands: Aspects of Autonomous Vehicle Control
13 Arlette van Wissen (VU), Agent-Based Support for Behavior Change: Models and Applications in Health and Safety Domains
14 Yangyang Shi (TUD), Language Models With Meta-information
15 Natalya Mogles (VU), Agent-Based Analysis and Support of Human Functioning in Complex Socio-Technical Systems: Applications in Safety and Healthcare
16 Krystyna Milian (VU), Supporting trial recruitment and design by automatically interpreting eligibility criteria
17 Kathrin Dentler (VU), Computing healthcare quality indicators automatically: Secondary Use of Patient Data and Semantic Interoperability
18 Mattijs Ghijsen (UVA), Methods and Models for the Design and Study of Dynamic Agent Organizations
19 Vinicius Ramos (TUE), Adaptive Hypermedia Courses: Qualitative and Quantitative Evaluation and Tool Support
20 Mena Habib (UT), Named Entity Extraction and Disambiguation for Informal Text: The Missing Link
21 Kassidy Clark (TUD), Negotiation and Monitoring in Open Environments
22 Marieke Peeters (UU), Personalized Educational Games - Developing agent-supported scenario-based training
23 Eleftherios Sidirourgos (UvA/CWI), Space Efficient Indexes for the Big Data Era
24 Davide Ceolin (VU), Trusting Semi-structured Web Data
25 Martijn Lappenschaar (RUN), New network models for the analysis of disease interaction
26 Tim Baarslag (TUD), What to Bid and When to Stop
27 Rui Jorge Almeida (EUR), Conditional Density Models Integrating Fuzzy and Probabilistic Representations of Uncertainty
28 Anna Chmielowiec (VU), Decentralized k-Clique Matching
29 Jaap Kabbedijk (UU), Variability in Multi-Tenant Enterprise Software
30 Peter de Cock (UvT), Anticipating Criminal Behaviour
31 Leo van Moergestel (UU), Agent Technology in Agile Multiparallel Manufacturing and Product Support
32 Naser Ayat (UvA), On Entity Resolution in Probabilistic Data
33 Tesfa Tegegne (RUN), Service Discovery in eHealth
34 Christina Manteli (VU), The Effect of Governance in Global Software Development: Analyzing Transactive Memory Systems.
35 Joost van Ooijen (UU), Cognitive Agents in Virtual Worlds: A Middleware Design Approach
36 Joos Buijs (TUE), Flexible Evolutionary Algorithms for Mining Structured Process Models
37 Maral Dadvar (UT), Experts and Machines United Against Cyberbullying
38 Danny Plass-Oude Bos (UT), Making brain-computer interfaces better: improving usability through post-processing.
39 Jasmina Maric (UvT), Web Communities, Immigration, and Social Capital
40 Walter Omona (RUN), A Framework for Knowledge Management Using ICT in Higher Education
41 Frederic Hogenboom (EUR), Automated Detection of Financial Events in News Text

42 Carsten Eijckhof (CWI/TUD), Contextual Multidimensional Relevance Models
43 Kevin Vlaanderen (UU), Supporting Process Improvement using Method Increments
44 Paulien Meesters (UvT), Intelligent Blauw. Met als ondertitel: Intelligence-gestuurde politiezorg in gebiedsgebonden eenheden.
45 Birgit Schmitz (OUN), Mobile Games for Learning: A Pattern-Based Approach
46 Ke Tao (TUD), Social Web Data Analytics: Relevance, Redundancy, Diversity
47 Shangsong Liang (UVA), Fusion and Diversification in Information Retrieval

---

2015 01 Niels Netten (UvA), Machine Learning for Relevance of Information in Crisis Response
02 Faiza Bukhsh (UvT), Smart auditing: Innovative Compliance Checking in Customs Controls
03 Twan van Laarhoven (RUN), Machine learning for network data
04 Howard Spoelstra (OUN), Collaborations in Open Learning Environments
05 Christoph Bösch (UT), Cryptographically Enforced Search Pattern Hiding
06 Farideh Heidari (TUD), Business Process Quality Computation - Computing Non-Functional Requirements to Improve Business Processes
07 Maria-Hendrike Peetz (UvA), Time-Aware Online Reputation Analysis
08 Jie Jiang (TUD), Organizational Compliance: An agent-based model for designing and evaluating organizational interactions
09 Randy Klaassen (UT), HCI Perspectives on Behavior Change Support Systems
10 Henry Hermans (OUN), OpenU: design of an integrated system to support lifelong learning
11 Yongming Luo (TUE), Designing algorithms for big graph datasets: A study of computing bisimulation and joins
12 Julie M. Birkholz (VU), Modi Operandi of Social Network Dynamics: The Effect of Context on Scientific Collaboration Networks
13 Giuseppe Procaccianti (VU), Energy-Efficient Software
14 Bart van Straalen (UT), A cognitive approach to modeling bad news conversations
15 Klaas Andries de Graaf (VU), Ontology-based Software Architecture Documentation
16 Changyun Wei (UT), Cognitive Coordination for Cooperative Multi-Robot Teamwork
17 André van Cleeff (UT), Physical and Digital Security Mechanisms: Properties, Combinations and Trade-offs
18 Holger Pirk (CWI), Waste Not, Want Not! - Managing Relational Data in Asymmetric Memories
19 Bernardo Tabuenca (OUN), Ubiquitous Technology for Lifelong Learners
20 Lois Vanhée (UU), Using Culture and Values to Support Flexible Coordination
21 Sibren Fetter (OUN), Using Peer-Support to Expand and Stabilize Online Learning
22 Zhemin Zhu (UT), Co-occurrence Rate Networks
23 Luit Gazendam (VU), Cataloguer Support in Cultural Heritage
24 Richard Berendsen (UVA), Finding People, Papers, and Posts: Vertical Search Algorithms and Evaluation
25 Steven Woudenberg (UU), Bayesian Tools for Early Disease Detection
26 Alexander Hogenboom (EUR), Sentiment Analysis of Text Guided by Semantics and Structure
27 Sándor Héman (CWI), Updating compressed colomn stores

28 Janet Bagorogoza (TiU), Knowledge Management and High Performance; The Uganda Financial Institutions Model for HPO
29 Hendrik Baier (UM), Monte-Carlo Tree Search Enhancements for One-Player and Two-Player Domains
30 Kiavash Bahreini (OU), Real-time Multimodal Emotion Recognition in E-Learning
31 Yakup Koç (TUD), On the robustness of Power Grids
32 Jerome Gard (UL), Corporate Venture Management in SMEs
33 Frederik Schadd (TUD), Ontology Mapping with Auxiliary Resources
34 Victor de Graaf (UT), Gesocial Recommender Systems
35 Jungxao Xu (TUD), Affective Body Language of Humanoid Robots: Perception and Effects in Human Robot Interaction

---

2016 01 Syed Saiden Abbas (RUN), Recognition of Shapes by Humans and Machines
02 Michiel Christiaan Meulendijk (UU), Optimizing medication reviews through decision support: prescribing a better pill to swallow
03 Maya Sappelli (RUN), Knowledge Work in Context: User Centered Knowledge Worker Support
04 Laurens Rietveld (VU), Publishing and Consuming Linked Data
05 Evgeny Sherkhonov (UVA), Expanded Acyclic Queries: Containment and an Application in Explaining Missing Answers
06 Michel Wilson (TUD), Robust scheduling in an uncertain environment
07 Jeroen de Man (VU), Measuring and modeling negative emotions for virtual training
08 Matje van de Camp (TiU), A Link to the Past: Constructing Historical Social Networks from Unstructured Data
09 Archana Nottamkandath (VU), Trusting Crowdsourced Information on Cultural Artefacts
10 George Karafotias (VUA), Parameter Control for Evolutionary Algorithms
11 Anne Schuth (UVA), Search Engines that Learn from Their Users
12 Max Knobbout (UU), Logics for Modelling and Verifying Normative Multi-Agent Systems
13 Nana Baah Gyan (VU), The Web, Speech Technologies and Rural Development in West Africa - An ICT4D Approach
14 Ravi Khadka (UU), Revisiting Legacy Software System Modernization
15 Steffen Michels (RUN), Hybrid Probabilistic Logics - Theoretical Aspects, Algorithms and Experiments
16 Guangliang Li (UVA), Socially Intelligent Autonomous Agents that Learn from Human Reward
17 Berend Weel (VU), Towards Embodied Evolution of Robot Organisms
18 Albert Meroño Peñuela (VU), Refining Statistical Data on the Web
19 Julia Efremova (Tu/e), Mining Social Structures from Genealogical Data
20 Daan Odijk (UVA), Context & Semantics in News & Web Search
21 Alejandro Moreno Célleri (UT), From Traditional to Interactive Playspaces: Automatic Analysis of Player Behavior in the Interactive Tag Playground
22 Grace Lewis (VU), Software Architecture Strategies for Cyber-Foraging Systems
23 Fei Cai (UVA), Query Auto Completion in Information Retrieval
24 Brend Wanders (UT), Repurposing and Probabilistic Integration of Data; An Iterative and data model independent approach
25 Julia Kiseleva (TU/e), Using Contextual Information to Understand Searching and Browsing Behavior

26 Dilhan Thilakarathne (VU), In or Out of Control: Exploring Computational Models to Study the Role of Human Awareness and Control in Behavioural Choices, with Applications in Aviation and Energy Management Domains
27 Wen Li (TUD), Understanding Geo-spatial Information on Social Media
28 Mingxin Zhang (TUD), Large-scale Agent-based Social Simulation - A study on epidemic prediction and control
29 Nicolas Höning (TUD), Peak reduction in decentralised electricity systems - Markets and prices for flexible planning
30 Ruud Mattheij (UvT), The Eyes Have It
31 Mohammad Khelghati (UT), Deep web content monitoring
32 Eelco Vriezekolk (UT), Assessing Telecommunication Service Availability Risks for Crisis Organisations
33 Peter Bloem (UVA), Single Sample Statistics, exercises in learning from just one example
34 Dennis Schunselaar (TUE), Configurable Process Trees: Elicitation, Analysis, and Enactment
35 Zhaochun Ren (UVA), Monitoring Social Media: Summarization, Classification and Recommendation
36 Daphne Karreman (UT), Beyond R2D2: The design of nonverbal interaction behavior optimized for robot-specific morphologies
37 Giovanni Sileno (UvA), Aligning Law and Action - a conceptual and computational inquiry
38 Andrea Minuto (UT), Materials that Matter - Smart Materials meet Art & Interaction Design
39 Merijn Bruijnes (UT), Believable Suspect Agents; Response and Interpersonal Style Selection for an Artificial Suspect
40 Christian Detweiler (TUD), Accounting for Values in Design
41 Thomas King (TUD), Governing Governance: A Formal Framework for Analysing Institutional Design and Enactment Governance
42 Spyros Martzoukos (UVA), Combinatorial and Compositional Aspects of Bilingual Aligned Corpora
43 Saskia Koldijk (RUN), Context-Aware Support for Stress Self-Management: From Theory to Practice
44 Thibault Sellam (UVA), Automatic Assistants for Database Exploration
45 Bram van de Laar (UT), Experiencing Brain-Computer Interface Control
46 Jorge Gallego Perez (UT), Robots to Make you Happy
47 Christina Weber (UL), Real-time foresight - Preparedness for dynamic innovation networks
48 Tanja Buttler (TUD), Collecting Lessons Learned
49 Gleb Polevoy (TUD), Participation and Interaction in Projects. A Game-Theoretic Analysis
50 Yan Wang (UVT), The Bridge of Dreams: Towards a Method for Operational Performance Alignment in IT-enabled Service Supply Chains

---

2017 01 Jan-Jaap Oerlemans (UL), Investigating Cybercrime
02 Sjoerd Timmer (UU), Designing and Understanding Forensic Bayesian Networks using Argumentation
03 Daniël Harold Telgen (UU), Grid Manufacturing; A Cyber-Physical Approach with Autonomous Products and Reconfigurable Manufacturing Machines
04 Mrunal Gawade (CWI), Multi-core Parallelism in a Column-store
05 Mahdieh Shadi (UVA), Collaboration Behavior

06 Damir Vandic (EUR), Intelligent Information Systems for Web Product Search
07 Roel Bertens (UU), Insight in Information: from Abstract to Anomaly
08 Rob Konijn (VU) , Detecting Interesting Differences:Data Mining in Health Insurance Data using Outlier Detection and Subgroup Discovery
09 Dong Nguyen (UT), Text as Social and Cultural Data: A Computational Perspective on Variation in Text
10 Robby van Delden (UT), (Steering) Interactive Play Behavior
11 Florian Kunneman (RUN), Modelling patterns of time and emotion in Twitter #anticipointment
12 Sander Leemans (TUE), Robust Process Mining with Guarantees
13 Gijs Huisman (UT), Social Touch Technology - Extending the reach of social touch through haptic technology
14 Shoshannah Tekofsky (UvT), You Are Who You Play You Are: Modelling Player Traits from Video Game Behavior
15 Peter Berck (RUN), Memory-Based Text Correction
16 Aleksandr Chuklin (UVA), Understanding and Modeling Users of Modern Search Engines
17 Daniel Dimov (UL), Crowdsourced Online Dispute Resolution
18 Ridho Reinanda (UVA), Entity Associations for Search
19 Jeroen Vuurens (UT), Proximity of Terms, Texts and Semantic Vectors in Information Retrieval
20 Mohammadbashir Sedighi (TUD), Fostering Engagement in Knowledge Sharing: The Role of Perceived Benefits, Costs and Visibility
21 Jeroen Linssen (UT), Meta Matters in Interactive Storytelling and Serious Gaming (A Play on Worlds)
22 Sara Magliacane (VU), Logics for causal inference under uncertainty
23 David Graus (UVA), Entities of Interest — Discovery in Digital Traces
24 Chang Wang (TUD), Use of Affordances for Efficient Robot Learning
25 Veruska Zamborlini (VU), Knowledge Representation for Clinical Guidelines, with applications to Multimorbidity Analysis and Literature Search
26 Merel Jung (UT), Socially intelligent robots that understand and respond to human touch
27 Michiel Joosse (UT), Investigating Positioning and Gaze Behaviors of Social Robots: People's Preferences, Perceptions and Behaviors
28 John Klein (VU), Architecture Practices for Complex Contexts
29 Adel Alhuraibi (UvT), From IT-BusinessStrategic Alignment to Performance: A Moderated Mediation Model of Social Innovation, and Enterprise Governance of IT"
30 Wilma Latuny (UvT), The Power of Facial Expressions
31 Ben Ruijl (UL), Advances in computational methods for QFT calculations
32 Thaer Samar (RUN), Access to and Retrievability of Content in Web Archives
33 Brigit van Loggem (OU), Towards a Design Rationale for Software Documentation: A Model of Computer-Mediated Activity
34 Maren Scheffel (OU), The Evaluation Framework for Learning Analytics
35 Martine de Vos (VU), Interpreting natural science spreadsheets
36 Yuanhao Guo (UL), Shape Analysis for Phenotype Characterisation from High-throughput Imaging
37 Alejandro Montes Garcia (TUE), WiBAF: A Within Browser Adaptation Framework that Enables Control over Privacy
38 Alex Kayal (TUD), Normative Social Applications

39 Sara Ahmadi (RUN), Exploiting properties of the human auditory system and compressive sensing methods to increase noise robustness in ASR

40 Altaf Hussain Abro (VUA), Steer your Mind: Computational Exploration of Human Control in Relation to Emotions, Desires and Social Support For applications in human-aware support systems

41 Adnan Manzoor (VUA), Minding a Healthy Lifestyle: An Exploration of Mental Processes and a Smart Environment to Provide Support for a Healthy Lifestyle

42 Elena Sokolova (RUN), Causal discovery from mixed and missing data with applications on ADHD datasets

43 Maaike de Boer (RUN), Semantic Mapping in Video Retrieval

44 Garm Lucassen (UU), Understanding User Stories - Computational Linguistics in Agile Requirements Engineering

45 Bas Testerink (UU), Decentralized Runtime Norm Enforcement

46 Jan Schneider (OU), Sensor-based Learning Support

47 Jie Yang (TUD), Crowd Knowledge Creation Acceleration

48 Angel Suarez (OU), Collaborative inquiry-based learning

---

2018 01 Han van der Aa (VUA), Comparing and Aligning Process Representations

02 Felix Mannhardt (TUE), Multi-perspective Process Mining

03 Steven Bosems (UT), Causal Models For Well-Being: Knowledge Modeling, Model-Driven Development of Context-Aware Applications, and Behavior Prediction

04 Jordan Janeiro (TUD), Flexible Coordination Support for Diagnosis Teams in Data-Centric Engineering Tasks

05 Hugo Huurdeman (UVA), Supporting the Complex Dynamics of the Information Seeking Process

06 Dan Ionita (UT), Model-Driven Information Security Risk Assessment of Socio-Technical Systems

07 Jieting Luo (UU), A formal account of opportunism in multi-agent systems

08 Rick Smetsers (RUN), Advances in Model Learning for Software Systems

09 Xu Xie (TUD), Data Assimilation in Discrete Event Simulations

10 Julienka Mollee (VUA), Moving forward: supporting physical activity behavior change through intelligent technology

11 Mahdi Sargolzaei (UVA), Enabling Framework for Service-oriented Collaborative Networks

12 Xixi Lu (TUE), Using behavioral context in process mining

13 Seyed Amin Tabatabaei (VUA), Computing a Sustainable Future

14 Bart Joosten (UVT), Detecting Social Signals with Spatiotemporal Gabor Filters

15 Naser Davarzani (UM), Biomarker discovery in heart failure

16 Jaebok Kim (UT), Automatic recognition of engagement and emotion in a group of children

17 Jianpeng Zhang (TUE), On Graph Sample Clustering

18 Henriette Nakad (UL), De Notaris en Private Rechtspraak

19 Minh Duc Pham (VUA), Emergent relational schemas for RDF

20 Manxia Liu (RUN), Time and Bayesian Networks

21 Aad Slootmaker (OUN), EMERGO: a generic platform for authoring and playing scenario-based serious games

22 Eric Fernandes de Mello Araujo (VUA), Contagious: Modeling the Spread of Behaviours, Perceptions and Emotions in Social Networks

23 Kim Schouten (EUR), Semantics-driven Aspect-Based Sentiment Analysis

24 Jered Vroon (UT), Responsive Social Positioning Behaviour for Semi-Autonomous Telepresence Robots

25 Riste Gligorov (VUA), Serious Games in Audio-Visual Collections

26 Roelof Anne Jelle de Vries (UT),Theory-Based and Tailor-Made: Motivational Messages for Behavior Change Technology

27 Maikel Leemans (TUE), Hierarchical Process Mining for Scalable Software Analysis

28 Christian Willemse (UT), Social Touch Technologies: How they feel and how they make you feel

29 Yu Gu (UVT), Emotion Recognition from Mandarin Speech

30 Wouter Beek, The "K" in "semantic web" stands for "knowledge": scaling semantics to the web

---

2019 01 Rob van Eijk (UL),Web privacy measurement in real-time bidding systems. A graph-based approach to RTB system classification

02 Emmanuelle Beauxis Aussalet (CWI, UU), Statistics and Visualizations for Assessing Class Size Uncertainty

03 Eduardo Gonzalez Lopez de Murillas (TUE), Process Mining on Databases: Extracting Event Data from Real Life Data Sources

04 Ridho Rahmadi (RUN), Finding stable causal structures from clinical data

05 Sebastiaan van Zelst (TUE), Process Mining with Streaming Data

06 Chris Dijkshoorn (VU), Nichesourcing for Improving Access to Linked Cultural Heritage Datasets

07 Soude Fazeli (TUD), Recommender Systems in Social Learning Platforms

08 Frits de Nijs (TUD), Resource-constrained Multi-agent Markov Decision Processes

09 Fahimeh Alizadeh Moghaddam (UVA), Self-adaptation for energy efficiency in software systems

10 Qing Chuan Ye (EUR), Multi-objective Optimization Methods for Allocation and Prediction

11 Yue Zhao (TUD), Learning Analytics Technology to Understand Learner Behavioral Engagement in MOOCs

12 Jacqueline Heinerman (VU), Better Together

13 Guanliang Chen (TUD), MOOC Analytics: Learner Modeling and Content Generation

14 Daniel Davis (TUD), Large-Scale Learning Analytics: Modeling Learner Behavior & Improving Learning Outcomes in Massive Open Online Courses

15 Erwin Walraven (TUD), Planning under Uncertainty in Constrained and Partially Observable Environments

16 Guangming Li (TUE), Process Mining based on Object-Centric Behavioral Constraint (OCBC) Models

17 Ali Hurriyetoglu (RUN),Extracting actionable information from microtexts

18 Gerard Wagenaar (UU), Artefacts in Agile Team Communication

19 Vincent Koeman (TUD), Tools for Developing Cognitive Agents

20 Chide Groenouwe (UU), Fostering technically augmented human collective intelligence

21 Cong Liu (TUE), Software Data Analytics: Architectural Model Discovery and Design Pattern Detection

22 Martin van den Berg (VU),Improving IT Decisions with Enterprise Architecture

23 Qin Liu (TUD), Intelligent Control Systems: Learning, Interpreting, Verification

24 Anca Dumitrache (VU), Truth in Disagreement - Crowdsourcing Labeled Data for Natural Language Processing

25 Emiel van Miltenburg (VU), Pragmatic factors in (automatic) image description
26 Prince Singh (UT), An Integration Platform for Synchromodal Transport
27 Alessandra Antonaci (OUN), The Gamification Design Process applied to (Massive) Open Online Courses
28 Esther Kuindersma (UL), Cleared for take-off: Game-based learning to prepare airline pilots for critical situations
29 Daniel Formolo (VU), Using virtual agents for simulation and training of social skills in safety-critical circumstances
30 Vahid Yazdanpanah (UT), Multiagent Industrial Symbiosis Systems
31 Milan Jelisavcic (VU), Alive and Kicking: Baby Steps in Robotics
32 Chiara Sironi (UM), Monte-Carlo Tree Search for Artificial General Intelligence in Games
33 Anil Yaman (TUE), Evolution of Biologically Inspired Learning in Artificial Neural Networks
34 Negar Ahmadi (TUE), EEG Microstate and Functional Brain Network Features for Classification of Epilepsy and PNES
35 Lisa Facey-Shaw (OUN), Gamification with digital badges in learning programming
36 Kevin Ackermans (OUN), Designing Video-Enhanced Rubrics to Master Complex Skills
37 Jian Fang (TUD), Database Acceleration on FPGAs
38 Akos Kadar (OUN), Learning visually grounded and multilingual representations

---

2020 01 Armon Toubman (UL), Calculated Moves: Generating Air Combat Behaviour
02 Marcos de Paula Bueno (UL), Unraveling Temporal Processes using Probabilistic Graphical Models
03 Mostafa Deghani (UvA), Learning with Imperfect Supervision for Language Understanding
04 Maarten van Gompel (RUN), Context as Linguistic Bridges
05 Yulong Pei (TUE), On local and global structure mining
06 Preethu Rose Anish (UT), Stimulation Architectural Thinking during Requirements Elicitation - An Approach and Tool Support
07 Wim van der Vegt (OUN), Towards a software architecture for reusable game components
08 Ali Mirsoleimani (UL),Structured Parallel Programming for Monte Carlo Tree Search
09 Myriam Traub (UU), Measuring Tool Bias and Improving Data Quality for Digital Humanities Research
10 Alifah Syamsiyah (TUE), In-database Preprocessing for Process Mining
11 Sepideh Mesbah (TUD), Semantic-Enhanced Training Data AugmentationMethods for Long-Tail Entity Recognition Models
12 Ward van Breda (VU), Predictive Modeling in E-Mental Health: Exploring Applicability in Personalised Depression Treatment
13 Marco Virgolin (CWI), Design and Application of Gene-pool Optimal Mixing Evolutionary Algorithms for Genetic Programming
14 Mark Raasveldt (CWI/UL), Integrating Analytics with Relational Databases
15 Konstantinos Georgiadis (OUN), Smart CAT: Machine Learning for Configurable Assessments in Serious Games
16 Ilona Wilmont (RUN), Cognitive Aspects of Conceptual Modelling
17 Daniele Di Mitri (OUN), The Multimodal Tutor: Adaptive Feedback from Multimodal Experiences

# Curriculum Vitae

Georgios Methenitis graduated from the Technical University of Crete, in Greece, with a Diploma in Electronic and Computer Engineering in September 2012. In December 2014 he graduated with a Master's degree in Artificial Intelligence from the University of Amsterdam, The Netherlands. As a student at the University of Amsterdam, he was a member of the Dutch Nao Team and participated in two RoboCup competitions. He conducted his MSc thesis in collaboration with the Advanced Concepts Team at the European Space Agency. In February 2015, he started his PhD research at the Centrum Wiskunde & Informatica in Amsterdam and the Delft University of Technology. The results of his PhD research are presented in this dissertation.