

DISS. ETH NO. 26972

FOSTERING SUSTAINABILITY USING INFORMATION TECHNOLOGY:
REAL-TIME FEEDBACK, INCENTIVES, AND SMART MARKETS

A thesis submitted to attain the degree of
DOCTOR OF SCIENCES of ETH ZURICH
(Dr. sc. ETH Zurich)

Presented by

ANSELMA MERET WÖRNER
M.Sc., Karlsruhe Institute of Technology
Born on January 21, 1992
Citizen of Germany

Accepted on the recommendation of

Prof. Dr. Elgar Fleisch
Prof. Dr. Wolfgang Ketter
Prof. Dr. Verena Tiefenbeck

2020

Für meinen Lieblingspapa.

“I’ve always felt that the human-centered approach to computer science leads to more interesting, more exotic, more wild, and more heroic adventures than the machine-supremacy approach where information is the highest goal.”

- Jaron Lanier, 2010

Acknowledgements

Most PhD students will agree that doing a PhD can seem like a roller coaster ride. The Highs and Lows you go through in figuring out your research path can be overwhelming, but the good thing is, you are not alone in it. Scientific research – as I have gotten to know it – is always a team effort, and any success is the product of the knowledge, of the experiences, and of the interactions of multiple people.

This dissertation is the result of my work from 2017 to 2020, conducted at ETH Zurich and at Stanford University. Throughout this period, I have had the great privilege of meeting and collaborating with many inspiring individuals, and I believe this is the greatest value I will take away from these past few years. I would like to express my gratitude to the many people who have influenced the work presented in this dissertation and mention some of them here. First of all, I would like to thank Prof. Elgar Fleisch for giving me the chance to work in this amazing environment and meet so many interesting people along the way. It has been an honor and a joy to work in your team. You are an inspiration to all of us. (Also, thank you for bringing me to Zurich – my favorite city in the world.) I also thank my co-advisor Prof. Wolf Ketter for providing great ideas, constructive feedback and encouragement for my research. I very much appreciate the time and effort you put into our joint work. It has been a pleasure working with you and your team in Cologne – here, special thanks also goes out to Philipp Kienscherf. Furthermore, I thank Prof. Inês Azevedo for hosting me in Stanford, for your valuable input to our joint work, and for being such a great role model; it was a truly inspiring time for me. In addition, I would also like to thank Prof. Clemens Puppe for first guiding me on the academic path and encouraging me to pursue it.

Considering the last few years of my work, however, I owe most gratitude to Prof. Verena Tiefenbeck, my supervisor and team leader in the Bits to Energy Lab: You have continuously encouraged and supported me not only on a professional level, but also on a personal level. I have learned so much from you and from your ever constructive feedback, and I am truly happy and grateful that we got to work together so closely.

The same holds for our team – Lili, Arne and Sandro. It was a pleasure working with you! I am really proud of what we achieved together, and I have learned way more from you three than you probably realize. In particular, I want to thank Lili for being there for me when times were tough, I cannot express how much that meant to me. This also brings me to our ‘Office Girls’ group: Cristina, Iris, Lili and Raquel – you are living proof that *empowered women empower women* and I am so happy that you have accompanied me

throughout this research endeavor. I further thank everyone I got to meet at the Fleisch-Chair and I am happy that many of you have grown into friends through the experiences we share: I love our team spirit and the interesting and inspiring discussions we had over lunch, coffee, BBQ, or during an after hour at Limmat.

All of you have made my time at ETH seem more like an interesting conversation and exchange of ideas than just a professional duty. Here, Liz and Monica also deserve huge credit for the magic that they do to keep the whole group up and running – we could not do it without you.

In addition to all of these amazing people I met during my PhD studies, I was lucky enough to already have a great support system to start out with. I would like to thank my mother and my father (1955–2018) for teaching me the importance of education; for making me stand up for myself; and, most of all, for always believing in me.

I also thank Claudia and Marc, as well as Brigitte and Wolfgang, for always being there for me and supporting me unconditionally. Furthermore, I am deeply grateful to have the greatest friends who are an ever-steadfast presence in my life despite growing geographical distances – since our high school days in Heidenheim, or since our joint studies in Karlsruhe. In particular, I would like to thank Caro, Lena, Lori and Susa – I would probably not be here without you, and if I was, I would for sure not love it this much.

Finally and most importantly, I feel incredibly lucky and am most grateful that I have gotten to enjoy almost every minute of this ride because I have learned to cherish the Highs that I am so fortunate to experience, rather than pondering over inevitable Lows. This gift, I owe to Max – my one and only.

Zurich, July 2020

Anselma Wörner

Contents

- Abstract** **i**
- Disclaimer** **vii**
- Abbreviations** **x**
- 1 Introduction** **1**
 - 1.1 Motivation 1
 - 1.2 Objectives & Contribution 3
 - 1.3 Approach 4
 - 1.4 Thesis Outline 6
- 2 Overview of the Related Work** **7**
 - 2.1 Decision Support Systems 7
 - 2.2 Smart Markets 12
 - 2.3 Green IS 16
 - 2.4 Research Gaps 18
- 3 Article A) Real-Time Feedback in the Absence of Volunteer-Selection Bias and Monetary Incentives** **23**
 - 3.1 Motivation 23
 - 3.2 Method 27
 - 3.3 Results 31
 - 3.4 Additional Analyses 33
 - 3.5 Discussion 36
- 4 Article B) Self-Set Goals in IS-Enabled Behavior Change** **39**
 - 4.1 Introduction 39
 - 4.2 Related Work 41
 - 4.3 Methodology 46
 - 4.4 Results 48
 - 4.5 Discussion & Conclusion 54
- 5 Article C) Blockchain Technology as Enabler for P2P Markets & Excerpt from Article D) as Use Case Analysis for the Energy Sector** **58**
 - 5.1 Introduction 58
 - 5.2 Background on Blockchain Technology 60

5.3	Methodology	62
5.4	Analytical Framework & Evaluation	64
5.5	Use Case Analysis: P2P Energy Exchange	74
5.6	Conclusion and Future Work	78
6	Article E) Evaluation of a P2P Energy Market in the Real World	80
6.1	Introduction	80
6.2	Background & Related Work	82
6.3	Method	87
6.4	Results	91
6.5	Discussion & Conclusion	97
7	Article F) Bidding Behavior on a P2P Energy Market	102
7.1	Motivation & Introduction	102
7.2	Related Literature and Theoretical Background	105
7.3	Study Design	108
7.4	Experimental Results	115
7.5	Further Analyses	122
7.6	Discussion & Conclusion	125
8	Intelligent Agents for Smart Load Scheduling	129
8.1	Motivation	129
8.2	Reinforcement Learning	131
8.3	Summary of Article G) Multi-Agent Reinforcement Learning for Electric Vehicle Charging	132
8.4	Outlook	137
9	General Discussion & Conclusion	139
9.1	Synopsis of Findings	139
9.2	Contributions & Implications	143
9.3	Limitations	148
9.4	Conclusion	150
	Bibliography	152
	Appendix	184
A	Literature Review on Blockchain-Based Energy Markets	185
B	Quartierstrom Webapp	187
C	Additional Analyses on Bidding Behavior	191
D	Reinforcement Learning	194
	List of Figures	197
	List of Tables	199

Abstract

Despite growing public attention and policy efforts for environmental sustainability, worldwide energy consumption and greenhouse gas emissions are still increasing (International Energy Agency, 2020a). Technological advances have enabled improvements in energy efficiency, and hundreds of billions of US dollars are being invested in renewable energy generation every year (International Energy Agency, 2020a), however, actually implementing these technologies requires changes to commercial practices, regulatory frameworks, and market structures. The replacement of conventional energy generation by renewable resources is advancing rather slowly (International Energy Agency, 2020b), as their integration requires a fundamental transition in the energy sector. Energy that was traditionally supplied by few power plants is now generated in smaller, distributed renewable generators that are not centrally controlled. The resulting growing number of stakeholders and increasing volatility in supply challenges existing market structures, as well as the grid infrastructure. Beyond technological and structural aspects, human behavior is what ultimately drives energy consumption and the adoption of renewable technologies. Consumer choices have a massive impact on resource use, as residential households consume more than 20% of the worldwide total final energy consumption (International Energy Agency, 2020a). Likewise, personal transportation makes up roughly the same amount in most countries (eurostat, 2017; International Energy Agency, 2020a). However, empirical data and research studies reveal a persistent gap between individuals' intentions for environmentally-friendly behavior and their actual energy consumption and associated emissions.

Information and communication technology can play a pivotal role in advancing the energy transition. Ubiquitous connected devices and the data they capture can be useful tools to identify inefficient consumption patterns or the potential for investments in new technologies. By capturing, analyzing, and evaluating energy data in high granularity, information technology can support sustainable practices – not only on a macro or organizational level, but also on the individual level. Yet, given the recency and fast pace of technological progress in this area, existing research has mostly focused on the

technical capabilities of information and communication technologies in environmental contexts. There is a lack of empirical and applicable knowledge on the impact of such ‘green’ information systems in the real world.

To tackle this issue, this thesis examines different ways in which information technology can foster sustainability in the real world, a) among individual consumers and b) in integrating renewable energy resources into the energy market. To that end, state-of-the-art information systems designed in conceptual studies are deployed in field experiments. In the studies presented, a smart metering device, a blockchain system, and an autonomous intelligent agent are designed based on recent work from the computer science discipline. To understand the behavioral effects in real-world settings, field experiments are conducted. The design and implementation of these field experiments builds on theories from psychology and economics research. Herein, the work presented in this thesis complements conceptual research on information systems (IS) with a social-science perspective and empirical validations.

More precisely, Chapters 3 and 4 present data collected in two large field experiments on feedback interventions for resource conservation during an energy-intensive activity – namely showering. The results reveal significant savings effects induced by activity-specific real-time feedback, even in the absence of monetary incentives. The presented findings verify the effectiveness of behavioral interventions for resource conservation. They thus provide robust and unique empirical evidence for real-time feedback as a scalable and cost-efficient policy instrument for fostering resource conservation among the broader public. More so, the results highlight the importance of understanding motivational drivers in the design of behavioral interventions – in general, and in environmental contexts in particular. For practitioners, as well as policy makers, these results represent applicable and highly relevant insights for the design of conservation programs.

In a second set of studies (Chapters 5–7), this thesis examines smart energy markets for the integration of distributed energy resources like rooftop solar systems. Based on market design theory, a peer-to-peer (P2P) market in which households bid prices for local solar energy is designed and implemented on a blockchain infrastructure. The system is deployed in a field experiment and tested for the duration of an entire year. The collected data is the first empirical evidence on a P2P energy market and on individual bidding behavior in this context. The results show that P2P energy markets are technologically feasible and that they can provide dynamic price incentives for balancing the grid, while actively engaging consumers in the decision-making process on the energy market. These studies generate first-hand empirical insights that scrutinize the (so far mostly theoretical) proposals on smart, consumer-centric energy markets. The observed behavior reveals an intention–behavior gap in participants’ willingness to pay for local solar energy and indicates learning effects on the market dynamics over time. These unique findings shed light on the prices that can be expected on local energy markets and elucidate avenues for future research on the trade-off between automation and consumer engagement. The findings further serve as a decision basis for policy makers for creating a regulatory framework for future energy markets. Capturing and processing high resolution energy data is essential for coordinating decentralized energy resources and, thus, for delivering the

energy transition. Furthermore, semi-autonomous agents, which elicit consumer preferences and then act on their behalf, emerge as a desirable system, combining consumer engagement and efficiency in the future energy market.

A further simulation study addresses the need for such computational tools in smart energy markets due to the increasing complexity of high electricity demand and distributed energy resources. Software agents that use reinforcement learning are employed to schedule electricity loads created by electric vehicles. The results show that intelligent agents are successful in coordinating loads to relieve the grid infrastructure in this simulation, thus paving the way to facilitate the electrification of transportation in the existing grid.

All in all, this dissertation demonstrates that information systems can indeed spur the energy transition and foster sustainability, both on the individual consumer level and on the market level. The findings from large-scale field experiments generate novel empirical insights and contribute to the impact-oriented work on green information systems. A multi-disciplinary approach leveraging state-of-the-art technologies (blockchain technology, agent-based simulation, and reinforcement learning) and testing them in the field expands existing conceptual knowledge to a more holistic understanding. The research presented shows that, by incorporating behavioral factors and economic incentives, information systems can induce energy conservation in the real world and encourage renewable energy generation – and thus tackle some of the wicked problems entailed in mitigating climate change.

Kurzfassung

Trotz öffentlicher Aufmerksamkeit und politischen Bemühungen für ökologische Nachhaltigkeit steigen Treibhausgasemissionen und der weltweite Energieverbrauch unablässig an (International Energy Agency, 2020a). Obwohl in den vergangenen Jahren deutliche technische Fortschritte im Bereich der Energieeffizienz gemacht wurden und jährlich Hunderte Milliarden US-Dollar in erneuerbare Energieerzeugung investiert werden (International Energy Agency, 2020a), erscheint die tatsächliche Umsetzung von Einsparungsmassnahmen in der Praxis schwer. Die Ablösung konventioneller Energieerzeugung durch erneuerbare Ressourcen schreitet nur langsam voran (International Energy Agency, 2020b), da deren Einbindung gleichzeitig eine Anpassung von Geschäftspraktiken, von rechtlichen Rahmenbedingungen und von Marktstrukturen erfordert: Energie, die traditionell in grossen Kraftwerken zentral erzeugt wurde, soll in der Zukunft von kleineren und örtlich verteilten erneuerbaren Generatoren geliefert werden. Die dadurch wachsende Zahl von Stakeholdern im Energiemarkt stellt bestehende Marktstrukturen in Frage, und die hohe Volatilität erneuerbarer Energie belastet die Netzinfrastruktur. Neben technologischer und struktureller Aspekte sind Energieverbrauch und Emissionen allerdings auch durch individuelles Verhalten getrieben. Private Haushalte verbrauchen mehr als 20% des weltweiten Energiebedarfs (International Energy Agency, 2020a) und der Personenverkehr beläuft sich in den meisten Ländern auf weitere 20% (eurostat, 2017; International Energy Agency, 2020a). Empirische Daten und Studien aus der Verhaltensforschung zeigen auf, dass eine anhaltende Kluft zwischen der Absicht zu umweltfreundlichem Verhalten und tatsächlichem Handeln besteht – was die Energiewende auch auf dieser Ebene verlangsamt.

Informations- und Kommunikationstechnologie (IKT) kann hier Abhilfe schaffen: Vernetzte Geräte und die von ihnen erfassten Daten können verwendet werden, um ineffiziente Verbrauchsmuster oder Potential für Investitionen in erneuerbare Technologien zu erkennen. Durch die Erfassung, Analyse und Auswertung von hochaufgelösten Energiedaten kann IKT ökologisch nachhaltiges Verhalten vorantreiben – nicht nur auf makro-ökonomischer oder organisationaler, sondern auch auf individueller Ebene. Angesichts der rasanten Fortschritte im Bereich IKT konzentriert sich die Forschung zu Informationssystemen für Nachhaltigkeit bislang jedoch hauptsächlich auf deren rein technische Möglichkeiten; es mangelt an empirischer und anwendungsorientierter Forschung zu den Effekten solcher ‘grünen’ Informationssysteme in der Praxis.

Die vorliegende Dissertation untersucht verschiedene Ansätze, um Nachhaltigkeit in der Praxis mithilfe von IKT zu fördern. Zu diesem Zweck werden modernste Informationssysteme in konzeptionellen Studien entworfen und, insbesondere, in Feldexperimenten auch

auf ihre Effekte hin empirisch untersucht. Auf Grundlage aktueller Forschung aus der Informatik, werden für die vorgestellten Studien ein intelligentes Messgerät, ein Blockchain-basierter Markt und ein intelligenter Software-Agent entwickelt. Diese Systeme werden in Feldexperimenten eingesetzt, um die aus ihrer Nutzung resultierenden Verhaltensseffekte in der realen Welt zu analysieren. Die Konzeption und Implementierung der Feldexperimente baut dabei auf Theorien aus der Psychologie und der Ökonomie auf. So ergänzt diese Dissertation die bisher hauptsächlich konzeptionelle Forschung zu IKT im Nachhaltigkeitsbereich durch eine sozialwissenschaftliche Perspektive und empirische Validierungen.

Konkret stellen die Kapitel 3 und 4 zwei umfangreiche Feldexperimente zu Verhaltensinterventionen vor. Hierbei werden Individuen während einer energieintensiven Aktivität – des Duschens – mit Echtzeit-Feedback zu ihrem Ressourcenverbrauch konfrontiert. Die Experimentalergebnisse zeigen deutliche und statistisch signifikante Einsparungseffekte dieser Feedback-Intervention. Beide Studien liefern robuste empirische Belege dafür, dass Echtzeit-Feedback eine skalierbare und kosteneffiziente Maßnahme zur Förderung von Ressourceneinsparungen in der breiten Öffentlichkeit darstellt.

In einer zweiten Reihe von Studien (Kapitel 5–7) untersucht die vorliegende Arbeit digitale Energiemärkte für die Integration von verteilten Generatoren wie Solaranlagen. Basierend auf der Markt-Design-Theorie wird eine Handelsplattform entworfen und entwickelt, über die private Haushalte online Solarenergie, die auf ihren Dächern produziert wird, an Nachbarn verkaufen können. Das System wird auf einer Blockchain-Infrastruktur implementiert und in einem Feldexperiment eingesetzt, in dem die Plattform über die Dauer eines Jahres hinweg mit realen Nutzern getestet wird. Die gesammelten Daten liefern die ersten empirischen Erkenntnisse zum Handel von Solarenergie unter privaten Haushalten (auch ‘peer-to-peer’ Markt genannt). Die hier gezeigten Ergebnisse belegen, dass solche lokalen, Nutzer-zentrierten Energiemärkte bereits technologisch umsetzbar sind und dass ein solcher Marktplatz dynamische Preisanreize für den Ausgleich des Stromnetzes bieten kann. Die Ergebnisse legen zudem eine Diskrepanz zwischen von den Teilnehmern zuvor angegebener und während des Experiments tatsächlich beobachteter Zahlungsbereitschaft für Solarenergie offen. Die hier enthaltenen Studien zu lokalen Energiemärkten tragen hierin einzigartige empirische Erkenntnisse zur bisher meist theoretischen Erforschung solcher verbraucherzentrierten, digitalen Energiemärkte bei. Die vorgestellten Erkenntnisse über individuelles Verhalten in diesem Kontext geben Einblicke in Verbraucherpräferenzen für Solarenergie. Darüber hinaus erweisen sich Softwareagenten, die Präferenzen von Verbrauchern erlernen und in deren Namen handeln, als ein vielversprechendes Instrument, das Verbraucherengagement und Effizienz verbindet. Die vorgestellten Studienergebnisse dienen als Grundlage für politische Entscheidungsträger bei der Umsetzung der Energiewende und der Schaffung von verbraucherorientierten Energiemärkten. Die Erfassung und Verarbeitung hochaufgelöster Energiedaten erweist sich dabei für die Koordination dezentraler Energieressourcen als wesentlich.

Zum Abschluss befasst sich eine zusätzliche Simulationsstudie (Kapitel 8) mit dem Einsatz von künstlicher Intelligenz im Energiemarkt, um die zunehmenden Komplexität verteilter Energieressourcen beherrschbar zu machen. In einer Simulation werden

Software-Agenten untersucht, die neuste Techniken des Maschinellen Lernens nutzen, um Ladelasten von Elektrofahrzeugen zu koordinieren. Die Ergebnisse zeigen, dass es intelligenten Software-Agenten gelingt, Lasten bei dynamischen Marktpreisen so zu koordinieren, dass die Netzinfrastruktur entlastet wird. Dies gelingt in der vorgestellten Modellierung, ohne dass Daten über individuelle Fahrprofile oder deren Energiebedarf an einen zentralen Aggregator kommuniziert werden müssen. Dieser Ansatz kann zur Elektrifizierung des Verkehrs in bestehender Netzinfrastruktur beitragen, die notwendig ist, um eine deutlichen Emissionsreduktion zu erreichen.

Alles in allem belegt diese Dissertation, dass IKT die Energiewende vorantreiben und Nachhaltigkeit sowohl auf der Ebene des einzelnen Haushalts, als auch auf kollektiver Marktebene fördern kann. Die vorgestellte Forschung veranschaulicht wie Informationssysteme Energieeinsparungen induzieren und die Erzeugung erneuerbarer Energie vorantreiben können, wenn bei deren Gestaltung verhaltenswissenschaftliche Erkenntnisse beachtet, und effektive Anreize gesetzt werden. So können Informationssysteme einige der vielschichtigen Probleme lösen, die mit der Eindämmung des Klimawandels verbunden sind.

Disclaimer

Major parts of this thesis have already been published in peer-reviewed conference proceedings and journals, or are currently under review. As a result, many sections of this thesis correspond to these articles and papers (status as of August 26, 2020):

- A) Tiefenbeck, V., **Wörner, A.**, Schöb, S., Fleisch, E., and Staake, T. (2019) “Real-time feedback promotes energy conservation in the absence of volunteer selection bias and monetary incentives”. Published in *Nature Energy*.
<https://doi.org/10.1038/s41560-018-0282-1>
- B) **Wörner, A.**, Tiefenbeck, V. (2018) “The role of self-set goals in IS-enabled behavior change”. Published in *Proceedings of the European Conference on Information Systems (ECIS)*. (Award for ‘Runner Up Best Research Paper’)
<https://doi.org/10.3929/ethz-b-000291478>
- C) **Wörner, A.**, and Tiefenbeck, V., Ketter, W. “Blockchain-Based Markets – A Literature Review”. Under review at *Decision Support Systems*.
- D) **Wörner, A.**, Meeuw, A., Ableitner, L., Wortmann, F., Schopfer, S., and Tiefenbeck, V. (2019) “Trading Solar Energy within the Neighborhood: Field Implementation of a Blockchain-Based Electricity Market”. Published in *Energy Informatics*.
<https://doi.org/10.1186/s42162-019-0092-0>
- E) **Wörner, A.**, Ableitner, L., Meeuw, A., Wortmann, F., and Tiefenbeck, V. (2019) “Peer-to-Peer Energy Trading in the Real World: Market Design and Evaluation of the User Value Proposition”. Published in *Proceedings of the International Conference of Information Systems (ICIS)*.
<https://doi.org/10.3929/ethz-b-000395174>
- F) **Wörner, A.**, Tiefenbeck, V., Ableitner, L., Meeuw, A., Fleisch, E., Wortmann, F., and Azevedo, I. (2020) “Peer-to-Peer Energy Markets”. Under review at *Information Systems Research*.

-
- G) Kienscherf, P. A., **Wörner, A.**, Ketter, W., and Tiefenbeck, V. “Multi-Agent Reinforcement Learning for Electric Vehicle Charging Guided by Nodal Pricing”. To be submitted to the *Journal of Artificial Intelligence*.

Chapter 3 of the present thesis corresponds to a slightly modified version of paper A), which was published in *Nature Energy* in 2019. I declare that I have conducted the data analysis, and that I co-wrote the manuscript with Verena Tiefenbeck. Verena Tiefenbeck and Thorsten Staake designed the study, and Samuel Schöb wrote the software of the study devices that enabled the technical side of the data collection. Thorsten Staake and Elgar Fleisch provided valuable feedback for the article and secured funding for the study.

Chapter 4 corresponds to a slightly modified version of article B), which was awarded the Best Paper Runner-Up Award at the *European Conference on Information Systems (ECIS) 2018* and was published in the conference proceedings. I confirm that I conducted the data analysis, and I am the lead author of this article. Verena Tiefenbeck conducted the data collection, and reviewed and edited the article.

Chapter 5 corresponds to a modified version of article C), which was submitted to the journal *Decision Support Systems (DSS)* and is currently under review. I hereby confirm that I am the lead author of this article, including study design and analysis. The article benefited immensely from the edits and reviews of my co-authors Verena Tiefenbeck and Wolf Ketter. In addition, the chapter contains a literature review on blockchain-based energy markets included in article D), which was published in *Energy Informatics*. I hereby confirm that I conducted this literature review, and I am the lead author of the paper.

Chapter 6 corresponds to a slightly modified version of article E), which has been published in the *Proceedings of the International Conference of Information Systems (ICIS) 2019*. I hereby confirm that I am the lead author of the article. I drafted the article, designed the research question and theoretical embedding, and conducted the quantitative analyses. The data examined in the article was collected in a joint research project conducted together with Liliane Ableitner, Arne Meeuw, Sandro Schopfer, Verena Tiefenbeck, Felix Wortmann, and industry partners including the WEW Walenstadt. Arne Meeuw was the lead developer of the information system deployed in the project, Liliane Ableitner, Sandro Schopfer and myself have designed and developed individual modules of this information system. Article E) was drafted by myself and benefited immensely from edits and reviews by Liliane Ableitner and, in particular, by Verena Tiefenbeck.

Chapter 7 corresponds to a slightly modified version of article F), which was submitted to the journal *Information Systems Research* in early July 2020 and is now under review. I hereby confirm that I am the lead author of the article. I drafted the article, designed the research question and theoretical embedding, and have conducted the quantitative analyses. The data examined in the article was collected in the same research project described above for article E). Furthermore, Verena Tiefenbeck, Felix Wortmann, and Inês Azevedo have provided valuable input and edits to this manuscript.

Section 8.3, of Chapter 8 presents a summary of a simulation study included in manuscript G), which will be submitted to the *Journal of Artificial Intelligence* in August 2020.

Philipp Kienscherf and I have developed and written that manuscript together, and both have designed parts of this simulation. Philipp Kienscherf implemented most of the code for the simulation. The article further benefited immensely from the input and edits of Wolfgang Ketter and Verena Tiefenbeck.

In addition to the above-named documents, the following manuscripts were part of my PhD research, but are outside the scope of this dissertation:

- H) Tiefenbeck, V., **Wörner, A.**, Schöb, S., Fleisch, E., and Staake, T. (2018) “Real-time feedback reduces energy consumption among the broader public without financial incentives”. Policy brief. Published in *Nature Energy*.
<https://doi.org/10.1038/s41560-019-0480-5>
- I) Brenzikofer, A., Meeuw, A., Schopfer, S., **Wörner, A.**, and Dürr, C. (2019) “Quartierstrom: P2P Energiemarkt auf der Blockchain”. Published in *CIREN*.
<http://dx.doi.org/10.34890/360>
- J) Ableitner, L., Meeuw, A., Schopfer, S., Tiefenbeck, V., Wortmann, F., and **Wörner, A.** (2019) Quartierstrom Whitepaper. Available on *arXiv*. [arXiv:1905.07242](https://arxiv.org/abs/1905.07242)
- K) Meeuw, A., Schopfer, S., **Wörner, A.**, Tiefenbeck, V., Ableitner, L., Fleisch, E., and Wortmann, F. (2020) “Implementing a blockchain-based local energy market: Insights on communication and scalability”. Published in *Computer Communications*. <https://doi.org/10.1016/j.comcom.2020.04.038>
- L) Ableitner, L., Tiefenbeck, V., Meeuw, A., **Wörner, A.**, and Wortmann, F. (2020) “User Behavior in a Real-World Peer-to-Peer Electricity Market”. Published in *Applied Energy*. <https://doi.org/10.1016/j.apenergy.2020.115061>
- M) Ableitner, L., Tiefenbeck, V., **Wörner, A.**, and Fleisch, E., “Designing a peer-to-peer energy market from the user perspective”. Submitted to *Environmental Innovation and Societal Transformations* in June 2019.
- N) Yan, X., **Wörner, A.**, Awater, P., Monscheidt, J., Gemsjäger, B., and Tiefenbeck, V. (2020) “Long-Term Electric Load Forecast for Urban Areas with an Energy Decomposition Model” (Extended Abstract). Forthcoming at *DACH+ Conference on Energy Informatics*.

Abbreviations

AI	artificial intelligence
CHF	Swiss Francs
CO ₂	carbon dioxide
DER	distributed energy resources
DLT	distributed ledger technology
EIA	US Energy Information Administration
EUR	Euro
EV	electric vehicle
GHG	greenhouse gas
IEA	International Energy Agency
IS	information systems
MWh	Megawatt hour
MDP	Markov Decision Process
kWh	Kilowatt hour
P2P	peer-to-peer
PV	photovoltaics
RL	reinforcement learning
SCR	self-consumption rate
sd	standard deviation
SFOE	Swiss Federal Office of Energy
SSR	self-sufficiency rate
USD	US Dollars

1. Introduction

This chapter outlines the the general motivation and objectives of this dissertation and provides an overview of the articles contained.

1.1 Motivation

Resource depletion and climate change are among the main threats to the well-being of human society and the environment (Davis et al., 2018; Stanley, 2000; United Nations, 2019). Environmental concerns that seemed abstract in the past have started to materialize in tangible effects, directly affecting the life of human individuals. Human-caused climate change contributes to effects such as increased wildfire activity (Abatzoglou and Williams, 2016), global spread of mosquito-borne infectious diseases (Altizer et al., 2013; Hales et al., 2002; Rocklöv and Dubrow, 2020) or urban freshwater shortages (Gober et al., 2010; McDonald et al., 2011). As a result, public awareness for environmental concerns has increased. In a historic landmark agreement in 2015, 195 countries have pledged to take measures to limit climate change by reducing greenhouse gas (GHG) emissions (United Nations Conference of the Parties (COP), 2015). Policy makers around the globe have thus set out ambitious goals to reduce energy consumption and increase the share of renewable energy generation, also captured in the United Nations Sustainable Development Goal Nr. 7 (United Nations, 2019).¹

¹ In addition, between 2017 – the year I started my PhD research – and today, the number of Google searches for the terms ‘climate change’ and ‘sustainability’ have roughly doubled (Google Trends, 2020). Children and teenagers around the globe have started school strikes for climate change (The Economist, 2019) demanding responsible resource use, and this year’s World Economic Forum featured sustainability on the top of its agenda (World Economic Forum, 2020).

Yet, statistics on ever-increasing energy consumption (International Energy Agency, 2020a) and emissions (International Energy Agency, 2020b) reveal that implementing effective measures to actually put these intentions into practice is difficult. Consuming energy resources is deeply ingrained in human life, and resource use is determined by a myriad of conscious and unconscious individual choices. Environmental sustainability – the “conservation, deployment, and reuse of resources in responsible ways” (Malhotra et al., 2013, p. 1265) – necessitates changing established usage patterns or establishing new ones. This is difficult from many perspectives: For companies, technological innovations that improve energy efficiency require considerable investments or adjustments in their supply chains. From a structural perspective, the integration of renewable energy resources adds complexity to the energy market. Established energy markets are strongly centralized and hierarchical, with few power plants distributing energy to thousands of households. Wind and solar energy generation, in contrast, is geographically distributed and is dependent on local weather conditions; it cannot be simply switched on or off according to power demand (Ketter et al., 2018; Ramchurn et al., 2012). To further increase the share of renewable resources in total energy generation, electrification of transportation is required (Williams et al., 2012). This creates technical challenges, as the electricity grid infrastructure is not built to cover the high loads that would result from uncoordinated charging of large numbers of electric vehicles (Robu et al., 2013; Rogers et al., 2012). Beyond technological and structural challenges, mitigating climate change also requires changes in individual consumption behavior as residential households consume more than 20% of the worldwide total energy consumption (International Energy Agency, 2020a). However, energy consumption is abstract to individuals, and the effects of one’s own consumption behavior or sourcing choices seem minor, and they get lost in everyday business (Tiefenbeck et al., 2018a).

Given the complexity of these issues, consumers, as well as decision makers, need support to actually be able to ‘walk the talk’ and reduce energy consumption and emissions. Fortunately, help is around the corner. Information and communication technology provides powerful tools to empower individual consumers a) by actively supporting them in formerly abstract decisions, or, in contrast, b) by automating complex processes to disburden individuals in their everyday lives. Both of these approaches, ‘human-in-the-loop’ (Mattern et al., 2010) and ‘human-out-of-the-loop’ (Fleisch, 2010) computing, have their place in the diverse and multi-layered problem settings relevant to foster sustainability (Ketter et al., 2018; Tiefenbeck, 2017).

For instance, designing effective measures addressing individuals’ consumption behavior

was difficult and costly in the past. Using information systems (IS), “data is democratized from scientific practices and made universal and meaningful for use by all individuals” (Swan, 2013, p. 95). Real-time information captured by smart meters and other connected devices can be used to better manage energy resources (Gupta, 2017) and to provide cost-efficient, individualized behavioral interventions (Loock et al., 2013; Melville, 2010; Tiefenbeck et al., 2018a). Consumption feedback during specific activities can make resource consumption more salient and thus encourage conservation on an individual level (Tiefenbeck et al., 2018a).

In other instances, IS can automate processes to save energy or utilize renewable energy using computational tools. In particular, in the dynamic energy market, data on energy demand and supply can be used to balance the grid and prevent technical failures without overburdening individuals (Bichler et al., 2010; Gholami et al., 2016; Ketter et al., 2018). In that, “[s]mart technologies will help consumers and energy service companies working for them to reap the opportunities available on the energy market by taking control of their energy consumption (and possible self-production).” (European Commission, 2015, p. 11)

1.2 Objectives & Contribution

The objective of this thesis is to examine different ways in which information technology can foster sustainability a) among individual consumers and b) in integrating renewable energy resources into the energy market. The research studies presented employ, and examine the effects of a smart metering device, a blockchain-based energy market, and an intelligent software agent. Herein, the incentives that drive consumer behavior are in the focus. The goal is to contribute to scientific theory and to create actionable implications for practitioners on bringing effective and scalable ‘green’ IS solutions to the real world that can help mitigate climate change.

The complexity of sustainability issues requires a multi-perspective approach examining not solely technological aspects, but also integrating social sciences (Melville, 2010; Tiefenbeck, 2017). Goes (2013) argues that “[Recommendation and Personal Information Systems] are designed for consumers to overcome their cognitive limitations when making choices in IT-mediated environments. Ironically, by and large [these systems] have ignored the cognitive biases that come into play in decision-making.”, p. v. And in the same vein, “E-marketplaces continue to evolve, and behavioral economics can inform the design of such platforms.” (Goes, 2013, p. v). Addressing potential cognitive biases and evaluating the motivators for individual behavior seems particularly relevant in the sustainability

context, as (long-term) effects of resource consumption or energy-efficiency investments are not immediately perceptible by consumers (Tiefenbeck, 2017), and misconceptions about effective conservation measures prevail (Lesic et al., 2018).

Furthermore, in this type of policy-relevant research, it is particularly important to validate solution concepts empirically (Editorial, 2017). Research in various domains has shown that theoretical predictions or laboratory studies may not provide a full picture of behavioral effects that materialize in practice (Allcott et al., 2012b; List et al., 2006). (Natural) field experiments can provide additional insights to garner a more nuanced understanding of their real-world impact (Malhotra et al., 2013).

To address these issues, the research presented in this thesis draws on theoretical background from different disciplines, such as (behavioral) economics, and psychology, as well as IS research. Although the latest technological innovations in blockchain technology and machine learning are employed in the work presented here, this thesis does not focus on *either* the technical aspects of an information system *or* user behavior in interaction with it, but takes an integrative approach in the evaluation of the systems. Three unique field experiments allow to investigate the empirical effects of the presented IS and contribute novel insights on their performance and resulting user behavior in practice.

1.3 Approach

This dissertation comprises six research studies, which cover three distinct field experiments. These field experiments are at the core of this thesis, creating unique insights on consumer behavior and decision making in the energy market. They are complimented by simulations and a conceptual framework that increase the background knowledge and conceptual understanding of the technologies employed.

The first study (Chapter 3) investigates the effect of real-time feedback in the absence of volunteer-selection bias and monetary incentives. In a natural field experiment, an uninformed sample of hotel guests was presented with a feedback intervention on resource consumption during showering. Showering was chosen as an example of an energy-intensive, habitual, low-involvement target activity. The individuals in the hotel setting do not incur monetary savings from reducing their resource consumption; in addition, these individuals have not self-selected into taking part in an energy-related research study, but rather are merely guests of one of six hotels that partnered with us to conduct the experiment. The data collected comprises a sample of $n = 19,596$ showers, measured in 265 rooms in six hotels.

In a second field experiment (Chapter 4), $n = 413$ residential households were presented with the same feedback intervention on resource consumption in the shower for two months. The study examined the role of self-set goals in IS-enabled feedback interventions. Based on the consumption data collected and on supplementary surveys, the formation of self-set conservation goals in response to the feedback intervention and their relationship to incurred savings was evaluated. Both studies on feedback interventions used the same IS artefact – a smart shower meter produced by the ETH spin-off company Amphiro AG.

The second part of this dissertation focuses on the creation of smart electronic markets that can handle the complexity and volatility of distributed renewable energy resources. This research includes a field experiment conducted within the research project Quartierstrom. The project was a lighthouse project funded by the Swiss Federal Office of Energy, and was conducted in cooperation with a consortium of industry and academic partners, as well as a local utility provider (Wasser- und Elektrizitätswerk Walenstadt). In this project, our research group had the opportunity to design, implement, and deploy the first blockchain-based energy market in Switzerland in which households engaged in peer-to-peer trading of solar energy via a blockchain-based information system. For the duration of an entire year, $n = 37$ households interacted with the market using a web application. Participants bid prices for local solar energy via an auction mechanism; the prices settled in the auction had direct impact on participants' real electricity bills. The three articles presented in Chapters 5, 6 and 7 examine P2P energy markets, from conceptualization to empirical evaluation: Chapter 5 presents a conceptual framework on blockchain-based markets and a systematic literature review. Chapters 6 and 7 evaluate the unique data collected in the field experiment in the SFOE lighthouse project to evaluate the market design and the user behavior of the deployed P2P energy market in the real world.

As an outlook, Chapter 8 evaluates the possibilities of machine learning for scheduling loads in a smart energy market. Using a simulation based on real consumption data and synthetic driving profiles, autonomous intelligent agents coordinate charging schedules for electric vehicles to reduce peak demands in the electricity grid. This project represents the first step in a new research direction; it is summarized in Chapter 8, providing an outlook on future research on artificial intelligence in smart energy markets.

1.4 Thesis Outline

The remainder of this dissertation is structured as follows: Chapter 2 presents an overview of the related research relevant to this thesis and highlights research gaps, providing a framework for the subsequent research articles. Chapters 3 - 7 represent separate research articles that have already been published in, are currently under review at, or are about to be submitted to peer-reviewed publications (see also Disclaimer). Chapter 8 presents a summary of recent work on a multi-agent simulation of intelligent agents for load scheduling and provides an outlook on future research questions in this area. The thesis closes with a synopsis of the findings of the different articles and a discussion of its contributions and its limitations.

2. Overview of the Related Work

This thesis examines the application of information systems (IS) to foster resource conservation and to create marketplaces for distributed energy resources. While each of the subsequent chapters contains a more in-depth review of the related work for each respective article, this chapter provides an overview of the relevant literature to put the different articles contained in the thesis into context. This overview starts with an introduction of research streams on IS-enabled decision support in form of behavioral feedback and intelligent agents (Section 2.1). Section 2.2 presents research on electronic smart markets, explaining the basics of market design theory and P2P markets. Section 2.3 gives an overview of the existing research in Green IS, as the overarching domain of the studies in this thesis.

2.1 Decision Support Systems

In our lives, we are increasingly surrounded by sensors and connected devices that measure our every-day activities like how many steps we take a day, or how much energy we consume, but also capture our environment, e.g. weather or locational data. As a result of ubiquitous data collection and communication technologies, IS can provide decision support to individuals in many domains and in different ways.¹

¹Note that in this thesis, the term ‘decision support system’ is used in a broader definition to describe information systems that support individuals in different types of choice situations (Banker and Kauffman, 2004; Bichler et al., 2010; Ketter et al., 2018). This can mean explicitly by providing feedback to inform educated decision making, or, on the other end of the spectrum, automating specific choices for an individual based on algorithmic rules.

The term is used in a more narrow meaning in some of the computer science literature, namely referring to recommender systems, which is not intended here.

On the one hand, IS can “help people collect personally relevant information for the purpose of self-reflection and gaining self-knowledge” (Li et al., 2010, p.558), thus supporting individuals’ own insightful reflection. Personalized feedback can influence individuals’ behavior in habitual activities without actively interfering in their decisions. As an example, according to Fox and Duggan (2012), already in 2012, 19% of the smartphone owners in the US had downloaded an app to specifically track and improve their health. This phenomenon is called ‘self-tracking’, ‘quantified self’, or ‘living by numbers’ (Consolvo et al., 2009; Li et al., 2010; Lupton, 2014). As Swan (2012) states, the overarching goal of such continuous monitoring and immediate feedback enabled by connected devices is lasting behavior change.

At the same time, IS can also provide data-driven decision support in form of concrete advice or even automate decision processes for individuals (Bichler et al., 2010; Ketter et al., 2018; Wooldridge and Jennings, 2018). Using distributed sensors, as well as distributed computational power (Rogers et al., 2012), intelligent agents can process information in real time and act accordingly in specific contexts.

2.1.1 Feedback Interventions

Feedback, i.e., informing individuals about their behavior in order to reinforce and/or modify certain actions (Karlin et al., 2015), has proven to be successful in changing individual behavior, not only in work-related, but also in private contexts (Kluger and DeNisi, 1996; Moon and Sproull, 2008). Studies in various areas provide evidence that individuals respond to feedback, in particular for habitual or ‘low-involvement’ activities, which individuals perform in their everyday lives without extensive deliberation, such as feedback on physical activity (Bravata et al., 2007), nutrition (Vandelanotte et al., 2005), or energy consumption (Abrahamse et al., 2005; Allcott, 2011).

Originating in psychology research, ‘Feedback Intervention Theory’ (Kluger and DeNisi, 1996, p. 259) states that “feedback interventions change the locus of attention”, but its effect depends on various moderators. While many studies find positive results for feedback interventions, the observed effect sizes vary with the target activity, the general setting, and the way feedback is presented (Kluger and DeNisi, 1996). A key finding is that in particular for habitual activities, the more closely (in time and place) the feedback is linked to a particular target activity, the better (Ehrhardt-Martinez et al., 2010; Froehlich et al., 2010). This emphasizes the importance of the medium by which the intervention is provided; ideally, it should be accessible during the targeted activity and process information in real time (Tiefenbeck et al., 2018a). Personal IS increasingly

make this possible at population level: Ubiquitous connected devices offer more and more possibilities to deliver personalized feedback (Froehlich et al., 2010; Hermsen et al., 2016). With smartphones, fitness trackers and other personal IS, feedback interventions are scalable to the broad population and can be integrated easily into individuals' everyday activities (Consolvo et al., 2009; Hermsen et al., 2016; Saha and Mukherjee, 2003).

In the energy context, a number of field studies have found that feedback on energy consumption is effective and reduces misconceptions in many cases (Abrahamse et al., 2007; Allcott, 2011; Tiefenbeck et al., 2018a). For example, consumers are often unaware of the energy efficiency of their homes and devices (Allcott et al., 2012b), or underestimate the long-term benefits of potential investments into energy efficiency upgrades (Allcott and Mullainathan, 2010). Feedback interventions can help mitigate these issues, and technological advances make it increasingly feasible to monitor the energy consumption of households, individual appliances or even specific activities in real time (Tiefenbeck, 2017). A series of studies using interventions such as monthly energy reports or real-time feedback based on smart meter data have provided evidence that digital interventions can be a scalable and cost-effective instrument for fostering energy savings (Loock et al., 2013; Tiefenbeck et al., 2018a). Other programs send out "home energy reports": periodic mailings that compare the electricity use of individual households with similar homes in the neighbourhood, thus tapping into social norms (Allcott and Mullainathan, 2010). Yet, most existing feedback studies do not discuss the motivation or personal incentives for users to change their behavior once they are presented with consumption feedback. Several scholars point out that research on the drivers of effective use of IS-enabled feedback is still sparse (Burton-Jones and Grange, 2013; Gimpel et al., 2013; Sjoeklint et al., 2015). Based on a meta-study on digital technologies for disrupting and changing behavior, Hermsen et al. (2016) conclude that more research should study the factors that drive and sustain behavioral effects in order to design more effective interventions and necessary IS.

2.1.2 Intelligent Agents

While feedback interventions have been successful for changing consumption behavior in specific activities, many consumption patterns are hard to change for us human beings (Tiefenbeck et al., 2018a). In particular, the energy consumption of appliances which run in the background, like heating or battery charging, seems abstract and its optimization can easily overburden individuals (Ehrhardt-Martinez et al., 2010; Faruqui et al., 2010). In environments in which resource consumption is non-tangible or supply is volatile (as is the case with home energy control (Mattern et al., 2010; Strueker and Dinther, 2012) or

renewable energy generation), the support of algorithms for their decision making can be useful for humans for bare capacity reasons (Strbac, 2008; Valogianni and Ketter, 2016). Here, IS cannot only serve as feedback instruments to prompt behavioral reactions, but can also act in itself, automating processes for the user (Ketter et al., 2018; Melville, 2010; Valogianni et al., 2019). In an agenda for “autonomous agents and multi-agent systems research within the smart grid”, Rogers et al. (2012), p. 2167, highlight “Supporting Consumers in the Transition to a Smart Grid” as a key research area. Intelligent agents can provide (semi-)autonomous decision support for individuals by using algorithms which analyze observed data and adhere to preferences defined by the users up front (Bapna et al., 2004; Bollinger and Hartmann, 2020; Ketter et al., 2018; Rogers et al., 2012).

While it is hard to find a formal definition of ‘intelligent agents’ in the literature – as it is for the term ‘artificial intelligence’ – Wooldridge and Jennings (2018) dare an attempt: They define an intelligent agent as a computer system that operates without the direct intervention of humans, interacts with others and acts and reacts depending on information on the environment it perceives. Moreover, an agent may be “conceptualised or implemented using concepts that are more usually applied to humans”, such as “knowledge, belief, intention, and obligation” (Wooldridge and Jennings, 2018, p. 117). The related research in the computer science discipline focuses on the optimization and learning algorithms employed by these agents to adequately react to their environment (Valogianni et al., 2014; Wooldridge and Jennings, 2018). Most notably the evolution of reinforcement learning in the past decade has led to the development of very sophisticated adaptive strategies. Learning models allow agents to learn from past situations to improve their reactions to new information they observe (Mnih et al., 2015; Peters et al., 2013; Reddy and Veloso, 2011b; Valogianni et al., 2013; Vinyals et al., 2019; Yang et al., 2018). Rogers et al. (2012) point out that the robustness of these machine learning approaches to real-world deployment is key to their effectiveness.

Still, the success of intelligent agents in practice is not solely determined by their inner machine learning mechanisms, but also by their acceptance and interactions with the human user (Rahwan et al., 2019). Scholarly work has – rather recently – started to investigate questions around trust in AI and the degree of automation which humans prefer in different contexts (Bollinger and Hartmann, 2020; Komiak and Benbasat, 2006; Logg et al., 2019; Wang and Benbasat, 2005). For instance, research on ‘algorithm appreciation’ challenges the widespread belief that individuals generally do not like to rely on algorithms (i.e. that they exhibit ‘algorithm aversion’), and finds that they in fact appreciate advice even from black box algorithms in some cases (Logg et al., 2019). Findings from another

recent lab study, on the other hand, indicate that people prefer to use algorithms when they can modify them slightly, and perform better overall then. Individuals seem to want at least some control over an outcome rather than being out of the loop in the decision-making process (Dietvorst et al., 2016). Rahwan et al. (2019) propose to establish ‘machine behavior’ as a new research field with the objective to study how intelligent agents shape human behavior and vice versa, as well as collective vs. individual machine behavior. Along the same lines, Bichler et al. (2010) argue that “the degree of autonomy that an agent or automated decision support system should have is a completely open research question especially in dynamic and complex market environments [...]”, p. 697.

Practical applications for intelligent software agents range across domains, from algorithmic trading (Ketter et al., 2013; Rahwan et al., 2019), to digital health applications (Barata et al., 2019; Tinschert et al., 2017) or driver assistance (Gahr et al., 2018; Sharon and Stone, 2017). As touched upon above, in the energy sector alone, there are numerous areas in which intelligent agents can be employed to reduce complexity for individuals. One of the use cases which is gaining increasing importance with the electrification of transportation (Williams et al., 2012) and the diffusion of distributed energy resources (Ramchurn et al., 2012), is the smart scheduling of loads (also called autonomous demand response (Ketter et al., 2018)). Based on (real-time) data and dynamic price signals, intelligent agents can be used to schedule flexible electricity loads, like charging of electric vehicles to times when electricity generation is cheapest or cleanest (Ketter et al., 2018; Peters et al., 2013; Tiefenbeck, 2017; Vázquez-Canteli and Nagy, 2019). However, these efforts need to be coordinated on a collective level in order to prevent high demand peaks which strain the electricity grid (Flath et al., 2014; Valogianni et al., 2020). One way to coordinate supply and demand patterns is to establish a central entity which aggregates information from all parties and determines a distribution (or ‘allocation’) of resources (‘top-down approach’) (Valogianni et al., 2015). Another, more decentralized way, is to let a market govern the distribution among individual consumers based (‘bottom-up approach’) (Dietz et al., 2003; Slavova and Constantinides, 2017). Such so-called smart markets can create incentives for efficient resource use while keeping the grid in balance (Ketter et al., 2018).

2.2 Smart Markets

Markets are institutions that govern resource allocation among a group of individuals (Dietz et al., 2003). Basic research on the interactions of individuals on markets and the incentives they face originates from micro-economics, and has developed into its own field of ‘market design theory’, an overview of which is provided in the following Section 2.2.1. While the research on market mechanisms is rooted in economic theory, most new, emerging markets are enabled by computational tools, advanced user interfaces, and smart devices. The term ‘smart markets’ relates to computer-assisted markets relying on algorithms and data to determine prices and allocation of goods (Bichler et al., 2010; McCabe et al., 1991). Understanding their impact in practice requires insights from economic theory, computer science, operations research, and information systems (Bichler et al., 2010) as both the facilitating technology, as well as the market design, strongly influence their adoption and efficiency in practice (Bapna et al., 2004; Goes et al., 2012; Lampinen and Brown, 2017). With the advent of easily accessible online platforms, smart markets have attracted many private individuals and have thus enabled a move away from pipeline value chains to P2P markets in which a variety of professional, but also private stakeholders (consumers and producers) interact directly (Bichler et al., 2010; Constantinides et al., 2018; McCabe et al., 1991; Slavova and Constantinides, 2017).

In the sustainability context, smart energy markets have gained importance with the transition to renewable energy resources. Numerous examples illustrate, however, that individual sustainability efforts, such as investments in infrastructure or demand side management, are often most efficient, if coordinated on the collective level (Slavova and Constantinides, 2017). In combination with intelligent agents, smart energy markets will be required to balance supply and demand in the highly complex environment with distributed, volatile generators in the electricity grid (Ketter et al., 2018). In addition, consumer-centric P2P markets for renewable energy are considered a promising vehicle for integrating residential consumers in the energy transition (Morstyn et al., 2018; Slavova and Constantinides, 2017; Weinhardt et al., 2019). Section 2.2.2 presents an overview of the theoretical background on P2P markets. Large parts of this thesis will be dedicated to P2P energy markets for integrating and incentivizing distributed renewable energy generators.

2.2.1 Market Design Theory

The term ‘market design’ was coined most prominently by Roth (2000) who defines it in its simplest form as “the creation of a venue for buyers and sellers, and a format

for transactions.”, p. 8. Economic theory provides a whole set of tools to assess the functioning of markets. The theoretical literature is mainly based in an area of game theory called ‘mechanism design’ and was adapted to a more applied vocabulary in the market design literature initiated by Roth. Mechanism design theory formally assumes that a market mechanism is a function or algorithm that takes individual consumption preferences as input and computes an outcome, which consists of an allocation of the traded goods and the respective prices (Mas-Colell et al., 1995). In sum, allocation and prices translate into transactions that must be settled once the outcome is determined. The market design refers to the entire ecosystem around the market mechanism (Mas-Colell et al., 1995; Roth, 2000, 2008): market participants, the bidding language, i.e. in which way participants can express their preferences, a format for transactions and settlements, the market mechanism itself, and finally the integration into the regulatory environment.

While the central function of a market is allocation and pricing of the resources that are to be sold, one should not underestimate the fact that markets essentially are social systems that govern interactions between individuals and thus influence human behavior depending on the defined rules (Dietz et al., 2003; Lampinen and Brown, 2017). Hence, the incentives market participants face when interacting depend on every aspect of the market design and must be well understood (Borenstein, 2005). If done thoroughly, the market design itself can make sure to “align social goals with the profit motivated interests of private parties by defining an appropriate set of rules and incentives” (Ketter et al., 2013, p. 264). Roth (2008) defines four general conditions for a well-functioning market:

- Providing thickness: sufficient proportion of buyers and sellers on the market
- Avoiding congestion: transactions must be processed fast enough
- Providing safety: transactions must be simple and without risk of data loss
- Avoiding repugnance: ban transactions that are considered unethical because they lead to undesired incentive structures

In addition, the domain and practical application also introduce several requirements and limitations to the market design problem. Although strongly rooted in game theory and sub-fields like auction theory, market design is not only relevant for economic theory, but also from a practical (Levitt and List, 2008; Roth, 1991) and behavioral perspective (List, 2011; List et al., 2006). In practice, market performance is strongly linked to the form, frequency, accuracy, and distribution of information provided to participants. Hence, the information systems implementing smart markets entail implications on a

markets' ability to achieve the conditions laid out by Roth (2008). Lampinen and Brown (2017) thus suggest: "Since markets are often instantiated in a technological form, we see an opportunity for our community to take an active role in designing markets and intervening critically where they do not work fairly or effectively." (Lampinen and Brown, 2017, p. 4331).

In particular, involving consumers in price-setting procedures in electronic markets has become increasingly relevant in the digital economy, as market platforms for anything are being created online and, oftentimes, accessible to almost anyone. Online auctions exemplify how the internet enables individual consumers to actively engage in pricing decisions (Bapna et al., 2003). Bapna et al. (2003) argue that the influence of the internet and the ubiquity of information are not fully understood yet in this context, and that classical economic theory may not incorporate all of the newly arising phenomena. Assumptions of rational, risk neutral bidders with perfect information, which are used in economic theory to derive analytical models of markets, usually do not hold in practical applications (Adomavicius et al., 2009; Klemperer, 2002). Additionally, even if they did, analytical models for interaction of agents under uncertainty, with a large number of bidders, different types of goods and/or in a repeated setting are highly complex and computationally hard to solve (Bapna et al., 2003; Bichler et al., 2019; Cramton et al., 2006). Bichler et al. (2010) thus argue that smart market design should, in a first step, focus on preference elicitation and provide (agent-assisted) decision support thereupon. This is the reason why, for decades now, IS research has been strongly involved in the creation and interactive design of smart markets (Bichler et al., 2010; Ketter et al., 2018) and the analysis of individual behavior in markets in practice (Goes et al., 2010; Lu et al., 2016).

2.2.2 P2P Markets

In recent years, research on electronic markets was particularly spurred by the rise of online platforms like Ebay, Amazon Marketplace, Uber, and AirBnb (Constantinides et al., 2018; Einav et al., 2016; Zimmermann et al., 2018). From an economic perspective, these online platforms technically facilitate 'P2P markets', two-sided markets in which neither side of the market is exclusively made up of professional companies (Einav et al., 2016). P2P markets generally differ in many aspects from traditional markets. The selling agents are usually not professionals producing a flexible quantity of goods for sale, but private individuals selling excess capacities (Kiesling et al., 2017) like free labor time, spare housing space, unused clothing, or excess electricity production. Moreover, both buyers and sellers may be very heterogeneous in demand patterns and production

capacities. Characteristics that make a certain domain attractive for peer production are (Einav et al., 2016): variability in demand, low scalability of production, and existence of accessible, well-functioning markets.

Peer-to-peer markets can be designed as mediated or as bilateral markets (Bichler and Segev, 2001; Malinova and Park, 2016). In mediated markets, a central intermediary collects demand and supply and determines the pricing. Purely bilateral peer-to-peer markets, by contrast, are less organized, and trades are arranged directly between individual buyers and sellers, which means that there is not a unique point of information aggregation. Various combinations or hybrid models of mediated and bilateral market designs can be implemented on digital platform (Tiwana, 2003). Airbnb for example hosts a platform on which individuals can offer housing capacities, and users are free to determine individual prices, but the monetary transaction runs through the platform (Lampinen and Brown, 2017). Airbnb charges a fee on this transaction for the information aggregation, the provision of the interface and the billing service. By contrast, Uber intervenes more strongly in the participants' interaction: It defines the prices for rides offered on their platform using proprietary algorithms (Subramanian, 2017) and also takes care of accounting and billing for the drivers. As opposed to centralized markets with mediating resellers that aggregate supply from one side of the market, P2P markets often use auction mechanisms in which buyers and sellers post their respective willingness to pay or sell for their supply and demand (Fontoura et al., 2005). Auctions allow for price discovery in settings in which there is high uncertainty about preferences or volatility of prices by letting buyers and sellers post their respective willingness to pay or sell for their supply and demand (Lu et al., 2016). Auctions thus are appealing for markets with many small, individual parties and iterative trading of similar items because auctions allow prices to respond to market regimes (Ketter et al., 2012).

In a P2P market, information about offered goods and individuals' preferences is dispersed over many individuals, so information aggregation is key for defining an efficient market allocation (Einav et al., 2016). One intuitive possibility to aggregate information is to centralize the process using a common communication platform, which is the approach companies like Uber take (Einav et al., 2016). The necessity to elicit and process distributed information inherently creates transaction costs in form of search and processing costs. From an economic perspective, transaction costs are the search, negotiation, and enforcement costs involved in market exchange (Coase, 1937; Williamson, 1979), so they include more than the settlement costs for processing a transaction – they include the costs that were incurred to identify and build trust in the counterparty and to define

and negotiate the transaction in question (Kiesling et al., 2017). In his seminal paper, Coase (1937) argues that the search and structuring of information on the supply, quality, and providers of goods on the market is too costly for individual agents, so they rely on firms as intermediaries to aggregate relevant information based on which they can make a buying decision. In smart markets without this type of intermediation, information must thus be provided in an accessible and processable manner for individual participants to make reasonable, possibly even automated, decisions.

2.3 Green IS

In recent years, it has become increasingly clear that information and communication technology plays a crucial role in fostering sustainability (Seidel et al., 2017; vom Brocke et al., 2012). As Watson et al. (2010) argue, information provision to enable and motivate sustainable “economic and behaviorally driven solutions”, p.24, is vital to enable a more efficient use of resources.

With this aim, Green IS has emerged as a new research area concerned with the application of information systems to sustainability problems. It is a multi-disciplinary research field tying together concepts known from research in the economics, psychology and engineering disciplines. Malhotra et al. (2013) conducted a meta-review in which they analyzed the existing literature on Green IS published in the top IS journals, and Gholami et al. (2016) updated the review with literature published additionally until 2016.² Both meta-reviews find that existing Green IS literature published in top IS journals contains mostly conceptual and analytical studies like reviews or case studies. In that, the authors identify a lack of design-oriented and impact-oriented research that can really make a difference in sustainability issues like reducing resource depletion and tackling climate change. Figure 2.1 shows a further update of these two reviews, which was compiled for this thesis. It contains the Green IS studies published since 2016, in addition to the studies already identified in Malhotra et al. (2013) and Gholami et al. (2016).

The update for this dissertation was compiled using the same methodology used in Malhotra et al. (2013) and Gholami et al. (2016), searching for the keywords ‘green’ and ‘environmental sustainability’ in the IS senior scholars’ basket of eight journals.³ The identified

²Note that, obviously, there exist further articles that are published in other journals and still discuss Green IS solutions, for instance in the managerial literature or in dedicated environmental journals, e.g., Gottwalt et al. (2011); Hopf et al. (2017); Ketter et al. (2018); Orlov et al. (2020); Tiefenbeck et al. (2018a, 2019)

³i.e. European Journal of Information Systems, Information Systems Journal, Information Systems Research, Journal of AIS, Journal of Information Technology, Journal of MIS, Journal of Strategic Information Systems, and MIS Quarterly

articles were categorized as proposed by Malhotra et al. (2013), p.1266, in: “conceptualize (review papers, conceptual frameworks, etc.); analyze (case studies, ethnographic analyses, quantitative empirical analyses, hermeneutics, etc.); design oriented (design science); or impact oriented (implementation and sustainability impacts using action research, in vivo real-time approaches, etc.)”.

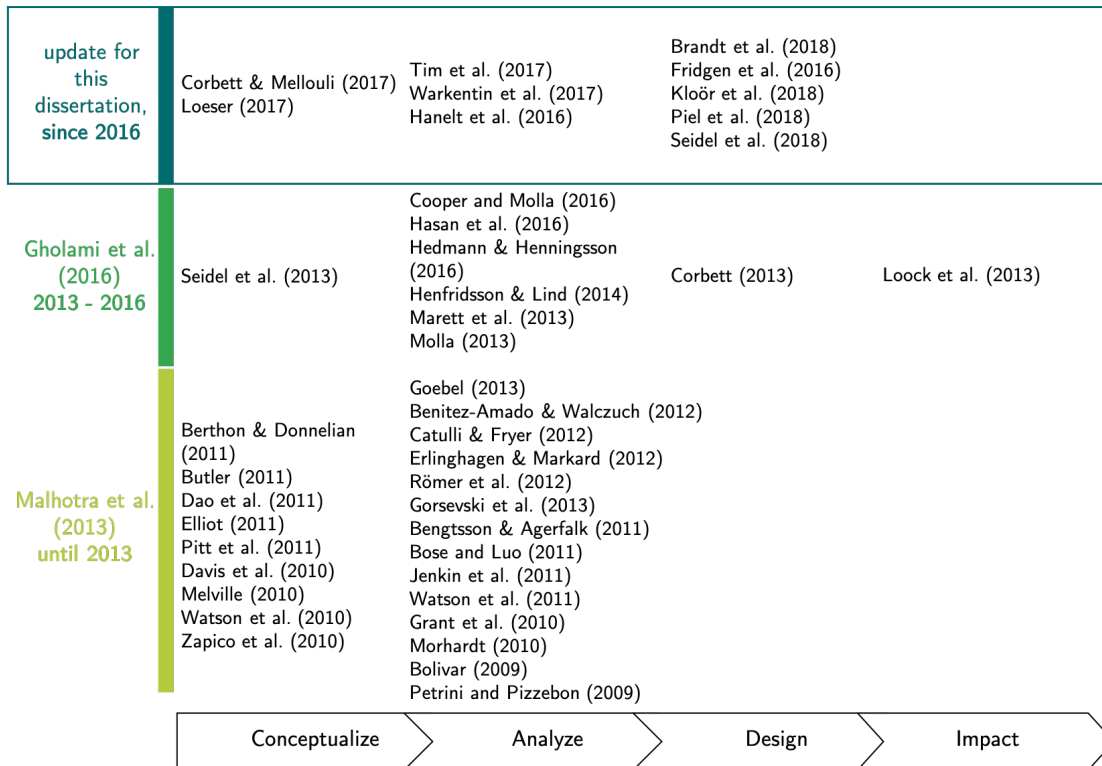


Figure 2.1: Updated literature review in the Green IS framework created by Malhotra et al. (2013, in light green). The dark green box contains articles published since the updated meta-review by Gholami et al. (2016, in green) . Existing research mostly focuses on analyzing and designing Green IS solutions, neglecting the practical design and impact of these.

In contrast to the strong focus on conceptual and analytic research that the earlier meta-reviews found, the updated landscape shows that there has been a much stronger focus on design-oriented research in the last five years. This is partially due to an EJIS special issue, which contained three design science studies (Brandt et al., 2018; Kloör et al., 2018; Seidel et al., 2018). However, the lack of impact-oriented research still persists.

In addition, the meta-reviews reveal that most existing work on energy-related topics in the IS discipline focuses on the use of IS to foster sustainable business practices in organizations (Malhotra et al., 2013; Watson et al., 2010). Just recently, few studies have

started to follow a more consumer- rather than purely organization-centered approach, e.g. Loock et al. (2013); Tim et al. (2018). However, keeping the human in the loop and integrating individual, private consumers in the sustainability movement is just as crucial. Residential households consume a large share of the worldwide total final energy consumption (International Energy Agency, 2020a) and their individual consumption patterns and sourcing preferences impact the sector as a whole. As the European Commission (2015, p.1) argues: “We have to empower consumers through providing them with information, choice and through creating flexibility to manage demand as well as supply.”

2.4 Research Gaps

Reviewing the Green IS literature revealed that more impact-oriented empirical research is needed for the IS discipline to truly contribute to a more sustainable society and to tackling climate change. The present thesis aims at the empirical validation of consumer-centric approaches for more sustainable resource use (an overview of the research studies is provided in Figure 2.2). Leveraging the existing research on decision support systems and smart markets (Sections 2.1 and 2.2), the articles in the following chapters address individual consumption behavior by examining feedback interventions on the one hand, and coordinating consumption patterns on the collective level using smart energy markets on the other. In that, the incentives consumers face in their resource consumption and sourcing choices will be a recurring theme throughout this thesis. Moreover, the tension between empowering conscious consumer choices by providing information versus automating decision processes using computational tools is discussed from different perspectives. Based on the presented topics and literature streams, this thesis addresses three research gaps in six research studies (Figure 2.2).

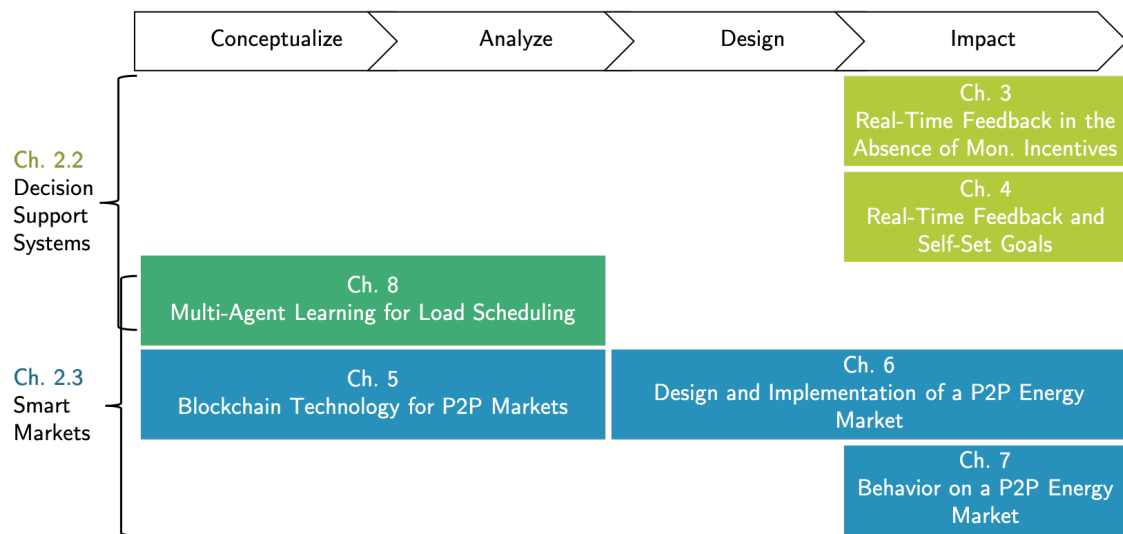


Figure 2.2: Overview of the studies presented in this thesis within the Green IS framework by Malhotra et al. (2013). The two major topics examined are feedback interventions (light-green) and P2P energy markets (blue), aiming at practical impact of Green IS with the human in the loop. Finally, Chapter 8 (green) combines research on algorithmic decision support and smart markets.

The following paragraphs provide a short description of these research gaps and the overarching research questions targeted in the subsequent articles (Chapters 3 – 8):

Real-Time Feedback & Incentives for Resource Conservation, Chapters 3 & 4

Technological advances enable the cost-effective provision of behavioral interventions to large and diverse groups of people at virtually any place in real time (Saha and Mukherjee, 2003; Swan, 2012). In the sustainability context, previous studies have found that real-time feedback on specific activities yields large and significant conservation effects in resource consumption (Tiefenbeck et al., 2018a, 2016). Yet, as existing programs often emphasize both financial and environmental benefits of energy savings, and consumers usually incur monetary savings from reduced energy consumption, it is unclear which motive primarily drives behavior change. Moreover, recent related work on feedback indicates that participant recruitment in framed field experiments strongly influences the motivation for and effects of feedback (Frederiks et al., 2016; Kelly and Knottenbelt, 2016). The lack of external validity of self-selected samples limits the generalizability of reported results when it comes to estimating the savings effects that can be expected among the general population (Allcott and Rogers, 2012; McKerracher and Torriti, 2013). This is problematic as the efficacy of conservation measures among the broader public and an understanding of motivational drivers for consumption behavior is crucial for policy

makers.

The article presented in Chapter 3 addresses this issue by confronting an uninformed sample of hotel guests with an IS-enabled feedback intervention on resource consumption. The large-scale, natural field experiment takes place in six hotels (total of 265 rooms, $n = 19,596$ observations), to tackle the research question: *Does real-time feedback induce conservation effects even among a sample of uninformed individuals in the absence of monetary incentives for resource conservation?*

The article presented in Chapter 4 investigates self-set goals as further driver for behavior change in response to feedback, examining the overarching question: *Do individuals define consumption goals when presented with real-time feedback on resource consumption?* In this second, large-scale field experiment, individuals in 413 residential households were presented with real-time feedback on resource consumption during showering. Complementary surveys provide evidence on individuals' self-set goals and their effects on consumption behavior.

P2P Energy Markets in the Real World, Chapters 5 – 7

In the energy sector, P2P markets could provide an opportunity for smaller renewable generators to participate in the energy market, and thus foster the integration of renewable resources (Andoni et al., 2018; Burger et al., 2016; Mengelkamp et al., 2017a; Morstyn et al., 2018; Ramchurn et al., 2012). Yet, while the integration of renewable energies and distributed generation is an urgent practical problem, and the number of articles on the topic has steeply grown over the past few years, existing studies on P2P energy markets mostly remain on a conceptual or, at best, analytical level (Mengelkamp et al., 2017b) – as is the case for a lot of Green IS research in general, see Section 2.3. This thesis tackles this issue by examining P2P energy markets systematically on every level of the Green IS framework (Malhotra et al., 2013), from the conceptualization through to an empirical field experiment with a P2P energy market (see Figure 2.2).

The article presented in Chapter 5 investigates the characteristics of blockchain technology, to understand: *What are the benefits and risks of implementing P2P markets on a blockchain infrastructure?* An analytical framework is created to characterize P2P markets based on economic theory and identify advantages and disadvantages of blockchain technology in this context. Applied to the use case of a P2P energy market for renewable energy, this framework condenses the specific features of the technology for this application and derives implications for the market design.

While market design theory provides a theoretical foundation (Bichler et al., 2010; Roth, 2008), there have been no empirical studies that examined the performance of P2P energy markets in the real world. Chapter 6 and 7 present a unique one-year field experiment with a P2P energy market in a local community in Switzerland: Chapter 6 describes the concrete market design and examines the user value proposition of this P2P energy market based on the data collected in the first three months of the experiment. This chapter presents the design of the system and the impact in terms of energy allocation and prices achieved to answer the question: *Which value propositions do P2P markets create from the user perspective and to what extent are they an effective measure to empower once passive consumers to assume a more active role in these markets?*

Empirical studies in various domains have shown that game-theoretical predictions on bidding behavior or their underlying assumptions often fail in practice (Adomavicius et al., 2009; Kahneman et al., 1990; Shogren et al., 2001). Chapter 7 presents a detailed analysis of the prices bid for locally generated solar energy during the field experiment, examining the questions: *Does the bidding behavior observed in the field deviate from cost-minimizing behavior? And how does it evolve over time?* The bidding for solar energy observed in the field and its evolution over time create novel insights on consumer preferences, market outcomes and expected prices in P2P energy markets. The study further inquires participants' preferences for algorithmic decision support.

Multi-Agent Learning for Load Scheduling, Chapter 8

Smart energy markets promise to provide incentives to balance electricity supply and demand in the grid, even incorporating the volatility of renewable and distributed energy resources. Yet, these dynamic mechanisms are often too complex and energy demand of larger appliances too abstract for individuals to adjust their consumption patterns accordingly (Reddy and Veloso, 2011a; Tiefenbeck et al., 2018a; Vytelingum et al., 2010). While the potential of decision support in form of providing feedback or recommendations on consumption patterns is thus limited in this context, intelligent agents can be a useful tool to manage the energy demand of appliances (semi-)autonomously (Ketter et al., 2018). Recent studies show that agents employing reinforcement learning can control electric loads to shift demand to times of renewable energy generation and to relieve pressure off the grid infrastructure (Valogianni et al., 2014; Vázquez-Canteli and Nagy, 2019), see also Section 2.1.2. Yet, autonomous load control has been studied mostly for individual agents with a focus on the employed machine learning techniques. So far, little is known on the interaction of intelligent, learning agents on a collective level (Vázquez-Canteli and Nagy, 2019).

Chapter 8 first explains reinforcement learning as a technique and then presents early research on reinforcement learning for load scheduling in a multi-agent setting: A simulation study is summarized, in which an auction mechanism coordinates charging demand for electric vehicles with very precise price signals for each consumer household (i.e., agent), tackling the question: *Is multi-agent learning in a smart energy market effective in reducing demand peaks while still respecting individual preferences?* This study thus connects the two research streams described in the previous sections, the application of intelligent agents (Section 2.1.2) interacting within a smart market (Section 2.2).

3. Article A) Real-Time Feedback in the Absence of Volunteer-Selection Bias and Monetary Incentives

3.1 Motivation

Individuals' choices and behavior are a key lever influencing energy consumption, along with technical energy efficiency of the products and infrastructure used (Sovacool, 2014). To tackle environmental challenges, it is important to put people at the centre of energy research, and to empirically validate how to promote sustainable decision-making among individual consumers. Energy consumption is a low-involvement topic for most people; many consumers are unaware of the energy efficiency of their homes and devices (Allcott et al., 2012b), or underestimate the long-term benefits of potential investments into energy efficiency upgrades (Allcott and Mullainathan, 2010).

As digitization advances, it becomes increasingly feasible to monitor the energy consumption of households, specific appliances, or activities in real time (Tiefenbeck, 2017). As a result, digitally enabled behavioral interventions can be deployed at population scale and become more powerful through personalization and context specificity. Beyond that, it becomes increasingly possible to systematically evaluate the impact of behavioral interventions with large and diverse samples of participants. The availability of high-resolution consumption data enables more and more personalized and flexible interventions (Tiefen-

beck, 2017). These developments may open up new avenues towards more powerful digital strategies for behavior change.

Yet, while early feedback intervention studies in which people were provided with information about their energy consumption reported large savings effects of 5-15% (Darby, 2006; Ehrhardt-Martinez et al., 2010) with small convenience samples, spurring the large-scale roll-out of smart meters in many countries, those savings have not materialized in larger field trials (Davis et al., 2013; Kelly and Knottenbelt, 2016; McKerracher and Torriti, 2013). The most widespread form of feedback intervention are ‘home energy reports’: periodic mailings that compare the electricity use of individual households with similar homes in the neighbourhood, thus tapping into social norms (Allcott and Mullainathan, 2010). Deployed at population-level (households can opt out, but few do), those programs typically yield electricity savings of 2% (Allcott et al., 2012a; Allcott and Mullainathan, 2010). Other programs use digital technologies, delivering feedback on electricity use via web portals or in-home displays; studies with large opt-in samples report electricity savings in the range of 1-5% (Buchanan et al., 2015; Delmas et al., 2013; McKerracher and Torriti, 2013; Schleich et al., 2013) – far less than the savings reported by earlier studies with smaller samples and a higher degree of involvement from study administrators (Abrahamse et al., 2005; Ehrhardt-Martinez et al., 2010).

Early studies were subject to several methodological issues that compromised the internal and external validity of the results, overestimating the savings potential of these feedback interventions (Davis et al., 2013; Delmas et al., 2013; McKerracher and Torriti, 2013). For instance, a meta-analysis of 156 field trials on energy conservation found substantially smaller savings effects of 1.99% for high-quality studies with adequate controls, compared to studies without such controls (9.57%) (Delmas et al., 2013). Likewise, a meta-analysis of 33 field trials on in-home displays found weighted mean conservation effects of 2.61% for high-quality (‘class A’) studies using representative sampling techniques, compared to 8.21% for ‘class C’ studies characterized by small samples of volunteers and a high degree of involvement from study administrators (McKerracher and Torriti, 2013).

Although randomized controlled trials eliminate most threats to the internal validity of studies (Campbell, 1969; Haynes et al., 2012; Vine et al., 2014), the external validity of the results may still be compromised if the people who choose to participate in a study differ from the study’s intended population (Davis et al., 2013). The vast majority of feedback programs on energy consumption use opt-in recruitment strategies, where participants actively register for taking part in those programs (Davis et al., 2013; Kelly and Knottenbelt, 2016). There is increasing evidence that individuals who sign up for

energy efficiency studies or demand-side-management programs are indeed different from the general population: Participation rates are higher amongst households with high levels of education and income (Clark et al., 2003; Sulyma et al., 2008), amongst more altruistic and environmentally concerned individuals (Clark et al., 2003), and amongst those with a higher interest and expertise in energy topics than the general population (Baladi et al., 1998; Kelly and Knottenbelt, 2016). Most behavioral programs do not even provide information about the number of households initially contacted, and those that do report participation rates in the range of 4-8% (Baladi et al., 1998; Ehrhardt-Martinez et al., 2010; Herter et al., 2013; Lossin et al., 2016; Schleich et al., 2013). These numbers have raised concerns that the results of opt-in studies may be largely biased by an already-motivated sub-group of the general population ('energy enthusiasts' or 'positive greens'), who represent only a small fraction of the population (Behaviors Union, 2008; Kelly and Knottenbelt, 2016). The response of these volunteers to the treatments may not be very indicative for the response of the general population (volunteer-selection bias). On the one hand, it is conceivable that those individuals are already more aware of effective energy conservation measures and have already taken action prior to the intervention, making it more difficult for them to realize additional savings in those studies (Tiefenbeck et al., 2018a). On the other hand, it is likely that they are particularly open and receptive to these interventions, thus inflating estimates of intervention effectiveness (Frederiks et al., 2016).

One commonality that the majority of larger energy-feedback trials share is that they provide aggregate consumption information at the household level. This makes it difficult for the individual to establish a link between the current action and its impact on energy consumption (Faruqui et al., 2010). The results of a recent randomized controlled field trial suggests that real-time feedback on a specific, energy-intensive activity may induce much larger savings (Tiefenbeck et al., 2018a). In a two-month study with an opt-in sample of 697 Swiss households, the treatment group received real-time feedback on the environmental impact of specific, energy-intensive activity (showering), while they could directly take action. The intervention yielded large and stable energy savings of 22% on the target behavior over the duration of the study. At the household level, this reduction led to much larger conservation gains – also in absolute terms – than aggregate feedback on energy use among the same pool of households. From a technology and cost perspective, the large-scale rollout of focused real-time interventions is increasingly feasible (Tiefenbeck et al., 2018a). Yet, given the decline in the effect size and the resulting wave of disillusionment once smart metering trials with aggregate feedback moved from small convenience samples to a broader population, the key question is whether the promis-

ing large savings effects of activity-specific real-time feedback will also materialize among individuals who do not self-select into a research study.

Another controversial issue is that the communication strategies of most energy- conservation programs focus on the financial benefits for the consumer as incentives for behavior change (Allcott and Sweeney, 2017; German Federal Ministry for Economic Affairs and Energy, 2016; Schwartz et al., 2015; United States Environmental Protection Agency, 2010). From a standard economics perspective, this approach makes sense, as rational consumers should respond to monetary incentives in their resource-consumption decisions (Goette et al., 2010; Karlan and List, 2007; Landry et al., 2006). Consequently, monetary incentives play a key role in demand side management (Borenstein, 2005; Olmstead and Stavins, 2009); they have the potential to break established consumer patterns and to initiate the development of new patterns of behavior by making an alternative behavior more attractive (Bamberg, 2006; Maki et al., 2016). However, in many contexts, financial motives are not a viable strategy to promote energy conservation: employees, tenants whose rents include utilities, or hotel guests do not pay the marginal cost of their energy consumption. While the provision of large, persuasive monetary benefits does not scale well to the wider population, monetary incentives may also crowd out the intrinsic motivation for pro-social behavior (Frey and Oberholzer-Gee, 1997; Sandel, 2012; Schwartz et al., 2015) and generate adverse effects (Gneezy and Rustichini, 2000): As individuals tend to internalize the logic of reward systems easily, monetary incentives can lead to the deterioration of morals and reduce intrinsic motivation (Sandel, 2012; Thøgersen, 1994).

The present study evaluates whether the large savings effects from digital activity-specific feedback (Tiefenbeck et al., 2018a) are also realistic in settings where a volunteer-selection bias can be ruled out, and where study subjects have no financial incentives for resource conservation. A smart device provides activity-specific feedback on resource consumption to uninformed hotel guests during a habitual resource-intensive activity: showering. The findings show that even in this setting, the digital behavioral intervention creates large and significant conservation effects of 11.4% or 0.215 kWh per shower. Given that most people take a daily shower, scaling up this kind of intervention could produce substantial energy (and water) savings. More importantly, the results suggest that activity-specific real-time feedback – and possibly other digital interventions – have the potential to transform behavioral interventions into a highly relevant policy instrument for fostering energy conservation and behavior change at population level (Allcott and Mullainathan, 2010; Delmas et al., 2013; Karlin et al., 2015).

3.2 Method

3.2.1 Experimental Setup

This article presents a natural field experiment which targets showering as an example of a resource-intensive, low-involvement activity. Participants were not recruited as individuals; instead, a collaboration with six hotels enabled the experiment to take place in hotel rooms without the guests being informed about the experiment upfront. Shower data was collected in batches after several weeks and without a time stamp, which guaranteed complete anonymity of the guests' identity. The study was approved by the Internal Review Board of the University of Bamberg. Similar natural field experiments in hotel contexts have been conducted to investigate the impact of other behavioral interventions such as commitment strategies (Baca-Motes et al., 2013) or social comparisons (Goldstein et al., 2008; Schultz et al., 2008).

Overall, guests staying in 265 different hotel rooms took part in the experiment. Hotel guests encountered smart shower meters as part of their rooms' bathroom equipment. The guests who stayed in rooms assigned to the treatment group received real-time feedback on how much energy and water they consumed over the course of their shower (details below). This activity-specific consumption feedback was displayed by a shower meter that had previously been used in framed field experiments in private households (Tiefenbeck et al., 2018a). Room assignment was randomized over floor levels and room categories to minimize confounding factors from differences in infrastructure (e.g., water pressure) and did not change throughout the study. Approximately 40% of the rooms were assigned to the control group. They serve as reference group to calculate the treatment effect. In those rooms, the same device was installed, but it displayed only water temperature. While the temperature reading does not convey information about the resource use and remains relatively static over the course of a shower, it indicates that the device measures data, thus reducing potential differences between the treatment and control group due to Hawthorne effects (Schwartz et al., 2013; Tiefenbeck, 2016).

As the smart shower meters are powered by the water flow via a small internal generator, the screen displaying the feedback switches on as soon as water flows through the device and remains active for up to three minutes after the end of a shower. Thus, short interruptions to the water flow (for instance while soaping) still result in a single shower being recorded. Before the device switches off, it stores the final data in its internal memory, which was read out at the end of the study.

The experiment took place in six different hotels in Switzerland, recruited based on ex-

isting contacts. Data was collected between February and April 2016. Four of the hotels focus on business customers and the other two on private tourists. The categorization into business and tourism was defined based on information provided by the hotels’ management. Of course, it cannot be ruled out that business hotels also had guests on private holidays, or that the tourism hotels hosted some business guests during the course of the study. Depending on their size, the participating hotels allowed the installation of shower meters in 10 to 96 of their guest rooms, respectively (see Table 3.1).

Hotel	Category	# Participating Rooms	# Observations
Hotel 1	Business, four-star	96	7,923
Hotel 2	Business, four-star	67	6,123
Hotel 3	Business, four-star	43	2,789
Hotel 4	Business, three-star	11	1,494
Hotel 5	Tourism, four-star	42	814
Hotel 6	Tourism	10	453
Total		269	19,596

Table 3.1: Overview of the participating hotels.

Display content in the treatment group

Most feedback devices display a bundle of elements rather than a single numeric metric in order to put the measurement data into context (Ableitner et al., 2017); frequently used elements include historic comparisons, peer comparisons, analogies, and energy savings tips. Likewise, the smart shower meters in rooms assigned to the treatment condition displayed water consumption in litres (one decimal), energy use in (k)Wh, current water temperature, a dynamic rating of the current energy-efficiency class (A-G) and a four-stage animation of a polar bear standing on a melting ice floe with stage transitions at predefined energy use thresholds. This is the same intervention with the same device and display elements as the treatment group of the opt-in household sample in Tiefenbeck et al. (2018a). The energy consumption displayed on the screen represents the lower bound of the energy used (without losses), and is calculated using the standard engineering formula for heat energy ($Q = m * c_p * \Delta t$, with heat energy Q , mass of water m , heat capacity c_p , and Δt the difference between the measured water temperature of the ongoing shower and the average cold-water temperature). In the analysis of energy savings, the same average heating efficiency and losses as in Tiefenbeck et al. (2018a) are taken into account.

3.2. Method

The energy-efficiency class displayed was inspired by the (static) energy-efficiency class scale indicated on household appliances in Europe. The smart shower meter dynamically indicates the energy efficiency of the current shower based on the energy use in the ongoing shower, starting in energy efficiency class A and progressing to B, C etc. at predefined kWh-thresholds; the thresholds were defined based on the distribution of energy use per shower in a pilot study. The four stages of the polar-bear animation are tied to the energy-efficiency class, and change with the transitions from B to C, D to E and E to F, respectively. While the polar bear may be an eye-catching and memorable display element, it does not seem to drive savings effects. A related study specifically examined the effect of variations of the design choices of the feedback elements; the results indicate that if the polar bear animation makes any difference, it reduces rather than increases the effectiveness of the display (Ableitner et al., 2017).




Figure 3.1: Smart shower meter. a) The smart shower meter for displaying real-time feedback on resource consumption to hotel guests was installed between the showerhead and shower hose; b) two snapshots of the treatment group's display; c) control group display.

3.2.2 Data

For each water extraction, the smart shower meter recorded energy and water consumption, average water temperature, interruptions and the duration. In addition, in 168 of the rooms, the average flow rate per shower was also measured. Based on the data stored on the device, energy consumption can be converted to water consumption and vice versa. Given the high correlation between water and energy consumption per shower (0.989), the choice of the unit of analysis does not change the results in any meaningful way (Tiefenbeck et al., 2018a). This article focuses on resource consumption in units of

energy in kWh. The raw data set included observations of 25,647 measured showers from 269 hotel rooms at six different hotels (Tiefenbeck et al., 2018b). In a first pre-processing step, the data was cleaned by removing outliers from malfunctioning devices; to this end, observations that deviated by over 3 standard deviations from the mean of the energy consumed or water volume per shower were removed from the sample – i.e., only observations in the interval $[(x - 3 * sd), (x + 3 * sd)]$ were retained. Furthermore, data points which most likely did not represent showers were removed – e.g. water extractions of volumes below 6.5 litres and observations deviating over 2 standard deviations from average temperature, which probably represent cleaning or other procedures. A member of the research team accompanied cleaning personnel at one hotel for several hours to gather information on cleaning practices to identify water extractions for cleaning. The specific choice of 6.5 litres was based on this assessment; robustness checks with other threshold values (5 litres or 10 litres) generated very similar results.

After this pre-processing step, the final data set included 19,596 showers from 265 hotel rooms (11,384 observations in the treatment group and 8,212 in the control group). Since the study is a natural field experiment with uninformed participants, there is no socio-demographic data about the guests who stayed in the rooms with the smart shower meters during the study.

3.2.3 Data analysis

The data points observed in the control group quantify the energy use per shower in the participating hotels without feedback. A simple linear regression model (and a log-linear transformation) estimates the treatment effect of the feedback intervention:

$$y_i = \beta_0 + \beta_1 x_i + \epsilon_i \tag{3.1}$$

$$\ln y_i = \beta_0 + \beta_1 x_i + \epsilon_i \tag{3.2}$$

where the dependent variable y_i is the energy consumption in shower i . The variable x_i is binary, indicating treatment ($=1$) or no treatment ($=0$), and thus coefficient β_1 estimates the treatment effect. The intercept β_0 represents the control group mean in this model, as $x_i = 0$ for observations in the control group. The results of this analysis are reported in Table 3.2, column 1, and in Table 3.3, column 1 with the natural logarithm of the dependent variable.

For rooms in which the smart shower meter also measured the flow rate, an additional

model that controls for the centred flow rate in litres f_i is computed:

$$y_i = \beta_0 + \beta_1 x_i + \beta_2 f_i + \epsilon_i \quad (3.3)$$

$$\ln y_i = \beta_0 + \beta_1 x_i + \beta_2 f_i + \epsilon_i \quad (3.4)$$

The results for Model (3.3), which includes the flow rate in the regression, are reported in Table 3.2, column 2, and for Model (3.4) in Table 3.3, column 2, with log-linear transformation. The average flow rate per shower could be measured only for 168 rooms and 8,824 observations, so only these data points are included in the estimation of these models. In both model specifications, standard errors were clustered at the room level to account for infrastructural influences. Two-sided t-tests were conducted to test whether the coefficients were significantly different from zero.

Finally, a fixed effects model with dummy variables for the individual hotels provides more understanding of the effects of the six hotels with their different infrastructure and setting. In this model, the constant represents the estimates for the largest hotel (Hotel 1) and dummy variables are included for the other hotels. Results are reported in Table 3.4.

$$y_i = \beta_0 + \beta_1 x_i + \alpha_2 h_{2i} + \alpha_3 h_{3i} + \alpha_4 h_{4i} + \alpha_5 h_{5i} + \alpha_6 h_{6i} + \epsilon_i \quad (3.5)$$

3.3 Results

This article presents a natural field experiment in the context of an energy-intensive habitual activity: showering. In a randomized controlled trial, guests at six Swiss hotels (see Table 3.1) encountered a smart shower meter fitted to the shower in the bathroom of their hotel room. The devices measured the energy and water consumption of every shower taken, and displayed feedback on each ongoing shower in real time.

Hotel guests exposed to real-time feedback consumed significantly less energy per shower than the control group (Figure 3.2). The treatment effect of the intervention is large and significant: Guests in the treatment group used on average 0.215 kWh less energy per shower than the control group mean of 1.883 kWh (Table 3.2, Column 1). This represents a reduction of 11.4% ($t(19,594)=-4.88$, $p<0.001$). Controlling for flow rate (Column 2), the effect is still highly significant, with a reduction of 0.188 kWh ($t(8,822)=3.59$, $p<0.001$), or 10.0%. To determine whether subsampling for observations in which flow rate is available biases this results, a third model specification for this subsample is included without

3.3. Results

controlling for flow rate (Column 3). The treatment effect is significant in all three models and large (ranging between 10.0% and 13.2%).

These results illustrate that activity-specific real-time feedback can be an efficient measure to foster energy conservation, not only among a volunteer sample, but also among a random, uninformed sample of individuals.




Figure 3.2: Effect of consumption feedback. Group-wise distribution of energy use per shower in hotel rooms with energy consumption feedback (treatment group) and the control group, shown as boxplots ($n=19,596$). The line in the middle of the box represents the median, the diamond the mean energy use. The box spans the first quartile to the third quartile, and the whiskers extend up to 1.5 times the interquartile range from the top or bottom of the box.

	Energy use per shower [kWh]		
	Model (3.1)	Model (3.3)	Model (3.1')
Consumption feedback (treatment=1, control=0)	-0.215*** (0.044)	-0.188*** (0.050)	-0.252*** (0.056)
Flow rate (mean-centred, l/min)	-	0.098*** (0.010)	-
Constant	1.884*** (0.032)	1.881*** (0.036)	1.902*** (0.039)
Observations	19,596	8,824	8,824
R ²	0.008	0.047	0.011

Table 3.2: Main treatment effect. Notes: Standard errors are in parentheses, adjusted for clustering at the room level; *, ** and *** indicate significance at the 5%, 1% and 0.1% levels respectively.

3.4 Additional Analyses

A log-linear regression model provides an alternative functional form of the estimation. The results (reported in Table 3.3) are consistent with the results of the non-transformed version reported above and show a strong and significant treatment effect of the real-time feedback. To further corroborate the reported results, the same models are run with varying filter thresholds, reducing the data pre-processing to an absolute minimum, with very similar results: Removing only observations deviating over five standard deviations from mean energy or water consumption, and mean average temperature, yields a sample of 25,490 out of the initial 25,647 observations. Running Model (3.1) on this sample yields a slightly smaller, but still highly significant treatment effect of -0.182 kWh (standard error of the mean 0.044, $p < 0.001$). To get an understanding of the effects of the six hotels with their different infrastructure and setting, a fixed effects model with dummy variables for the individual hotels is computed. The results are presented in Figure 3.3 and Table 3.4 and show that the treatment effect is highly significant, albeit slightly smaller than in Models (3.1) and (3.2). Only in Hotel 5, the energy use per shower differs significantly from the other hotels, which may be due to different infrastructure (e.g., more low-flow showerheads) or guest characteristics. Otherwise, the impact on energy use per shower is very similar between the different hotels. Regardless of the model specification, the treatment effect is large and significant; thus, non-self-selected participants also respond to real-time feedback in the complete absence of monetary incentives.

Furthermore, a cost-benefit analysis for installing the metering device in the hotels'

3.4. Additional Analyses

showers based on the treatment effect estimated in Model (3.1) provides insights on economic effects. Tiefenbeck et al. (2018a) assumed a retail price of 40 CHF for the smart shower meter and fuel cost for water heating of 0.128 CHF/kWh and water cost of 3.8 CHF/m³. Extrapolating from the treatment effect of 0.21 kWh and 3.56 litres per shower and assume on average 1.2 showers per day per room, as observed during the period of the present study, this results in an amortization time of 2.2 years.

	Energy use per shower (log) [kWh]		
	Model (3.2)	Model (3.4)	Model (3.2')
Consumption feedback (treatment=1, control=0)	-0.124*** (0.024)	-0.112*** (0.028)	-0.148*** (0.032)
Flow rate (mean-centred, l/min)	-	0.056*** (0.006)	-
Constant	0.413*** (0.017)	0.410*** (0.020)	0.422*** (0.022)
Observations	19,596	8,824	8,824
R ²	0.008	0.047	0.011

Table 3.3: Main treatment effect with log transformation of the dependent variable. Notes: Standard errors are in parentheses, adjusted for clustering at the room level; *, ** and *** indicate significance at the 5%, 1% and 0.1% levels respectively.

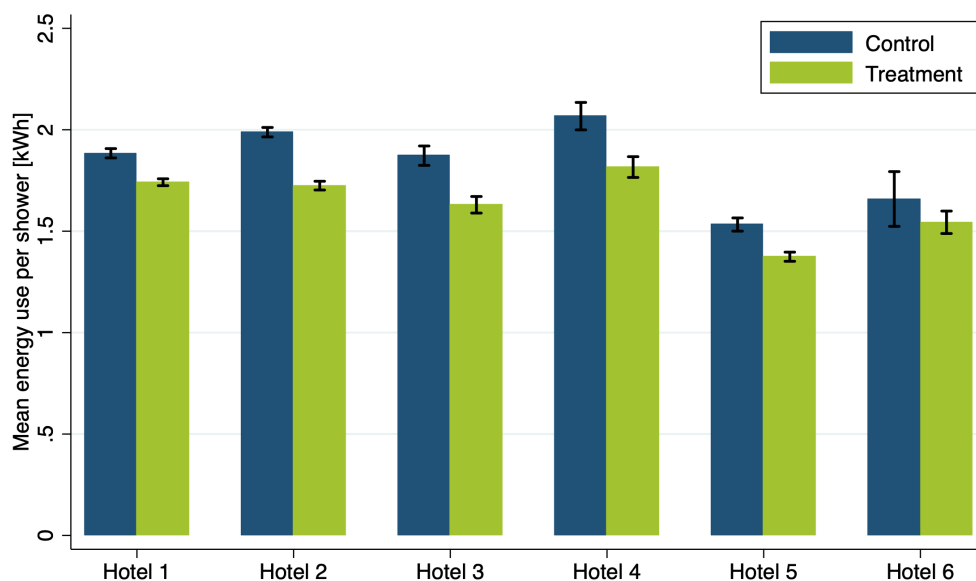


Figure 3.3: Effect of consumption feedback in different hotels.

	Energy use per shower [kWh], Model (3.5)
Consumption feedback (treatment=1, control=0)	-0.197* ** (0.041)
Hotel 2	0.037 (0.051)
Hotel 3	-0.072 (0.091)
Hotel 4	0.117 (0.090)
Hotel 5	-0.361* ** * (0.058)
Hotel 6	-0.191 (0.127)
Constant (Hotel 1)	1.918* ** * (0.045)
Observations	19,596
R ²	0.021

Table 3.4: Treatment effect of consumption feedback and fixed effects for different hotels. Notes: Standard errors are in parentheses, adjusted for clustering per room; *, ** and *** indicate significance at the 5%, 1% and 0.1% levels respectively

Comparisons with a volunteer-household sample

In line with the earlier findings on volunteer-selection bias (Davis et al., 2013; Kelly and Knottenbelt, 2016; McKerracher and Torriti, 2013), the treatment effect in the previous study with a volunteer-household sample (Tiefenbeck et al., 2018a, with 0.592 kWh, or 22%) was larger than the effect observed in the hotel setting (0.215 kWh, or 11.4%). However, it is important to note that this study does not seek to quantify the self-selection effect. The present hotel setting and the study context in (Tiefenbeck et al., 2018a) differ in multiple aspects other than the two key variables of interest. First, this study examines the behavioral response to feedback in the short term (one or a few nights spent per guest at the hotel). Consequently, any attempt to compare the two studies would need to focus on the short-term behavior of the household volunteer sample. Indeed, an analysis of the first three showers only in the volunteer sample also yielded a smaller treatment effect (17.8% or 0.46 kWh) than in the full two-month evaluation. Robustness checks using the first two, four or five (instead of three) showers were conducted, with very similar results. Second, individuals may react differently to feedback in their familiar environment at home vs. in a hotel room, or guests may perceive the mere presence of the shower meter as a signal that the hotel management cares about environmental issues and pays attention

to how much energy and water their guests are using.

Another remarkable difference relates to participants' energy use per shower in the absence of feedback, measured in the respective control groups. Hotel guests in the control group consumed 28% less (!) energy per shower, namely $m = 1.88$ kWh, $sd = 1.25$ kWh, than the control group in the household setting, with a mean of $m = 2.62$ kWh, $sd = 1.67$ kWh (Tiefenbeck et al., 2018a), $t(20,236)=-11.1$, $p < 0.001$. This difference is noteworthy for two reasons. First, according to standard economic theory, one would expect individuals to take longer showers at a hotel than at home, as they do pay a marginal cost for every kWh of consumed energy. Yet, the results suggest that the hotel guests did not exploit the zero marginal cost of consumption. The lower consumption in the control group may be largely attributed to differences in the technical infrastructure between the hotels and households: Lower flow rates in the hotel rooms are likely caused by a higher share of water-saving showerheads installed in the hotel rooms. Second, this difference may also partially explain the smaller treatment effect in the hotels. The study in the household setting had revealed a strong positive interaction between the treatment effect and baseline consumption: In the household study, a 1-kWh increase in baseline consumption led to a 0.32 kWh increase in the savings effect (Tiefenbeck et al., 2018a). To put it simply, it is far easier to cut a 20-minute shower short by a few minutes (and kWh) than to realize substantial reductions on a one-minute shower. The control-group mean of energy consumption per shower in the hotel sample is 1.88 kWh, compared to 2.62 kWh in the household sample (0.74 kWh difference). Interpreting the control group mean as a proxy for baseline consumption, an increase in baseline consumption by 0.74 kWh would increase the savings effect by 0.24 kWh, which almost exactly matches the difference in the observed savings effect. Thus, controlling for the lower consumption at the hotels in the absence of feedback, the savings effects among the hotel guests and among the volunteer household sample are in fact comparably large.

3.5 Discussion

In the case of aggregate feedback, most energy efficiency studies yielded much smaller savings effects once those interventions were evaluated with large, non-self-selected samples (Davis et al., 2013; Kelly and Knottenbelt, 2016; McKerracher and Torriti, 2013). In other words, those interventions resonated much less with broader, non-self-selected audiences than with those individuals who had opted to participate. By contrast, with a highly significant treatment effect of 11.4% among uninformed hotel guests, this study provides empirical evidence that activity-specific real-time feedback can induce substantial behav-

ior change among a broader population, even in a setting without monetary incentives for resource conservation. Thus, providing real-time feedback on a specific energy-intensive activity may not only generate large and persistent savings effects (Tiefenbeck et al., 2018a) among the small percentage of the population who tend to opt into energy efficiency studies (Baladi et al., 1998; Herter et al., 2013; Lossin et al., 2016; Schleich et al., 2013); the results indicate that the intervention successfully induces substantial behavior change and resource conservation among broader audiences.

Regarding the cost-effectiveness of the intervention in the hotel context studied, based on the savings effects observed, the device pays itself off in a hotel within 2.2 years on average, which is a very low amortization time as compared to other energy efficiency investments (Allcott et al., 2012b; Schopfer et al., 2018). Thus, the results suggest that even in settings where third parties pay for the marginal cost of resource consumption, activity-specific feedback can be a cost-effective and scalable strategy to foster energy conservation.

Despite all best efforts, there are limitations to this study. While effects for 100% of the hotel guests are measured, these may not be representative of the general population. Despite the diverse sample which includes different types of hotels, with different comfort categories, room rates and primary target customers (business vs. tourism), and in different locations, additional studies in other settings and other countries would be valuable. Furthermore, despite the efforts to limit differences in potential Hawthorne effects by displaying real-time water temperature on the control group devices here, it is conceivable that the treatment with real-time feedback on resource consumption draws more attention to the fact that the smart shower meter measures data than does real-time information on water temperature, conveying a stronger feeling of being monitored among its users. Moreover, due to the short duration of guests' hotel stays, this study is not able to examine effects over time. While several studies (lasting between two and 16 months (Tiefenbeck et al., 2018a, 2016)) document the mid-term effect stability of activity-specific real-time feedback with opt-in samples, further research needs to investigate whether the large savings effects also persist over time among non-self-selected participants.

This article provides robust empirical evidence that activity-specific real-time feedback can induce substantial behavior change and resource conservation – even for a sample of individuals who neither volunteered to participate in an environmental study, nor reaped financial benefits from energy conservation. Given the debate on volunteer self-selection (Davis et al., 2013) and the dwindling treatment effects of other feedback interventions once they are deployed among broader samples (Buchanan et al., 2015; Davis et al., 2013;

Kelly and Knottenbelt, 2016; McKerracher and Torriti, 2013), this empirical validation is critical to provide solid recommendations for the design of future energy conservation programs (Editorial, 2017). Information technology increasingly makes it possible to monitor behavior in real time, to provide individuals with feedback on their ongoing activities and to collect granular data on the real-world impact of interventions from millions of individuals in the field (Tiefenbeck, 2017) at rapidly declining costs. The results of this study highlight the potential of digital interventions to transform behavior in energy-intensive activities, which can be implemented and monitored at the population level.

4. Article B) Self-Set Goals in IS-Enabled Behavior Change

4.1 Introduction

Feedback (providing individuals with information regarding their actions/performance) has proven to be successful in changing individual behavior, not only in work-related contexts, but even in private contexts (Kluger and DeNisi, 1996; Moon and Sproull, 2008). In particular for habitual or low-involvement activities, which individuals perform in their everyday lives without extensive deliberation (Loock et al., 2013), studies in various areas provide evidence that individuals respond to feedback, including feedback on physical activity (Bravata et al., 2007) or energy consumption (Abrahamse et al., 2005; Allcott, 2011).

A vast body of literature in social psychology and behavioral economics has studied personal motivation and behavior change, and integrating the principles and insights of those fields can be beneficial to the IS discipline (Goes, 2013). One of the key mechanisms governing behavior is goal setting. After a “35-year Odyssey” of empirical research on goal setting, Locke and Latham (2002) attribute motivation for actions largely to task-related goals. They argue that virtually “all action is the result of cognition and motivation” (p. 707) and can thus be influenced by goal setting. Yet, the majority of their insights is based on laboratory experiments, not on observable outcomes in the real world. This article makes an attempt to link behavior change in the real world, induced by personal IS, to the existing knowledge on goal-setting theory.

In many situations, individuals set goals by themselves: They strive to run a marathon

in less than three hours (Allen et al., 2016), make an effort to attain a certain grade (Levy and Baumgardner, 1991), or define a specific body weight they want to achieve or maintain (Lupton, 2014). Goals can also be successful if assigned by an external party like a principal (Latham and Locke, 1991), or a software artefact (McCalley and Midden, 2002). While the motivational effect of assigned goals has been observed in many different settings, they can also be rejected or even backfire if applied in a wrong manner (Locke and Latham, 2002; Ordóñez et al., 2009). Several scholars emphasize that assigned goals can create unexpected, adverse reactions if they are not suited for the individual and suggest that future research should test goal-setting theory not only in the lab, but especially in natural settings (Loock et al., 2013; Lupton, 2014; Ordóñez et al., 2009). One difficulty with assigning goals is how to define appropriate ones. Given individuals' heterogeneous preferences and living contexts, they may themselves be best qualified to judge their specific situations and capabilities (Hinsz et al., 1997). Moreover, although externally assigned goals have proven to be efficient in organizational contexts (Locke and Latham 2002), they may be perceived as intrusive in private activities such as consumer choice (Camerer et al., 2003).

This raises the question whether in the design of feedback technologies it is necessary to specify and communicate goals (e.g., 10,000 steps per day), or whether users will formulate adequate goals by themselves even for habitual activities when exposed to feedback on their behavior. More generally: *Does IS-enabled feedback prompt individuals to self-set goals for the measured behavior by themselves and if so, is there a relationship between the difficulty of the self-set goals and the measured behavior?*

To answer this question, this study seeks to understand the goal-setting behavior of individuals in response to personalized IS-enabled real-time feedback on a habitual activity. In a framed field experiment, the goal-setting behavior of 413 participating households is examined to see whether individuals set goals by themselves; how ambitious these goals are; and what effect they have. To that end, individuals receive feedback on residential energy consumption, more precisely a personal IS displays real-time feedback on the resource consumption of the ongoing shower. That particular activity was chosen for four reasons: a) showering is an energy-intensive daily behavior that accounts for 12-16% of residential energy use (Bertrand et al., 2017), b) the amount of energy and water consumed is largely influenced by the individual's daily decision-making (unlike the energy use of most appliances, which is largely determined by technical characteristics), c) it is a low-involvement activity (Loock et al., 2013), and d) it is an activity (typically) carried out by an individual in isolation who is not being monitored by others and which is not

subject to clear social norms (e.g., how many liters of water or kWh of energy would be considered as a ‘normal’ shower).

The article is structured as follows: The first section presents related literature on IS-enabled feedback interventions and on goal-setting theory along with the hypotheses which will be examined, followed by a description of the experimental design and results. The article closes with a discussion of these results and implications for future research.

4.2 Related Work

4.2.1 Feedback and Personal Information Systems

Feedback interventions have proven to successfully induce behavior change with and without the use of personal IS in a variety of different applications. In the context of resource consumption, Allcott (2011) analyzes a series of programs in which households receive periodic Home Energy Reports that compare their electricity consumption to similar homes in the neighborhood. The author shows that the program cost (dollars per kWh saved) is substantially lower than the marginal cost of electricity production, which implies that these programs yield net benefits, not costs. Abrahamse et al. (2005) conduct a meta-study of over thirty experiments on different behavioral interventions for resource conservation. They find that feedback is effective in most cases, yielding resource savings of 2-28%, depending on the setting, yet the drivers of the observed behavior change are unclear. Other applications of feedback interventions include health, nutrition and transportation behavior. For instance, Vandelanotte et al. (2005) and Bravata et al. (2007) find that feedback improves nutrition behavior and physical activity and is most effective if the information is personalized. Froehlich et al. (2009) create an application that provides visual feedback on transportation behavior and find evidence that participants’ subsequent transportation behavior is more eco-friendly.

While all of the aforementioned studies report positive results for feedback interventions, the observed effect sizes vary with the target activity, the general setting, and the way feedback is presented (Kluger and DeNisi, 1996). A key finding is that in particular for habitual activities, the more closely (in time and place) the feedback is linked to a particular target activity, the better (Ehrhardt-Martinez et al., 2010; Froehlich et al., 2010). This emphasizes the importance of the medium by which the intervention is provided; ideally, it should be accessible during the targeted activity and process information in real time (Tiefenbeck, 2016). Personal IS increasingly make this possible at population level: With the ubiquity of smartphones, sensors, and networks, IS-enabled behavioral

interventions offer more and more possibilities to deliver personalized feedback (Froehlich et al., 2010; Hermsen et al., 2016).

Still, although many self-tracking devices aim at enabling the user to change her behavior in a desired way (Sjoeklint et al., 2015), a profound understanding of the motivations and effects on the user is missing (Gimpel et al., 2013). Most of the feedback studies presented above do not discuss the motivation or personal incentives for users to change their behavior once they are presented with IS-enabled feedback. Based on a meta-study on digital technologies for disrupting and changing behavior, Hermsen et al. (2016) conclude that more research should study the factors that drive and sustain behavioral effects in order to design more effective IS. Sjoeklint et al. (2015) argue that personal IS providing personalized feedback can represent an “instrument supporting the user’s willpower to reach a specific daily goal”, p. 4. They conduct an explorative study on wearable devices based on semi-structured interviews and identify activity-specific goals, partly assigned and partly self-set, as motivators to change undesired habits. Yet, several interview participants stated that that they used the feedback merely to learn about themselves and not to actually change their behavior, or that the goals assigned by the devices were inadequate or undesirable.

4.2.2 Goal Setting

The impact of goals on performance has been studied by psychologists and behavioral economists for decades. The founding fathers of the empirically based ‘goal-setting theory’ that integrates insights from both disciplines are Edwin Locke and Gary Latham. While they have focused on goals and performance in work-related tasks (Locke and Latham 2002), their studies are so extensive that many findings can be transferred to private settings (Locke, 1996). Goals can be defined “as the object or aim of an action” (Locke, 1996, p.181), or as “internal standards that specify the conditional requirements for positive self-evaluation, which provides incentive for action” (Williams et al., 2000, p. 161). According to goal-setting theory, goals regulate behavior by representing a conscious reference point that guides subsequent activities (Locke and Latham, 2006). They serve as motivators because the achievement of a goal leads to satisfaction of the individual, whereas goal failure leads to dissatisfaction (Locke, 1996). Outcomes of an activity are thus evaluated as gains or losses against that reference point. In behavioral economics, a reference point can be formally represented as a jump or discontinuity in an individual’s utility function; or as kink or discontinuity in the first or the second derivative of the utility function, depending on the model (Allen et al., 2016; Kahneman, 1992; Kahneman

and Tversky, 1979). This implies that satisfaction and dissatisfaction are typically perceived unsymmetrically: Failing a goal has a qualitatively stronger effect than achieving it. This gives theoretical support for the motivational power of goals, as it disproportionately increases the gain from achieving a goal compared to just missing it. In psychological terminology, the interpretation as reference point translates to: Goals involve discrepancy production ('feedforward control', i.e. initial motivation by setting goals above current levels) and discrepancy reduction ('feedback control', i.e. adjustments of effort to achieve desired goal) (Williams et al., 2000). For this mechanism to work, explicit feedback showing progress on task-performance is crucial (Kluger and DeNisi, 1996; Locke, 1996).

Key moderators of goal setting are thus task complexity, situational constraints, commitment to a goal and task-related feedback (Locke and Latham, 2006). In many daily activities in private settings, task complexity is low and situational constraints often cannot be circumvented for practical reasons. As there is no external control mechanism to incite commitment, the individual's own commitment to the goal is critical. Gaining this commitment is easier for self-set goals (Hinsz et al., 1997; Locke, 1996), as they have been chosen willingly and consciously by the individual. Moreover, self-set goals do not run the risk of being rejected for being too difficult. Ordóñez et al. (2009) even suggest that assigned goals can 'go wild' and create unexpected adverse reactions, for example if the same goal is applied to a set of different people and does not match the individual situation. This may explain why self-set goals have been more successful than assigned ones in some experiments involving private activities: In a series of lab experiments, McCalley and Midden (2002) provided feedback to $n=100$ participants on the energy consumption of simulated washing cycles. The authors find that participants who were explicitly prompted to self-set a goal by typing it into a user interface consumed less energy than both, participants who were not asked to set a goal and those who were assigned a fixed one. In a field study with 1,791 participants, Loock et al. (2013) presented subjects with smart meter readings of their household energy consumption. While some (randomly assigned) subjects were asked to freely set a goal for their energy consumption, others received a predefined default goal. The results indicate that default goals can have positive and negative influence on individual behavior, depending on goal difficulty; medium-level default goals are most effective, whereas too low or too high defaults are outperformed by self-set goals. In addition to being efficient, self-set goals do not have to be derived in a complicated manner, but are often chosen by the individual by herself (Locke, 1996, p.120): "When provided with feedback on their own performance or that of others, people often spontaneously set goals to improve over their previous best or beat the performance of others simply as a way of challenging themselves [...] The effect of

performance feedback (knowledge of score) depends on the goals set in response to it.” However, Locke (1996) states that one possible downfall is that individuals who self-set goals may choose less ambitious, easier-to-reach goals than goals that are assigned by a third party.

4.2.3 Research Gap & Hypotheses Development

Although goals have been shown to be important motivators of human behavior, the role of implicit, self-set goals as motivators for behavior change is not well understood yet, especially in the context of IS-enabled feedback that enables constant self-tracking (Sjoeklint et al., 2015). It is not clear whether personal IS can best support the desired behavior change by deriving and displaying explicit goals for the user or whether individuals will set a goal by themselves that is ‘adequate’ (ambitious, but achievable) in their specific situation in response to IS-enabled feedback. If the latter is the case, then IS should not assign explicit goals, to avoid adverse reactions (e.g., of individuals rejecting an assigned goal they perceive as too ambitious or as too easy to reach in their particular situation). Therefore, this article aims to understand which role self-set goals play in the behavior change triggered by activity-specific real-time feedback provided by personal IS. More precisely, five hypotheses are developed to understand whether individuals set goals by themselves (H1), whether self-set goals are ambitious (H2, H3), and what is the relationship between self-set goals and actual energy savings (H4, H5). The hypotheses apply knowledge from goal-setting theory and previous studies to IS-enabled behavioral interventions. The study is placed in the context of a low-involvement behavior which users typically are not highly interested in. Individuals receive real-time feedback on their resource consumption (water and energy) in the shower. First, the study investigates whether individuals choose a goal by themselves without being nudged to when presented with consumption feedback, as suggested by Locke (1996) and Hermsen et al. (2016):

H1: When exposed to real-time feedback on their resource consumption, individuals are likely to choose a conservation goal by themselves.

After having established whether individuals choose self-set goals when exposed to real-time feedback on their resource consumption, the difficulty of those goals are examined. As described earlier, self-set goals have several advantages over externally set goals: They do not have to be derived by a second party for every individual and they do not run the risk of being rejected for being inappropriate or too difficult. Yet, the disadvantage commonly brought forward is that self-set goals will not be ambitious enough, as Locke (1996) argues that individuals will choose their self-set goals below what they could actually reach. Since

goal failure creates dissatisfaction, a tendency to choose goals which are easily achievable seems intuitive. Based on this, the hypothesis that individuals will not choose ambitious goals is tested. As it is hard to define what is ‘ambitious’ in the present context, goal attainment is interpreted as a (negative) proxy for goal difficulty:

H2: Individuals’ self-set conservation goals are not too ambitious, i.e. they reach them most of the time.

As several studies discuss whether gender plays a role in goal-directed behavior, differences among genders are examined. For instance, Venkatesh et al. (2000b) present several arguments suggesting that men are more motivated by needs of achievement and more directed towards goals than women. Likewise, Levy and Baumgardner (1991) show in an experimental study that men choose more ambitious goals than women. Based on this literature, H3 reads:

H3: Men choose more ambitious conservation goals than women.

Finally and most importantly, this study examines the relationship between goal setting and resource conservation (i.e. task performance in the present setting). The theoretical explanations for the motivational effect of goals presented earlier (Allen et al., 2016; Locke, 1996; Locke and Latham, 2002; Sjoeklint et al., 2015), and empirical studies on the provision of goals for resource conservation (Loock et al., 2013; McCalley and Midden, 2002) suggest the hypothesis:

H4: Individuals who set a conservation goal for themselves consume less resources during the intervention phase than those who did not.

Locke and Latham (2002) further claim that the higher the goal, the better the individual’s task performance. One of the core findings of their work is that there exists a positive linear relationship between goal difficulty and performance. Against this backdrop, this study analyzes the relationship between goal difficulty and resource conservation:

H5: There is a positive linear relationship between goal difficulty and resource conservation.

4.3 Methodology

4.3.1 Experimental Setup

This paper presents a two-month framed field experiment which targeted showering as a resource-intensive, low-involvement activity. Participants received a smart shower meter, which displayed real-time feedback on their energy and water in the ongoing shower (see Figure 3.1). Participants installed the IS artifact themselves (simple process, no tools required). The device recorded energy and water consumption, average water temperature, interruptions, and duration of each shower. To collect data on the participants' behavior in the absence of feedback, all devices displayed only water temperature during the first ten showers. That period serves as baseline measurement; afterwards, the feedback intervention started. From then on, the device of two thirds of the households displayed energy and water consumption in real time (see Figure 3.1). A third of the households was assigned to the control group, which served as a reference group: Their shower meter continued to display only water temperature. Participants were recruited among a sample of 5,919 residential customers of the Swiss utility company ewz who had participated in an electricity smart metering study. Only one- and two-person households were admitted due to technical constraints in the storage capacity of the device. Individuals needed to opt in by filling out an online survey and agree to share their shower data with the researchers. Among the 1,348 households who registered, 700 were chosen on a first-come-first-served basis due to cost and logistics limitations. For further details on the experimental set-up and the randomization checks of the sample of participants, please refer to (Tiefenbeck et al., 2018a).

In a pre-experimental survey, participants disclosed socio-demographic information and answered several questions on personality and environmental attitudes. The questions on individuals' personality factors were based on the HEXACO Personality Inventory (Lee and Ashton, 2004) and the questions on environmental attitudes had the same wording and scales as the nationally representative Swiss Environmental Survey (Diekmann et al., 2009). Comparisons to Swiss national statistics and a representative Swiss environmental survey (Diekmann et al., 2009) indicate that the present sample is younger, more urban, and more educated, but slightly less (!) environmentally friendly than an average Swiss person. Among other questions in the pre-experimental survey, participants were asked how often they compared their own performance with others individuals' (five-point Likert-scale, 1 = never, 5 = often) and whether they acted environmentally friendly even if this incurred costs and efforts (1 = do not agree, 5 = agree). After the experiment, participants were asked to fill out a survey that included Likert scales assessing their per-

ception of the shower meter and on their goal-setting behavior. The question regarding goal setting was formulated as follows: “The smart shower meter displays information on your water and energy consumption since the baseline phase has ended. Have you set yourself a goal per shower that you try not to exceed (e.g., max. water volume or energy efficiency class)?” The question could be answered by checking yes or no. Those who responded yes were asked to specify that goal in a text box. A free text box was chosen in order not to prime participants on particular numbers or metrics.

4.3.2 Data

A complete data set (both surveys and shower data for the entire study duration) is available for 621 households. Among these, 208 had been assigned to the control group and 413 to the treatment group. After the experiment, participants in the treatment group received the post-experimental survey with the question on their goal-setting behavior. The analyses will thus focus on these 413 households.

Since participants indicated their goal in a free text box, some of them stated multiple goals or a range (e.g., “30-50 liters”). For the analyses, ranges were converted to the mean value, (40 liters in the example), except for the analysis presented in Figure 4.1, for which the upper and the lower bound of those responses were weighted with 0.5 each (otherwise participants who indicated a range would be counted twice). As most of the participants stated a goal related to water consumption in liters, the following analyses will focus on water consumption in liters rather than on energy in kWh in the following analyses. In this study, energy consumption can be easily converted to water consumption and vice versa based on the data stored on the device. Given the high correlation between water and energy consumption per shower (0.989), the choice of the unit of analysis does not change the results in any meaningful way. If a participant stated multiple different types of goals, like “below 39 °C and below 50 liters” (which happened in 16 cases), the water volume-related goal or the one that was easiest to convert to water volume was chosen (to have one common metric). Whereas the shower data recorded in two-person households include observations from both household members, the survey was completed by one person per household. As a result, answers to questions on attitudes or goals set thus solely reflect the respondent’s perspective. Therefore, for data consistency, 2-person households were excluded in the main numerical analyses that involve the shower measurement data, i.e. for the evaluations of hypotheses H2, H4, and H5. The remaining subsample includes 196 individuals and 10,878 measured data points. (For the sake of completeness and as additional sensitivity analysis, all analyses were also conducted with the full sample including the 2-person households; the results are very similar.)

4.4 Results

The real-time feedback provided by the smart shower meters results in substantial behavior change and resource savings. As soon as the device starts displaying feedback from shower 11 on, resource consumption per shower in the treatment group drops considerably. Overall, participants in the treatment groups reduce their water consumption per shower by 9.5 liters compared to the control group. This amounts to 22% savings both in water and energy consumption; the savings effects are stable over the duration of the study. For a detailed description of the analysis and of the effects induced on consumption behavior, please see Tiefenbeck et al. (2018a). While that article focused on the savings effects of the intervention, their stability, and the cost-effectiveness of the intervention, the present article digs deeper into the underlying psychological mechanisms. In particular, this article assesses whether the large savings were mediated by the participants setting goals for themselves in response to the feedback.

Survey statements are examined to answer the first hypothesis (H1), whether the participants set themselves a goal regarding their maximum resource consumption per shower in response to the real-time feedback. (Note that they had not been exhorted or encouraged to do so at any point of the study.) The post-experimental survey shows that, indeed, 221 of the 413 questionnaire participants confirm that they have set a goal by themselves, which implies: When exposed to real-time feedback on their resource consumption, many individuals (54%) in the sample did indeed set a conservation goal by themselves. To get a better understanding of the goals individuals chose, Figure 4.1 depicts the absolute frequency of goals stated as maximum water consumption goal in liters. This subset includes 154 surveys. As the figure illustrates, most participants specified round numbers (i.e., multiples of tens), which is in line with existing studies on self-set goals (Allen et al., 2016). A total of 94 of the 154 liter-goals (61%) were round numbers; overall, 127 goals (83%) were multiples of five, whereas other numbers were hardly ever chosen. The most popular goal was 30 liters, which was chosen by 23 participants.




Figure 4.1: Self-set goals related to water consumption. Participants exhibit a clear propensity for choosing multiples of 10.

The aim of H2 was to examine how ambitious the self-set goals were. The results show that most individuals chose a goal well below their average water consumption in the baseline period. This is consistent with other studies on resource consumption behavior, which show that individuals try to conserve resources if presented with relevant information (Abrahamse et al., 2005; Allcott, 2011). Yet, the relative difficulty of the chosen goals is surprising: On average, people chose a goal 12.1 liters (sd 21.2) below their average baseline consumption, which corresponds to a 20.1% reduction in their water consumption (sd 26.7). Note that the high standard deviation indicates that the goals vary substantially between individuals and that goals were not defined as percentage savings relative to the baseline by the participants, but as absolute value or range they tried not to exceed. Figure 4.2 depicts to what extent the participants who stated having set themselves a maximum consumption goal reached that goal. The histogram displays deviations between their self-set goal and their actual resource use in the intervention phase (when the feedback was visible). The subset of observations from those individuals includes 10,878 showers. A negative deviation from zero indicates that the individual used less water than specified in the individual goal, i.e. the individual met her conservation goal. The deviations are roughly normally distributed around the zero marker. The

interpretation of this distribution is that the self-set goals serve as reference points for the users.

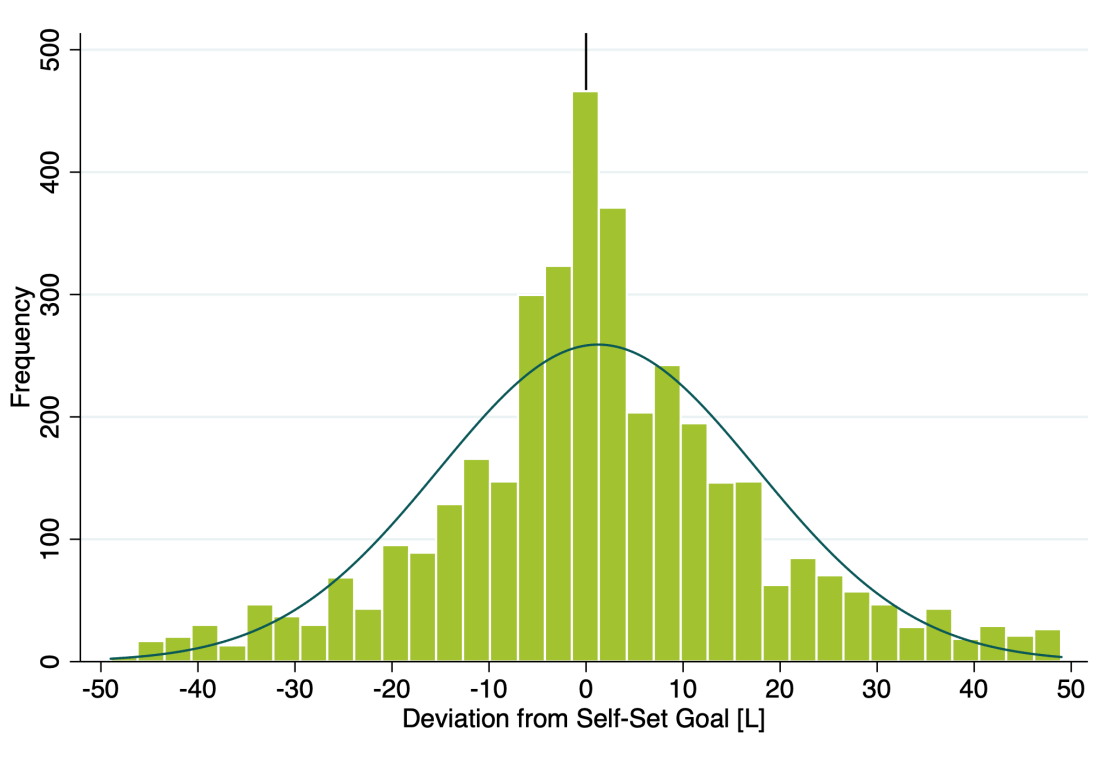


Figure 4.2: Distribution of deviations from self-set goal. The bulging around (in particular just below) 0 indicates that self-set goals serve as reference points for the individuals.

Still, remarkably, participants on average fail their goal by 2.7 liters (sd 25.3). Thus, the prediction of H2 that individuals set themselves unambitious goals which they can achieve easily is not confirmed. This is in clear contradiction to the literature on goal-setting theory in which lacking goal difficulty is used as a key argument why goals should be assigned externally, as opposed to encouraging self-set goals. Possible explanations are that situational constraints did not allow for the desired, ambitious conservation, or that the self-set goals were viewed as rough guide rather than as strict limits. For robustness, additional analyses reconfirmed that the fact that the average shower does not meet the goal is not biased by outliers: The median deviation is 1 liter above the individual goal and participants fail their self-set goals in 50.2% of the showers. Regarding H3 (assessing gender differences in goal difficulty), men on average set goals 22.4% below their baseline water consumption, whereas the average goal reported by women was 16.8% below their baseline water consumption. This difference, however, is not statistically significant ($t(90)=-0.67$, $p=0.50$), so the hypothesis that men choose more ambitious goals cannot be

confirmed. Moreover, these numbers are confounded by the fact that men in this sample started with a considerably higher average baseline consumption, 48.4 liters (sd 34.6) than women, 40.3 liters (sd 22.5). Previous analyses (Tiefenbeck et al., 2018a) show that there is a strong positive interaction between resource conservation and baseline consumption – simply put, it is easier for high consumers to conserve resources.

To deduct practical implications for the design of IS-enabled behavioral interventions, the crucial question is how self-set goals actually affect behavioral outcomes. In other words, did participants who set themselves a goal use less resources than those who did not? A linear regression model (4.1) is used to test the relationship between self-set goals and resource conservation in this sample. Regression analysis is a very established technique in behavioral research (Loock et al., 2013; Moon and Sproull, 2008; Venkatesh et al., 2000a). In this linear regression, dependent variable y_i is water consumption per shower of household i in the intervention period. The independent variable x_{1i} is a binary variable indicating whether individual i reported a self-set goal (=1) or not (=0). The model controls for baseline consumption x_{bi} , which is the average water use of household i during the first ten showers.

$$y_i = \beta_1 x_{1i} + \beta_b x_{bi} + \epsilon_i \quad (4.1)$$

The first column in Table 4.1 shows the regression results for all treatment participants in 1-person households. The average water consumption per shower during the intervention phase was significantly (4.8 liters per shower) lower among those individuals who self-set a goal. Thus, H4 is confirmed. The numbers also reveal the strong correlation between baseline consumption and savings effects. Mean water consumption during the intervention phase was 37.6 liters (sd 29.2) for participants who chose a goal as opposed to 40.8 liters (sd 40.6) for those who did not. Thus, not only are participants likely to set a goal by themselves, but those who do consume significantly less resources than those who do not (H4).

To evaluate the results for H5, the relationship between goal difficulty and conservation, a second regression is conducted. The dependent variable, y_i , in Model (4.2) is again water consumption per shower for household i , but goal difficulty in liters now serves as independent variable x_{2i} instead of the binary value of having set a goal or not. As described above, goal difficulty x_{2i} is defined, for the households who chose a goal, by the difference between stated goal and measured baseline consumption in liters.

$$y_i = \beta_2 x_{2i} + \beta_b x_{bi} + \epsilon_i \quad (4.2)$$

The results for Model (4.2) can be found in the second column of Table 4.1. Coefficient $\beta_2 = -0.56$ is negative and highly significant, which implies that the more ambitious a goal is, the more individuals conserve on average (Table 4.1, Column 2). Each liter the self-set goal was below the individual baseline consumption yields 0.56 liters in average water conservation per shower. This means that higher goal difficulty increased the conservation effect, which confirms H5. These findings are in line with one of the key propositions of Locke and Latham (2002), who argue that there is a positive linear relationship between goal difficulty and task performance. Analyses of the consumption medians (instead of means) that were conducted as robustness checks provide very similar results.

Note that the findings do not establish a causal relationship in H4 – based on the regression results, it is not possible to say whether self-set goals cause the savings, or whether the kind of individuals who set themselves a goal are also the ones who conserve more resources. Therefore, confounding effects were assessed to evaluate whether individuals who are generally more concerned about the environment were more likely to set themselves a conservation goal and to conserve more resources. If that were the case, then goal setting might be merely another manifestation of interest in the topic, rather than serve as mediator for behavior change. To evaluate whether goal-setting behavior can be explained by other latent variables collected, like individuals' personality traits, their tendency to compare themselves to others, or environmental awareness, the results of the pre-experimental survey served as input to estimate Model (4.3). Column 3 in Table 4.1 contains the results of a third regression which included the same variables as in Model (4.1) in addition to control variables for the self-reported environmental awareness, tendency to compare oneself to others, and the six HEXACO-dimensions (honesty, emotionality, extraversion, agreeableness, conscientiousness, openness). The results do not detect any significant influence of these variables on resource conservation, whereas self-set goals and baseline consumption are still highly significant in this extended model. In particular, the estimates for goal-setting are barely affected by environmental attitudes and personality traits. While confounding effects of other, unobserved variables on the correlation of resource savings and goal setting cannot be ruled out, these results are indicative of self-set goals acting as mediators of resource conservation.

4.4. Results

	Model (4.1)	Model (4.2)	Model (4.3)
Goal set? (0=no, 1=yes)	-4.75** (1.73)	-	-5.11* (1.99)
Goal difficulty (i.e. baseline - goal [L])	-	-0.56*** (0.16)	-
Baseline consumption [L]	0.72*** (0.06)	0.86*** (0.07)	0.68*** (0.09)
Environmental awareness	-	-	-3.13 (1.90)
Tendency to compare one-self to others	-	-	-0.34 (0.94)
Honesty	-	-	-0.21 (1.76)
Emotionality	-	-	3.16 (1.85)
Extraversion	-	-	-0.28 (1.67)
Agreeableness	-	-	-0.46 (1.48)
Conscientiousness	-	-	-0.34 (1.74)
Openness	-	-	-0.97 (1.43)
Constant	8.48** (2.58)	-3.71 (2.02)	19.08 (14.07)
Observations	196	90	153
R ²	0.76	0.76	0.76

Table 4.1: Relationship between goal setting (Models (4.1) and (4.3)) / goal difficulty (Model (4.2)) and water consumption in liters. Standard errors are in parentheses, adjusted for clustering at the household level. (*, **, and*** indicate significance at the 5%, 1%, and <0.1% levels, respectively.)

4.5 Discussion & Conclusion

4.5.1 Discussion

This paper evaluates goal-setting behavior in response to IS-enabled real-time feedback on the individual's resource consumption in the ongoing shower. As a first result, the majority (54%) of individuals set a conservation goal by themselves without being exhorted or encouraged to do so in any way. While goal-directed behavior has been studied in the context of various high-involvement activities (Allen et al., 2016; Levy and Baumgardner, 1991; Locke and Latham, 2002), it is remarkable that it also occurs in the context of a low-involvement activity like showering and with respect to the consumption of the low-involvement good energy, which individuals typically show little interest in (Attari, 2010). It is important to note that there was no mentioning of the study investigating personal motivation or goal-setting behavior at any previous point in the study, nor any exhortation that participants should set themselves a goal – they simply did so by themselves in response to the feedback intervention. Moreover, there is no evidence for significant influence of any of the personality factors that were assessed in the pre-experimental survey on goal setting in analyses that were conducted for completeness.

The results show that individuals chose ambitious conservation goals which they were not able to meet on many occasions. Based on the results for H2, the concern raised by Locke (1996) that individuals may not choose ambitious goals by themselves to avoid the negative sensations associated with failure, does not seem warranted in the present setting. Moreover, since goal difficulty has a positive impact on conservation in this study (Table 4.1), goal failure for an ambitious goal can still imply substantial resource conservation, i.e. success of the intervention. Future research should thus investigate whether motivating individuals to choose even more ambitious goals can further increase the effects.

The results on H4 indicate that individuals who self-set a goal used significantly less water in response to the feedback intervention than those who did not. Moreover, for H5, the data shows a positive correlation of goal difficulty and resource conservation: Individuals who set themselves an ambitious goal (far below their resource consumption in the baseline period) exhibit larger conservation effects. It is important to note, however, that the research design does not allow for establishing a cause-effect relationship between self-setting of goals and conservation behavior: It is both conceivable that a) the act of self-setting a goal mediates conservation efforts and the associated savings, as McCalley and Midden (2002) found, or that b) the kind of individuals who self-set a goal were systematically different from those who did not in the first place (in a dimension that

was not assessed). If the former is true (i.e., that the act of self-setting a goal induces larger conservation effects by increasing task motivation), then personal IS should actively prompt individuals to set themselves a goal. On the other hand, if the latter applies, then the act of goal-setting could be just another manifestation of the individual's interest and commitment to the conservation task, which at the same time produces larger conservation effects. Thus, it cannot be ruled out that the goal-setters in this study exhibited some unobserved characteristic (e.g., general interest in technology or numbers) that caused both the goal-setting and the higher conservation effects among that subset of participants. However, the analyses did not reveal any indication that those who set a goal by themselves were more environmentally aware, nor that they had a higher general tendency to compare themselves to others. These results are indicators that self-set goals act as mediators of resource conservation. Further IS-enabled field studies could improve the identification and establishment of a causal relationship by random assignment, whereby one treatment group is systematically prompted to self-set a goal, while another treatment group is not.

4.5.2 Limitations

Against all best efforts in the study design, there are some limitations to the results of this study. Participants who indicated in the survey at the end of the study that they had set themselves a maximum consumption goal were asked to state that goal in a free text box. They answered in different metrics, some set multiple goals, and some defined target ranges rather than maximum values. This introduces some blurriness in the data that the study design cannot fully control for. However, the open way to answer the question enables a less biased view of how individuals chose their personal goals. Wherever possible, established validated scales like the HEXACO inventory were used (Lee and Ashton, 2004). In order to compare the environmental attitudes of the sample with the general population, the same wording as in the Swiss environmental survey (Diekmann et al., 2009) was applied. Nevertheless, like any self-reported data, there is a risk of social desirability bias. For the reported self-set goals, it is impossible to say whether the participants set their goals at the beginning of the intervention phase or if they adjusted their goals over time, as their task-related knowledge increased. Future research should investigate how self-set goals evolve over time.

One limitation is the fact that participation in the study was voluntary. Despite the best efforts to mitigate those issues by comparing socio-demographic data and environmental attitudes with the general Swiss population, caution may be warranted with the external validity of the results, as in any study with an opt-in sample (see Chapter 3). While there

is no ‘baseline’ measure as to whether individuals also had maximum consumption goals for the target behavior prior the feedback intervention, the survey asked participants how much water they thought they used per shower, both before and after the study. The results indicate that prior to the intervention, most individuals had a poor sense for their actual water consumption, which improved significantly with the intervention (Tiefenbeck et al., 2018a). Thus, it is not possible to establish a baseline for goal setting as individuals cannot be expected to come up with meaningful goals for their shower-related resource consumption by themselves in the absence of the feedback intervention.

Another aspect that should be tackled in future research is a systematic comparison of self-set goals to externally assigned goals. On the one hand, the results suggest that IS-enabled real-time feedback successfully induces many individuals to set themselves ambitious goals even for low-involvement activities, and that individuals who set themselves ambitious goals in response to the feedback intervention conserve more resources. On the other hand, future research should directly compare the impact of self-set and externally assigned goals to determine in which situations which strategy is more effective and in line with individuals’ preferences.

4.5.3 Conclusion

Personal information systems are becoming more and more pervasive and make it possible to reach mass audiences (almost) in real time. The combination of feedback interventions with these technologies has the potential to support individuals in changing their behavior into healthier, more sustainable, or socially desirable habits (Consolvo et al., 2009; Hermsen et al., 2016; Li et al., 2010). It is thus in the interest of society to understand the mechanisms that drive behavior change in response to that kind of increasingly ubiquitous information.

Based on goal-setting theory, this study investigates self-set goals as mediators of individual behavior change in response to real-time feedback in the context of a low-involvement activity. The findings are based on real-world measurement data (over 10.000 observations) from a framed field experiment with 413 participating households. To the study authors’ best knowledge, this is the first field study that investigates the formation of self-set goals in response to IS-enabled feedback. In particular, the results show that even in the context of a habitual activity and regarding the consumption of a low-involvement good, many individuals set themselves a goal in response to real-time feedback, without ever being encouraged to do so. Against the predictions of Locke (1996), the data illustrates that individuals tend to set themselves ambitious goals that they do not achieve

easily. In line with goal-setting theory, the evidence reveals a positive relationship between goal setting and resource conservation and, in particular, between goal difficulty and resource conservation, meaning the more ambitious the chosen goal is, the more resources are conserved.

Since the majority of participants in this study set a goal by themselves even for the consumption of a low-involvement good like energy, it is likely that most individuals will set themselves goals for other activities tracked by IS as well. Combined with the strong relationship of goal setting and effect size found, self-set goals may explain the strong behavioral effects observed for existing self-tracking applications (Consolvo et al., 2009; Froehlich et al., 2010; Lupton, 2014). Moreover, given the difficulty of defining adequate goals externally, and the risk of goal rejection and adverse reactions associated with external goals, these findings question whether IS artifacts should assign goals to users. The results rather suggest that IS artifacts should encourage users to self-set goals and provide functionalities to store and display them during the target activity.

5. Article C) Blockchain Technology as Enabler for P2P Markets & Excerpt from Article D) as Use Case Analysis for the Energy Sector

5.1 Introduction

Various electronic markets which enable the exchange of goods or services between peers have emerged in the past decades and enjoy increasing popularity among consumers (Bichler et al., 2019; Einav et al., 2016; Parker and Van Alstyne, 2005). Technology-enabled markets empower consumers to also become producers of goods or services that they sell to others online in more and more sectors, so-called ‘prosumers’ (Ramchurn et al., 2012). Airbnb, Amazon, Etsy, Uber, and other sharing platforms provide the possibility to exchange goods, or service capacities among peers, or share knowledge, often in real time (Einav et al., 2016; O’Reilly and Finnegan, 2010; Van Alstyne et al., 2016).

While these platforms experience great popularity, some critical issues persist: Platform operators represent intermediaries that control market mechanisms and act as trusted third parties to reduce friction on the market (O’Reilly and Finnegan, 2010). Consumers and regulators are increasingly worried about increasing power of the few companies that operate these markets and question their integrity (Constantinides et al.,

2018; Roger Aitken, 2017; Subramanian, 2017). In various sectors, a small number of internet companies have gained almost monopolistic power (Slavova and Constantinides, 2017). Blockchain-based markets promise to enable the same kind of peer-to-peer (P2P) interactions as other market platforms, but can be operated by the market participants themselves instead of being centrally controlled by one company hosting the platform (Avital et al., 2016; Subramanian, 2017). This idea of decentralized marketplaces appeals to many as it may offer an alternative way to enable exchange of goods and services between peers — which seems particularly attractive at a time in which trust in political and corporate institutions has been shaken by the financial crisis and privacy concerns (Beck et al., 2017b). While there is a growing body of literature on P2P markets instantiated by digital platforms in the IS discipline (Constantinides et al., 2018; Parker and Van Alstyne, 2005; Zimmermann et al., 2018), research on blockchain-based markets is still in its infancy. There is still little understanding of the economic impact and of the socio-technical systems and markets that are being created using blockchain technology (Beck et al., 2017a; Malinova and Park, 2016) and a differentiation to existing, centrally governed platforms is missing. Although multiple articles promote blockchain technology for its potential of revolutionizing, for instance, the financial and the energy sector by enabling decentralized P2P exchanges (Hasse et al., 2016; Kastrati and Weissbart, 2016; Laszka et al., 2017; Mengelkamp et al., 2017a), only few provide concrete details on how a successful market design should look like and what the challenges are. It is unclear to what extent blockchain-based markets can meet the requirements that existing market platforms successfully address and how they can create an added benefit compared to a system with centralized governance (Constantinides et al., 2018). This article combines technical aspects and an economic perspective to investigate what differentiates a mediated, centralized market platform from a blockchain-based market, thus tackling the following research question:

RQ1: Can blockchain infrastructure meet the requirements that other digital platforms successfully address and how can it create an added benefit compared to a system with centralized governance?

An analytical framework is created that lays out the characteristics and requirements of P2P markets. To that end, this study derives eight dimensions from economic theory which contribute to the market performance and which provide a comprehensive guiding structure to analyze the characteristics of different P2P markets and instantiations of those. The framework is rooted in the four general market tasks defined by Roth (2008): providing thickness, avoiding congestion, providing safety & simplicity, and avoiding re-

pugnance. Seminal work on economic theory of market design (Akerlof, 1970; Roth, 2008; Williamson, 1979) and on existing P2P markets (Ba and Pavlou, 2002; Constantinides et al., 2018; Einav et al., 2016) identifies the characteristics of market instantiations that feed into these tasks. On that basis, the characteristics of centrally mediated digital platforms are compared to blockchain-based decentralized markets. This enables a differentiated evaluation of the characteristics and potential benefits of blockchain-based markets. To examine the implications of this assessment more concretely, one real-world use case of societal relevance is investigated further: a local P2P energy market.

RQ2: Can blockchain technology facilitate the energy transition in local energy markets?

Providing access to resources and energy security is one of the most important challenges for our society (Hentschel et al., 2018; Ketter et al., 2018). Due to environmental and resiliency benefits, distributed energy resources (DER) are gaining importance for energy supply as the costs of solar and wind power systems decrease (Hentschel et al., 2018; Khalilpour and Vassallo, 2015). Given this decentralization of the energy sector, several scholars and entrepreneurs suggest that the energy domain could strongly benefit from blockchain technology in the creation of new, local marketplaces for renewable energy (Basden and Cottrell, 2017; Creyts and Tribovich, 2018; Kastrati and Weissbart, 2016).

5.2 Background on Blockchain Technology

While transfer of information was simplified and scaled up by the internet, the transfer of value over the internet is not inherently secure and often includes interactions with unknown parties. The lack of trust, or transparency, is a key reason why third parties or big, well-known platform operators that are able to provide institutional trust usually organize and control online transactions (Pavlou and Gefen, 2004). The rise of large platform operators is thus not only due to network effects, but also due to them guaranteeing a certain level of trustworthiness and security of transactions (Coase, 1937; Einav et al., 2016). Blockchain technology, on the other hand, promises to enable transparent and secure transfer and enforcement of property rights between peers, without a central point of authority (Catalini and Gans, 2016; Malinova and Park, 2016).

A ‘blockchain’ is a communication protocol that allows for “secure transfer and enforcement of property rights” among participants of the system network (Catalini and Gans, 2016): It is a ledger, i.e. an account book that tracks the balance of all accounts in the system by registering all transactions that take place. More specifically, the ledger contains sets of transactions, the ‘blocks’, that are organized in a directed tree in chronological order. All nodes in the blockchain network keep a local replica of the ledger, so it is actually

a ‘distributed ledger’. For any new block to be added to the ledger and all its distributed replica, all nodes in the system must come to an agreement that the transactions contained in this block are valid, which is ensured by a ‘consensus mechanism’. Consensus mechanisms rely heavily on cryptographic methods to ensure tamper-proofness of the ledger entries, anonymity of the network participants and to prevent attacks by compromised nodes in the system to add invalid transactions to the ledger (Decker and Wattenhofer, 2013).

Different blockchain protocols differ in the degree of centralization, efficiency, and consensus mechanisms (Schweizer et al., 2017). When designing a specific blockchain system, different architecture choices are possible, e.g. choosing a private or public, and permissioned or permissionless system architecture, and choosing between different blockchain protocols, which have different implications for their fit to the requirements of the particular use case (Wüst and Gervais, 2017). Blockchains can be extended from recording mere transactions to execute some code that does not only test whether a transaction is valid based on the current account balances, but that includes other contractual clauses or market rules that apply to the transaction of goods. The code containing such clauses is called a ‘smart contract’ (Szabo, 1996). There are blockchains that provide Turing-complete programming languages for smart contracts, which means they allow the scripting of any arbitrary logic in the contract, which increases their universality of use (Buterin, 2014). Smart contracts allow the implementation of a market mechanism in a way that is transparent to every participant of a network without giving any single party the possibility to change these rules (Mattila et al., 2016). A smart-contract-based market logic thus signifies a digital institution that is controlled by the distributed network and not by a single intermediary (Beck et al., 2018), as illustrated in Figure 5.1. Former intermediaries are either removed or turned into just one of many market participants. Blockchain-based markets also do not require a central infrastructure from a technical perspective, as the blockchain is operated in a distributed ledger on the computing devices on all nodes in the network.

In a broader sense, blockchains can thus be seen as ‘cryptographic economic systems’ (Beck et al., 2016), as an information system enabling economic interactions based on cryptographic methods. Blockchains allow for the creation of new marketplaces and business models for the transfer of value in domains “where there exists a need for a reliable record of transactions in a decentralized environment where not all parties, whether humans or machines, can be fully trusted” (Beck et al., 2017b, p. 99). Some scholars even argue that blockchain represents a general purpose technology (Catalini and Gans,

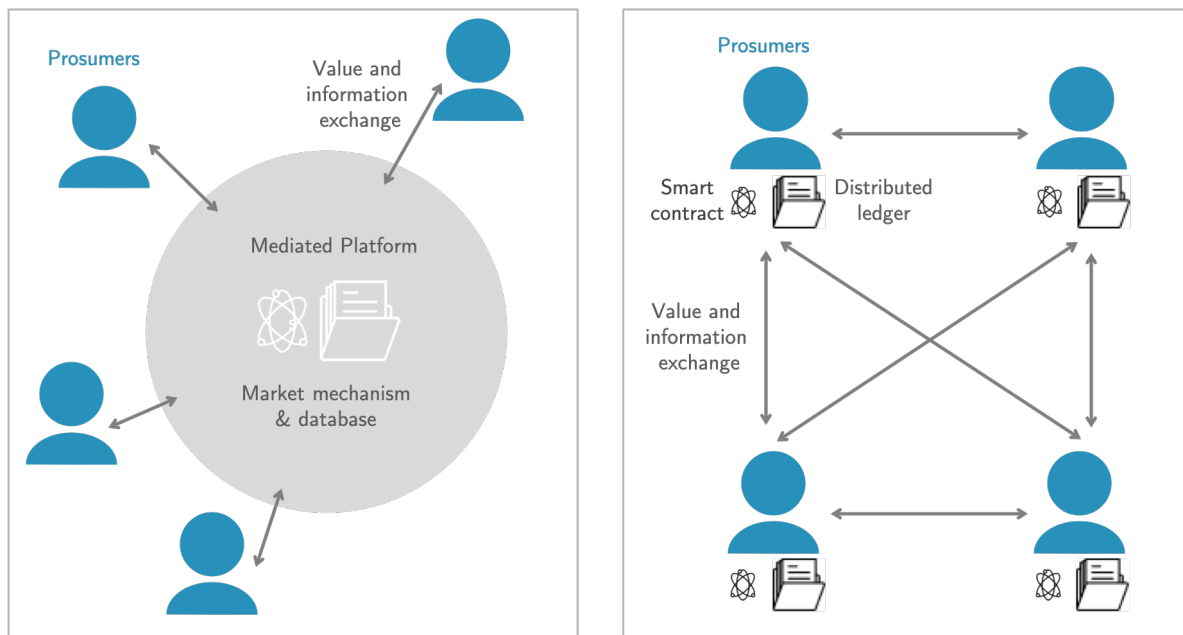


Figure 5.1: P2P market, instantiated as mediated platform (left graph) or as blockchain-based market (right graph). The market mechanism and transaction records are a black box for users of mediated platforms. In a blockchain-based system, the market mechanism is implemented as smart contract that is run by all network participants and all transactions are stored in the distributed ledger.

2016), as it has the potential to disrupt different sectors (Beck et al., 2017b; Wörner et al., 2016) by replacing platform intermediaries with smart contracts and creating decentralized marketplaces for P2P exchange (Figure 5.1). This article aims to analyse such blockchain-based markets through an economic lens to provide insights on the value the technology can bring to such applications.

5.3 Methodology

As the recent publication dates of most of the blockchain-related references indicate, research on blockchain-based applications is only in its infancy (Lindman et al., 2017) and a solid general understanding of emerging marketplaces built on a decentralized infrastructure is missing (Constantinides et al., 2018). At the same time, many of the concepts coming together in decentralized P2P markets have been extensively studied before (Lindman et al., 2017): market design and causes for market failure by economists (Akerlof, 1970; Einav et al., 2016; Roth, 2008), and digital platforms and infrastructure by information systems researchers (Ba and Pavlou, 2002; Bichler et al., 2010; Constantinides et al., 2018; Yoo et al., 2010). Avital et al. (2016) argue that to design viable

blockchain-based applications, it “might require experts like ‘coding economists’. Information systems researchers with economic understanding should develop the systems and theories to support them, serving as boundary spanners and driving forces.” (p. 4).

5.3.1 Analytical Framework for P2P Markets

To address the lack of understanding of blockchain-based markets from an economic lens, an analytical framework is created to characterize P2P markets. Roth (2008) defines the four market tasks as providing thickness, avoiding congestion, providing safety & simplicity, and avoiding repugnance. Based on the extensive literature from the economics, as well as IS discipline, eight dimensions are identified which influence to what extent a market can fulfill these tasks. Following the approach by Lee (2001), the framework allows the comparison of different types of marketplaces (in his case, physical and virtual marketplaces). In contrast to existing work on market design from the economics discipline, this study takes a more comprehensive set of dimensions: The framework also integrates aspects influenced by the digital infrastructure a market is implemented on, such as processing time and system security, which are crucial for the operation of electronic markets (Bichler et al., 2010; Constantinides et al., 2018). Using related literature from the IS discipline, blockchain-based markets were compared to centrally mediated market platforms along this framework.

5.3.2 Use Case Analysis: P2P Energy Market

As suggested by Risius and Spohrer (2017), designers of a blockchain system need to be aware of the features that are relevant to the respective industry branch. Given its societal relevance, the energy sector was chosen for a more detailed analysis as a concrete use case. A systematic literature search provides an overview on the status quo of blockchain-based electricity exchanges. The findings highlight potential success factors and implications for a blockchain-based electricity markets on the dimensions of the analytical framework. Relevant academic literature was identified in a broad search on Google Scholar and Web of Science, as well as a forward and backward search using Google Scholar. The terms ‘blockchain renewable energy’, ‘blockchain microgrid’, and ‘blockchain electricity trading’ were used as search keywords. Given the early stage of the technology, the picture of the status quo of the blockchain applications would be incomplete when focusing on academic publications alone. Due to duration of the reviewing process, academic publications have an inherent time lag; consequently, information on many existing projects so far has only been published in industry white papers. Therefore, the literature search also covered

a list of companies described in two industry guides on blockchain for the energy sector by Montemayor and Boersma (2018) and Hasse et al. (2016), and related posts on the publishing platform Medium.com, which is very popular in the blockchain community.

5.4 Analytical Framework & Evaluation

5.4.1 Analytical Framework

In contrast to many existing market platforms, blockchain-based markets are not run by an intermediary. They remove the central platform from P2P markets, often simply dubbed market platform or marketplace. The present study investigates, which consequences this implies for the economics and the design of P2P markets. The four market design tasks defined by Roth (2008) introduced in Section 2.2.1 serve as basis to further derive dimensions that describe P2P markets. The resulting framework with the eight dimensions derived is illustrated in Figure 5.2 and described in Table 5.1.

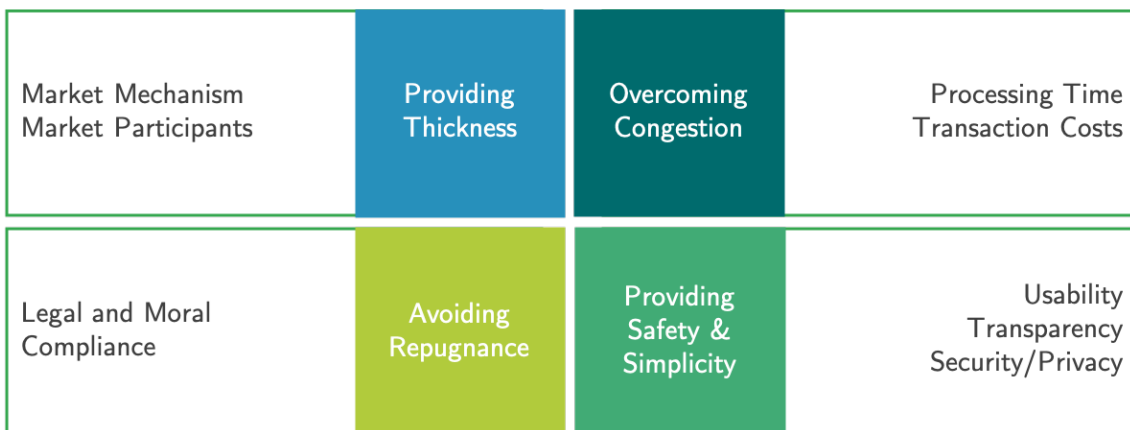


Figure 5.2: Analytical Framework for P2P Markets based on Market Design Theory (Roth, 2008).

Providing thickness: market participants and market mechanism

Providing thickness relates to the presence of sufficient supply and demand (also known as liquidity) on a market. Peer-to-peer markets typically benefit from network effects or positive externalities when there is a large number of users on both sides of the market (Bakos and Katsamakas, 2008; Slavova and Constantinides, 2017; Van Alstyne et al., 2016). Compared to other markets, the **participants** of peer-to-peer markets are usually private individuals, so the marketplace needs to be accessible and attractive to achieve the necessary amount of thickness on the market (Basu et al., 2019). These participants may

switch from merely consuming a good or a service to also producing it depending on their individual requirements (Ramchurn et al., 2011), so-called ‘prosumers’. Moreover, market thickness is also influenced by the efficiency of matching supply and demand. Einav et al. (2016) argue that peer-to-peer exchange is flourishing in applications with variability in demand, low scalability of production, and existence of well-functioning **market mechanisms**. A peer-to-peer exchange should thus establish a market mechanism that aggregates information and determines market outcomes on that basis. The definition of a specific market mechanism is a game-theoretical problem and should meet different economic properties that are important in the application domain, for instance incentive compatibility to reduce collusive behavior and efficiency to reduce welfare losses. Peer-to-peer markets can be cleared using central market mechanisms, or bilateral negotiation protocols (Bichler, 2001). On purely bilateral peer-to-peer markets, trades are arranged directly between individual buyers and sellers, which means that there is no unique point of information aggregation. However, purely bilateral governance is complex (Morstyn et al., 2019) and may be inefficient on peer-to-peer markets due to the high communication necessary for arranging individual trades (Tiwana, 2003). Some scholars thus argue this is not suitable for markets with many participants or recurrent trades (Williamson, 1979). In mediated markets, by contrast, a central intermediary collects demand and supply (Lu et al., 2016), and sometimes also determines the pricing. This intermediary is represented by a professional reseller in a wholesale market or a platform provider for the provision of an infrastructure, on which users can interact, and of algorithms, which implement a market mechanism. Many existing peer-to-peer platforms are currently implemented via auction mechanisms in which buyers and sellers place bids for their goods at distinct times (Rafaeli et al., 2002). On eBay for instance, individuals bid their willingness to pay for the offered goods, whereas on Uber, drivers bid their service at a specific time and place at fixed prices prescribed by the platform (Einav et al., 2016). For the digital infrastructure, this implies that the information system should allow for a variety of market mechanisms to be implemented. Therefore, it needs to provide an accessible programming language and environment to implement the market mechanism. Airbnb, for instance hosts a platform on which individuals can offer housing capacities, and users are free to bilaterally determine prices, but the monetary transaction runs through the platform (Lampinen and Brown, 2017). Airbnb charges a fee on this transaction for the information aggregation, the provision of the interface, search algorithm, and billing service. By contrast, Uber intervenes more strongly in participants’ interaction: It defines the prices for rides offered on their platform using proprietary algorithms (Subramanian, 2017) and also takes care of accounting and billing for the drivers.

Avoiding congestion: processing time and transaction costs

Avoiding congestion means that a market must be capable of **processing** the desired thickness in supply and demand in a reasonable amount of **time**, and at limited effort (Bichler et al., 2010). Many electronic markets today operate close to real time, and users' willingness to interact depend on that speed (Xia et al., 2012). While processing input from an entire network of participants instantaneously is not necessary for all use cases, it may be beneficial or even mandatory for some time-critical application domains in which prices are highly volatile, such as the financial or electricity sector (Rosen and Madlener, 2013). However, 'congestion' cannot only be caused by technical latency, but also by costs for settling transactions between many individuals rather than few institutional entities: The volume of each transaction on a peer-to-peer market will be much lower than the volumes of goods or assets sold in traditional business interactions. **Transaction costs** are thus another key dimension: The costs participants incur for aggregating information on supply and demand, and on their trading partner must be low, so that costs do not outweigh the benefit of the trade. Transaction costs hence include more than the settlement costs for processing a transaction — they include costs that were incurred to identify and build trust in the counterparty and to define, negotiate, and settle the transaction in question. In his seminal paper, Coase (1937) argues that transaction costs are the essential reason for the existence of firms. The search and structuring of information on the supply, quality, and providers of goods on all of the different markets an individual participates in, is too costly for individual agents. For that reason, we rely on firms as intermediaries to aggregate relevant information based on which they can make a buying decision (Coase, 1937). In peer-to-peer markets, information is dispersed over many individuals, so information aggregation is key for defining an efficient market allocation (Einav et al., 2016). The necessity to elicit and process distributed information inherently creates transaction costs in form of search and processing costs.

Providing safety and simplicity: usability, transparency and security

By Roth's definition, safety and simplicity of use are necessary to engage participants in interacting on any kind of market. In particular, for private individuals in electronic market, who are not professionals in trading or may not be IT-savvy, good **usability** and a low (perceived) risk are crucial, otherwise these individuals will not take part in the market (Basu et al., 2019; Lampinen and Brown, 2017). The user interface must be simple to understand and to access, so that the participants can state their preferences in an easy way (Bichler et al., 2010; Lampinen and Brown, 2017).

5.4. Analytical Framework & Evaluation

Market Task (Roth, 2008)	Dimension	Explanation
Providing Thickness	Market Mechanism (Bichler, 2001; Roth, 2008; Tiwana, 2003)	A market mechanism implements contractual clauses which establish rules for market interactions. The main task is to elicit truthful information on the participants' preferences and to compute efficient allocations.
	Market Participants (Basu et al., 2019; Lu et al., 2016; Roth, 2008)	Thickness, or liquidity, to create network effects on a P2P market requires a sufficient number of market participants, which are brought together on one platform. The market thus needs to be accessible for the relevant audience.
Overcoming Congestion	Processing Time (Lu et al., 2016; Rosen and Madlener, 2013; Xia et al., 2012)	Transactions need to be processed fast enough, so that a marketplace can handle the required thickness and allow participants to trade sufficiently fast.
	Transaction Costs (Coase, 1937; Einav et al., 2016)	To keep the user experience convenient and reduce time for search and for defining and settling transactions, transaction costs must be low. Costs of information aggregation and of defining and settling transactions must be low also due to the high number and lower volume of trades occurring between peers.
Providing Safety & Simplicity	Usability (Basu et al., 2019; Lampinen and Brown, 2017; Roth, 2008)	Interaction on a P2P market must be easily understandable and accessible for private individuals, the user interface must be simple and clearly state the relevant information. Some applications may require software agents to support participants in trading based on their preferences.
	Transparency (Akerlof, 1970; Pavlou and Gefen, 2004; Weber, 2016)	Information asymmetry may lead to opportunistic or even fraudulent behavior, which reduces the attractiveness of a market and can lead to market failure. Information on quality of goods, pricing mechanisms, and reliability of participants should thus be transparent and accessible to all participants in the same way.
	Security/Privacy (Clemons et al., 2017; Roth, 2018)	Reliable operation of the system must be guaranteed, and malicious attacks prevented to avoid market failure due to technical issues and to support participants' trust in system integrity. Privacy concerns about personal data need to be addressed.
Avoiding Repugnance	Legal and Moral Compliance (Clemons et al., 2017; Roth, 2008)	Moral and legal compliance of all transactions on a market should be warranted by installed rules, in order not to create adverse reactions or encourage undesired behavior (e.g. corruption, collusion, pollution).

Table 5.1: Analytical framework for P2P markets.

For a market to be economically efficient, it is further important for each party to trust in the others' integrity, e.g. to reliably execute the agreed transactions, to provide the agreed upon quality, or to adhere to other conditions of the trade (Akerlof, 1970; Weber, 2016). **Transparency** reduces information asymmetry and is thus key on a market, on which peers interact without quality control and reinforcement by a mediating institution (Einav et al., 2016). However, the degree to which information asymmetry is an issue varies with the traded goods (Ba et al., 2005), the market mechanism, and whether the interacting peers know each other or not. One way to reduce information asymmetry is the introduction of mediating institutions on markets who aggregate information for all participants to provide some quality control (Akerlof, 1970; Ba et al., 2005; Pavlou and Gefen, 2004). Pavlou and Gefen (2004) argue that due to the anonymity of the participants, institution-based trust in a well-known platform provider is particularly suitable to build peer-to-peer markets. As an alternative option, many electronic markets reduce information asymmetry by implementing feedback mechanisms in which everyone can share their experiences with another party, which creates transparency on other participants' market behavior (Ba and Pavlou, 2002; Wu et al., 2013). These mechanisms, however, still underlie the control of the platform operator and might be manipulated for their needs (Zhang et al., 2018b). Furthermore, safety in the context of electronic markets also relates to data **privacy** and **system security**. When interacting with other entities on a market, individuals reveal information about their preferences concerning pricing, personal schedule, their identity, or their connections (Clemons and Row, 1988; Roth, 2018). Security of the digital infrastructure against malicious attacks or events like electricity outages is critical to guarantee reliable processing of all transactions.

Avoiding repugnance: legal and moral compliance

The fourth and final task for any market design is avoiding repugnance. A marketplace should provide and enforce a clear set of rules to prevent activities or transactions that are morally questionable or even illegal (Roth, 2018). **Compliance** to **moral** and **legal** standards should be given on any market; yet, the large number of participants and interactions within a peer-to-peer market made up of private individuals, and the anonymous nature of many electronic markets may particularly promote opportunistic behavior or collusion (Clemons and Row, 1988).

5.4.2 Evaluation

Table 5.2 compares characteristics of blockchain-based markets to centralized markets along the presented analytical framework based on existing literature on blockchain technology.

Most existing research on blockchain-based markets concern topics around **market mechanisms**, **transaction costs**, and **transparency**: In a blockchain system, the implementation of a market logic that is controlled by a distributed network and coded into software thus leads to a reduction of search and administrative costs and an increase in transparency by design. In a way, smart contracts can act as intermediary collecting information and computing the resulting market outcome. However, contractual clauses that can be implemented in smart contracts are limited to deterministic and objectively testable conditions. It is thus difficult to include qualitative assessments of physical assets or services into smart contract. Using the trading of real-world assets, namely cars, as an example: There are possibilities to securely track data on mileage records (Chanson et al., 2017), which could be used on a blockchain-based marketplace for real-world assets, but it will be more difficult to check for rust spots or the cleanliness of a car’s interior using a smart contract (unless there is a substantial upgrade in sensors collecting information on this) (Notheisen et al., 2017). Hence, a blockchain-based market may not be able to increase transparency with regard to characteristics that are difficult to quantify as verification issues arise for information about the physical world. However, transparency of transactions on the blockchain ledger and distributed control over the market mechanism reduces potential information asymmetry between the participants on a marketplace, and thus creates value. Low information asymmetry or transparency is closely related to the supposed trust created by blockchain technology, which is frequently brought up in the broader media (The Economist, 2015). As pointed out by Milkau in Beck et al. (2017b), Satoshi Nakamoto’s paper (Nakamoto, 2008) describing Bitcoin as the first application of blockchain technology appeared right on the peak of the global financial crisis and benefited from a momentum of distrust in official institutions and decentralization movements. Markets in which the integrity of existing institutions is questionable – due to political reasons or monopolistic market power for instance – are thus another major use case for blockchain-based markets. On the flipside, as Carvalho (2020) puts it and others agree, “if a central, third party is trusted, then there is no need to use a blockchain”, p. 3. If an existing platform provider is trusted and puts an effort into providing transparency about its governance practices and employed market mechanisms, the argument for removing the intermediary loses its strength.

5.4. Analytical Framework & Evaluation

Market Task	Dimension	Blockchain-Based Exchange	Mediated Platform
Providing Thickness	Market Mechanism	Turing-complete programming language for smart contracts allows for any kind of contract or deterministic market mechanism to be implemented, like on a centralized system. The mechanism is transparent to all participants and underlying network consensus which may foster welfare maximization.	Most algorithms can be implemented on a central server, even non-deterministic ones. Market mechanisms are usually inaccessible to the user, and platform providers can execute and optimize the mechanism to its own benefit, and may be willing to incur losses in total welfare.
	Market Participants	Market participants need to become part of the blockchain network (e.g. use interface connecting to blockchain or even run blockchain node), which may be challenging to users that are not tech-savvy at this stage, but can be circumvented with a good user interface.	Everyone can sign up online to market platforms. Platform operators define terms of participation. Sign up procedures are similar on most platforms and familiar to most users, may however require specific payment channels.
Overcoming Congestion	Processing Time	A limited number of transactions per time can be validated, depending on the blockchain protocol and the distributed computational power. Transactions involving a lot of data are thus subject to high latency.	Centralized server computes transactions, often in real time. Yet, settling transactions involving multiple institutions can take very long, as they may involve cross-system communication or offline settlement.
	Transaction Costs	Reduced validation costs for transactions may give rise to smaller players/smaller market environments. Conversely, costs of coordination and (re)negotiation of smart contracts can also be high depending on their complexity and the number of involved parties. Direct transaction fees to run the network and incentivize consensus vary strongly and depend on the blockchain protocol.	Participants remunerate intermediary for operation of the market and thus pay for information aggregation and matching. Providers often charge fees for providing ancillary services (e.g., recommender systems, payment insurances). However, direct costs for central servers are usually lower than operating a distributed network.

Providing Safety & Simplicity	Usability	Simple user interface is needed to allow users to safely interact with the system. This is not part of a blockchain protocol, but needs to connect to it and to address existing user concerns about the technology.	User interface and supporting user interface/trading software can be implemented on top of the infrastructure, provided by the intermediary.
	Transparency	Immutable storage of transactions provides transparency. Market mechanism implemented in smart contract is accessible to all users.	Information on transaction history is not available or provided indirectly in form of feedback mechanisms. Participants cannot access or control market mechanism. Intermediary is focused on establishing institutional trust.
	Security/ Privacy	Distributed blockchain system has no single point of failure and is thus resilient in operation of the market at all times. However, privacy risks of linking public keys to real identities are perceived as critical. Perceived security is further reduced by lack of understanding and trust in complex technology and code.	Market data is stored on central servers to which the identities of the participants may be known, which also raises privacy concerns. In terms of system security, centralized servers are more prone to server downtimes or malicious attacks than a distributed system. Yet, perceived security of established, successful platforms is often high.
Avoiding Repugnance	Legal and Moral Compliance	Legal and moral compliance is governed by the P2P network as a whole and are subject to network consensus. Legal compliance of smart contracts governed by the network is difficult to ensure. In addition, energy consumption and hardware use – i.e. environmental costs for the operation of the system – can be high, depending on the blockchain protocol.	Intermediary has responsibility for legal compliance as legal entity. Stakeholders may pressure intermediary to limit external costs; environmental costs for platform hosting are limited to the operation of one central server – but can be comparably high to a blockchain network.

Table 5.2: Blockchain vs. centrally operated electricity exchange platform

Many articles on blockchain-based markets also argue that blockchain systems enable low transaction costs, which were identified as key enablers of P2P markets (Beck et al., 2017b; Catalini and Gans, 2016; Malinova and Park, 2016). Administrative transaction

costs for search and negotiation, as well as for accounting and billing are reduced if a mediating company is replaced by a market mechanism implemented in smart contracts. Smart contracts can aggregate information algorithmically, determine market allocation and prices automatically, and can include tests whether a transaction has been processed on the blockchain. The smart contract logic can reduce the need for an institution to mediate between the participants within a network (Sikorski et al., 2017; Wüst and Gervais, 2017). This closely relates to Coase’s theory of the firm, which argues that firms are not necessary if a market already provides transparency and aggregates relevant information in a reliable manner (Coase, 1937). Blockchain-based markets may thus enable the market entrance for small-scale producers in various industries like music or energy (Torbensen and Ciriello, 2019), similar to the open source environment (Chong et al., 2019).

However, it is important to note that at this stage, low transaction costs and instantaneous processing are an idealistic aim of blockchain technology, which is not a reality (yet) when it comes to the settlement of transactions on public blockchains (Beck et al., 2018; Wüst and Gervais, 2017). The currently most widely used consensus algorithm, proof of work, causes high external costs in form of energy consumption; furthermore, participants incur high transaction fees and latency for the settlement of transactions on the chain. The external costs in electricity consumption are, however, not an inherent problem to all blockchains and can be overcome by alternative consensus mechanisms (e.g. proof of stake) (Buterin, 2014) or alternative system architectures that leverage state channels or permissioned blockchains. The ecosystem is not developed far enough for verification costs on public blockchains to approach zero at this moment. Yet, extrapolating from the developments in the past years, it is realistic that further improvements will solve these issues in the near future (Yli-Huumo et al., 2016). Moreover, depending on the use case, permissioned blockchains may be a suitable option to reduce transaction costs and processing times while still meeting the use-case specific requirements for some marketplaces (Wüst and Gervais, 2017).

Similarly, the transparency of blockchain-based systems is still controversial. On the one hand, the distributed nature of a blockchain system helps to reduce information asymmetry on the matching logic and transaction history; in the case of mediated platforms, trading rules are determined by intermediaries and are not transparent to individual market participants. On the other hand, this transparency is not necessarily perceived by users, especially if they are not very tech-savvy. It seems like a strong assumption that trust in institutions can be replaced by trust in algorithms, as is underlined by research on ‘algorithm aversion’ from other domains, as well (Dietvorst et al., 2016).

Dimensions that have so far received less attention in the literature on blockchain based-markets are **processing time**, **usability**, and **legal and moral compliance**. Only few studies discuss how users are attracted, or which users are targeted although this is essential to provide liquidity (Basu et al., 2019; Roth, 2008) and to create network effects on any exchange (Parker and Van Alstyne, 2005). While it may seem that this is not an issue related to the technological infrastructure of a marketplace and thus not specific to blockchain technology, missing understanding among nontechnical users and entry barriers may represent an obstacle for applications to reach a minimum network size (Kazan et al., 2018). Miscione et al. (2019) point out that although having the ideal to enable peers to directly interact, this claim is reversed in practice, and organizations are the more prominent users of blockchain. Cai et al. (2019) find that blockchain adoption is highly influenced by peer influence. Furthermore, providing safety and security on a market is not only a technical issue, but also relates to the perceived security and ease of use for the participant. A large scale survey conducted by Abramova and Böhme (2016) shows that there is a strong positive interaction between perceived ease of use and usage behavior of Bitcoin. Other authors agree that it is crucial to improve usability of blockchain systems and to reduce complexity and entry barriers for the user in order to attract participants and make them feel confident in interacting with the system (Janze, 2017; Kazan et al., 2018). It is striking that many studies neglect the role of the actual user in practice, which is somewhat inconsistent with the notion of democratization by decentralization (Buterin, 2014; Kranz et al., 2019), and central to economic decision making. Anecdotal evidence from the blockchain industry emphasizes the lack of focus on the user, as many of the early blockchain-based peer-to-peer markets never reach an effective network size and suffer from distrust in the underlying market mechanisms and code integrity (Sun Yin et al., 2019; Wu, 2019).

The last dimension in the framework, legal and moral compliance to avoid repugnant actions on the market, may pose a further challenge to a blockchain-based system. A platform operator may employ a legal expert to monitor transactions to ensure compliance and can be held accountable by regulators. In the absence of a central intermediary, this legal integrity is hard to ensure. It is difficult to define responsibilities without linking legal entities to participants in a blockchain network (Beck et al., 2018). Sun Yin et al. (2019) argue that pseudonymity in the Bitcoin ecosystem, for instance, leads to intransparency and may foster criminal activities, which reduces trustworthiness and adoption of the system. This can create difficulties in many use cases, for instance in the financial sector, or illegal sharing of data or files, and will require further research and attention of policy makers.

5.5 Use Case Analysis: P2P Energy Exchange

Of the many opportunities for blockchain technology to create new marketplaces and to enable the secure transfer of value (Constantinides et al., 2018), the energy domain is among the most relevant for society. Energy provision is a highly complex challenge that is of political, economic, and societal importance (Ketter et al., 2018). Substantially reduced cost and improved technology in the smart grid are turning DER (Akorede et al., 2010), such as wind and solar energy, into the key levers to change the electricity market from a vertical structure to a decentralized, bottom-up landscape (Green and Newman, 2017; Hentschel et al., 2018; Khalilpour and Vassallo, 2015). To account for that, energy markets in several countries are being liberalized, allowing smaller players to enter the market (Ketter et al., 2018; Rosen and Madlener, 2013). Consequently, new concepts for energy markets are being developed both by academic scholars and industry researchers: virtual power plants, cooperatives or local P2P markets for DER (Hentschel et al., 2018; Morstyn et al., 2018) – possibly enabled by blockchain technology (Basden and Cottrell, 2017; Hentschel et al., 2018; Meeuw et al., 2018; Mengelkamp et al., 2017a). P2P exchange of electricity signifies a shift to a decentralized bottom-up market (illustrated in Figure 5.3, right side), in which individual consumers and prosumers can directly trade electricity without the mediation of a central utility provider acting as reseller. Prosumers can sell excess electricity to other consumers within local communities on the low-voltage distribution grid level (Morstyn et al., 2018). This puts small generators in the focus and creates a competitive environment for distributed generation (Basden and Cottrell, 2017).

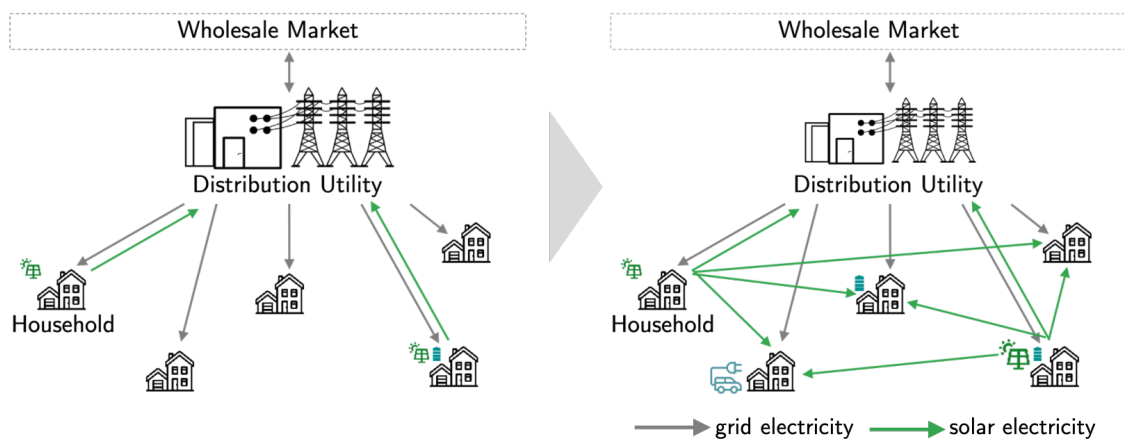


Figure 5.3: Traditional hierarchical electricity market (left graph) versus decentralized electricity market enabling P2P exchange of solar energy (right graph). Local electricity sourcing from prosumers to consumers enables individual households to take a more active role in electricity sourcing and pricing.

To get an understanding of the status quo of the research on blockchain-based energy markets, a systematic literature review on the relevant academic articles, as well as the industry projects, was conducted. Tables A.2 and A.1 in the appendix section provide an overview of identified existing studies on P2P energy markets using blockchain technology. All in all, academic and industry research on decentralized energy markets enabled by blockchain technology is still in its infancy. While start-ups in particular were quick to publish whitepapers, the majority of those articles provide only vague ideas. The academic publications identified give a more refined view of the economic implications of the created P2P markets; nevertheless, most of the research is still conceptual or presents a small proof-of-concept. Most articles either focus on the practical reasoning for a decentralized electricity market, e.g. (Mengelkamp et al., 2017a), or provide a technical description of a planned blockchain system (Kang et al., 2017). Yet, the connection of the two, i.e. a link between technical feasibility and implied practical value to the electricity market, is missing. Most articles merely describe a proof of concept that focuses on the technical feasibility of electricity trading, ignoring economic considerations or user-related aspects of creating a novel energy market. However, the analytical framework and the evaluation of blockchain-based P2P markets provided in Table 5.2 can put some of the arguments made in the literature in context. Table 5.3 provides a concise overview of the literature analysis for this use case, which have direct implications for the design of blockchain-based energy markets in practice. The main implications are discussed below.

From a business perspective, small, local energy markets may not be very attractive to operate for a platform company as they do not (and are not supposed to) scale up (Slavova and Constantinides, 2017). In contrast, if a P2P market is operated in a decentralized manner by the participants themselves, central operation is no longer necessary. This may be one factor why ambitions to sell electricity from prosumers to consumers are only gaining momentum now with the development of blockchain technology. Additionally, the transaction volume for electricity consumed by a household in short time intervals is very low (typically amounting to less than 0.05 USD in 2016 in the US in 15 minutes (US Energy Information Administration, 2016)). For such low-volume transactions, minimal costs for transaction matching, settlement, and billing are vital (Rosen and Madlener, 2013). This again reduces the incentive for competitive companies to host P2P electricity markets – and makes a case for a decentralized, blockchain-based system owned by the market participants. On the other hand, it excludes blockchain protocols with high transaction fees or low throughput and suggests the use of a permissioned blockchain (Wüst and Gervais, 2017). In fact, most academic studies on electricity exchange focus on permissioned blockchains (Table A.2). By contrast, most of the industry projects plan

to employ public blockchains in combination with payment channels to reduce settlement costs and external costs.

Due to its infrastructural importance and environmental consequences of its sourcing, a market mechanism for the allocation of electricity should adhere to social welfare, allocative efficiency, and incentive compatibility as much as possible (Rivola et al., 2018). As in most other market applications for blockchain systems, a smart contract protocol providing the tools to implement contractual clauses subject to consensus of the network is thus necessary. The review of whitepapers in the energy sector (see Table A.1) indicates that industry research seems to almost entirely neglect the mechanism design or implied incentives.

In the energy sector, security and privacy have a high priority. Reliable operation of a marketplace must thus be guaranteed, and malicious attacks prevented (Aitzhan and Svetinovic, 2016; Ketter et al., 2013; Miller et al., 2017). The resilience of the distributed transaction ledger in a blockchain creates a clear benefit over a centralized market, as the market can keep running even if some nodes in the network are corrupted. On the other hand, high-resolution electricity consumption data is sensitive data which provides information on energy usage patterns and could even identify households within a local community (Laszka et al., 2017), as well as details about their way of living (Hopf et al., 2017). Production and consumption data will be traceable on a blockchain to some extent, so transactions between nodes must be settled in an privacy-preserving manner (Aitzhan and Svetinovic, 2016; Ketter et al., 2018), which still requires further research. In addition, usability should not be underestimated in a market in which private individuals are not used to take active decisions, but traditionally always have been price takers in a retail market (Slavova and Constantinides, 2017). As the user interface of many blockchain-based applications is still very complex to grasp for the ordinary citizen, and since there are still hardly any applications out there for monitoring household energy sourcing, this is an area with a lot of open questions for future research.

The final dimension is the risk of immoral or illegal activities on the market, which could prevent individuals from wanting to trade on an electricity market. If regulation generally allows for P2P energy trading, the opportunities for illegal or immoral behavior may seem limited at first, but one aspect may create adverse reactions: External costs in the form of environmental costs generated by the blockchain system must not exceed the savings created on the local energy market (Ketter et al., 2013). The high energy consumption of blockchains that rely on proof-of-work consensus (Beck et al., 2017a), for instance, are thus not suitable for this use case from an environmental perspective.

5.5. Use Case Analysis: P2P Energy Exchange

Market Task	Dimension	Implementation in Blockchain System
Providing Thickness	Market Mechanism	<ul style="list-style-type: none"> • Blockchain protocol with smart contract language, e.g. Ethereum • Auction mechanism implemented in smart contract, allowing for preference elicitation on energy sourcing and incentivizing local electricity consumption (Mengelkamp et al., 2017a)
	Market Participants	<ul style="list-style-type: none"> • Consumers, Producers, utility provider • Permissioned blockchain to limit participation to local area
Overcoming Congestion	Processing Time	<ul style="list-style-type: none"> • Online mechanism with iterative market clearing intervals of short time intervals 24/7 • Permissioned blockchain to limit number of participants
	Transaction Costs	<ul style="list-style-type: none"> • Network consensus to ensure transaction integrity • Automated matching and auction clearing on blockchain system • Automated billing in longer intervals
Providing Safety & Simplicity	Usability	<ul style="list-style-type: none"> • Simple and understandable user interface, abstracted from blockchain technology • Smart bidding agent to quote bids in real time for the user might be necessary
	Transparency	<ul style="list-style-type: none"> • Standardized, calibrated smart meters • Distributed ledger of transactions • Market mechanism in smart contract, changes underlying community consensus
	Security/Privacy	<ul style="list-style-type: none"> • Consensus-based market mechanism • Signed messaging and data encryption • Private blockchain
Avoiding Repugnance	Legal and Moral Compliance	<ul style="list-style-type: none"> • Consensus protocol with low external costs (i.e., no proof of work)

Table 5.3: Implementation of market tasks in blockchain-based energy market

5.6 Conclusion and Future Work

The power of P2P markets that connect individual users for sharing and trading has been demonstrated by the steep success of digital market platforms like eBay, Uber, Airbnb or Etsy. High hopes are put in blockchain technology to facilitate decentralization of such markets to put the individual participant in a more active role and to reduce the power of a single intermediary (Beck et al., 2016; Subramanian, 2017). Despite the ongoing wave of enthusiasm for the technology itself, there is still little understanding of the added value of running a P2P market on a blockchain infrastructure. This article provides novel insights on these socio-technical systems, in particular, with the following three contributions.

First, the study presents an analytical framework for characterizing P2P markets. Based on Roth's theory of market design and related literature from the economics and information systems disciplines, eight dimensions are used to characterize a P2P market (Table 5.1). The framework can be used to analyze different P2P markets and to derive managerial implications for the implementation and market design of a viable marketplace. As a second contribution, the framework highlights dimensions in which a blockchain system can add value to a marketplace as compared to when it is run by a central platform (Table 5.2). One key benefit of blockchain technology is the automated execution of contractual clauses and of billing procedures. This is particularly important in markets with frequent, low volume trades and in which there is a lack of trust in a central intermediary. The distributed nature of a blockchain-based market helps to reduce information asymmetry on the matching logic and pricing; in the case of mediated platforms, these rules are determined by intermediaries and are not transparent to the individual participants. At the same time, the characterization shows that with regard to markets in which quality or physical aspects are hard to measure digitally (e.g. markets for used cars, housing), it is difficult to establish transparency in a blockchain-based system or to include such aspects in smart contracts. The evaluation further identifies processing time, privacy concerns and legal compliance as potential challenges for a decentralized system to successfully fulfill the key tasks of a market. Overall, the evaluation provides focal points for both academic scholars and practitioners for assessing whether a specific use case can benefit from blockchain technology. Third, an exemplary use case illustrates this kind of analysis, namely for P2P energy market, in Section 5.5. The electricity sector is currently undergoing a paradigm shift that involves the integration of an increasing share of volatile renewable energies and decentralization of energy sourcing (Basden and Cottrell, 2017; Hasse et al., 2016; Kastrati and Weissbart, 2016; Morstyn et al., 2018).

In the context of this ongoing systemic decentralization, the findings of this study show that many aspects of blockchain-based marketplaces suit the characteristics of the electricity domain well. Blockchain technology can add value when it comes to resilience of the system and increased transparency of energy accounting, when there is distrust in platforms handling sensitive energy data. However, these benefits can only be realized if the system implements an efficient market mechanism and actively integrates participants in the P2P market to realize the vision of a ‘democratized’ decision process. Moreover, the metering of the physical energy flows provides a technological challenge and may raise privacy concerns in a local community.

Given the current hype around blockchain technology, the technical focus (and technology enthusiasm) of many developers outweighs the expected value propositions in practice (Beck et al., 2017b). There is vast potential for further scientific research regarding the examination of digital P2P markets and the role of intermediaries. Future work in the information systems domain should investigate user adoption, how well markets governed by smart contracts work in practice (Beck et al., 2018), and by whom these smart contracts will be implemented in the real world. In addition, it might be fruitful to study potential business models for the implementation of blockchain-based markets. Despite the apparent contradiction between the aim to create a profitable business and decentralization, digital market platforms have brought businesses to life that had been unthinkable only a few years earlier (Lee, 2001). Finally, scholars and policy makers will need to examine which regulatory frameworks have to be in place to regulate blockchain-based markets without central providers who can be held accountable for legal compliance of the market activities.

6. Article E) Evaluation of a P2P Energy Market in the Real World

6.1 Introduction

Renewable energy generation plays an increasingly important role in meeting future electricity demand and in reducing greenhouse gas emissions (Gholami et al., 2016; Ramchurn et al., 2012). Yet, the integration of distributed energy resources (DER) creates a challenge for the existing market structures (Koolen et al., 2017). Today's established power markets are strongly centralized and hierarchical with electricity distribution from a few power plants down to thousands of households. Wind and solar energy generation, in contrast, is geographically distributed, strongly volatile, and cannot be simply switched on or off according to the demand (Andoni et al., 2018; Ramchurn et al., 2012). Moreover, the novel, more active role of consumers who own solar panels and produce energy by themselves ('prosumers') creates challenges at different fronts, in particular for industry incumbents and traditional electricity markets.

Information technology can play a key role in this transformation of electricity markets, as it provides tools to control distributed networks and enable bidirectional communication with the user (Ramchurn et al., 2012; Seidel et al., 2017). Electronic markets and digital platforms have revolutionized a variety of industries by enabling a shift from traditional pipeline markets to P2P platforms, which now shape these industries (Van Alstyne et al., 2016). Information systems can provide personalized information (Tiefenbeck, 2017) and

create electronic marketplaces that can handle stochastic supply and demand in real time (Bichler et al., 2010; Gholami et al., 2016). Green IS research is thus in the ideal position to study innovative platforms which seize the possibilities of technological advances to market DER and foster sustainability among the broader public (Ketter et al., 2018; Seidel et al., 2017; Slavova and Constantinides, 2017).

Recently, advances in distributed ledger technologies and the simultaneous decentralization of energy supply and have led to ambitions to create decentralized energy markets in which prosumers can directly sell excess renewable energy from peer to peer (Burger et al., 2016; Mengelkamp et al., 2017a; Morstyn et al., 2018). Using a digital platform, electricity from solar panels could be traded locally among neighbors, without a central utility provider serving as intermediary for these transactions. P2P energy markets have the potential to generate value on multiple levels: They allow for local matching of supply and demand for renewable energy, enable consumers to actively influence energy sourcing, and provide incentives for investments in renewable generation (Morstyn et al., 2018). Overall, this may reduce depletion of natural resources and greenhouse gas emissions in the long run, thus fostering sustainability (Andoni et al., 2018; Morstyn et al., 2018). However, the performance of P2P energy markets has not been studied in practice yet. While there are several conceptual articles on decentralized energy markets (Andoni et al., 2018; Mengelkamp et al., 2017a; Morstyn et al., 2018), empirical evidence for the feasibility and impact in the real world is still missing. This is not only due to the early stage of the technology or regulatory challenges. More importantly, like in other domains, energy trading on P2P markets implies a fundamental shift regarding the role of the participating citizens (Fridgen et al., 2018; Morstyn et al., 2018). This raises the question: *Which value proposition do P2P markets create from the user perspective, and to what extent are they an effective measure to empower once passive consumers to assume a more active role in these markets?*

This article presents a framed field study to explore the impact of P2P energy markets in the real world. More precisely, the research team implemented a platform for trading solar energy among peers in a local community in Switzerland with 37 participating households. After an intensive period of selecting, developing, testing and deploying the information system, an active market phase with collection of trading data started in January 2019. Based on market design theory, a double auction mechanism allocates solar energy and determines prices on this market. Each household explicitly states their willingness to pay for local solar energy and prosumers additionally define the conditions under which they are willing to sell energy produced by their solar panels. By analyzing the three-month

data set collected between January and March, this study examines energy matching, preference satisfaction and resulting benefits in a real-world instance of a P2P electricity market.

To the best of the study authors' knowledge, this article presents the first empirical evidence of this extent on a P2P electricity market in the field. This study contributes early empirical research on a novel approach to tackle the energy transition using electronic markets, thus addressing one of the most pressing societal problems (United Nations, 2019). By designing and implementing P2P trading in a local electricity market and by evaluating its impact from the user perspective, this article goes beyond the stage of merely conceptual or analytical research that characterizes the majority of research projects in the Green IS area (Gholami et al., 2016; Malhotra et al., 2013). The findings contribute to Green IS research and research on the sharing economy by testing an innovative solution concept to design electronic markets (Bichler et al., 2010) for fostering sustainability in the field (Seidel et al., 2017). In particular, real price preferences for local solar energy are collected in a real-world setting from the users' input to the market mechanism. These findings can serve as input to design local energy markets on a larger scale and possibly, to create personalized trading agents for electricity trading. The data thus provides meaningful information for policy makers to address the challenges of incorporating DER and to design future electricity markets. Furthermore, the results provide empirical evidence for the user value proposition created on an electronic P2P market enabled by a distributed information system.

6.2 Background & Related Work

6.2.1 Designing Smart Energy Markets

Different electronic peer-to-peer or 'platform' markets, which have emerged in recent years, have been subject to research both by economists and information systems scholars (Bichler et al., 2010; Slavova and Constantinides, 2017). While market design is rooted in economic theory, most new, emerging markets are enabled by computational tools and smart devices which in turn strongly influence the efficiency of and human interaction with these 'smart markets' (Bichler et al., 2010; Zimmermann et al., 2018). The peer-to-peer exchange of goods and services on Airbnb, Uber and other sharing platforms represent economic systems, but these systems are strongly influenced by the information systems that support them (Glaser, 2017; Lampinen and Brown, 2017). This interdisciplinary nature makes smart markets a subject of interest for information systems research (Bichler

et al., 2010; Melville, 2010). Smart market design is concerned with the question how information systems can be leveraged to design well-functioning markets and how the user can be supported in the decision-making process without being overburdened with information (Bichler et al., 2010). Bichler et al. (2010) argue that the first step in designing smart markets is preference elicitation to understand and to model user behavior and preferences. The user perspective is necessary to make the right design choices on market mechanisms, input format and information provision. Furthermore, based on the user preferences elicited, real-time decision support systems can be created that adapt to the individual user and dynamic market conditions and that provide personalized recommendations. Herein, market designers should strive to align participants' incentives with the social goals respecting the specific characteristics and requirements of the domain (Ketter et al., 2013). Energy markets represent some of the most information-intensive instantiations of markets due to the volatility in supply and demand and its dependency on environmental conditions (Koolen et al., 2017). Given that providing sustainable energy supply is one of the most critical societal tasks (United Nations 2015, 2019), several calls in recent years have encouraged research on Green IS and smart markets for sustainable energy (Gholami et al., 2016; Melville, 2010; Seidel et al., 2017; Watson et al., 2010). Yet, the task of creating smarter energy markets is a wicked one, as the development of solution concepts viable in the real world involves engineering problems as well as active integration of the user (Seidel et al., 2017). Due to the complexity of impact-oriented Green IS research that examines the actual "in-field" impact of such systems" (Malhotra et al., 2013), p. 1270, is very scarce (Gholami et al., 2016). Based on the assumption that energy is considered a homogeneous commodity, user preferences have been largely ignored in this sector for a long time.

6.2.2 P2P Energy Markets

In recent years, the energy market is undergoing substantial changes, not only on the physical, but also on the digital layer (Ketter et al., 2018). The deployment of smart meters is enabling monitoring of the consumption of individual market participants in real time (Gholami et al., 2016). Information systems can further support algorithmic control within energy networks and bidirectional communication between consumers and prosumers, making it possible to implement a market mechanism that matches supply and demand based on real-time data and to provide decision support systems that individual consumers can interact with (Bichler et al., 2010; Ramchurn et al., 2012; Watson et al., 2010). Consequently, new markets for DER are being developed by both academic scholars and industry research. One approach to better mirror the decentralization of energy

supply in the energy market is to form local microgrids, i.e. electricity distribution systems which attempt to balance supply and demand on a local level (Brandt et al., 2014; Slavova and Constantinides, 2017). Due to an increasing share of generation assets that is operated by private consumers and with the aim of creating a more consumer-centered market, the concept of P2P trading of local electricity in such microgrids has attracted interest among practitioners and scholars alike (Basden and Cottrell, 2017; Mengelkamp et al., 2017a; Wörner et al., 2019b). P2P exchange of electricity signifies a paradigm shift to a decentralized bottom-up market in which individual consumers and prosumers can directly trade electricity on demand without the mediation of a central utility provider acting as reseller. Morstyn et al. (2018) argue that, from the user perspective, the value proposition of P2P trading of renewable energy is threefold, p. 95: energy matching, preference satisfaction, and uncertainty reduction.

- **Energy matching:** The efficient coordination of supply and demand of energy requires a market mechanism that incorporates the specific characteristics of electricity as well as prosumer preferences. Ideally, the market mechanism incentivizes local production and storage capacities according to local demand in real time, thus reducing transactions with the central utility provider and required generation from centralized power plants (Ketter et al., 2013; Morstyn et al., 2018). Existing literature on P2P energy markets (Mengelkamp et al., 2017a; Morstyn et al., 2018) mostly proposes some type of online, double auction as market mechanism, as market-based prices reflect supply and demand on a market in real time while allowing to engage the participant in the decision-making process at the same time.
- **Preference satisfaction:** Allowing consumers to state preferences on energy sourcing and different resources to be traded according to these preferences. Several studies (Capstick et al., 2015; Lee et al., 2015) as well as media reports (Aljazeera, 2019; The Economist, 2019) suggest that in many countries public awareness for climate change and energy-related sustainability issues is rising. More and more individuals do not perceive electricity as a homogeneous commodity anymore and increasingly display preferences for local energy supply (Silva et al., 2012; Tabi et al., 2014). Hence, the integration of renewable energy drives more user-centric approaches (Andoni et al., 2018; Koolen et al., 2017). This trend is also reflected in recent statements of the European Consumer Organisation (2016) and the European Commission (2015), p. 1, who highlighted the need “to empower consumers through providing them with information, choice and through creating flexibility to manage demand as well as supply”. In a choice experiment with consumers in Germany,

Tabi et al. (2014) find that a large majority of consumers displays a preference for renewable energy supply and one quarter of them deem the location of electricity generation an important attribute. Likewise, results from an online survey by Ecker et al. (2018) suggest that consumers are willing to incur a price premium of 20% on average for renewable energy produced in their own homes. Yet, all these findings are based on self-reported survey data and an investigation of individual preferences and social behavior in a real market setting to develop efficient markets is still missing (Andoni et al., 2018).

- **Uncertainty reduction:** As prices for residential photovoltaic systems have been falling over the past years, the number of small generators has been increasing. This has granted more and more consumers a new and more active role as prosumers who “both produce and consume electricity depending on their local requirements” (Ramchurn et al., 2012), p. 88. Yet, it is unclear for prosumers whether and how they can market energy produced from their generators in the long run: Recently, investments in renewable generation are highly uncertain, as subsidized feed-in-tariffs are declining or even abolished in many countries (Morstyn et al., 2018). Ideally, trading energy within a local P2P market increases revenues for prosumers and creates incentives for investments in renewable energy, hence reducing uncertainty of investments in DER. In turn, this may lead to investment spillovers (Bakos and Katsamakas, 2008) increasing the adoption of renewable generators or smart load scheduling solutions as has been observed in P2P markets in other domains (Van Alstyne et al., 2016).

The Brooklyn Microgrid was the first running electricity exchange deployed in the field, in which locally produced energy from solar systems was sold within a neighborhood, and participants were directly involved in the trading (Mengelkamp et al., 2017a). The Brooklyn Microgrid currently applies a uniform double auction (Lacity, 2018; Mengelkamp et al., 2017a), but so far, there is little information on the reasoning for the chosen market mechanism and there is no empirical data on the observed market outcomes available (yet). There are some simulation studies that examine individual aspects of peer-to-peer energy markets: In a simulation study based on load profiles from 4,190 households in Ireland, Griego et al. (2019) compare different compositions of load profiles for P2P microgrids. They find that communities of at least 10 households and a share of 40-60% perform best in terms of self-sufficiency rate (SSR), i.e. the share of energy consumption the local market can cover with local electricity production, and self-consumption rate (SCR), i.e. the share of locally produced electricity that can be consumed locally. Mengelkamp et al.

(2018) conduct a simulation with load profiles from Germany, in which they implement a time-discrete double auction and use artificial bidding data. They find an overall SSR of max. 42%, but point out that the prices and allocation they find need to be validated with real field data on price preferences and bidding behavior. Block et al. (2008) design a combinatorial auction mechanism for electricity and heat trading in a microgrid, but provide no quantitative evaluation or field testing. Several industry publications and whitepapers (Hasse et al., 2016; LO3 Energy, 2017; Miller et al., 2017) present case studies with a stronger focus on the individual consumers. Yet, conceptual market designs presented in these publications are not empirically validated either. Moreover, these publications do not consider a market environment that allows users to actively take part in the pricing and allocation mechanism in the microgrid. To this date, there has not been any real-world data reported on trading within a P2P energy market that provides empirical evidence to what extent the value propositions conjectured by Morstyn et al. (2018) and supported by other proponents of decentralized energy markets translate to the real world.

6.2.3 Energy auctions

The performance of a market depends on the interaction of the individual market participants with the market mechanism, on the input language and on the settlement process of triggered transactions (Ketter et al., 2013). Energy markets are complex multi-agent systems with diverse market participants exhibiting different individual preferences and trading strategies (Bichler et al., 2010; Ketter et al., 2013). Both supply and demand are volatile; to avoid blackouts, it is critical that supply and demand match at all times. Furthermore, electricity markets are vulnerable to strategic behavior, as participants have abundant opportunities for (implicit) collusion (Klemperer, 2002). Consequently, the design of P2P energy markets – and electricity markets in general – needs to mitigate these risks and constraints adequately.

Most existing electricity markets and concepts for P2P energy markets employ an auction mechanism (Dauer et al., 2015; Koolen et al., 2017; Mengelkamp et al., 2017a). This means that the participants express their preferences as bids containing a price and a quantity of electricity, which they want to purchase or sell. To balance supply and demand, all bids are collected in an order book and matched according to specific rules, similar to the operation of stock markets (Andoni et al., 2018; Fridgen et al., 2016). The rules of an auction have strategic implications on how the market participants formulate their bids to maximize their expected utility. Different types of auctions exhibit different

properties, such as: Pareto efficiency, which means that individuals with higher willingness to pay should be prioritized higher in the allocation of goods; incentive compatibility, which demands that agents' never have an incentive to misrepresent their true preferences in their bids; or the expected prices on the market (Mas-Colell et al., 1995). Regarding pricing rules, a main differentiation is whether prices are uniform (i.e. all bidders pay the same market clearing price) or discriminatory (i.e. bidders prices differ depending on their bids) (Fabra et al., 2002). Several studies show that discriminatory price auctions for electricity foster a more competitive environment; uniform price auctions, on the other hand, are more prone to collusive behavior on one side of the market (Fabra et al., 2002; Klemperer, 2002). While a uniform price regime yields slightly lower average prices according to several studies based on simulations and lab experiments, discriminatory-price auctions can reduce volatility of prices (Rassenti et al., 2003). As market failure on a local electricity market may reduce efforts to create new, innovative market structures for the energy sector, it is crucial to study the design of electricity markets. The Californian electricity market in 2000 and 2001 serves as a negative real-world example that illustrates the potential real-world implications of a poor market design. During that period, the Californian market experienced tremendous volatility in prices and even some blackouts caused by poor auction design and strategic behavior of several market participants (including utility providers and generation plant operators), among other factors (Borenstein et al., 2002).

6.3 Method

6.3.1 Study Site and Setup

In a real-world market platform deployed for this study, prosumers can sell the (surplus) energy produced by their solar panels directly to consumers within their neighborhood using an auction mechanism. The study sample comprises 37 participants in a local municipality in Switzerland, 31 of which are prosumers who own (part of) a PV panel. Most of the participants are private households, with the exception of one flower shop and one nursery home for elderly people with approximately 50 residents. All participants are customers of the local utility provider (blinded for review), whose support and active role has been vital to the launch and success of the project. Together with the academic researchers, they selected and recruited the participants from a neighborhood with a high penetration of residential PV panels and served as a trusted local point of contact.

The research team deployed smart meters, which measure electricity loads in time intervals of 15 minutes, in every participating household. Each household received one

device that measures electricity consumption. Prosumers received another smart meter for measuring electricity production from their PV panels and participants who own a battery storage system received a third smart meter for measuring battery loads. Altogether, the research team installed 75 smart meters in the participating households; all devices are connected to the internet.

6.3.2 Design and Implementation of the Market Mechanism

Based on experimental and simulation studies on auction mechanisms for electricity trading (Klemperer, 2002; Nicolaisen et al., 2001; Rassenti et al., 2003; Rosen and Madlener, 2013), a market mechanism that takes into account the specific setting of P2P exchange between private households was implemented: a time-discrete, discriminative double auction. A double auction was identified as the most suitable archetype of an auction mechanism for the present setting to enable prosumers as well as consumers to decide for which conditions they are willing to sell or buy sustainable electricity (Rosen and Madlener, 2013). This double auction takes the prices defined by the participants, as well as their consumption and production loads measured by the smart meters as input. Due to the propensity of electricity markets for collusive behavior described above, discriminatory pricing seemed more advantageous than uniform pricing (Klemperer, 2002). In addition, the limited volatility observed for discriminatory pricing (Rassenti et al., 2003) is critical to provide affordable and calculable costs for energy supply and not to alienate households taking part in energy trading for the first time (Rosen and Madlener, 2013). Although prices are expected to be slightly higher in a discriminatory auction than with a uniform pricing scheme (Fabra et al., 2002; Rassenti et al., 2003), this is not necessarily a disadvantage in the local electricity exchange, as it benefits the prosumers generating renewable energy and may thus foster the profitability and diffusion of DER. Moreover, related literature has shown that many consumers state that they are willing to incur higher costs for renewable or local energy (Ecker et al., 2018; Tabi et al., 2014).

The participants' buy and sell orders for local electricity are collected over a 'clearing period' of 15 minutes. After the orders are collected, the auction mechanism is run to clear the market and determine the resulting electricity trades. The discriminative double auction matches the highest buy order with the lowest sell order (in terms of price) and progresses like this through the entire order book. The price for each matched trade is the mean between the sell and the buy price of the respective orders ('discriminative/midpoint pricing'). A sample orderbook is depicted in Figure 6.1: The blue and green curve show buy and sell orders, respectively. The dashed grey line indicates the realized prices resulting from the auction mechanism. As electricity supply and demand need to be balanced

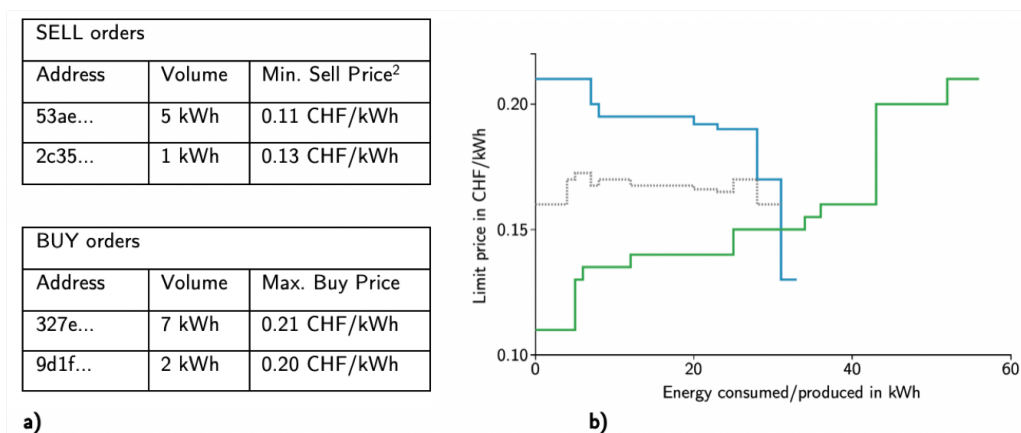


Figure 6.1: a) Extracts from a sample orderbook, containing energy loads and prices b) chart of a sample order book of one time slot during the day.

at all times, the local utility provider serves as backup for the microgrid: When there is not enough solar energy in the local market, or when there is more production than demand within the local market, the local utility provider covers (and absorbs, respectively) these undersupplied (and excess, respectively) capacities at its standard tariffs. In the study location, the standard electricity tariff incurred by residential consumers is 20.75×10^{-2} CHF/kWh and the feed-in-tariff for local production that is fed back into the grid is 9.79×10^{-2} CHF/kWh (including network charges that have to be reimbursed to the utility company).

6.3.3 Design and Implementation of the P2P Trading Platform

The software enabling the P2P trading and communication within the microgrid is running in a decentralized manner on smart meters. Each of the participating households was equipped with a smart meter (a Raspberry Pi with expansion modules to measure voltage and current) which ran a permissioned blockchain system based on the Tendermint consensus protocol (Kwon 2014) (an overview of the system is provided in Figure 6.2). The auction mechanism was implemented on the application layer of this blockchain and is running as a smart contract without using a central server. All bids were handled in a pseudo-anonymous way, as each smart meter received an address which only the research team knew. Hence, participants did not know with which of their neighbors they were trading electricity with, or who asked for which price. More details on the technical details of the system architecture are available in the technical report (Ableitner et al., 2019).

In order to encourage an active participation of the households and to elicit their price preferences (Bichler et al., 2010) regarding local solar energy, the participants received

access to a personalized web application for the P2P trading. The application allows them to monitor real-time data on their energy consumption (and production, if applicable), on their past trading behavior and, in particular, to place price bids: By moving a slider element, they can state their willingness to pay for solar electricity produced by their neighbors (Ableitner et al., 2019). Prosumers can further define their minimum ask price for selling energy from their solar panels to their neighbors. The participants are free to define their price bids just once or adjust them as often as they wish. The application provided them with a concise overview of their energy data and their trading outcomes on the local market in real-time at their discretion, as earlier research indicates that participants may be interested in the local origin of the energy they buy (Ecker et al., 2018; Meeuw et al., 2018). (Note that the details on the development of the user interface and analyses related to user experience and system usage are beyond the scope of this manuscript.)

While in theory, the participants have the possibility to adjust their maximum buy and minimum sell price as often as they want, the research team did, obviously, not expect them to continuously monitor the auction execution or to take action on a daily or (sub)hourly basis. Once they have set their price bids in the web application, orders are posted by the smart meters every 15 minutes. The auction is executed every 15 minutes to clear the market so that prices reflect availability of solar energy in near real time (Rosen and Madlener, 2013). Participants received a monthly report summarizing the information available on the web application. It included their energy consumption and production, resulting expenses, share of local energy supply and the average price they incurred for local energy. The report was sent out at the end of each month via email.

A key feature that sets this field study apart from prior research is that the participants are in fact charged according to the prices defined by the participants on the P2P market and that the electricity trades computed by the described auction mechanism occur in reality. Consequently, the price preferences elicited from the participants are not merely responses to a hypothetical scenario in a survey, but they influence the actual electricity costs participants incur. Participants have been made aware of this fact in an information event prior to the study (attended by 29 out of 37 households) and all participants have signed a letter of consent in advance.

6.3.4 Data Collection and Analyses

Data is collected in the P2P market for a duration of three months, from January 7, 2019 to March 31, 2019. In addition to the trading data, a pre-experimental sur-

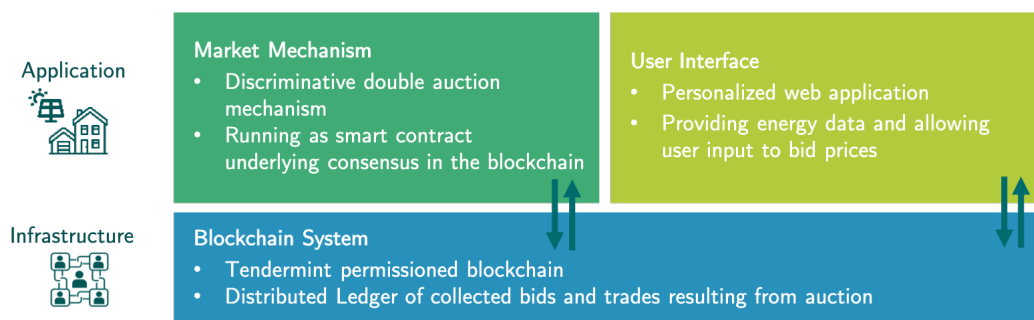


Figure 6.2: Schematic representation of the main system components of the P2P trading platform: blockchain infrastructure, market mechanism, and user interface

vey gathered supplementary information on participants' preferences and their socio-demographics. The study sample comprises 37 participating households, including 31 prosumer households. The prosumers' total peak production capacity exceeds 280 kWp. Over the duration of the study, the solar panels have produced 48,981 kWh and the participating households have consumed a total of 130,378 kWh. Over the study period of three months (8,024 clearing periods), a total of 292,316 orders posted on the market were collected. The time-discrete, discriminative double auction matched 424,049 trades from these orders, which were stored on the blockchain.

Based on this data, this article examines to what extent the implemented market realizes the value propositions laid out by Morstyn et al. (2018). To that end, this study analyzes the energy allocation and market efficiency achieved during the study period of three months, and examine the preferences elicited in form of prices bid by the participants. Furthermore, prices and resulting savings and revenues for consumers and prosumers, respectively, serve as performance indicators of the market. Most results are reported as relative values, as the absolute values depend strongly on absolute prices of the local utility provider and the absolute energy demand and production in the specific microgrid.

6.4 Results

Based on the data collected, this study analyzes the efficiency of the P2P market, participants' price preferences, and realized prices for local electricity to empirically evaluate the value proposition of P2P trading from the user perspective along the three dimensions described by Morstyn et al. (2018).

Energy Matching

To evaluate the efficient coordination of supply and demand in the P2P market, the following section examines how energy was allocated within the microgrid. More precisely, the analyses reveal whether the P2P market for solar energy provided incentives for renewable energy consumption and production and led to energy matching on a local level. Figure 6.3 depicts the weighted mean price to pay for energy by participants. The green line depicts the average day from the study period of January to March, the light green line the day with most solar production and the light blue line the day with least solar energy. As the chart illustrates, the price for energy decreases over the course of each day when orders could be matched, depending on the availability of solar energy and demand within the local market (green bars depict solar production, blue bars consumption). This type of price curve was also observed in related simulations (Mengelkamp et al., 2017a). The fact that the market price represents the relation of supply and demand in this way is desirable, as it incentivizes electricity consumption when there is most renewable production and highlights the very idea of a functioning market (Ketter et al., 2013). The lowest average market price was achieved at middays on the sunniest days of the study period. In almost all clearing periods, average prices were between the feed-in tariff (9.79×10^{-2} CHF/kWh) and the residential retail tariff (20.75×10^{-2} CHF/kWh), except for very few periods, in which consumers paid a price premium for local energy.

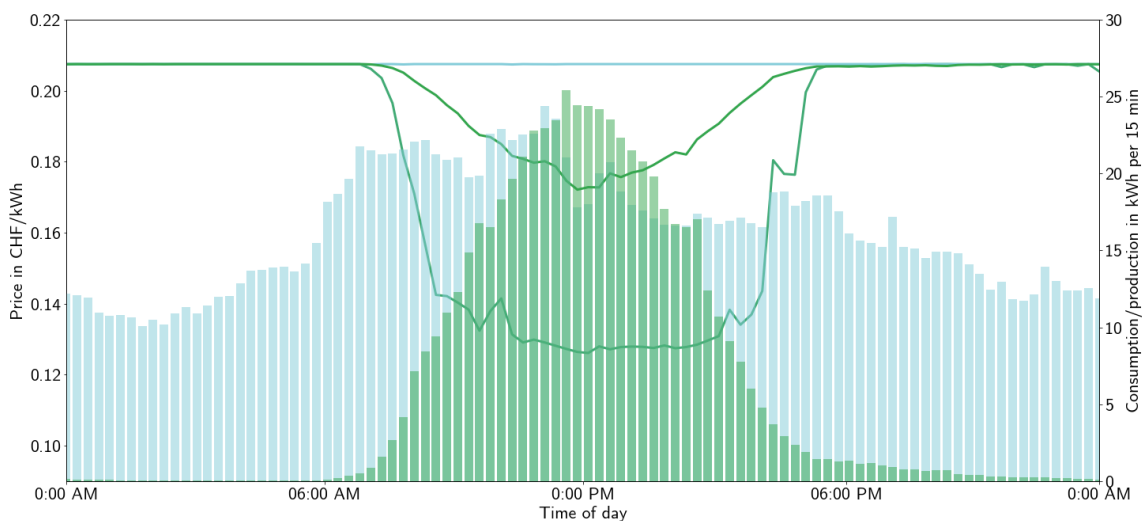


Figure 6.3: Price evolution over time of day during the three-month study period in winter: average day (dark green line graph), sunniest day (light green line), least sunny day (light blue line). The bar chart depicts average production and consumption loads in kWh. Prices within the P2P exchange reflect availability of local energy and range between feed-in and residential retail tariff.

The influence of sunny hours on the average energy price depicted in Figure 6.3 within the local market indicates that a considerable share of energy was indeed traded among peers and not at the fixed tariff defined by the utility provider (20.75×10^{-2} CHF/kWh). As all participating households are located within the same neighborhood, consumption profiles – and even more so production profiles – exhibit a high correlation between households (Griego et al., 2019). Given that, it is striking to what extent the local P2P trading increased local consumption of solar production: Without P2P trading, the overall SSR at community level corresponds to the share of electricity demand covered by the prosumers consuming their own solar energy. Over the duration of the study, the participant’s total SSR without P2P trading would have been 15.5%. With the P2P trading system enabled, the collective SSR almost doubled to 26.3%. Similarly, in the absence of P2P trading, the SCR, i.e. the share of produced solar energy that is consumed by the prosumers in their own houses, would have been 41.2% over the duration of the study. With the P2P trading system, the collective SCR of the market participants increased to 70.0%. These results are remarkable given that the data comprises three winter months in a community with a higher prosumer share than recommended in Griego et al. (2019) and assumed in Mengelkamp et al. (2017a). Overall, by enabling P2P trading, transactions of 14,092 kWh that would normally have involved the utility company were replaced by transactions among households within the microgrid. This implies that the load profiles and preferences stated by the participants could be matched for transactions of this volume.

To get an understanding of the efficiency of the market, this volume is now compared to the volume of energy that could have mathematically been traded within the P2P market given local supply and demand – in other words, the local solar production that occurred concurrently with consumption within the microgrid. Over the period of three months, the collective self-sufficiency rate reaches 26.3%, and self-consumption rate 70.0% with P2P trading. Local supply and demand actually concurred for 16,439 kWh and could thus technically have been matched within the microgrid. This means that during the three-month period of the study, 2,347 kWh of locally produced solar energy were not sold within the P2P market although there was local demand for it, due to a mismatch in participants’ bid prices. Since orders which cannot be filled within the P2P market have to be settled with the utility provider in any case, these transactions represent an inefficiency. If participants’ price bids had not been taken into account, the collective SSR would have been 1.8 percentage points higher (i.e., 12.6% of energy consumed could have been bought from the local market instead of 10.8%); yet, that fraction was supplied by the utility company. The freedom of decision-making granted to participants by actively including them in the pricing process thus comes at a trade-off of this decrease in SSR.

Preference Satisfaction

To evaluate price preferences for local electricity and their satisfaction on the P2P market, preferences stated in the pre-experimental survey are compared to prices bid in the market setting. As a first step, participants were asked prior to the field experiment whether they would be willing to incur higher costs for solar energy or for local energy supply: In this survey, 13 out of 31 participants who filled out the survey stated that they were willing to incur a price premium for solar energy and 17 that they would for local energy. As a second step, the prices bid by the participants in the market environment are put in relation individual preferences stated on the market: The histogram in Figure 6.4 displays all bids made on the P2P market. These bids reveal several interesting insights on the preferences elicited from the study participants. First, 27 of the 37 participants chose to define price bids other than the default prices, at least at some point during the study which indicates their willingness to engage on the market. Consumers offered 19.23×10^{-2} CHF/kWh (sd=2.37) on average for solar energy. The average sell price that prosumers wanted to earn was 13.67×10^{-2} CHF (sd=3.85). This implies that in general, many transactions among the peers can be matched within the microgrid. However, a lot of buy prices as well as sell prices bid intersect in the interval between 12.5 and 18×10^{-2} CHF/kWh. This indicates that there may occur cases in which sell orders ask for a higher price than offered in the buy orders – which explains the inefficiencies identified above. Moreover, the high standard deviation and broad distribution of bids indicates that participants have heterogeneous preferences and that many participants did seize the opportunity to influence the decision making process on the market. While some prosumers asked a price premium be paid by the consumers (sell bids $>20.75 \times 10^{-2}$ CHF/kWh), none of them is willing to incur opportunity costs for selling their energy locally by offering their solar energy below the feed-in tariff ($< 9.79 \times 10^{-2}$ CHF/kWh). Hence, the bids by prosumers in this study do not display other-regarding preferences or prosocial behavior for selling electricity locally. On the consumer side, 11% of the buy orders are higher than the utility tariff; these orders were posted by 6 different participants. While these 6 participants (temporarily) offered to incur a slight price premium for solar energy from the microgrid, overall, the participants' real-world price settings in the field study considerably deviate from their self-reported preferences indicated in the pre-experimental survey. In other words, once their choices were consequential for their real-world income, they were less willing to pay a price premium for local solar energy (and to incur opportunity costs for selling their energy locally, respectively) than their responses to the hypothetical scenario in the survey prior to the field study had suggested. These findings call into question the results of survey-based evaluations of individuals' willingness to pay for renewable ener-

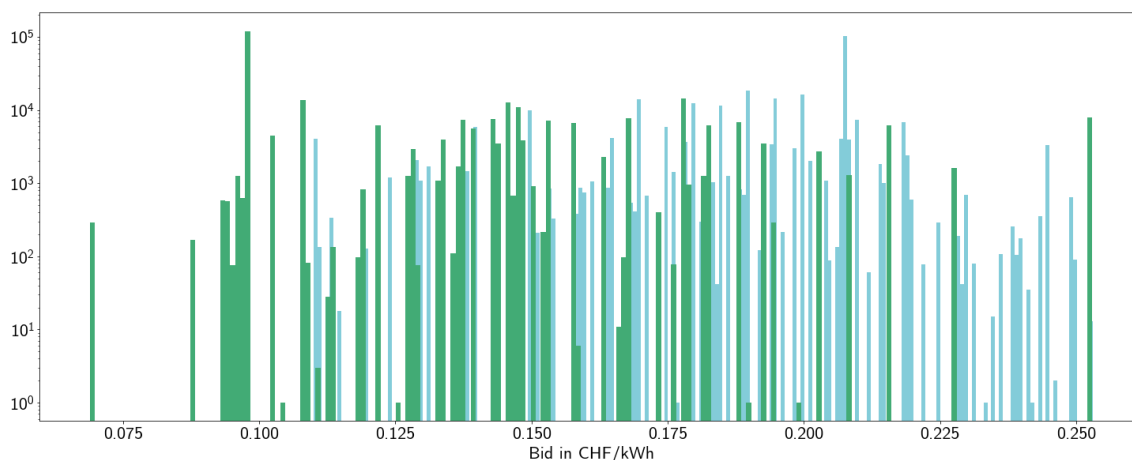


Figure 6.4: Histogram of prices bid for local solar energy on log-scale: Sell prices bid are displayed in green, buy prices bid in blue (default tariffs defined by the utility provider: 9.79×10^{-2} CHF/kWh and 20.75×10^{-2} CHF/kWh).

gies (Ecker et al., 2018; Tabi et al., 2014). Participants’ self-reported inclination towards renewable and local energy (as stated in the pre-experimental survey), which is in line with preferences reported by other survey-based studies in the existing literature, does not translate into their behavior in the market setting in which participants’ bids determine the actual costs they incur.

Uncertainty Reduction

Having examined the preferences elicited, the prices realized are now examined to understand to what extent P2P trading may help to reduce uncertainty for the prosumers (Morstyn et al., 2018). To that end, the transactions realized and their implications for the users are assessed. The mean price per kWh for transactions among peers is 16.80×10^{-2} CHF (sd 1.78). Except for a few cases, prices for almost all transactions fall within the limits of the fixed feed-in tariff of 9.79 (as lower bound) and the residential retail tariff of 20.75×10^{-2} CHF (as upper bound). As illustrated in Figure 6.3, prices for local solar energy vary over the course of the day, depending on the availability of solar energy. This has two implications, both of which are in line with the results above: 1) On average, both sellers and buyers benefit from the P2P transaction, as they trade at a price that is below the price that the consumer would have to pay to the utility company and above the revenue that the prosumer would earn for feeding into the grid. 2) The average prices realized do not include a price premium over the grid tariffs.

In fact, the results imply that all users have benefited from the P2P trading in the field study: Their incurred electricity costs in the P2P market are below the costs users would

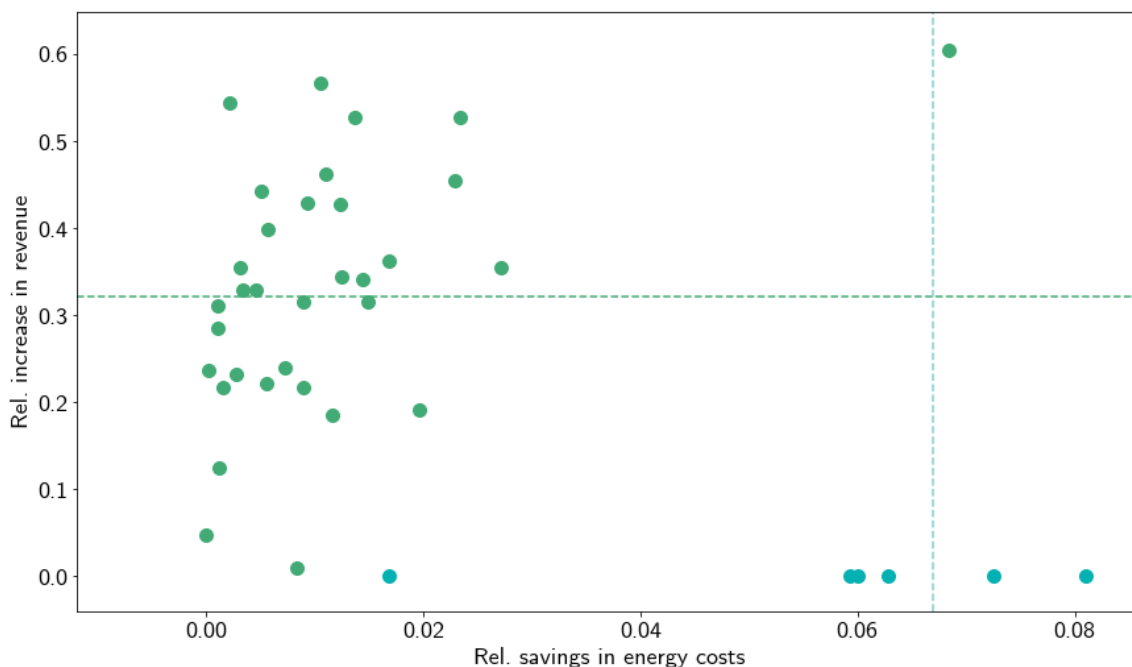


Figure 6.5: Savings and additional revenue incurred by each participant in the P2P market relative to their expenses/revenue for trading electricity with traditional tariffs.

have incurred if they had not been part of the local market and had bought from and sold to the electricity provider. Summing up the transactions for each user in this way, each of the participants either saves electricity costs, earns more for the solar electricity she produces, or both. The scatter plot in Figure 6.5 shows the relative increase in revenues for sold electricity (on the y-axis) and relative savings on electricity expenses (on the x-axis) by each of the users. Pure consumers are depicted in blue (no electricity sold, hence no revenue increase), prosumers in green. On average, users earned 32.2% ($sd=0.15$) more for the electricity they sell, and saved 1.8% ($sd=0.022$) of their electricity expenses. At first glance, the relative savings on electricity purchased seem very small. Upon closer inspection, the numbers are not surprising, as prosumers cannot save much in buying solar energy, as they mostly consume solar energy from their own roofs during sunny hours – they benefit from the peer-to-peer market on the seller side. Focusing on pure consumers alone (who do not own a solar panel), these saved an average of 6.7% ($sd=0.008$) of their electricity bill. Moreover, savings may likely increase in summer months with more excess supply and prolonged hours of sunlight. Taken together, the results indicate that the market design is supporting the overall goal of providing a profitable market for renewable energy produced by small prosumers and reducing uncertainty for prosumers' investments.

6.5 Discussion & Conclusion

6.5.1 Discussion

This article investigates a P2P market for solar energy in the field and collects high-resolution empirical data of 37 participating households over the duration of three months. The article thus contributes to the discussion on smart markets for renewable energy (Bichler et al., 2010; Ketter et al., 2013) and green IS (Gholami et al., 2016; Malhotra et al., 2013; Melville, 2010). To the best of the study authors' knowledge, this is the first scientific evidence on trading conducted on a P2P energy market in the real world. This study analyzes the data collected with respect to the three value propositions of P2P trading proposed by Morstyn et al. (2018): energy matching, preference satisfaction, and uncertainty reduction. It is important to note that the quantitative results achieved in the study sample are not generalizable to the broader public. Given the novelty of this research area and the complexity of the energy market (Ketter et al., 2013), this study examines the value propositions of P2P markets that have been theorized in the literature in the field and provide a first benchmark for the real-world impact of peer-to-peer energy trading among users in the field. With this impact-oriented approach (Gholami et al., 2016), the present study tackles the first stage of smart market design of understanding the user value and eliciting user preferences, as described in Bichler et al. (2010).

Despite the local proximity of the participating households, the findings indicate that by matching supply and demand within the P2P market, the share of self-sufficiency of the P2P community can be increased by 70% (from SSR of 15.5% to 26.3%), even during the winter months of January to March. These numbers will likely increase over the summer months with longer hours of sunlight. Moreover, given that prosumers still sell around one third of the solar energy to the utility provider, SSR and SCR could be increased by shifting flexible loads or by deploying more storage capacities in the microgrid. The double auction employs price limits stated by the users to match trades. Yet, there is a tradeoff between computing an efficient energy matching on the one hand, and on the other hand enabling individual preference satisfaction (Morstyn et al., 2018) of the users by letting them bid prices: Energy that cannot be matched on the P2P market needs to be supplied from the utility provider at the residential retail tariff, as security of supply needs to be guaranteed at any time.

The inefficiency observed in this field study reduced the technically possible SSR by 1.8% percentage points – with a decreasing trend over time. This seems like an acceptable trade-off; in exchange, the market design implemented in this study allowed a greater

influence of the participants, as they could directly state their willingness to pay for local solar power and thus actively influence prices. Aside from that, the real-time prices achieved in the market reflect the relation of supply and demand on the market very well, as is shown in Figure 6.3. The pricing achieved by the matching mechanism selected thus also incentivizes shifting consumption loads to periods in which local solar energy is available, which could be achieved using smart appliances or storage capacities in the future (Fridgen et al., 2016). This result is also interesting beyond the context of the energy domain, as it shows that with the right market design, a P2P market can be relatively efficient while, at the same time, enabling individuals to participate in the decision making on a market (Lampinen and Brown, 2017). The double auction mechanism can handle manually defined, heterogeneous individual preferences and still run autonomously in real-time.

Furthermore, in this field study, the vast majority of residential solar energy was sold within the P2P market and increased revenues from renewable generation, thus reducing uncertainty of returns on investments for prosumers. As argued above, the auction mechanism manages to provide incentives for local generation, which in turn creates incentives for investments in renewable generation and reduces insecurity of investment. It is important to note that the feed-in tariff granted in this field study is relatively low compared to current, subsidized tariffs in European countries. However, this projects the future market structure, as feed-in tariffs and their financial support schemes have been reduced in the past years and might even disappear in some countries (Karneyeva and Wüstenhagen, 2017), which illustrates the importance of studying novel market structures integrating distributed prosumers. The direct preference elicitation from the consumers (by letting them bid a price per kWh of solar energy) may seem like an extreme approach to involve the user in a rather abstract decision making process. Nevertheless, the results indicate that direct involvement of consumers is indeed crucial to understand the heterogeneity of preference profiles and consumer behavior in a real market environment, as this may differ strongly from their statements made in surveys. Moreover, there may be additional societal benefits of engaging consumers directly in the energy market (European Consumer Organisation 2016): By allowing consumers to influence the energy sources they use or even the prices they pay, they assume a more active role.

This empowerment may increase the salience and the understanding of energy supply. The double auction represents an extreme approach that directly allowed the users to bid prices for different sources of energy. The results can now serve as empirical starting point for designing decision support systems which automatize smart trading strategies adapting

to the consumer type (Bichler et al. 2010) or which provide consumer analytics for energy consumption. Also beyond the energy sector, information systems provide various avenues to support users in decision processes both in their professional and private lives. Many of these systems include autonomous agents and regardless of the specific application context (Bichler et al., 2010; Gholami et al., 2016), a key question will be how to make sure that these systems act according to the users' preferences. The discrepancy between the participants' self-reported price preferences in a hypothetical scenario and their actual price settings in the field study highlight the importance of empirical research to better align the strategies of autonomous agents with the individuals' actual preferences in the real world.

Overall, the results presented here confirm the value propositions of P2P markets that were theorized in the related literature (Andoni et al., 2018; Mengelkamp et al., 2017a; Morstyn et al., 2018) and have, partially, been observed in electronic peer-to-peer markets in other domains (Einav et al., 2016; Zimmermann et al., 2018). Trading energy directly between private households may become part of a future energy landscape, since thus field study shows that the technology to put such platforms into practice already exists. Yet, future research needs to investigate whether these benefits are actually perceived and appreciated by the individual user and whether they justify the costs, time, and efforts involved in the deployment of a distributed information system. This article takes a step to address the dearth of impact-oriented research in the field of green IS (Gholami et al., 2016); yet, fostering sustainability is a wicked problem with many interrelated aspects and consequently, the design of smart energy markets for the future requires further research.

6.5.2 Limitations & Outlook

Despite all best efforts, this study is not without limitations. The very complex technical setting in the field and the criticality of energy supply for all users imposes some natural restrictions to the study design. Due to the complexity of the study implementation, the explorative approach on a critical infrastructure, and to the associated costs, the sample was limited to 37 participating households. Furthermore, the sample recruited features a high share of prosumers; as early adopters, they may be more interested in energy or sustainability topics than the general public, therefore the results may be subject to volunteer-selection bias (Tiefenbeck et al., 2019). Future research needs to investigate to what extent a broader population is receptive to P2P energy markets and how these markets and the user interfaces need to be designed not to overwhelm citizens who so far did not have any active role in the electricity market. This being said, it is all the more

remarkable that in the field study, the high willingness to pay for local solar power which the same participants had stated in the pre-experimental survey and which is in line with previous survey-based studies (Ecker et al., 2018; Tabi et al., 2014) is not reflected in the bidding behavior.

From an economic perspective, a particular feature of the application context is that the utility provider backs up every order that could not be matched within the microgrid. If that was not the case, the strategic incentives in the market would have been reduced and Pareto efficiency would have been fulfilled. However, it is a necessity to keep the electricity grid in balance and to provide reliable electricity supply at all times, so the trade-off between respecting individual preferences and accepting inefficiencies is a natural property of market design in this domain.

One reason for the lack of empirical data on P2P electricity trading is that technological advances in communication technology and distributed ledger technologies have spurred the interest in decentralized platforms only in recent years (Albrecht et al., 2018; Basden and Cottrell, 2017; Buterin, 2014; Hasse et al., 2016). Naturally, the field implementation of such a complex socio-technical system raises various interesting questions in different research areas, including human-computer interaction, technical aspects, and regulatory issues. For instance, the choice of the technical system architecture is beyond the scope of the present article. In particular, this article does not aim to evaluate the advantages or disadvantages of the blockchain infrastructure implemented in this field test. Another aspect requiring further investigation relates to the design choices of the user interface implemented on this P2P market and its influence on the trading behavior and understanding of the users. In the course of the research project, several interventions will be implemented and qualitative data collected to assess these questions in detail. Going forward, it would be also interesting to investigate other market designs incorporating forecasts and include decision support systems for the user (Bichler et al., 2010), e.g. active control of flexible loads and storage capacities, or autonomous agents taking part in the auction mechanism based on user input. For that purpose, it would be interesting to examine spillover effects on renewable adoption and possible shifts in load profiles caused by the real-time pricing and additional information provided to the users. Such effects have been observed in other studies on P2P platforms (Bakos and Katsamakas, 2008). Finally, the deployment of P2P energy markets on a larger scale will have implications on the grid infrastructure, grid costs, and demand schedules which need to be carefully investigated from an engineering perspective on a systemic level. In this context, an obvious question relates to alternative models for grid fees and pricing schemes to recover

the costs for (super)regional transmission lines if the diffusion of P2P markets picks up and consequently, the share of locally produced and consumed energy increases (“Who pays for the grid?”). In sum, while the variety and breath of important questions arising in detail cannot all be addressed in this article, the empirical data collected provide a concrete starting point to foster the debate across disciplines.

6.5.3 Conclusion

In recent years, advances in personal information systems and in blockchain technology have enabled the creation of new marketplaces, in particular for trading or sharing of goods among private consumers. Given the increase of distributed energy resources, the energy sector can benefit from this evolution if a market design can be established that is beneficial to the user (Bichler et al., 2010; Morstyn et al., 2018). This article presents a framed field study to test a P2P energy market in the real world and provides early empirical evidence on the impact of this novel market platform from the user perspective. To that end, a P2P electricity exchange for solar energy was set up in a local community in Switzerland. A time-discrete, iterative double auction with discriminative pricing was implemented based on existing literature on P2P energy markets and market design theory. The established utility pricing serves as benchmark for trading data observed in the field study. Furthermore, the preferences displayed by the users on the real-world market are compared to survey-based findings on consumer preferences for local or solar energy, both elicited from the same participants, and reported in prior literature. The results suggest that the value propositions theorized in the literature can actually be realized for the user in P2P energy markets. If the regulatory framework allows, information systems can be a viable option for prosumers to sell their excess production locally and directly instead of being dependent on feed-in tariffs determined by regulators and utility companies. Policy makers should facilitate the creation of user-centric market structures that allow for local energy matching and provide the possibility to reflect heterogeneous consumer profiles. When addressing the user needs and employing efficient market mechanisms, information systems have the potential to create smart energy markets that foster sustainability on its three levels: socially, economically, and ecologically.

7. Article F) Bidding Behavior on a P2P Energy Market

7.1 Motivation & Introduction

To mitigate climate change and achieve sustainability, energy systems need to move towards low-carbon, affordable, and socially equitable energy services. Such aims have been underscored by the Sustainable Development Goals of the United Nations (2019) which call for clean and affordable energy supply, as greenhouse gas emitted by energy generation from fossil fuels is likely the dominant cause of climate change (IPCC et al., 2014). Distributed renewable electricity generators, together with an electrification of transport, hold a pivotal role in cutting greenhouse gas emissions (International Energy Agency, 2018; Siler-Evans et al., 2013; Williams et al., 2012). In this context, rooftop solar photovoltaics (PV) have been promoted by policy incentives and subsidies in many countries. However, owning solar panels is not yet popular among the masses: It is still often an unsatisfying endeavor to invest in renewable generators, as solar production does not usually coincide with residential energy demand, and consequently a considerable share of solar generation is fed back into the grid (Schopfer et al., 2018) at tariffs defined by regulators or utility providers. These feed-in tariffs, in turn, have been falling for the past few years, prolonging amortisation periods (Hoppmann et al., 2014; Karneyeva and Wüstenhagen, 2017). Moreover, such fixed-price mechanisms generally do not reflect real-time market conditions (Morstyn et al., 2018; Ramchurn et al., 2012), and fail to provide efficient incentives for generating or consuming renewable energy when it is available – even less so on a local level. Beyond that, many people live in rented homes where they cannot make long-term hardware investments (Andrews and Sánchez, 2011).

Simultaneously, in many applications, the use of information systems has led to “on-line auction markets for resource allocation in distributed systems” (Guo et al., 2012, p. 823). Two-sided platform, or ‘peer-to-peer’ (P2P) markets, like AirBnB, Uber and eBay, have revolutionized the way goods are marketed online (Adomavicius et al., 2009; Bapna et al., 2004; Einav et al., 2016). In the urgent ongoing discussion about climate change and the transition to renewable energy, smart P2P markets in which owners of solar panels can sell their energy to neighbors seem like a potential mechanism to facilitate the integration of renewable energy generators (Gholami et al., 2016). P2P energy markets are proposed as coordination mechanism for distributed energy resources (Morstyn et al., 2018; Ramchurn et al., 2012; Zhang et al., 2018a), to provide an alternative to the often rather crude top-down coordination of energy resources by regulatory frameworks or subsidies which is currently in place (Siler-Evans et al., 2013). In combination with a surge in household expenses for electricity that will result from electrification of transport, accessible online platforms may increase consumer interest in electricity sourcing and renewable generation.¹ Given the opportunity to sell their solar energy on such platforms, private consumers may start leading the way into a radically transformed energy market and, possibly, more efficient resource use, just like they did in other industries (e.g. real-estate, transportation).

Yet, to date, there is no study that has examined trading on a P2P market for solar generation between neighboring households in practice. In particular, there is a lack of empirical research and of consumer focus in this field, as most existing studies on the topic either (1) are of conceptual nature or (2) are empirical but focusing on hypothetical situations, or (3) focus on the technical implementation of information systems to manage solar PV (Andoni et al., 2018; Mengelkamp et al., 2017a; Morstyn et al., 2018). Research from other application contexts illustrates that individual behavior in online markets often deviates from theoretical predictions and can be very heterogeneous (Bapna et al., 2004; Lu et al., 2016), which may impact the efficiency of market mechanisms. As standard assumptions of rational, risk-neutral bidders with perfect information usually do not hold true in practical applications, Bapna et al. (2003) argue that behavior of individuals in online auctions is not fully understood yet, and that classical economic theory may not incorporate all of the newly arising phenomena in electronic markets. All the more so, it is crucial to carefully investigate novel market designs and arising behavior (Adomavicius et al., 2009) in a domain as societally important (and as complex) as the energy sector

¹Whereas electricity expenses currently represent roughly 1.4% of consumer expenditures by the average Swiss household, this volume is likely to double or triple with the electrification of transport and heating at current tariffs (Bundesamt für Statistik, 2020) – thus increasing its relevance.

(Gupta, 2017; United Nations, 2019).

This article presents the first real-world evidence on behavior in an IS-enabled marketplace for P2P energy trading. In a year-long framed field study, residential households were trading solar energy in an online auction market. 37 residential customers of a utility provider in a town in Switzerland formed a local P2P market. A majority of the participants already owned solar panels (these will be called ‘prosumers’) and had hitherto sold their excess production to their utility provider at fixed feed-in tariffs. This study examines the bidding behavior and market outcomes of this real-world implementation of a P2P market for solar energy.

My colleagues and myself designed, implemented, and deployed a trading platform using a blockchain system that ran on smart meters installed in the participants’ homes. Participants interacted with the P2P market using a personalized web application. Notably, they directly determined their maximal willingness to pay for locally produced solar energy, as well as the minimum price at which prosumers were willing to sell, using the application. Participants’ bidding behavior in this study did not merely represent intentional statements, but directly influenced their electricity bill. The buy and sell prices they stated on the web application were thus not a merely theoretical exercise, but had actual consequences on their expenses during the entire year of the field study. Leveraging the empirical evidence collected, this field study investigates the following question: *Does the individual behavior observed deviate from cost-minimizing behavior?* The experimental period of an entire year further provides insights on effects such as learning behavior or reactions to seasonal changes: *Does bidding behavior evolve over time?* The analysis of this data, draws on the existing body of literature on bidding in online auction markets (Bapna et al., 2003, 2004; Goes et al., 2012; Guo et al., 2012; Lu et al., 2016). Given that bidding behavior translated to actual, incurred costs in the present study, these findings complement survey statements observed in the related literature. Third, this article discusses intelligent agents as a tool to incorporate individual user preferences, while disburdening humans from making operative trading decisions (Bichler et al., 2010; Gupta, 2017). While autonomous agents could provide intelligence in this abstract domain of electricity supply in which individuals tend to lack long term thinking, it is crucial to understand individuals’ preferences for keeping ‘humans in the loop’. As Bichler et al. (2010) argue, “the degree of autonomy that an agent or automated decision support system should have is a completely open research question especially in dynamic and complex market environments [...]”, p. 697. To that end, participants’ reaction to an alternative automated pricing mechanism that supplements the interactive auction mechanism for a

period of four weeks is tested in the present field study. The within-subject design allows capturing the reaction of participants who actually have experienced both regimes in the real world: *Do participants prefer an automated dynamic pricing system rather than bidding prices themselves?* A tentative explanatory model further deepens the understanding of the data observed in the field and an agent-based simulation models different bidding strategies to put the observed behavior into context (Ketter et al., 2013).

The present study complements conceptual reviews, theoretical analyses, and empirical work with hypothetical situations of P2P energy markets by the human element and brings this novel market design to the real world. Allowing actual human participants to interact adds a behavioral perspective on such market structures in practice and evaluate whether online market mechanisms can efficiently coordinate distributed energy resources on the household level. To the best of the study authors' knowledge, this is the first study to collect empirical data bidding behavior on a P2P energy market. It provides first evidence on a pure market mechanism for the direct allocation of renewable energy among residential households. The implications of the findings may be instrumental for creating smart sustainable energy markets (Bichler et al., 2010) and, further, for designing autonomous agents providing algorithmic decision support in this context (Lu et al., 2016). This study thus contributes to the work on Green IS (Gholami et al., 2016; Gupta, 2017), collecting real-world evidence on the use of an electronic market for renewable energy. The article thereby goes "beyond conceptualizing, analyzing, and even designing" by conducting research "with demonstrable impact on mitigating the threat of climate change" (Malhotra et al., 2013, p. 1266), tackling one of the 'wicked problems' of our time (Creutzig et al., 2018; Gupta, 2017; Ketter et al., 2015).

7.2 Related Literature and Theoretical Background

7.2.1 P2P Energy Markets

Despite growing public attention for environmental sustainability and tremendously reduced costs for solar generation (International Energy Agency, 2018), the diffusion of renewable energy generation is still advancing too slowly (International Energy Agency, 2020b). Researchers as well as practitioners hence examine different ways for fostering the integration of individual solar panels or storage systems in the energy market (Parag and Sovacool, 2016) and ways to actively engage individual consumers in the energy transition (European Consumer Organisation, 2016). These activities have led to efforts to create decentralized market models in which electricity is sold directly from peer to peer (Basden and Cottrell, 2017; Morstyn et al., 2018; Parag and Sovacool, 2016). P2P exchange

of electricity is a paradigm shift from a centralized energy landscape to a decentralized bottom-up market model in which individual consumers and prosumers (households that produce energy themselves using solar panels or other electricity-generation assets) can directly trade electricity on demand without the mediation of a central utility provider acting as reseller and defining tariffs (Mengelkamp et al., 2017a). Prosumers can sell excess electricity they produce directly to other consumers within local communities on the local low-voltage distribution grid level. Thus, a higher share of the electricity demand of a community can be covered locally, which may help avoid or delay the need for investments in centralized generation infrastructure and in transmission (Jain et al., 2017). In addition, it grants the owners of renewable generators more choices; allows consumers who do not have the knowledge or possibility to invest in own infrastructure to support renewable generation; and it is hoped to reduce overhead costs levied by utility companies. However, given the recent developments in distributed energy resources and distributed communication, existing research on decentralized electricity exchange is still in an early stage (Mengelkamp et al., 2017a; Morstyn et al., 2018; Zhang et al., 2018a); and there is a lack of empirical data on P2P energy markets with real participants.

In a simulated P2P market with standardized consumption profiles, in which they implement a time-discrete double auction for solar energy, Mengelkamp et al. (2017b) construct artificial bidding data. They compare a learning strategy to a zero-intelligence bidding strategy in a uniform double auction market. They find that the self-consumption rate, as well as market prices achieved, vary with the bidding strategies employed. Financial outflow out of the P2P market is lower for the learning strategy. The authors point out, however, that their findings need to be validated in the field.

Hahnel et al. (2019) conducted an online study in which a group of participants were instructed to imagine they were owners of PV panels and were part of a P2P trading community. The authors identified three clusters of prosumers in these scenarios: price sensitive, autarky-focused, and heuristic prosumers. Participants' decisions thus mainly depended on electricity market prices (which they could not influence in this study) and on their own storage charging states. For the largest cluster of participants, the authors found a high price elasticity, highlighting the importance of financial factors and price preferences in this trading.

To understand what generally drives consumers to subscribe to already existing green electricity products, Tabi et al. (2014) conducted a conjoint analysis on a consumer survey among German households. They found that although a majority of consumers state a preference for renewable energy sources, only a small fraction had actually purchased a

‘green’ electricity product. The authors identified different groups of consumers, which differ in their sensitivity to prices and local sourcing, and in their likelihood of purchasing green electricity. In their study, the majority of ‘potential adopters’ of green electricity products do not state a high sensitivity to prices. In a survey conducted by Ecker et al. (2018), consumers state a willingness to pay a price premium of 20% on average for renewable energy produced in their own homes. In addition, they find that participants propose equal buying and selling prices for trading electricity on P2P markets in this survey.

All these results are based on hypothetical scenarios tested in an online study. As pointed out by Andoni et al. (2018), there is a lack of empirical evidence on the bidding of private individuals in such energy markets, and survey results do not always provide a true representation of agents’ actions (Fishbein and Ajzen, 1975). Probably the first prototype of a P2P energy market deployed in the field was the Brooklyn Microgrid, in which locally produced power from solar systems was sold within the neighborhood. In a case study, Mengelkamp et al. (2017a) describe the system components of this market. Market clearing was implemented by a time-discrete double auction in 15-minute intervals, but no data on the participants or their behavior has been made available.

7.2.2 Bidding in Online Auctions

Involving private individuals in price-setting procedures has become increasingly popular in the digital economy (Bapna et al., 2003). Online auctions are some of the purest market design problems that arise in practice (Roth, 2008), and they provide an ideal setting to study individuals’ bidding behavior. Here, empirical studies in various domains have shown that game-theoretical predictions for bidding behavior or the underlying assumptions often fail in practice, even in iterative auctions in which individuals have the chance to learn from earlier clearing periods.

For instance, one of the concepts of classical economic theory relating to the bids for buying and selling a good, says that willingness to pay (WTP) and willingness to accept (WTA) should be similar on an active market (Shogren et al., 2001), as sellers can rebuy a good they sell for a similar price and vice versa. However, seminal work from the behavioral economics field has shown that in many contexts, individuals ask for a higher price when selling a good than they would be willing to pay to buy it. Kahneman et al. (1990) attribute this phenomenon to an ‘endowment effect’, which is observed even in settings where no sentimental attachment to the traded good is expected. This effect is closely related to the notion of loss aversion described by Kahneman (1992).

Another essential economic concept in this context is equilibrium analysis in auction markets. However, the computation of equilibria under imperfect information, especially in sequential or multi-unit auctions, is often computationally intractable (Bapna et al., 2004). All the more so, it is hard to capture bidder heterogeneity in such analyses. This complexity limits the applicability of theoretical literature on auction design (Adomavicius et al., 2009). In recent years, scholars in many domains have hence turned to empirical experiments, as well as simulation models of markets, to complement theoretical analyzes.

Bapna et al. (2004) examined different types of bidding strategies and their resulting economic welfare in online auctions. They identified heterogeneous bidding strategies based on bidders' time of entry, time of exit, and number of bids in an auction. The surplus achieved in the auctions varied among different clusters. In a tentative analysis using a subset of their data, they also investigated learning effects: with repetition, bidders improved their performance. Lu et al. (2016) examined bidder heterogeneity in business-to-business auctions and identified five statistically different types of bidders based on these attributes, who also achieve varying levels of economic surplus. In an attempt to identify drivers for these distinct strategic choices, they derived an explanatory model based on factors like transaction costs and budget constraints rather than psychological or social factors, as they are investigating professional bidders in a B2B context. Using a similar approach, Goes et al. (2012) examined bidding behavior in sequential auctions of a well-known online auction platform. They confirmed the set of strategy types identified by Bapna et al. (2004) and further showed that bidders' experience had a significant effect on their behavior. Moreover, they find indications for the 'declining price anomaly'. The term describes the phenomenon that market prices in sequential auctions of identical goods decline over time (McAfee and Vincent, 1993).

In a longitudinal field study, Goes et al. (2010) investigated drivers for willingness to pay in sequential online auctions. Based on their data collected in online retail auctions, the authors argue that willingness to pay is strongly influenced by heterogeneity in demand, and by a bidder's experience in previous instances of a sequential auction. In particular, bidders successively reduce their WTP after winning in an instance of the auction.

7.3 Study Design

This framed field study examines bidding in P2P energy markets in the field: A blockchain-based P2P market enables the trading of solar energy in the real world. The study with $n = 37$ participants took place in a town in Switzerland and lasted for one full calendar year. It is one of the first realizations of a peer-to-peer electricity exchange

worldwide in which households can engage in direct trading of solar energy using an information system. Participants actively bid price preferences for locally produced solar energy and collected these bids at 15-minute intervals throughout the experiment. The information system that enabled P2P trading and the implementation of the study are described in more detail in the next sections.

7.3.1 Information System

The devices and network, forming the infrastructure for the trading platform of the layered, modular architecture described in Constantinides et al. (2018), are represented by smart meters running a blockchain network. The auction mechanism determining transactions on the P2P market constitutes the application layer. Finally, participants interact with a user interface for this application, a web application displaying individualized information. Each of these layers will be described in more detailed in the following paragraphs. (An overview of the different layers of the information system deployed to run the P2P energy market is shown in Figure 6.2.)

Infrastructure

To enable P2P trading among participants in real time, a distributed information system that was developed for the purpose of this experiment was deployed: Each household was equipped with one to three smart metering devices. Each household received one device that measures electricity consumption. Prosumers received another smart meter for measuring electricity production from their PV panels and participants who own a battery storage system received a third smart meter for measuring battery loads. In total, 75 metering devices were deployed. All devices are connected to the internet; they measure electricity loads in time intervals of 15 minutes.

Application

In combination with the electricity loads e , measured by the smart meters, willingness to pay p_b and willingness to accept p_s , defined by participants are input to a discriminatory double auction (Borenstein et al., 2002; Fabra et al., 2002) for allocating electricity trades within the market. The participants' buy and sell orders, $b_B = (p_b, e^+)$ and $b_S = (p_s, e^-)$, for local electricity were collected over a clearing period of 15 minutes. In theory, price preferences could thus have been adjusted on a 15-minute basis, which is, however, virtually impossible in practice for a human user and is not to be expected. After the orders are collected, an auction mechanism is run to clear the market and determine the resulting

electricity trades. Prices thus reflect availability of solar energy in near real time (Rosen and Madlener, 2013).

To enable prosumers and consumers to decide at which conditions they are willing to sell or buy sustainable electricity, time-discrete, iterative double auction (Algorithm 1) was implemented. The auction implements a discriminatory pricing rule, which means that the price participants pay directly depends on their own bid (Fabra et al., 2002). This increases comprehensibility for human participants as compared with a uniform market clearing price. Furthermore, discriminatory price auctions for electricity tend to yield lower volatility in prices and reduce vulnerability to implicit collusion – at the cost of higher prices in off-peak periods (Fabra et al., 2002; Klemperer, 2002; Rassenti et al., 2003).

Algorithm 1 Iterative Discriminatory Double Auction

```

1: initialize  $e = (e_1, \dots, e_{37}) = (0, \dots, 0)$ ;
2: initialize  $p = ((p_{1,b}, p_{1,s}), \dots, (p_{37,b}, p_{37,s})) = ((20.75, 9.79) \dots, (20.57, 9.79))$ ;
3: loop
4:   smart meters measure  $e = (e_1, \dots, e_{37})$  for this timeslot
5:   participants (may) adjust their price bids  $p_{i,b}$  and  $p_{i,s}$ 
6:   if clearing time is reached then
7:     create buy and sell order books  $b_B, b_S \subset b$ :
8:     if  $e_i \geq 0$  then
9:        $b_i = (p_{i,b}, e_i)$ ,  $b_B \leftarrow b_i$ 
10:    else if  $e_i < 0$  then
11:       $b_i = (p_{i,s}, e_i)$ ,  $b_S \leftarrow b_i$ 
12:    end if
13:    sort order books  $desc_p(b_B)$  and  $asc_p(b_S)$ 
14:    match orders:  $(b_B, b_S) \rightarrow t$  with  $t$  being the resulting set of trades  $t_h = (p_{h,t}, e_{h,t}, i, j)$  with  $p_{h,t} = \frac{1}{2} \cdot (p_{i,b} + p_{i,s})$ 
15:    settle remaining bids  $b'$  with utility provider  $u$ :  $(b'_B) \rightarrow t'_B$  and  $(b'_S) \rightarrow t'_S$ , with  $t_i = (t_u, e'_i, i, u)$  and  $t_j = (f_u, e'_j, u, j)$  respectively
16:  end if
17: end loop

```

The auction matches the highest buy offer with the sell order with the lowest ask price (in terms of price) and progresses in descending/ascending order through the entire order book (ties are resolved by a random draw). The price for each matched trade is the mean between the sell and the buy price of the respective orders (‘discriminatory/midpoint pricing’), $p = \frac{1}{2} \cdot (p_{\text{buyer},b} + p_{\text{seller},s})$.

All participants are customers of the local utility provider whose electricity tariff was $t_u = 20.75 \times 10^{-2}$ CHF/kWh² and the feed-in-tariff granted to prosumers was $f_u = 9.79 \times 10^{-2}$ CHF/kWh, including network fees. When electricity supply and demand within the microgrid was not balanced, the utility provider bought or sold excess capacities at these tariffs. The P2P market thus operated in grid-connected mode (Halu et al., 2016).

User Interface

Participants received access to an individualized web application, which had been designed for the purpose of this study. The application allowed them (1) to monitor real-time data on their energy consumption (and production, if applicable), and data on their past trading behavior and, in particular, (2) to place price bids: By moving a slider element, they could express p_b , their willingness to pay (WTP) for solar electricity produced by their neighbors in the microgrid.³ Prosumers were also able to define their willingness to accept (WTA): p_s , their minimum ask prices for selling energy from their solar panels to other households, as opposed to selling it to the utility provider at feed-in tariff (f_u). The slider element is depicted in Figure 7.1 and thoroughly described in Ableitner et al. (2020).




Figure 7.1: Slider elements included in the web application for prosumers of the P2P market. The consumer version only includes the buying side on the left.

²Electricity prices are indicated in Swiss Francs (CHF) per kWh, as the study took place in Switzerland.

³The superscript ^ denotes datapoints collected in this field study throughout this paper.

Feed-in-tariff and retail tariff of the utility provider were indicated on the slider elements for orientation, but the slider range also allowed participants to set their purchase bid above retail price (expressing a preference for green/local electricity) or to sell solar energy for free within the neighborhood. Furthermore, the application provided a concise overview of their energy data and their trading outcomes on the local market in real time at their discretion, as earlier research indicates that participants may be interested in the local origin of the energy they buy (Ecker et al., 2018; Meeuw et al., 2018). A screenshot of the web application’s energy data visualization is also given in Figure B.4, with further impressions in Appendix B.

7.3.2 Implementation of the Study and Participants

The field study took place from Jan. 7, 2019 until Jan. 6, 2020. During that year, participants had access to the user interface and could adjust their bid price bids; the latter was disabled during the month of April, as described below. An overview of the course of the study is given in Figure 7.2. The study was conducted in collaboration with the local utility provider. Together with the academic researchers, they selected and recruited the participants from a neighborhood with a high penetration of residential PV panels and served as a trusted local point of contact. The participant sample ($n = 37$) included 36 residential households and one retirement home for elderly people. The majority of the participating households ($n_p = 31$) either already had solar panels on their own roof or owned a share of a solar panel on the roof of their apartment building prior to the study; these participants were hence considered prosumers. The aggregate peak PV capacity was around 280 kW. In addition, $n_b = 7$ prosumers owned (a share of) a home energy storage system.




Figure 7.2: Course of the field study during 2019. The bidding in the discriminatory double auction was interrupted by a one-month experiment in which participants faced a dynamic uniform price. Reports on energy consumption and expenses/revenues were sent out every month by email.

Among the 37 participants, 28 registered on the web application to actively participate in the market. The remaining 9 households signed up for the study, but never registered

on the web application; thus, they were ‘passive’ traders whose price bids remained at default tariffs throughout the study. In addition to the real-time information on demand and supply on the web application, participants received a monthly report summarizing the information available on the web application. It included their energy consumption and production, resulting expenses, share of local energy supply and the average price they incurred for local energy for the last month. This report also stated the participant’s average price bids, as well as the average price bid by all of the market participants. Participants who had often failed to purchase electricity from the neighborhood when there was local supply on the market in the past month received the note stating, “X% of your electricity demand could have been met by local electricity. A local trade did not take place, as your bid price was too low/the producers’ bid price was too high to be matched.” Likewise, participants who defined a very high ask price received the respective message if they had often not sold on the P2P market even when there was electricity demand on the local market. The report was sent out at the end of each month via email. After six months and twelve months (i.e. in July 2019 and January 2020), participants received a bill sent out by the utility provider, which contained a financial summary of their local trading activities based on the data collected by the research team. The bill contained the same information as the monthly reports aggregated for six months, along with a payment slip to settle the resulting amount due.

The electricity bill was based on the prices and trades arranged based on their bids in the double auction. Their bidding behavior thus had actual monetary consequences, which had been communicated to participants at an information event prior to the study (attended by 29 out of 37 households) and on the user interface. In addition, all participants had signed a letter of consent in advance. This sets the present study apart from prior research, which oftentimes relied on survey results or laboratory studies on consumers’ willingness to pay for renewable or local energy. The study design allows for eliciting price preferences from actual trading activity over the duration of an entire year.

An in-subject experiment with a different pricing mechanism examined participants’ preferences and the pricing mechanism more in detail. In April 2019, the price-setting function of the sliders on the web application was disabled for one month. It was substituted by a simple pricing function that determined a dynamic, uniform market price for local solar energy in real time (every 15 minutes) based on current supply and demand of (solar) electricity. This pricing function was derived based on the trading prices observed during the active bidding phase (i.e. Jan–March). That way, the average price levels did not change during the experiment, hence avoiding the risk of biasing participants with new price levels. After one month, the system was switched back to the auction mech-

anisms, and the price sliders were enabled again. Prior to the one-month experiment, participants were informed that the research team was testing another way to determine prices on the P2P market during the month of April; while prices for local electricity would be determined centrally during that period, participants would still be able to see their resulting income and costs of local electricity trades on the web application. After one month, participants were informed that the manual bidding function with the price sliders would be reactivated.

7.3.3 Data

The one-year data collection resulted in an extensive data set containing cleaned 15-minute load profiles, auction clearing and transactions from more than 35,000 periods, and price preferences bid by the participants. The system measured electricity demand of 37 residential customers from a town in Switzerland and generation profiles from 31 of these households. Moreover, it collected the price preferences that participants revealed using the web application, as shown in Figure 7.1. In addition, demand profiles of residential customers of another utility provider served for additional analyses. A summary of the load profiles collected is provided in Table 7.1.

Sample	n	Time frame	Mean energy consumption/production p.a. (sd) [kWh]
Field Study Cons	37	full year 2019	11,737 (30,293)
Field Study Prod	31	full year 2019	6,173 (17,098)
Additional Cons	223	full year 2018	4,488 (2,725)

Table 7.1: Summary of the load profiles collected. All data collected in Swiss residential households at a 15-minute resolution.

Furthermore, pre- and post-experimental surveys gathered supplementary information on participants' preferences and their socio-demographics, complementing the bidding and trading data. Among the $n_{s1} = 32$ participants who filled out the pre-experimental survey, 30 were male, 2 were female, and the average age was 55.2 years (sd 12.9). Participating households mostly included couples and families with one or two children (in total 26/32). Average household size was 2.9 people (sd 1.19), and the nursery home had over 100 habitants, for which one representative interacted with the application and filled out the surveys. 21 of the survey respondents were employed, 10 were retired and 1 described herself as 'stay-at-home mum'. Most of the respondents had lived in the region for a long time; < 20% moved there only within the last 10 years.

7.4 Experimental Results

This study creates a unique opportunity to study a peer-to-peer electricity exchange in the field and to gather longitudinal data throughout all seasons, collecting real electricity consumption and production data, as well as real behavior of participating households in terms of their trading behavior and price preferences.

As Figure 7.3 shows, trading on the P2P market considerably almost doubled both the ‘self-sufficiency rate’ and the ‘self-consumption rate’ of the community: The self-sufficiency rate refers to the share of electricity consumed that originates from one of the PV panels in the community. Without the P2P market, the community would have had a self-sufficiency rate of 21%, resulting from prosumers who consume electricity from their own roof (green block in figure). With the P2P market, the community’s collective self-sufficiency rate increased to 39% – the delta resulting from electricity purchased from other members of the community (turquoise block on top). The self-consumption rate refers to the share of locally produced energy, which could be sold within the neighborhood instead being sold at feed-in tariff. A third of the community’s energy production was consumed within the household that produced it. With the P2P setting, the share of electricity that stayed within the community increased to 60%⁴. When buy price bids did not match prices asked by prosumers, both parties traded with the utility provider, which creates an inefficiency, as there is concurrent consumption and production on the P2P market that is not matched locally. Such inefficient trades make up 2% of the total consumption.

Overall, the self-sufficiency and self-consumption rates achieved in the field study are in line with simulation studies on P2P energy markets, which project numbers in the same range (Griego et al., 2019; Mengelkamp et al., 2017b). The fact that there is still a relatively high share of solar energy sold to the utility company indicates that the share of prosumers is very large, as compared with pure consumers in this sample of households; hence, in the majority of ‘sunny hours’ (i.e., the time period during which the market is active), the supply on the P2P market exceeds local demand.

7.4.1 Observed Bidding Behavior

In the pre-experimental survey, participants of the field study were asked whether they would be willing to pay a price premium for locally generated solar energy as opposed to their standard electricity tariff ($t_u = 20.75 \times 10^{-2}$ CHF/kWh). 32 of the 37 study participants replied to the survey; most of them stated that they would indeed be willing to pay

⁴A more detailed explanation and evaluation of these metrics is provided in the article presented in Chapter 6.




Figure 7.3: Energy sourcing during the field study (entire year 2019), rounded to MWh. Self-consumption in prosumers own houses in green, P2P trading in turquoise. P2P trading increased self-sufficiency rate by 18 percentage points and self-consumption rate by 27 p.p. on the collective level.

a price premium, and these participants indicated a WTP an average increase of roughly 13%. These findings are in line with existing literature in which survey participants state an increased willingness to pay for local/renewable energy (Ecker et al., 2018).

During the experiment, participants had the chance to define their willingness to pay/willingness to accept using price sliders in the web application. Most participants used this functionality and changed their price settings, thus overruling the default bids of $p_b = t_u = 20.75 \times 10^{-2}$ and $p_s = f_u = 9.79 \times 10^{-2}$ CHF/kWh. Figure 7.4 shows the averages prices bid by the participants over the entire year on the left and resulting prices paid by consumers on an average day on the right. The average price bid for buying electricity on the local market was 18.48×10^{-2} CHF/kWh (sd 2.70), and 13.01×10^{-2} CHF/kWh (sd 3.38) for selling.

The bidding data reveals a different impression than the survey responses: Participants defined almost exclusively prices below the utility’s electricity tariff in the web app; thus hardly any participants were offering a price premium in the real market. And conversely, prosumers usually defined ask prices above the feed-in tariff offered by the utility (f_u) for selling electricity on the local market. There is no evidence for a willingness to earn less than the feed-in tariff in order to supply other households of the community.

Based on these bids, trades among peers cleared at an average price of $\bar{p}_t = 15.65 \times 10^{-2}$

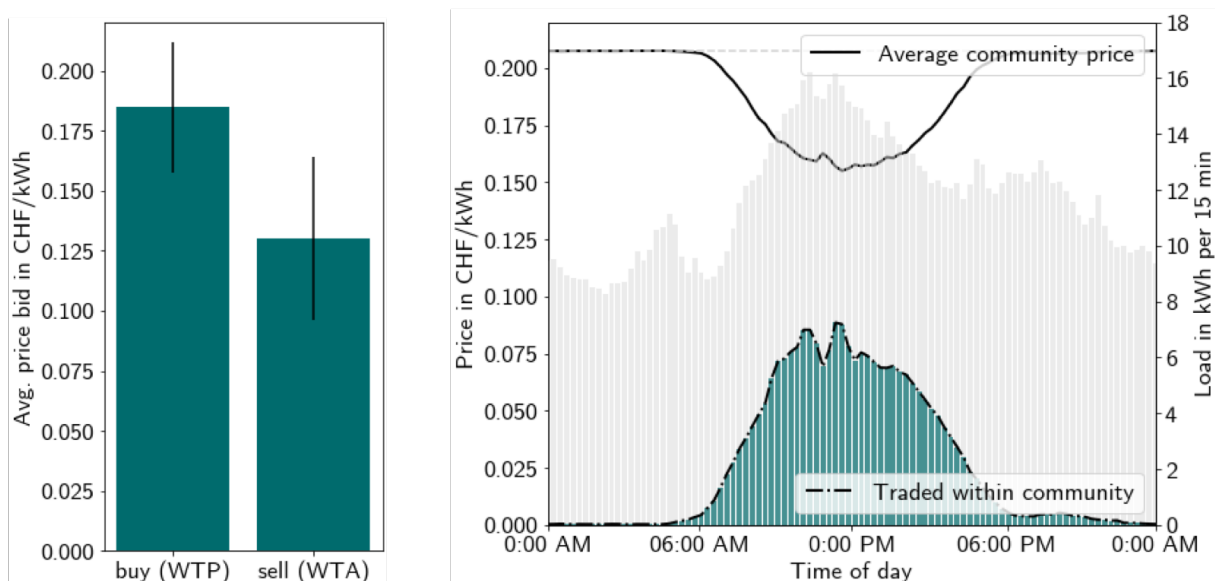


Figure 7.4: Bids and prices observed in the field. (Left graph:) Participants bid on average 18.48×10^{-2} CHF/kWh to buy and 13.01×10^{-2} CHF/kWh to sell energy on the P2P market. (Right graph:) Resulting market prices paid by consumers on a typical day in CHF/kWh. Prices drop during midday when solar energy is traded within the local market (solid line).

CHF/kWh (sd 1.86). The price for trading within the community was thus almost 25% lower than the utility tariff. The right hand side of Figure 7.4 depicts average daily electricity prices for consumers on the market. These prices depicted are the prices consumers in this market incur on average – namely a weighted average of the prices among peers resulting from the auction mechanism and utility tariff $t_u = 20.75 \times 10^{-2}$. The black curve shows that average prices dropped during the day when solar power is available hence creating an incentive for consuming solar energy when it is available. Participants incurred a total average price per kWh of $\bar{p}_c = 19.01 \times 10^{-2}$ CHF (sd 2.65).

Looking at participants individually, most prosumer households stated a buy price higher than the price to which they were willing to sell. As Figure 7.5 shows, most households' average bids are in the same range, yet they are not on the identity line. This indicates that $WTP \neq WTA$ in this data set, in contrast to the findings of Ecker et al. (2018), whose survey participants were willing to buy and to sell at roughly the same price. One plausible interpretation is that prosumer households will be buyers in time slots in which there is an undersupply of solar electricity on the local market, so one could argue that their buy and sell bids are not for identical goods: while their bids for selling electricity affect sunshine hours (at which time supply typically exceeds demand), their bids for buying electricity affect hours when supply is low. Second, the electricity




Figure 7.5: Average price bids observed in the field study. Each dot represents one participant. Most combinations of buy offer and ask price deviate strongly from the identity line, which represents $WTP = WTA$.

tariff charged by the utility is also considerably higher than the feed-in tariff prosumers are granted (electricity tariffs charged to consumers include grid fees, taxes and duties in addition to the energy price paid to producers). These tariffs were set as defaults in the web application and most likely serve as reference point for participants. With that in mind, it seems understandable for them to have a higher willingness to pay than to accept. The plot also reveals a few outliers that ask for a higher price than they would be willing to give. This seems counter-intuitive, but it might reflect an endowment effect (Kahneman et al., 1990; Shogren et al., 2001).

Again, Figure 7.5 shows that (also on the individual level) participants did not live up to their survey statements about offering a price premium for local solar energy. On the other hand, there is no clear evidence cost minimizing behavior in a competitive market sense either: Assuming individuals do not care about the origin of their energy supply, their utility function is solely defined by the cost for energy they pay, which results in prices equal to marginal costs of production in the long run. Thus, strictly profit-maximizing sellers would reduce their ask price p_s to the feed-in tariff, which represents the marginal cost in the present setting (i.e., the opportunity cost of not selling to the utility provider). (While initial investments may be high, the marginal costs of renewable energy generation are generally low or close to zero (Koolen et al., 2017). In the case of solar panels, prosumers do not incur marginal costs for an additional unit of electricity being produced on their roof.)

7.4.2 Evolution over time

Over the course of the study, there is a highly significant decreasing trend in buy price bids at a -0.08×10^{-2} CHF/kWh (sd 0.1×10^{-4}) change per month, and an even more pronounced, highly significant decrease in sell price bids: -0.16×10^{-2} CHF/kWh each month (sd 0.11×10^{-4}). Table 7.2 summarizes monthly average price bids (\bar{p}_b, \bar{p}_s) and illustrates an underlying decreasing trend in price bids. As the average bids illustrate, the decline in sell price bids is even more pronounced than in sell prices, as the average sell bid in December, $\bar{p}_{s,Dec} = 12.78 \times 10^{-2}$, is 12% below the average bid in January, $\bar{p}_{s,Jan} = 14.39 \times 10^{-2}$, whereas this decrease was only 4% on the buyer side.

Figure 7.6 depicts the average buy and sell bids over time aggregated over all participants as well as the daily average market prices for locally produced solar energy. The latter are not only subject to participants' price bids; they are also a function of the energy produced and consumed on the market. As to be expected by standard market logic, the matched prices for solar energy decrease over the summer months, and increase in the winter months, when supply of solar energy is limited and demand for electricity is higher⁵. Moreover, in summer, average prices drop below the average ask price of the prosumers. This can be explained by the large share of prosumer households in the participant sample, which leads to an over-supply of solar power during very sunny hours. In those situations, only the prosumers with the lowest ask prices are matched with the few net consumers who actually have electricity demand in these instances – namely consumers have a considerably higher market power during sunny hours.




Figure 7.6: Evolution of bids and resulting prices over the course of the study.

⁵In Switzerland, AC usage is less relevant than space and water heating.

7.4. Experimental Results

	Jan.	Feb.	Mar.	Apr.	May.	Jun.	Jul.	Aug.	Sep.	Oct.	Nov.	Dec.
Mean buy bid	19.37	19.27	18.72	18.12	18.53	17.79	17.85	17.88	18.04	18.35	18.59	18.66
\bar{p}_b (sd)	(1.92)	(2.31)	(2.91)	(3.28)	(2.83)	(3.27)	(3.15)	(2.66)	(2.71)	(2.43)	(2.28)	(2.16)
Mean sell bid	14.39	14.01	13.64	13.43	13.46	13.11	12.63	12.45	12.32	12.52	12.46	12.78
\bar{p}_s (sd)	(4.31)	(3.75)	(3.72)	(3.60)	(3.57)	(3.16)	(3.17)	(3.01)	(2.95)	(3.18)	(3.05)	(3.12)

Table 7.2: Average price bids per month in 10^{-2} CHF/kWh. Both buy and sell price bids decrease significantly over time.




Figure 7.7: Number of price changes per month. Inefficiency in the market (dotted line) is reduced as price changes decrease.

Figure 7.7 depicts the number of price changes as well as the inefficiency occurring due to mismatching price bids on the market. While it is not surprising that most price changes occur in the first months of the study, it is remarkable that there is an increased bidding activity over the summer months. At the same time, there is a peak in allocative inefficiency on the market due to mismatching bids. While increased bid changing in May and June is likely to have caused some of the inefficiency during this time, it is important to note that these are also months with high solar production: The more supply there is, the deeper the order books can be matched. This means that the buy price bids too low to be matched to sellers in June may have been this low all along, but before, there was not enough supply on the market to even ‘reach’ these buy offers and result in allocative inefficiency. Still, inefficiency decreases after June and from September on: it is lower than it was from January to March.

When analyzing participants’ individual bid curves (see Appendix, Figure C.5), an initial observation is that the bidding activity is relatively high. All the more so, as energy is a low-involvement good and the potential financial benefits are relatively low in absolute terms. Many participants did not stick to the default setting, but defined different price bids (p_b, p_s). Among the 28 participants who registered for the web app, 18 defined their own price bid at least once during the field study (64%). More surprisingly, many of these participants returned to the web application at least at some point to adjust their bids: In fact, 12 participants (43%) changed their bids at least twice, and 7 participants (25%) changed their bids more than 3 times. Some of them adjusted their bids frequently,

seemingly following a seasonal pattern – the two most active traders adjusted their price levels more than 50 times each over the one-year duration of the experiment. Overall, bidding activity in form of price changes decreased over the course of the study. A more detailed analysis of participants’ perception, and interaction with the user interface is available in (Ableitner et al., 2020).

A qualitative analysis of individual bidding behavior reveals that 10 (15) bidders reduced the buy (sell) prices they bid over time. Furthermore, among the buy (sell) prices bid, 9 (2) participants seemed to follow a seasonal pattern, meaning that they bid lower prices in the summer than in other months. Both of these patterns are incorporated in the aggregated average bid curves depicted in Figure 7.6. In addition, there are a handful of bidders who changed their price settings over time, but did not follow an obvious pattern, but seemed rather erratic. Interestingly, there is no evidence that show that the monthly reports or semi-annual bills, which included the average buy and sell bids among all participants, induced participants to adjust their bids.

7.4.3 Autonomous Control

The idea of P2P energy markets is, among other factors, driven by the aim to actively engage end-consumers in the energy market and raise understanding and interest in the topic of energy supply (Hahnel et al., 2019; Mengelkamp et al., 2017a). However, active bidding in an auction mechanism may overburden residential households over time. There is a trade-off between active integration of human decision making and handing off control (Dietvorst et al., 2016) to a software agent or even a central algorithm to clear the market. A within-subject experiment with an alternative market mechanism deployed during one month of the study (April 1 to April 30, 2019) evaluates how the participants perceived the possibility of defining prices for the P2P market themselves and thus of taking an active role in the market, as opposed to relying on an automatic system that determines a uniform price for all market participants. After disabling the price-setting function on the web application during that time (see also Section 7.3.2), participants were asked about their preferred pricing mechanism in a short survey in May and again a few months later in the post-experimental survey. In May, around half of the participants who had filled out the survey ($n_{s2} = 24$) preferred to manually bid prices for the auction to the automated system with a uniform price (see also Ableitner et al. (2020)). After an entire year had passed, in January 2020, only 21% favored the auction with manual price setting, whereas the majority of participants preferred an automatic pricing mechanism (Figure 7.8). In addition, it seems reasonable to assume that participants who did not reply to the

surveys were less engaged in the P2P energy market and would hence be likely to favor an automated pricing mechanism, as well. Further, participants were asked in the post-




Figure 7.8: Survey on price bidding vs. automated pricing. Many participants appreciated the possibility of bidding prices in the auction mechanism in the first months of the study, but at the end of the study, the majority expressed a preference for an automated pricing mechanism.

experimental survey whether they would like to use an autonomous agent to do the energy trading on their behalf (agreement on a 5-point Likert scale). The average answer was 4.16: they rather agree that they would like to use such a system ($n_{s3} = 19$). Moreover, the respondents stated that they would trust such a system to act in their best interest (4.11, rather agree). Still, on average they stated that they would still be using the web application of the platform just to check their own energy data (3.84, rather agree).

7.5 Further Analyses

Simulation of Benchmark Strategies:

A simulation of P2P energy trading was used, to put the results observed in the field study into further context and evaluate the sensitivity of the market outcomes of this market design to bidding behavior. Using simulations to understand market mechanics is often proposed in the recent literature on behavioral economics and smart markets (Bichler et al., 2010; Gode and Sunder, 1993; Ketter et al., 2013; List, 2004).

Leveraging a larger variety of real-world load profiles, the simulation model allows to examine how the observed bidding strategies relate to intuitive benchmark strategies and what effects these had on market outcomes. The simulation is based on electricity demand profiles from a distinct, additional set of Swiss households from another rural area, hence removing particularities in the load profiles of the study sample for this analysis. To put the bidding behavior observed in the field study p_b , p_s into context, two benchmark strategies are modelled that span a spectrum from cost-oriented to green preferences:

- a) **Cost-Minimizer Strategy:** As described above (see 7.2.2), the equilibrium price for solar energy under perfect competition would be opportunity cost.⁶ Hence, $p_{b_r} = f_u + \epsilon$ with $\epsilon = 0.001$ are defined as cost-minimizer strategy for the buying and $p_{s_r} = f_u$ for the selling side, i.e. assuming individuals have no preference for green electricity sourcing.
- b) **Green Strategy:** In the pre-experimental survey, as well as in the related literature, many consumers state a higher willingness to pay for local or renewable energy. Bid prices for buying in the ‘Green Strategy’ are defined to be $p_{b_g} = t_u + 0.10t$. $p_{s_g} = p_{s_r}$ remains as observed in the field.

Figure 7.9 shows the average prices for buying electricity on the P2P market on an average day in September. The figure illustrates that prices observed in the field study (not surprisingly) resulted in between what would result in scenarios a) (blue line) and b) (green line). However, the simulation also shows that the difference between the observed behavior in the field (black line) and the Green Strategy is smaller. In turn, there is a relatively large margin to the Cost-Minimizer strategy. In the simulated model, average prices for consumers in scenario a) result in 16.26×10^{-2} CHF/kWh (sd 5.29), for scenario b) in 18.83×10^{-2} CHF/kWh (sd 2.79) and in 18.16×10^{-2} CHF/kWh (sd 3.37) sampling from the bids observed in the field. Remarkably, there is hardly any difference in the volumes traded in different strategies which indicates the high level of efficiency achieved by participants in the field.

Explanatory Modelling:

While the sample size of this study is too small to solely focus on statistical significance testing, a number of explanatory models are computed to better understand the bid choices that participants made. In Appendix C.2, a fixed effects model explores the average prices bids per month per individual. While caution is warranted in interpreting these results, the analysis indicates some interesting trends: The average prices bid by other participants which were reported in monthly reports show a highly significant effect on individuals’ buy and sell price bids. In addition, the sell prices bid, reacted to the energy demand on the market, and display a significant relationship with the share of solar production sold on the P2P market in the previous month.

⁶While there is no perfect competition in the field study, in practice, the P2P market would be open for other prosumers to join; beyond that, even in the setting observed in the field study, there is affluence of solar energy during most days as soon as the sun is shining. It would thus be the economically rational strategy for consumers who do not care about the energy source to bid just above opportunity costs – feed-in tariff, during the vast majority of time slots when solar energy is available at all.




Figure 7.9: Simulation of different bidding strategies, aggregated to an average day in September. Average prices in the observed empirical behavior fall in between the two extremal cases, volumes traded are almost identical.

Overall, the results of the model reveal that (1) the effects of market information on sell prices were larger in size than for buy prices and (2) the explanatory power of these variables is much higher for sell prices. This matches the observation that participants changed their sell price bids more often and the fact that there are many prosumers in the sample of participants for whom the selling side of the market is more relevant. The fact that the share of previous P2P sales shows a highly significant effect on sell prices bid indicates that participants learned from their previous behavior and reduced their sell bids when they were successful in previous auctions. These findings are in line with related research on learning from experience in sequential bidding in online auctions (Bapna et al., 2004; Goes et al., 2010), but might not have been expected in a low-involvement, low-incentive domain like the present application.

A further examination of demographic and personality-related items did not reveal any significant effects on the observed bid prices. As already reported above, many participants had stated a willingness to pay a price premium for local solar electricity, which did not translate to bid prices above the utility tariff. Still, these participants bid an average buy price 2 Rp/kWh above those who did not state they would be willing to incur a premium on solar electricity. The difference is not statistically significant though, and neither were effects of other items related to participants' attitude towards environmental topics.

7.6 Discussion & Conclusion

7.6.1 Discussion

In multi-unit, online auctions with real humans, many theoretical predictions for bidding behavior have shown to be flawed or not to provide sufficient explanatory power for the way individuals behave (Bapna et al., 2004; Goes, 2013). Therefore, in line with Bapna et al. (2004) or Lu et al. (2016), this study takes an inductive approach and observes behavior in this societally relevant context in a real-world setting, and tries to make sense of these observations. Given the exploratory nature of the study, an ex-post analysis of the bidding behavior observed is conducted, rather than testing of a preconceived theory (Bapna et al., 2004).

Contrary to participants' statements in the pre-experimental survey and contrary to existing survey studies evaluating responses to hypothetical P2P scenarios (Ecker et al., 2018; Hahnel et al., 2019), participants did not offer a price premium for buying on the local market instead of from the utility provider. This is remarkable, as participants' responses in the pre-experimental survey had suggested preferences for green and local electricity generation and that they would perceive an increased utility in buying electricity from local and/or renewable resources. On the other hand, participants also did not act in a purely cost-minimizing way – bidding at a level of marginal costs (in this particular case, opportunity costs in form of the feed-in tariffs). The bidding behavior observed falls in between the spectrum ranging from acting upon pro-environmental preferences to a purely monetary oriented cost-minimizer – with a trend towards the cost-minimizing competitive equilibrium: Participants decreased their bids after winning an auction, hence learning from previous experience on the market (Appendix C.2). It is indeed remarkable that participants in this experiment seem to have understood the basic market logic and implications on electricity supply and demand quite well. Not only the decreasing bidding trend based on prior trading experience and efficiency provides evidence for this, but also the participants' basic market understanding revealed in the post-experimental survey. It is remarkable that the explanatory model with solely market information as independent variables reveals some effects that follow economic the profit-maximization rationale and explain a reasonable degree in the variation of sell bids. In contrast, survey items on sustainability-related topics were not reflected in the average bids observed. This null-finding seems relevant, as it illustrates that statements made in surveys specifically pointing out 'green' or 'sustainable' electricity schemes may deviate from individuals' behavior in an actual market setting.

It is difficult to fathom what ultimately caused the discrepancy between participants'

stated preferences for local solar energy and their bidding behavior: on the one hand, it is conceivable that participants' answers in the pre-experimental survey (and in other related research) were subject to social desirability bias, leading them to indicate a higher preference for local solar energy than they actually have. On the other hand, it is also plausible that they were caught by the competitive notion of known market logic and did not actually gauge that their bidding behavior was not leading to an active support of renewable generation in the long run. Either way, the decline in price bids and increasing efficiency could indicate the *tâtonnement* (gentle convergence) to the competitive market equilibrium (Gode and Sunder, 1993), in which the price for solar energy is close to the external feed-in tariff.

The intention–behavior gap is a known phenomenon established in several studies on resource conservation (Gatersleben et al., 2002; Tiefenbeck et al., 2018a). The presented evidence also shows that in the market context, statements made in surveys or in hypothetical markets may not be representative of the way in which individuals will act once they find themselves in a real market setting.

Overall, participants seemed to appreciate the possibility of influencing pricing on the P2P market, mostly in the first months of the study when they actively used the bidding function. However, after the year had passed, the survey results imply that participants would be willing to hand over control to an automated agent acting on their behalf (Bapna et al., 2004). In line with that, recent research on ‘algorithm appreciation’ challenges the widespread belief that individuals do not like to rely on algorithms and in fact appreciate advice, even if it comes from black-box algorithms (Logg et al., 2019). The decline in the initially high bidding activity over time also supports this tendency. Smart trading agents for P2P energy markets might be a valuable approach allowing for active preference elicitation from participants without relying on their continuous interaction. Autonomous agents pointing out the effect of price preferences for individuals and engaging them in smart energy markets may help individuals make choices they truly prefer in the long run. Using autonomous agents to point out environmental consequences of market interactions by factoring in long-term external costs of alternative energy production may be one way out of the zero-marginal-cost dilemma that renewable generators face, without going back to top-down control of externally defined prices.

7.6.2 Limitations

Despite all best efforts, given the technical and operational complexity in conducting a field study of this character, there are some limitations to the findings of this study. Most

critically, the sample of 37 participating households (including one elderly residence) is not representative of the larger population. All participants live in the same town and are customers of the local utility provider (blinded for review), whose support and active role has been vital to the launch and success of the project. Moreover, given that the experiment had an impact on participants' electricity bill, they had to opt into the study at the invitation of their utility provider. The results may thus be subject to volunteer-selection bias (Tiefenbeck et al., 2019). Hence, the sample size and recruitment does not provide a sufficient basis for drawing conclusions about explanatory variables for bidding behavior among the broader public. Additionally, the ratio of pure consumers and prosumers is not representative of most neighborhoods today, and the imbalance of electricity demand and solar supply during sunny hours may have biased participants' behavior.

Future studies need to evaluate the extent to which results can be replicated with larger, more diverse samples and in other locations. However, this study is probably the first to observe actual price-bidding behavior with real financial consequences for participants in a real-world local energy market. To control for special traits in the consumption profiles of the study sample, which influence local market prices as well, additional simulations leverage electricity consumption profiles from another, unrelated sample of 223 Swiss households. Further research should examine bidding behavior beyond the prices defined by participants, and also investigate whether load shifting activities result as a consequence of local real-time pricing (i.e., potential adjustment in bid demand). To that end, a considerably larger sample and a control group will be required to control for unrelated shifts in energy demand.

A further limitation of the findings is that the behavior observed is subject to the external market conditions – in particular, to the tariffs t_u and f_u participants face when buying from or selling to the utility provider. In particular, dynamic external pricing might impact bidding behavior in a more fundamental way than shifts in fixed tariffs. An important research topic for the development of P2P energy markets will thus be the economic environment in the form of external tariffs.

Moreover, it is important to note that the regulatory conditions for implementing P2P energy markets in the real world are outside of the scope of this article. This work focuses on the bidding behavior and potential application for intelligent agents to increase the understanding of potential effects of such novel market structures.

7.6.3 Conclusion & Outlook

This study presents a unique field experiment on a local P2P energy market, in which 37 participating households take part in an interactive double auction for solar power generated within the neighborhood. Price preferences are elicited by having participants bid their willingness to pay and their willingness to accept for trading solar power on the local market. The evolution and learning effects in participants' bidding behavior is observed throughout an entire year. This represents the first empirical evidence on participants bidding behavior in a P2P energy market of this kind.

The results presented indicate that participants' bidding behavior does not reflect their previously stated intentions regarding local solar energy, thus contributing to empirical studies in other contexts of pro-environmental behavior that have also reported an intention-behavior gap. Moreover, the behavior observed indicates that residential consumers do a reasonable job in approaching competitive market price and that they understand standard market logic relating supply and demand to resulting prices. The decreasing price bids, despite 'greener' intentions, suggest that decision support may help individuals to actually live up to their eco-friendly intentions. Still, the empowerment of consumers and prosumers through active market participation leads to a welcome engagement in the energy market. This will likely become increasingly relevant to individual consumers, given the imminent rise in household expenses on electricity that will result from the electrification of transport and heating.

These findings contribute to the research on behavior in online auctions and point out domain-specific outcomes for the energy sector. P2P energy markets emerge as a possibility to coordinate and incentivize renewable energy resource, providing a promising alternative to top-down regulations and government subsidies (Siler-Evans et al., 2013). Going forward, the design of smart agents acting on behalf of individual consumers should incorporate individuals' stated preferences. As pointed out by Bapna et al. (2004), the identification of different types of bidders and understanding their motivational drivers can improve the design of smart bidding agents that can act in line with the individuals' interests and preferences in the long run. Similar to 'Robo-advisors' in the financial industry (Logg et al., 2019), smart trading agents or advisors could be developed to make sure individuals are aware of and approve of the long-term effects of their energy sourcing choices and consumption decisions.

8. Intelligent Agents for Smart Load Scheduling

8.1 Motivation

Passenger cars are the number one CO₂-emitting end use in many industrialized countries (residential appliances plus space heating are number two) (International Energy Agency, 2019). Hence, the electrification of transportation in combination with renewable electricity generation is instrumental to reduce greenhouse gas emissions (Williams et al., 2012). Williams et al. (2012) name decarbonized electrification as *the* “technology path to deep greenhouse gas emissions cuts by 2050”, p.53. However, the electricity demand that results from increasing electrification and from the higher diffusion of distributed generation (see Chapter 7) is challenging the stable operation of power grids. Higher demand peaks and volatility of generation put a strain on the grid infrastructure and may require costly balancing interventions (Fridgen et al., 2016). In addition, there is often a mismatch between demand curves of, for instance, electric vehicles and the volatile supply of wind or solar power (Ramchurn et al., 2012; Valogianni et al., 2020). To optimally capture the value of renewable power when and where it is produced, flexible demands like charging loads for electric vehicles or heat pump cycles need to synchronize with the dynamics of renewable generation, while, at the same time, respecting physical grid limitations (Ketter et al., 2018).

To create incentives to shift flexible demand to times when electricity generation is cheapest, Schweppe et al. (1988) proposed spot prices providing feedback on electricity generation costs back in 1988. This has led to considerable research on the resulting

behavior on the demand side. Numerous studies have investigated ‘demand-side management’ or ‘demand response’ (Gillingham et al., 2015; Schmidt et al., 2015; Strbac, 2008; Valogianni and Ketter, 2016; Vázquez-Canteli and Nagy, 2019), and the intuition behind it is the same for taking advantage of carbon-neutral generation in real time. However, the success of manual demand response programs has been limited as they require consumers to actively change energy demand profiles (Bollinger and Hartmann, 2020; Vázquez-Canteli and Nagy, 2019). On the other hand, real-time pricing may not be precise enough in a future scenario with increasingly flexible electricity loads, as price drops which apply for a larger group of consumers may result in high demand peaks in their distribution grid (the electricity grid level to which residential households are connected) (Ramchurn et al., 2011). Existing research in this field indicates that the desired balance of energy supply and demand can only be achieved if demand-response mechanisms are automated and smarter market mechanisms are applied.

With the advances in distributed computation and intelligent devices, the coordination of loads using intelligent agents has gained increasing attention (Gottwalt et al., 2011; Peters et al., 2013; Robu et al., 2013; Valogianni et al., 2014). Complimentary to behavioral interventions to reduce individual resource consumption, information systems can also provide tools to enable the integration of renewable resources. The volatility in renewable generation and the interdependences of demand profiles on a collective level are too abstract and complex, and too dynamic for individuals to keep track of (Peters et al., 2013; Robu et al., 2013). Software agents using artificial intelligence (AI) can provide autonomous decision support to manage electricity consumption to electricity grid by flattening peak demands and by moving demand to times when renewable electricity is generated (Valogianni et al., 2020).

This chapter introduces reinforcement learning as a machine learning method for intelligent agents in the smart grid. Autonomous learning agents can serve as demand-side managers for utilizing renewable energy closest to when and where it is created, which reduces the need for additional infrastructure like storage systems or upgrades of distribution grids, and for costly grid balancing efforts (Fridgen et al., 2016). Intelligent agents that serve as aggregators for individual end-consumers or control flexible loads can learn from dynamic price incentives and plan load schedules. The following section presents the basic intuition behind reinforcement learning (the theory and formalization is described in more detail in Appendix D). Section 8.3 summarizes a simulation on multi-agent reinforcement learning for smart scheduling of EV charging, followed by an outlook on questions for future research.

8.2 Reinforcement Learning

The basic idea behind reinforcement learning is that an agent that interacts with an environment tries to maximize expected future rewards for her actions based on her experiences (Mnih et al., 2015).

Markov Decision Processes (MDP) are used as a formal framework for this type of decision making in AI as they provide a framework for simple representation of stochastic environments over time (Sutton and Barto, 1998). In an MDP, an agent interacts with an environment at discrete time steps (Sutton et al., 1999), Figure 8.1 illustrates this process on a high level. In each time step, the agent perceives the state of the environment and takes an action, which leads to a reward for the agent and influences the future state of the environment. If the agent is equipped with intelligence, she will try to learn if and in what way the environment reacts to her actions and which rewards result from each step she takes to improve her choices in future steps. The agent's intelligence can be represented by different machine learning models with varying complexity and performance, ranging from simple decision trees or regression models to deep neural networks. (Appendix D contains a formal and more detailed explanation of Markov Decision Processes, Multi-Agent Learning, and Deep Reinforcement Learning.)

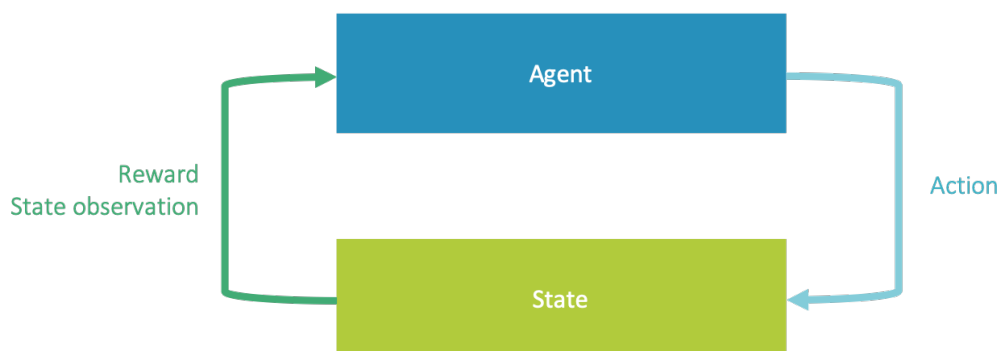


Figure 8.1: Simplified scheme of a Markov Decision Process (MDP). An agent observes a state, chooses an action, and observes the next state and reward she achieved with this action. On that basis, she decides for the next action while adapting her expected reward for her possible actions based on the past experiences (see Appendix D for more details).

Reinforcement learning is not a new concept, but computer scientists (Littman, 1994; Sutton and Barto, 1998), as well as economists (Erev and Roth, 1998), have formulated these learning processes inspired by the psychology of human learning already decades ago. However, the more recent improvements in machine learning and in computational power have increased the applicability of such models to real-world applications. Earlier studies

on reinforcement learning mostly focused on simple games, due to the clearly defined action and state spaces in (computer) games, and the resulting direct applicability of game theory (Erev and Roth, 1998). The difficulty of applying reinforcement learning in practice lies in the perception and representation of real situations (Mnih et al., 2015). To be able to learn from past experience, agents need to reduce the dimensionality of complex environments to a level which provides enough information to derive successful action policies, but is still computationally processable (Mnih et al., 2015). In recent years, machine learning models have improved significantly in that respect (Vinyals et al., 2019), driven, in particular, by the development of deep learning (Mnih et al., 2015). In deep reinforcement learning, the agent iteratively trains a deep neural network (in the blue box in Figure 8.1) based on the observed information about the environment and actions taken in the past. The neural net can process and abstract from a wider variety of input than simpler learning models, as it allows for non-linear function approximation (Mnih et al., 2015).

8.3 Summary of Article G) Multi-Agent Reinforcement Learning for Electric Vehicle Charging

This simulation addresses one of the use cases for intelligent agents in the smart grid, in particular, examining a multi-agent setting. Existing research shows that autonomous agents can learn from dynamic electricity prices and forecast demand to efficiently schedule electricity loads to reduce electricity costs for the users (Ketter et al., 2013; Ramchurn et al., 2011; Valogianni et al., 2020). However, if self-interested rational agents all receive the same price signal, real-time pricing may lead to ‘herding behavior’ creating high peak demands which can destabilize the grid (Fridgen et al., 2016; Valogianni et al., 2015). This issue illustrates that it is not sufficient to look at load shifting as an isolated problem for each individual consumer, but that the collective actions of load-shifting agents impact the electric loads in the grid. Electricity supply and demand need to be matched subject to grid constraints and to preferences of consumers and producers. Intelligent agents hence actually act in a multi-agent setting which needs to be coordinated efficiently.

Market mechanisms can be used to coordinate demand, in order to mitigate situations in which all agents defer their loads to the same time and thus create high peaks in power demand. For instance, Valogianni et al. (2020) present an adaptive pricing mechanism for EV charging. They propose a coordination mechanism in which a smart grid manager is learning a pricing strategy that incentivizes the desired collective demand profile. Given that actual costs for electricity consumption are not determined just by generation

costs, but also by distribution costs, Li et al. (2014) propose a distribution locational marginal pricing mechanism to alleviate congestion resulting from EV charging. Locational marginal prices for the next day are determined by a distribution system operator based on a power flow optimization. Charging of electric vehicles (EVs) is controlled by aggregators that minimize charging costs under this pricing mechanism and the constraint that EVs are charged within specific limits derived from driving patterns. Flath et al. (2014) argue that such nodal prices may be too complex and hard to understand for individual households. They thus propose area pricing zones instead, which follow the same logic in a lower granularity. Vytelingum et al. (2010) introduce trading agents for the electricity grid and compare different bidding strategies in a double auction with dynamic locational marginal pricing for electricity. In a simulation assuming fixed electricity demand, they find that agents bidding prices with an adaptive strategy outperform zero-intelligence agents that bid randomly in terms of resulting energy costs and that the grid is used more efficiently.

With the evolution of reinforcement learning in the past decade, adaptive strategies have become more sophisticated, as learning models allow agents to learn from past situations to improve their reactions to the information they observe. The simulation presented here employs learning agents for EV charging in a setting in which prices are determined by a locational marginal (or ‘nodal’) pricing mechanism, similar to the mechanism in Vytelingum et al. (2010), and uses real base-load data and simulated data for flexible EV demands. The mechanism creates incentives for agents to schedule the flexible charging loads such that it is most convenient for the grid while respecting user preferences.

8.3.1 Model & Simulation

An auction market allocates charging energy and determines nodal prices in the distribution grid. The auction is cleared every 15 minutes and agents submit a bid for electricity demand and valuation for every time step.

The agents in this environment are facing the task to minimize charging costs for the EV owner, while ensuring that the state of charge (SOC) of the EV batteries meet the driving needs at all times. Hence, an agent’s aim is to optimise the charging costs of their household by moving the charging process to cheaper times (i.e. more convenient for the grid), but still charging enough to fill the owner’s driving demand. In the present model, agents decide on these schedules individually without communicating driving patterns to a central aggregator. This represents a bottom-up coordination approach to the charging problem.

Market mechanism

In this model, multiple agents $i \in \{1, \dots, n\}$ in the same distribution grid face the task to minimize EV charging costs. The cost minimization problem is subject to the EVs' hourly availability to be charged at home and, naturally, to the battery capacities installed. In addition, agents incur a penalty if the SOC does not meet users' driving demand to prevent them from simply not charging. The optimization problem for an agent i is formulated as a maximization problem of aggregated rewards R_i (i.e., negative costs):

$$R_i = \sum_{t=1}^T r_{i,t} = - \sum_{t=1}^T \underbrace{p_{i,t} * P_{EV,i,t} * \Delta t}_{\text{Cost of electricity}} + \underbrace{\rho(1 - SoC_{i,t}) * \mathbb{1}_{t \in T_{D,i}}}_{\text{Penalty}} \quad (8.1)$$

$$SoC_{i,t+1} = SoC_{i,t} + \frac{P_{EV,i,t} * \Delta t}{Battery_i} \quad (8.2)$$

Here, $r_{i,t}$ represents the instantaneous reward at time t and is input to the reinforcement learning model. The price $p_{i,t}$ in the equation is determined by the auction mechanism that determines nodal prices based on optimal power flows in the distribution grid and on the bids entered by the agents. Prices are determined for each node in the distribution grid individually by the auction mechanism (details on this mechanism are contained in (Kienscherf et al., 2020)). In peak demand times, agents may have to pay a price premium on top of the standard electricity tariff, if they reduce electricity availability for other agents. The agents thus do not merely act as price takers like they do in traditional retail markets, but have a strategic choice in their bidding behavior.

Learning Model

In this simulation, agents leverage machine learning capabilities, as large parts of the environment are outside of the agents' own strategy space (other electricity loads, grid state), or even unobservable for them (others' charging demands). The MDP for the learning problem is defined by the state space \mathcal{S} , the action space \mathcal{A} , the transition function, $\delta(s, a)$, and the reward function $r(s, a)$ (Reddy and Veloso, 2011b):

$$M = \langle \mathcal{S}, \mathcal{A}, \delta, r \rangle \quad (8.3)$$

The action policy Π that is to be learned chooses the action that maximizes aggregate rewards in every state of the environment, i.e., $\Pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ (Sutton et al., 1999). The state space \mathcal{S} includes current market information like the real-time locational marginal

price at the agent's location, time of day, the current SOC, as well as the prices most recently paid.

The agents in this model have two decisions to make in each clearing period: Which price to bid and how fast to charge (i.e. how much energy to buy in the next 15 minute period)? In each auction clearing period, agents can choose among three price levels $\mathcal{A}_1 = \{low, medium, high\}$ and three charging powers $\mathcal{A}_2 = \{3.7, 11, 16.5\}$ kW to bid. The action space $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2$ thus consists of $|\mathcal{A}| = 9$ discrete actions. Depending on her charging decision, an agent is allocated energy (or not) by the auction mechanism and her consumption is priced. Accordingly, if the agent bids high during high demand times, she is allocated energy but has to pay a higher price as she imposes a charging reduction on other agents. The transition function $\delta(s, a)$ depends, on the one hand, on her actions taken, as these influence the charging state of her vehicle. On the other hand, numerous external interactions influence the market prices which are only observed through the resulting prices. Transition probabilities are thus non-stationary. The reward function $r(s, a)$ is determined by the market mechanism as described above, but not explicitly known to the agents.

To evaluate learning strategies to established benchmark strategies, this study compares the following strategies:

- a) **Benchmark Strategy:** No agent reschedules flexible loads, real-time prices are ignored (status quo). Agents start charging right when the EV is available for charging at home.
- b) **Learning Strategy:** The agent employs the deep Q-learning algorithm proposed by Mnih et al. (2015). Here, the optimal action policy $\Pi(s, a)$ is approximated by a small, one-layer deep neural network. The RL agent learns from the payments and penalty incurred in past periods and chooses bids on that basis (the learning algorithm is explained explained in Appendix D, Algorithm 2). The mechanisms by which prices in the distribution grid are computed (i.e., yielding the reward function) and the driving schedules are not explicitly known or modelled by the agent.

Data

Different data sets serve as base for this simulation: Residential electricity load data from Ireland is used to simulate the electricity demand of residential households (Commission for Energy Regulation, 2012). The original data was collected between July 2009 and December 2010 and comprises load profiles from 4232 households. The sampled load profiles are mapped to a distribution grid from an IEEE test case, which is representative

for a European residential distribution grid. Driving data is simulated, given the lack of empirical data on load profiles of electric vehicles. Arrival and departure times are sampled from normal distributions with 8 AM and 6 PM mean and 30 minute standard deviation, which caters to a pronounced, yet not overexaggerated demand peak in the evening and is in line with data from European Travel Surveys (Pasaoglu et al., 2012; UK Department for Transport, 2019; Wu et al., 2010). In addition, EVs are assumed to be mainly charged when at home.

The simulation spans one work week (five days) with 480 auction clearing periods and 100,000 training runs. Agents' performance is evaluated in on 1000 test runs.

8.3.2 Results

The results show that in this setting, agents using a neural network to learn reduce average charging costs by 68% as compared to the default strategy, in which they would just start charging as soon as the EV returns home. Especially in the evening hours, nodal prices observed in the grid for all participants are kept at the minimum, thus not indicating price premiums evoked by grid congestion. Moreover, many agents also reduce their charging demand to a medium rather than high charging speed. While the overall charging costs paid by the learning agents still exceed the absolute minimum by roughly 18%, the learning agents outperform the benchmark strategy currently employed for EV charging in practice by far (68%). Given that this is achieved in the absence of load control by a central aggregator and preserving the individual users' privacy on their driving schedules and demands, the application of intelligent agents in this type of market seems like a promising approach.

As related work illustrates, the application of intelligent agents in the smart grid may lead to undesired effects, if incentive structures are not fully thought through, e.g. herding effects as a result of learning real-time price. However, in an adequate market design, artificial intelligence can contribute to grid stability and optimize resource use, without the need for a central unit to control electricity loads. This aspect is particularly relevant in the light of increasing privacy concerns, as central load control requires individual households to communicate their electricity needs and EV usage schedules to a central point of authority. The combination of smart market design and intelligent agents provides an alternative solution for distributed load control, devoid of this need for information-sharing.

An extended version of this simulation and the results are described in more detail in Article G), (Kienscherf et al., 2020), an article that will be submitted to the Journal of

Artificial Intelligence in July 2020. As stated in the Disclaimer of this thesis, I am the second author of that article and have designed and drafted the study together with the first author, Philipp Arthur Kienscherf.

8.4 Outlook

The use case presented in this simulation is only one of the applications for intelligent agents in future electricity markets. As pointed out by Rogers et al. (2012) the transition the energy market is going through to adapt to increasing electricity demand and to more and more distributed energy resources is “one of the greatest engineering challenges of our day”, p. 2166. Intelligent agents can support this transition and help delivering the vision of a smart grid in a number of ways which entail fruitful avenues for future research: *Supporting consumers* in analyzing their own energy data and in interacting on smart energy markets is crucial. Individuals cannot follow the dynamics of renewable energy generation in real time. As the study presented in Chapter 7 showed, individuals are, however, generally interested in being integrated in the decision processes on the energy market. (Semi-)Autonomous software agents can elicit consumer preferences, and then act upon these in an autonomous way in the dynamic market environment, reducing complexity for the human consumer (Peters et al., 2013). As the simulation presented in this chapter (Section 8.3) shows, the tasks for such agents can go beyond bidding prices, to actively *coordinating flexible loads* in individual households. To take advantage of distributed resources, while respecting physical limitations of the grid requires smart mechanism design and real-time price incentives. Herein, it is important to understand the *interaction of intelligent agents in smart energy markets with dynamic pricing* also in the field. While AI capabilities can be examined in simulations at first, real-world studies for different use cases may also help to identify practical challenges in this context, which hamper the adoption of intelligent agents in the smart grid in practice (Ramchurn et al., 2012).

Furthermore, additional research for assessing the performance of different learning models for concrete use cases is needed. The simulation presented in Section 8.3 compares a very advanced learning model, namely a deep neural net, to the current default strategy. While the possibility of capturing input variables from the environment without having to pre-process them elaborately is an advantage of deep reinforcement learning, it is likely that a slightly simpler learning model could already achieve similar results in the proposed setting. As the complexity of learning models influences the requirements on distributed computational power, a careful evaluation of the learning models implemented by smart

grid agents is important. Further research will thus address this issue and assess the *trade-off between model complexity and performance improvements*. To that end, expanding the simulations presented above further, by using empirical data on different flexible loads like heat pumps or home battery storage, will also be relevant for practitioners.

9. General Discussion & Conclusion

The objective of this cumulative dissertation is to examine different ways in which information technology can foster sustainability with a focus on the individual consumer. The thesis started out by presenting the motivation and background for the studies presented in Chapters 3 – 8. The results of each individual study were explained and discussed in the corresponding chapters, highlighting implications for future research and for fostering sustainability, not only in theory, but in the real world. Chapter 9 now closes with a concise summary of those findings and synthesizes and discusses their implications for researchers, as well as for policy makers and practitioners.

9.1 Synopsis of Findings

The research gaps tackled in this thesis (introduced in Chapter 2.4) are threefold, and they introduce different approaches to leverage information systems for sustainability: Real-time feedback and incentives for resource conservation, P2P energy markets, and intelligent agents for load scheduling in smart energy markets. The findings regarding these research gaps are summarized in the following.

Real-Time Feedback & Incentives for Resource Conservation, Chapters 3 & 4

Previous research found personalized feedback to be effective in reducing resource consumption of individuals (Allcott and Mullainathan, 2010), in particular when pertaining to specific activities rather than consumption metrics aggregated to household level (Abrahamse et al., 2007; Tiefenbeck et al., 2018a). However, the drivers of these behavioral

changes induced by feedback interventions remained largely unclear, raising the following questions: *Does real-time feedback induce conservation effects even among a sample of uninformed individuals in the absence of monetary incentives? And do individuals define consumption goals when presented with real-time feedback on resource consumption?* Chapters 3 & 4 each present a large field experiment on a feedback intervention which focuses on these motivational drivers. Two distinct samples of participants were confronted with real-time feedback on resource consumption during showering. The main effects (energy conservation and the effect of self-set goals on energy conservation, respectively) are statistically significant in both studies. Robustness checks and additional analyses have increased the confidence in these results and underpin the quality of the study design.

The data collected thus provides robust empirical evidence showing that *a) individuals reduce their resource consumption while showering in reaction to activity-specific consumption feedback even if there are no monetary savings to be incurred, and b) these effects persist in a natural field experiment among a sample of hotel guests who did not self-select into a research study*, in contrast to many study samples in existing experiments on sustainable behavior (Chapter 3). Based on the results presented here, monetary incentives of the size of energy and water tariffs are not the main driver for resource savings effects and are not necessarily required in conservation programs – which increases the scalability and cost-efficiency of such measures.

In addition, the findings presented in Chapter 4 show that *c) individuals tend to set themselves consumption goals once they are provided with activity-specific real-time feedback. These self-set goals even were relatively ambitious and correlated with increased savings effects, which is in contrast to theoretical hypotheses proposed in goal-setting theory (Locke, 1996)*. Based on these findings, self-set goals could explain some of the effects of self-tracking enabled by personal IS which have been observed in many domains other than resource conservation, as well (Consolvo et al., 2009; Froehlich et al., 2010; Lupton, 2014). Given the difficulty of defining adequate goals externally and the risk of adverse reactions associated with external goals, these results question whether IS artifacts should assign goals to users and rather suggest that IS artifacts encourage users to set their own goals and provide functionality to store and display those.

P2P Energy Markets in the Real World, Chapters 5 – 7

The energy sector is undergoing a fundamental transformation to integrate renewable generators that are spatially distributed and often privately owned. To examine new mechanisms for coordinating these resources in the future, the second set of studies presented

in this thesis focuses on the design, implementation and impact of P2P energy markets, in which these smaller generators can interact and sell renewable energy to neighboring households. A lighthouse project funded by the Swiss Federal Office of Energy (SFOE) provided the unique opportunity to implement and study a prototype of such an energy market based on blockchain technology from the conceptual design, through to its impact during a 1-year long field experiment. The research studies around this project are of a more exploratory nature than the two preceding studies on feedback interventions, as the project is one of the worldwide first implementations of a P2P energy market.

First, the conceptual study on the characteristics of blockchain technology in Chapter 5 examines blockchain as an enabler for P2P markets to understand: *What are the benefits and risks of implementing P2P markets on a blockchain infrastructure?* A systematic literature review underlines that research on blockchain-based P2P markets is in a very early stage. *The identified advantages of blockchain technology as a platform for P2P markets are mainly the independence from a central entity, resilience of the distributed system, and – likely more relevant in the Swiss ecosystem – autonomous execution of market logic and settling of transactions in smart contracts. However, cost-efficient operation, usability, and governance of the decentralized system are still posing challenges to blockchain applications in practice.*

Chapter 6 presents the market design that was defined and implemented for the field experiment based on insights derived in the preceding chapter. Furthermore, the study examines data collected in the first three months of the experiment to understand: *Which value propositions do P2P markets create from the user perspective, and to what extent are they an effective measure to empower once passive consumers to assume a more active role in these markets?* This study provides unique empirical evidence on the user value propositions of P2P energy markets that had only been theorized in the literature thus far (Morstyn et al., 2018): *The auction mechanism implemented lead to real-time prices which accurately reflected the local availability of the solar energy, thus providing incentives for consuming locally generated solar energy. PV-owners sold a considerable share of their excess production locally, which illustrates that P2P energy markets can reduce uncertainty of returns for distributed renewable generators.* This first evidence further reveals that individuals interacted with the web application for the P2P market relatively frequently.

Focusing on behavioral aspects in more depth, the study presented in Chapter 7 examines the bidding behavior participants displayed during the one-year-long duration of the field experiment. Existing survey and lab studies have found that individuals stated a high willingness to pay for electricity if it is generated from renewable resources (Ecker

et al., 2018; Tabi et al., 2014). It was unclear, however, whether this willingness to pay will materialize in a real market setting, and whether individuals understand and are willing to deal with the complexity of a local energy market, which raises the questions: *Does the bidding behavior observed in the field deviate from cost-minimizing behavior? And how does it evolve over time?* While the pre-experimental survey gives the same impression as described in the literature (Chapter 7), i.e., participants stated a willingness to pay a price premium for buying locally produced electricity, the prices bid by participants during the field experiment with real consequences on their electricity bills conflicted with these statements: *The prices individuals offered for renewable energy on the P2P market were lower than the utility's standard electricity tariff. At the same time, PV-owners asked for higher remuneration than the feed-in tariff they were granted by the utility.* These results illustrate that individuals' behavior in a market context, in which bidding behavior translates to actual payments, may differ from survey statements on preferences for renewable energy. Furthermore, the data collected provides interesting insights from a behavioral economics perspective. *Examining the bidding behavior over time, the price bids indicate learning effects from previous periods (in line with Goes et al. (2010)) and reactions to seasonal changes over the course of the year. These resulted in an increase in efficiency on the market towards the second half of the experiment (see Figure 7.7). In addition, there is a decreasing trend in bids on both sides of the market (corroborating findings of McAfee and Vincent (1993)), which suggests an approach to the cost-minimizing equilibrium of a fully competitive market.*

Multi-Agent Learning for Load Scheduling, Chapter 8

As a final research project, the work presented in Chapter 8 provides an outlook on a promising future research direction: autonomous load scheduling by intelligent agents in a smart market. An agent-based simulation is used to examine the question: *Is multi-agent learning in a smart energy market effective in reducing demand peaks while still respecting individual preferences?* Chapter 8 introduces multi-agent reinforcement learning as a powerful computational tool to automate decision processes in energy markets that would otherwise overburden individuals. Dynamic and location specific market mechanisms can provide price signals to individual households (thus in an even higher resolution than in P2P markets) to relieve the grid infrastructure and capture the value of renewable energy. *The results of the agent-based simulation show that nodal prices provide precise enough signals to incentivize more balanced load profiles without requiring central load control by an aggregator collecting information from individual households. Using reinforcement learning to act upon these price incentives, agents shift loads for electric vehicle charging –*

which reduces peak loads in the grid. Given the forecasted surge in electric mobility, these insights are relevant to the design of smart energy markets that can sustain the resulting increase in electricity demand. Future work might integrate this type of intelligent load scheduling, for instance, in a P2P energy market to address both, consumer engagement, as well as grid stability and efficient infrastructure use – thus tying together some of the aspects of the studies presented before.

9.2 Contributions & Implications

Synthesizing the findings of the different studies presented in this thesis, several common denominators emerge. While the studies tackle different concrete research gaps addressing behavioral factors and the required technological facilitators, there are several recurring themes that shape the contributions of this dissertation. The following paragraphs summarize the key messages and contributions, and point out their implications.

Unique empirical evidence expands existing theoretical research and lab experiments on consumer behavior: Examining consumer behavior in the field is crucial for assessing the realistic impact of conservation programs or policy measures, as human behavior is influenced by a complex set of psychological, economic, as well as socio-cultural factors, and is highly context-dependent (Levitt and List, 2008). Experiments in behavioral economics research have revealed biases and inconsistencies in human behavior and decision making in various domains, which have primarily been studied in lab settings due to the controlled environment (Kahneman, 1992; Kahneman et al., 1990; Shogren et al., 2001), but are often even more complex in reality (Bapna et al., 2004; Levitt and List, 2007). Yet, evidence from real-world environments is rare, as field research is time consuming and costly (List, 2011).

Our research group had the exceptional opportunities to collect three unique data sets with great effort in large and elaborate field experiments. The field studies presented in this thesis (in Chapters 3, 4, 6 & 7) provide unique empirical evidence on the interaction of individuals with information systems and analyze their impact in the real world. The findings on consumer behavior and on bidding behavior derived thus represent highly relevant insights for policy makers. In particular, whereas scientific research often tends to focus mostly on counter-intuitive findings (e.g., preference reversals or cognitive biases in behavioral economics (Fehr and Gächter, 2000; Gigerenzer and Brighton, 2009)), policy makers need robust information on scalable conservation policies, regardless of whether the results are surprising or ‘as expected’ from a theoretical perspective (Editorial, 2017).

IS-enabled feedback represents a scalable conservation measure: Both studies on real-time feedback interventions (presented in Chapters 3 and 4) provide robust empirical evidence collected in large samples of participants. The findings imply that using ubiquitous connected devices, real-time feedback can be a cost-effective and scalable measure to reduce resource consumption during specific activities among the broader population. Given the rapidly declining costs of sensors and connected devices, governments could lead the way into more sustainable consumption practices by equipping public facilities such as schools or public transportation with metering devices, which can provide activity-specific feedback in real time. Similarly, companies could reduce their environmental footprint by installing such devices in office buildings, and manufacturers could integrate real-time consumption feedback in the design of new appliances.

As changing individuals' consumption behavior in their everyday lives has proven to be very difficult, real-time feedback stands out as a practicable and effective measure to foster more sustainable consumption patterns based on the results presented in this thesis. The societal relevance of these findings is illustrated by the fact that the study on real-time feedback among uninformed, non-incentivized hotel guests (Chapter 3) was featured not only in renowned scientific outlets, but also in large, international media outlets (e.g., Deutschlandfunk, GEO, Nature.com, Scientific American, Spiegel online)¹.

Smart energy markets are technologically feasible in practice, but are challenged by data availability and regulatory issues: While several similar projects in other countries have started in the last two years, the P2P energy market that my colleagues and myself designed, developed, and deployed within the SFOE lighthouse project is the first of its kind in Switzerland and one of the first worldwide. The research in this project tackles a crucial step in the energy transition, by creating a consumer-centric energy market that can address the challenges arising from the integration and coordination of distributed energy resources. In this context, it is important to note that the focus of the research presented in dissertation was on the market design, bidding behavior, and resulting market outcomes. Beyond that, further work on the design of the user interface and the performance of the blockchain infrastructure was conducted within the project and has been published by our research team, which is outside the scope of this thesis,

¹ www.nature.com/articles/d41586-018-07471-1, November 19, 2018,
www.scientificamerican.com/podcast/episode/smart-meters-speed-showers/, November 27, 2018
www.spiegel.de/wissenschaft/technik/energie-sparen-wie-wir-alle-beim-duschen-sparen-koennen-a-1239593.html, November 21, 2018
www.geo.de/natur/nachhaltigkeit/19962-rtkl-emissionen-forscher-entwickeln-geraet-das-beim-duschen-geld-und, November 21, 2018
www.deutschlandfunknova.de/nachrichten/wasserverbrauch-beim-duschen-echtzeitinfo-hilft-sparen, November 22, 2018

see Ableitner et al. (2020) and Meeuw et al. (2020).

The research presented here shows that leveraging information and communication technology, market mechanisms can be applied to coordinate resources on the individual household level. The studies illustrate that smart energy markets are technologically feasible and can grant end-consumers an empowered role as active decision makers in the energy market – if the infrastructure to measure electricity flows on household level is in place. Beyond the scientific research that was conducted in the project, the implementation and operation of the system uncovered many practical challenges and created insights on existing barriers for the diffusion of data-based, consumer-centric energy markets. These include a lack of communication infrastructure and of processing capabilities for high resolution energy data, and regulatory issues in the regulated Swiss energy market, which were discussed with the project partners and the SFOE. In April 2020, the Swiss government has decided on the cornerstones of a legislative amendment to the Swiss Electricity Supply Act, which enables the creation of local energy markets for renewable energy – and explicitly mentions ‘Quartierstrom’-approaches in this context (SFOE (2020), p.3 “Wer beispielsweise Solarenergie produziert, kann den überschüssigen Strom im Quartier verkaufen. Damit ermöglicht die Öffnung des Strommarkts lokale Lösungen wie Quartierstrom-Märkte und Energiegemeinschaften.”).

Again, the insights presented in this thesis have thus generated considerable attention from policy makers and other researchers internationally, as well as public media outlets (e.g., SRF aktuell, CNN Money Switzerland, World Economic Forum Blog, interview in SRF Kassensturz)², and resulted in several invitations for talks including at Stanford University and the research organization Pecan Street Inc. in Austin, TX (which unfortunately had to be cancelled due to the Covid-19 outbreak in March 2020).

Preference elicitation in real energy auction reveals active user involvement, but low price preferences for solar energy: The data set collected in the P2P energy market represents the first empirical evidence on individual trading behavior in this context that has been published in the scientific literature³. While the sample of participants in the P2P energy market is small in comparison to the two experiments on

²Among others, the research project was featured twice in SRF Schweiz aktuell, the major Swiss news show, www.srf.ch/play/tv/schweiz-aktuell/video/solarstrom-von-der-quartierstromboerse?id=84984c0a-37e7-484f-b8ae-036650329273, and other major media outlets: www.srf.ch/play/tv/kassensturz/video/warum-es-mit-dem-solarstrom-harzt?id=2b5c1e6f-5019-416d-a1c5-58b6d024a37f
www.weforum.org/agenda/2019/09/are-our-smart-meters-smart-enough/
www.youtube.com/watch?v=QHjAPEkrGlo.

³Chapter 6 has already been published (Wörner et al., 2019a), Chapter 7 is about to be submitted.

real-time feedback, the empirical bidding data collected goes beyond existing research that relies on surveys or lab experiments with hypothetical scenarios. Furthermore, the data was collected over a period of an entire year. This longitudinal data set hence provides the first ever evidence on a P2P energy market's performance, and participants' behavior over time and across different seasonal conditions. The insights on the evolution of bidding behavior are unique and complement existing research on individual behavior in online auctions.

Concretely, the empirical results indicate learning effects in their bidding behavior and that participants understand the auction mechanism surprisingly well and many adjust their bids several times in the first eight months of the study. These insights imply that integrating consumers in decisions on energy sourcing can be successful if they are provided with high resolution energy data in an accessible and personalized manner. Nevertheless, the findings reveal an intention-behavior gap in participants' willingness to pay for local solar energy – a known phenomenon, which is still often ignored in the design of sustainability programs and which is relevant for the expected prices on P2P energy markets. Contrary to participants' own prior statements in the pre-experimental survey and contrary to existing survey studies evaluating responses to hypothetical P2P scenarios (Ecker et al., 2018; Hahnel et al., 2019), participants did not offer a price premium for buying on the local market in the actual auction when bids impacted their real electricity bill. Intelligent agents pointing out environmental consequences of market interactions to consumers may be one way to drive interest and investments in renewable energy resources and to increase the salience of long-term external costs.

Field research and sound methodological approaches are key to derive meaningful knowledge on consumer behavior: Related to the previous point, the research presented highlights several instances in which survey statements were not sufficient to predict individuals' behavior or even contradicted their actual behavior. This phenomenon is not restricted to but particularly pronounced in environmental contexts, where many individuals feel moral obligations to favor environmental protection, but the implementation of these ideals is less vigorous (Allcott and Mullainathan, 2010; Gatersleben et al., 2002). Researchers should pay attention to this potential bias in their study design and treat survey or lab study results with the required caution.

In addition, the studies conducted discuss the importance of participant recruitment for the external validity of experimental results. In a unique setting, the field experiment conducted among hotel guests, Chapter 3, showed that real-time feedback on energy consumption during showering was effective among a population of individuals that did

not know they were part of a research study. Herein, IS-enabled feedback on resource consumption emerges as scalable and cost-efficient policy instrument for the broader population, which does not necessarily hold for interventions tested among smaller, volunteer samples. The savings effects observed among the uninformed sample of hotel guests was large and significant, but it was still smaller in size than the effects found in other studies with the same intervention among a sample of volunteering households (Tiefenbeck et al., 2018a; Wörner and Tiefenbeck, 2018). As discussed in Chapter 3, while the difference in effect sizes may be caused by the lower baseline consumption in the hotel study, volunteer-selection bias in the referenced studies cannot be ruled out. A transparent discussion of the limitations to such field studies on human behavior is crucial to warrant the value of scientific research.

Multi-disciplinary approach increases applicability and accessibility of scientific results on fostering sustainability: This thesis draws on a number of theories from the related research fields computer science, psychology, behavioral economics, and IS, e.g., goal-setting theory (Chapter 4), market design theory (Chapter 5 and 6), and multi-agent learning (Chapter 8). The integration different scientific perspectives contributes to a more holistic understanding of the impact of information technology in the sustainability context (Melville, 2010). While some of the aspects could have been analyzed in more depth if a less broad research scope was chosen, the different theoretical lenses inform the presented study designs and compliment in depth research conducted in the respective disciplines to mutually enrich each other (Bichler et al., 2010). Similarly, multiple methodological approaches are applied to generate the findings and to put them into context: Field experiments producing empirical evidence constitute the core of this thesis, but a systematic literature review providing an analytical framework, as well as agent-based simulations, derive further background knowledge on the studied topics. Herein, this thesis addresses calls for integrative and impact-oriented research on Green IS (Gholami et al., 2016; Ketter et al., 2018; Malhotra et al., 2013). More so, this integrated assessment contributes to the accessibility of study results to a broader community, as well as to its value for practitioners.

‘Human-in-the-loop’ approaches should compliment ‘human-out-of-the-loop’ approaches in smart energy markets: In the same vein, this thesis examines approaches in which information systems are used to actively engage individuals in consumption decisions on the one hand (Chapters 3 – 7) and points out meaningful use cases for autonomous decision making by intelligent software agents on the other (Chapters 7 – 8). While research in social sciences traditionally focuses on the one side and computer

science on the other, the combination of these approaches is crucial for tackling real-world problems, and denotes a distinct contribution of this thesis. Some behavior patterns, like showering behavior, have an immediate and measurable impact on resource consumption, but other questions regarding energy sourcing or investments in renewable technologies are way more complex and intangible. While all of the field experiments in this thesis are testament to the fact that real-time feedback raises salience of, and interest in energy consumption, the studies on smart energy markets also illustrate definite boundaries for active consumer engagement. Adapting charging schedules or price bids to the dynamically changing conditions of energy supply and demand cannot be done manually and require computational tools for capacity, as well as velocity reasons. The surveys conducted in the P2P energy market (Chapter 7) found that participants grew increasingly willing to replace the interactive auction with an automatic pricing mechanism, and that they were open to entrust control to an autonomous trading agent. In line with recent studies on algorithm appreciation vs. aversion from other domains, semi-autonomous agents eliciting and integrating consumer preferences and then acting on their behalf in smart energy markets emerge as most desirable system, combining consumer engagement and efficiency.

While the bottom line remains that automating processes using information technology is required to address the dynamic and highly stochastic conditions in future energy markets, the perspective that information systems can engage and empower individuals to contribute to the energy transition at the same time is an important one, which should be tackled in future research. Herein, human-in-the-loop and human-out-of-the-loop approaches do not have to be contradictory, but should be used as compliments to achieve an impact in protecting the environment.

9.3 Limitations

Despite all best efforts, certain limitations have to be considered in all of the presented findings. In particular, as argued by List (2011) “we must work carefully when drawing conclusions based on the results of field experiments.”, p.9. First of all, one of the main contributions of this thesis pinpoints a limitation to most of the other studies presented. As argued in the article on a natural field experiment in hotels (Chapter 3), volunteer-selection bias is likely to affect the results of many framed field experiments with volunteer samples. Notably, in environmental contexts, individuals signing up for such experiments often only represent a small fraction of the target population (Kelly and Knottenbelt, 2016). While the unique design of the hotel study strengthens the confidence in the

generalizability of its results, this level of external validity is not granted in the remaining studies in this thesis. As many research contexts inherently prohibit the recruitment of uninformed participants for studies, or increase the practical complexity, it is important to discuss this limitation in every scientific study on a field experiment and debate possible biases in the results. All field studies presented in this thesis hence transparently explain the recruitment process for study participants. The P2P energy market (Chapters 6 and 7), in particular, suffers from a small sample of volunteer participants. Here, it is important to note that the experiment was of exploratory nature, and statistical analyses of observed data has to be treated with caution.

External validity in field studies further often suffers from exogenous factors outside the control of the researchers that introduce noise or even systematic biases to the study results. One of the limitations to the findings presented in this thesis is that all concepts for fostering more efficient resource use presented are subject to external conditions. For instance, the behavior observed in the field studies may be impacted by energy tariffs, which are relatively low in Switzerland compared to the income level. The incentives for trading on the P2P energy market are particularly influenced by these tariffs, as well as by the regulatory framework for selling solar production and for smart metering infrastructure. This limits the direct applicability of the design of the P2P market presented in this thesis to other countries. In addition, all of the field experiments presented were conducted in Switzerland, and cultural factors may limit the transferability of behavioral results to other cultures.

One of the weaknesses of the studies on feedback interventions is the limited duration of the experiments that prohibited the analysis of long term effects. In turn, this was one of the learnings that could be addressed in the design of the P2P studies in which trading was observed throughout an entire year, thus capturing seasonal variations and learning effects over time.

Finally, it would have been interesting to combine the P2P market studied with intelligent agents directly in the field, by incorporating more automation for pricing and smart scheduling of flexible loads. While it was not possible to deploy more intelligence in the scope of the field experiment, participants were asked about their preference on automation and agent-assistance in post-experimental survey. The findings across the studies presented in this thesis illustrate that it is vital to understand the incentives that drive individual behavior to foster sustainable consumption patterns, but at the same time, they also highlight the boundaries of behavioral measures. Addressing all energy-intensive activities with real-time feedback or asking consumers to change their entire

electricity consumption profiles dynamically to capture the maximal benefit of renewable energy would clearly overburden individuals. Taking all things together, finding the right balance between directly addressing consumers in behavioral interventions and automating complex decisions on consumption patterns (always taking into account individual preferences) emerges as key to fostering sustainability on the large scale. The intersection of these efforts exhibits interesting avenues for future research, marrying machine learning methods with social engineering.

As a final remark, it is important to note that the approaches studied in this thesis present only a small selection of the many ways in which information technology can contribute to fostering sustainability. This thesis has a strong focus on targeting individuals, in providing decision support to consumers and creating consumer-centric smart markets for residential customers. However, there are other instances in which information systems can be used to support resource conservation or the transition to renewable energy generation, for example resource management in organisations or decision support for utility providers in the wholesale market (Fridgen et al., 2016). As one of the greatest challenges of our time, fostering sustainability requires the cooperation of individuals, organizations, and policy makers.

9.4 Conclusion

The tech-philosopher Jaron Lanier once pointed out that “[i]t is impossible to work in information technology without also engaging in social engineering” (Lanier, 2010, p. 4). This dissertation confirms this statement. Most information systems do not run in isolation, but shape individuals’ economic and social environment, or support decision processes and react to user input. The findings presented underscore the importance of a nuanced understanding of consumers’ interaction with technology and their motivational drivers in realistic settings, in order to create actual impact. In a domain as relevant as environmental sustainability, it is paramount to not only improve energy-efficient technologies or computational tools, but technological advances must be accompanied by a social–science perspective on how these technologies will be employed and integrated into consumers’ decision processes and consumption patterns.

This dissertation draws on theories from social sciences to integrate this perspective in applying information technology to sustainability problems. With three field experiments, complimented by conceptual studies and simulations, this work contributes to the impact-oriented work on Green IS and behavioral economics. The different studies examined ways in which information technology can be used to foster sustainability in the

real world a) by providing decision support for energy consumption and b) by creating smart energy markets that integrate distributed renewable energy resources. Based on the findings presented, behavioral interventions provided by smart devices at the point of action in real time can yield large resource conservation effects, even among the broader public. Furthermore, the studies illustrated that smart electronic markets can support the integration of distributed renewable energy resources and the electrification of transportation in established energy markets. The findings can support policy makers (and have already) in creating a legal framework that promotes the energy transition and in designing scalable conservation programs.

The study designs, the data sets collected and the implications drawn in this thesis are the result of collaborations with multiple different partners, including scientific researchers from other universities and countries, practitioners from the energy sector, as well as the SFOE. The different perspectives introduced by all of these stakeholders have shaped the work presented and have strengthened its value both for practitioners, as well as for policy makers and the scientific community. As already pointed out in the introduction to this dissertation, tackling a wicked problem like environmental sustainability requires multiple perspectives and a whole set of measures, including technological solutions, behavioral interventions, and economic programs to go hand in hand to compliment each other.

In conclusion, the presented research allows me to end this thesis on a hopeful note. While information technology has arguably enabled some behavioral patterns that have led to the environmental damage and resource depletion that we see today, it can also be an efficient vehicle for fostering sustainable consumption patterns and for integrating renewable resources – thus ultimately for mitigating climate change. I believe it is in our hands as learning human beings to use the possibilities technology provides us for the better and to adopt more sustainable behavior patterns. This dissertation takes one small step in this direction; and I am sure that there are many more to come.

Bibliography

- Abatzoglou, J. T. and Williams, A. P. (2016). Impact of anthropogenic climate change on wildfire across western US forests. *Proceedings of the National Academy of Sciences (PNAS)*, 113(42):11770–11775.
- Ableitner, L., Meeuw, A., Schopfer, S., Tiefenbeck, V., Wortmann, F., and Wörner, A. (2019). Quartierstrom - Implementation of a real world prosumer centric local energy market in Walenstadt, Switzerland. *arXiv preprint*, 1905.07242.
- Ableitner, L., Schöb, S., Tiefenbeck, V., and Fridgen, G. (2017). Real-World Impact of Information Systems : The Effect of Seemingly Small Design Choices. *Workshop on Information Technology and Systems (WITS)*.
- Ableitner, L., Tiefenbeck, V., Meeuw, A., Wörner, A., Fleisch, E., and Wortmann, F. (2020). User behavior in a real-world peer-to-peer electricity market. *Applied Energy*, 270:115061.
- Abrahamse, W., Steg, L., Vlek, C., and Rothengatter, T. (2005). A Review of Intervention Studies Aimed at Household Energy Conservation. *Journal of Environmental Psychology*, 25(3):273–291.
- Abrahamse, W., Steg, L., Vlek, C., and Rothengatter, T. (2007). The effect of tailored information, goal setting, and tailored feedback on household energy use, energy-related behaviors, and behavioral antecedents. *Journal of Environmental Psychology*, 27:265–276.
- Abramova, S. and Böhme, R. (2016). Perceived Benefit and Risk as Multidimensional Determinants of Bitcoin Use: A Quantitative Exploratory Study. In *Proceedings of the International Conference on Information Systems (ICIS)*.
- Adomavicius, G., Gupta, A., and Zhdanov, D. (2009). Designing Intelligent Software Agents for Auctions with Limited Information Feedback. *Information Systems Research*, 20(4):507–526.

- Aitzhan, N. Z. and Svetinovic, D. (2016). Security and Privacy in Decentralized Energy Trading Through Multi-Signatures, Blockchain and Anonymous Messaging Streams. *IEEE Transactions on Dependable and Secure Computing*, 15(5):840–852.
- Akerlof, G. A. (1970). The Market for 'Lemons': Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 84(3):488–500.
- Akorede, M. F., Hizam, H., Pouresmaeil, E., Akorede, M., Hizam, H., and Pouresmaeil, E. (2010). Distributed energy resources and benefits to the environment. *Renewable and Sustainable Energy Reviews*, 14:724–734.
- Albrecht, S., Reichert, S., Schmid, J., Strüker, J., Neumann, D., and Fridgen, G. (2018). Dynamics of Blockchain Implementation - A Case Study from the Energy Sector. *Proceedings of the 51st Hawaii International Conference on System Sciences*, pages 3527–3536.
- Aljazeera (2019). Young people across globe protest against climate change inaction. www.aljazeera.com/news/2019/03/young-people-globe-protest-climate-change-inaction-190315114657681.html, 2019-04-09.
- Allcott, H. (2011). Social norms and energy conservation. *Journal of Public Economics*, 95(9-10):1082–1095.
- Allcott, H., Angrist, J., Chandra, A., Davis, L., Deaton, A., Fowlie, M., Gine, X., Goldman, C., Harding, M., Hotz, J., Imbens, G., Katz, L., Levy, D., Ludwig, J., Pande, R., Rogers, T., and Vytlačil, E. (2012a). External Validity and Partner Selection Bias. *NBER Working Paper Series*.
- Allcott, H., Blasnik, M., Davis, L., Gillingham, K., Golden, M., Joskow, P., Langer, T., McMahon, J., Mullainathan, S., Nadel, S., Newell, R., Sallee, J., Sanstad, A., and Greenstone, M. (2012b). Is There an Energy Efficiency Gap? *The Journal of Economic Perspectives Journal of Economic Perspectives*, 26(1):3–28.
- Allcott, H. and Mullainathan, S. (2010). Behavior and Energy Policy. *Science*, 327(5970):1204–1205.
- Allcott, H. and Rogers, T. (2012). The Short-Run and Long-Run Effects of Behavioral Interventions : Experimental Evidence from Energy Conservation. *American Economic Review*, 104(10):3003–3037.
- Allcott, H. and Sweeney, R. L. (2017). The Role of Sales Agents in Information Disclosure: Evidence from a Field Experiment. *Management Science*, 63(1):21–39.

- Allen, E. J., Dechow, P. M., Pope, D. G., and Wu, G. (2016). Reference-Dependent Preferences: Evidence from Marathon Runners. *Management Science*, 63(6):1657–1672.
- Altizer, S., Ostfeld, R. S., Johnson, P. T. J., Kutz, S., and Harvell, C. D. (2013). Climate Change and Infectious Diseases: From Evidence to a Predictive Framework. *Science*, 341(6145):514–519.
- Andoni, M., Robu, V., Flynn, D., Abram, S., Geach, D., Jenkins, D., Mccallum, P., and Peacock, A. (2018). Blockchain technology in the energy sector: A systematic review of challenges and opportunities. *Renewable and Sustainable Energy Reviews*, 100:143–174.
- Andrews, D. and Sánchez, A. C. (2011). The Evolution of Homeownership Rates in Selected OECD Countries: Demographic and Public Policy Influences. *OECD Journal: Economic Studies*, 2011(1).
- Attari, S. (2010). Public perceptions of energy consumption and savings. *Proceedings of the National Academy of Science*, 107(37):1–16.
- Avital, M., Beck, R., King, J. L., Rossi, M., and Teigland, R. (2016). Jumping on the Blockchain Bandwagon: Lessons of the Past and Outlook to the Future. *Proceedings of the International Conference on Information Systems (ICIS)*.
- Ba, S. and Pavlou, P. A. (2002). Evidence of the Effect of Trust Building Technology in Electronic Markets: Price Premiums and Buyer Behavior. *MIS Quarterly*, 26(3):243–268.
- Ba, S., Stallaert, J., Whinston, A. B., and Zhang, H. (2005). Choice of Transaction Channels: The Effects of Product Characteristics on Market Evolution. *Journal of Management Information Systems*, 21(4):173–197.
- Baca-Motes, K., Brown, A., Gneezy, A., Keenan, E. A., and Nelson, L. D. (2013). Commitment and Behavior Change: Evidence from the Field. *Journal of Consumer Research*, 39(5):1070–1084.
- Bakos, Y. and Katsamakas, E. (2008). Design and Ownership of Two-Sided Networks: Implications for Internet Platforms. *Journal of Management Information Systems*, 25(2):171–202.
- Baladi, S. M., Herriges, J. A., and Sweeney, T. J. (1998). Residential response to voluntary time-of-use electricity rates. *Resource and Energy Economics*, 20(3):225–244.
- Bamberg, S. (2006). Is a residential relocation a good opportunity to change people’s travel behavior?: Results from a theory-driven intervention study. *Environment and Behavior*, 38(6):820–840.

- Banker, R. D. and Kauffman, R. J. (2004). The Evolution of Research on Information Systems: A Fiftieth-Year Survey of the Literature. *Management Science*, 50(3):281–298.
- Bapna, R., Goes, P., and Gupta, A. (2003). Replicating Online Yankee Auctions to Analyze Auctioneers’ and Bidders’ Strategies. *Information Systems Research*, 14(3):244–268.
- Bapna, R., Goes, P., Gupta, A., and Jin, Y. (2004). User Heterogeneity and Its Impact on Electronic Auction Market Design: An Empirical Exploration. *MIS Quarterly*, 28(1):21–43.
- Barata, F., Kipfer, K., Weber, M., Tinschert, P., Fleisch, E., and Kowatsch, T. (2019). Towards device-agnostic mobile cough detection with convolutional neural networks. In *IEEE International Conference on Healthcare Informatics*.
- Basden, J. and Cottrell, M. (2017). How Utilities Are Using Blockchain to Modernize the Grid. *Harvard Business Review Digital Articles*, 21:1–8.
- Basu, A., Bhaskaran, S., and Mukherjee, R. (2019). An Analysis of Search and Authentication Strategies for Online Matching Platforms. *Management Science*, 65(5):2412–2431.
- Beck, R., Avital, M., Rossi, M., and Thatcher, J. B. (2017a). Blockchain Technology in Business and Information Systems Research. *Business & Information Systems Engineering*, 59(6):381–384.
- Beck, R., Becker, C., Lindman, J., and Rossi, M. (2017b). Opportunities and risks of Blockchain Technologies – a research agenda. *Dagstuhl Reports*, 7(3):99–142.
- Beck, R., Müller-Bloch, C., and King, J. L. (2018). Governance in the Blockchain Economy: A Framework and Research Agenda. *Journal of the Association for Information Systems*, 19(10).
- Beck, R., Stenum Czepluch, J., Lollike, N., and Malone, S. (2016). Blockchain - The Gateway to Trust-Free Cryptographic Transactions. *Proceedings of the European Conference on Information Systems (ECIS)*.
- Behaviors Union (2008). A framework for pro-environmental behaviours. Technical report, Department for Environment Food and Rural Affairs.
- Bertrand, A., Mastrucci, A., Schüler, N., Aggoune, R., and Maréchal, F. (2017). Characterisation of domestic hot water end-uses for integrated urban thermal energy assessment and optimisation. *Applied Energy*, 186(2):152–166.

- Bichler, M. (2001). *The future of eMarkets : multi-dimensional market mechanisms*. Cambridge University Press.
- Bichler, M., Fux, V., and Goeree, J. K. (2019). Designing combinatorial exchanges for the reallocation of resource rights. *Proceedings of the National Academy of Science*, 116(3):786–791.
- Bichler, M., Gupta, A., and Ketter, W. (2010). Designing smart markets. *Information Systems Research*, 21(4):688–699.
- Bichler, M. and Segev, A. (2001). Methodologies for the design of negotiation protocols on E-markets. *Computer Networks*, 37(2):137–152.
- Block, C., Neumann, D., and Weinhardt, C. (2008). A Market Mechanism for Energy Allocation in Micro-CHP Grids. In *Proceedings of the Hawaii International Conference on System Sciences (HICSS)*.
- Bollinger, B. K. and Hartmann, W. R. (2020). Information vs. Automation and Implications on Pricing. *Management Science*, 66(1):1–25.
- Borenstein, S. (2005). The Long-Run Efficiency of Real Time Electricity Pricing. *The Energy Journal*, 26(3):93–116.
- Borenstein, S., Bushnell, J. B., and Wolak, F. A. (2002). Measuring Market Inefficiencies in California’s Restructured Wholesale Electricity Market. *The American Economic Review*, 92(5):1376–1405.
- Brandt, T., Feuerriegel, S., and Neumann, D. (2018). Modeling interferences in information systems design for cyberphysical systems: Insights from a smart grid application. *European Journal of Information Systems*, 27(2):207–220.
- Brandt, T., Stadler, M., and Neumann, D. (2014). Power Systems 2.0: Designing an Energy Information System for Microgrid Operation. *Proceedings of the International Conference on Information Systems (ICIS)*.
- Bravata, D., Smith-Spangler, C., Sundaram, V., Gienger, A. L., Lin, N., Lewis, R., Stave, C. D., and Olkin, I. (2007). Using Pedometers to Increase Physical Activity A Systematic Review. *Journal of the American Medical Association*, 298(19):2296–2304.
- Buchanan, K., Russo, R., and Anderson, B. (2015). The question of energy reduction: The problem(s) with feedback. *Energy Policy*, 77:89–96.
- Bundesamt für Statistik (2020). Haushaltseinkommen und -ausgaben in der

- Schweiz. www.bfs.admin.ch/bfs/de/home/statistiken/wirtschaftliche-soziale-situation-bevoelkerung/einkommen-verbrauch-vermoegen/haushaltsbudget.html, 2020-06-30.
- Burger, C., Kuhlmann, A., Richard, P., and Weinmann, J. (2016). Blockchain in the energy transition. A survey among decision-makers in the German energy industry. Technical report, DENA German Energy Agency.
- Burton-Jones, A. and Grange, C. (2013). From Use to Effective Use : A Representation Theory Perspective From Use to Effective Use : A Representation Theory Perspective. *Information Systems Research*, 24(3):632–658.
- Buterin, V. (2014). A Next Generation Smart Contract & Decentralized Application Platform. *Whitepaper*.
- Cai, X., Zhang, B., and Zhao, X. (2019). What Drives Adoption of Smart Contract?: Identifying Peer Influences in Blockchain User Network. In *Proceedings of the International Conference on Information Systems (ICIS)*.
- Camerer, C., Issacharoff, S., Loewenstein, G., O'Donoghue, T., and Rabin, M. (2003). Regulation for Conservatives: Behavioral Economics and the Case for "Asymmetric Paternalism". *University of Pennsylvania Law Review*, 151(3):1211–1254.
- Campbell, Donald, T. (1969). Reform as Experiments. *American Psychologist*, 24:409–429.
- Capstick, S., Whitmarsh, L., Poortinga, W., Pidgeon, N., and Upham, P. (2015). International trends in public perceptions of climate change over the past quarter century. *WIREs Clim Change*, 6:35–61.
- Carvalho, A. (2020). A permissioned blockchain-based implementation of LMSR prediction markets. *Decision Support Systems*, 130(113228).
- Catalini, C. and Gans, J. S. (2016). Some Simple Economics of the Blockchain. *National Bureau of Economic Research*, w22952.
- Chanson, M. ., Bogner, A. ., Wortmann, F. ., and Fleisch (2017). Blockchain as a privacy enabler: an odometer fraud prevention system ETH Library. *Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing*.
- Chong, A. Y. L., Lim, E. T. K., Hua, X., Zheng, S., Tan, C.-W., Beck, R., Rossi, M., and Thatcher, J. (2019). Business on Chain: A Comparative Case Study of Five Blockchain-Inspired Business Models. *Journal of the Association for Information Systems*, 20(9):1310–1339.

- Clark, C. F., Kotchen, M. J., and Moore, M. R. (2003). Internal and external influences on pro-environmental behavior: Participation in a green electricity program. *Journal of Environmental Psychology*, 23(3):237–246.
- Clemons, E. K., Dewan, R. M., Kauffman, R. J., and Weber, T. A. (2017). Understanding the Information-Based Transformation of Strategy and Society. *Journal of Management Information Systems*, 34(2):425–456.
- Clemons, E. K. and Row, M. (1988). McKesson Drug Company: A Case Study of Economost-A Strategic Information System McKesson Drug Company: A Case Study of Economost-A Strategic Information System. *Journal of Management Information Systems*, 5(1):141–149.
- Coase, R. H. (1937). The Nature of the Firm. *Economica*, 4(16):386–405.
- Commission for Energy Regulation (2012). CER Smart Metering Project - Electricity Customer Behaviour Trial, 2009-2010 [dataset]. 1st Edition. Irish Social Science Data Archive. Technical report.
- Consolvo, S., Mcdonald, D. W., and Landay, J. A. (2009). Theory-Driven Design Strategies for Technologies that Support Behavior Change in Everyday Life. *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*.
- Constantinides, P., Henfridsson, O., and Parker, G. G. (2018). Introduction-Platforms and Infrastructures in the Digital Age. *Information Systems Research*, 29(2):381–400.
- Corbett, J. and Mellouli, S. (2017). Winning the SDG battle in cities: how an integrated information ecosystem can contribute to the achievement of the 2030 sustainable development goals. *Information Systems Journal*, 27(4):427–461.
- Cramton, P., Shoham, Y., and Steinberg, R. (2006). *Combinatorial auctions*. MIT Press.
- Creutzig, F., Roy, J., Lamb, W. F., Azevedo, I. M., de Bruin, W. B., Dalkmann, H., and Hertwich, E. G. (2018). Towards demand-side solutions for mitigating climate change. *Nature Climate Change*, 8(4):268–271.
- Creyts, J. and Tribovich, A. (2018). Can blockchain help us to address the world’s energy issues? www.weforum.org/agenda/2018/01/how-can-blockchain-address-the-worlds-energy-issues/, World Econ.
- Darby, S. (2006). The effectiveness of feedback on energy consumption - a review for defra of the literature on metering, billing and direct displays. Technical report, DEFRA Review.

- Dauer, D., Karaenke, P., and Weinhardt, C. (2015). Load Balancing in the Smart Grid: A Package Auction and Compact Bidding Language. *Proceedings of the International Conference on Information Systems (ICIS)*.
- Davis, A. L., Krishnamurti, T., Fischhoff, B., and Bruine de Bruin, W. (2013). Setting a standard for electricity pilot studies. *Energy Policy*, 62:401–409.
- Davis, S., Lewis, N., Shaner, M., Aggarwal, S., Arent, D., Azevedo, I. L., Benson, S., Bradley, T., Brouwer, J., Chiang, Y.-M., Clack, C., Cohen, A., Doig, S., Edmonds, J., Fennell, P., Field, C., Hannegan, B., Hodge, B.-M., Hoffert, M., Ingersoll, E., Jaramillo, P., Lackner, K. S., Mach, K., Mastrandrea, M., Ogden, J., Peterson, F., Sanchez, D., Sperling, D., Stagner, J., Trancik, J., Yang, C.-J., and Caldeira, K. (2018). Net-zero emissions energy systems. *Science*, 360(eaas9793).
- Decker, C. and Wattenhofer, R. (2013). Information Propagation in the Bitcoin Network. *International Conference on Peer-to-Peer Computing*.
- Delmas, M., Fischlein, M., and Asensio, O. (2013). Information strategies and energy conservation behaviour: A meta-analysis of experimental studies from 1975-2011 (draft). *Energy Policy*, 61:729–739.
- Diekmann, A., Meyer, R., Mühlemann, C., and Diem, A. (2009). Schweizer Umweltsurvey 2007 - Analysen und Ergebnisse (Swiss environmental survey - Analyses and results). Report to the Swiss Federal Statistical Office (BFS) and to the Federal Office for the Environment (BAFU). *Technical report, ETH Zurich*.
- Dietvorst, B. J., Simmons, J. P., and Massey, C. (2016). Overcoming Algorithm Aversion: People Will Use Imperfect Algorithms If They Can (Even Slightly) Modify Them. *Management Science*, 64(3):1155–1170.
- Dietz, T., Ostrom, E., Stern, P. C., and Stern, P. C. (2003). The Struggle to Govern the Commons. *Science*, 302:1907–1912.
- Ecker, F., Spada, H., and Hahnel, U. J. (2018). Independence without control: Autarky outperforms autonomy benefits in the adoption of private energy storage systems. *Energy Policy*, 122:214–228.
- Editorial (2017). I’m Not Surprised. *Nature Energy*, 2:17101.
- Ehrhardt-Martinez, K., Donnelly, K. A., Laitner, S., York, D., Talbot, J., and Friedrich, K. (2010). Advanced Metering Initiatives and Residential Feedback Programs: A Meta-Review for Household Electricity-Saving Opportunities. *Washington, DC: American Council for an Energy-Efficient Economy*.

- Einav, L., Farronato, C., and Levin, J. (2016). Peer-to-Peer Markets. *Annual Review of Economics*, 8:615–635.
- Erev, I. and Roth, A. E. (1998). Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria. *The American Economic Review*, 88(4):848–881.
- European Commission (2015). A Framework Strategy for a Resilient Energy Union with a Forward-Looking Climate Change Policy - EUR-Lex - 52015DC0080. Technical report.
- European Consumer Organisation (2016). Current practices in consumer-driven renewable electricity markets, BEUC mapping report. Technical report.
- eurostat (2017). Energy consumption in households - Statistics Explained. ec.europa.eu/eurostat/statistics-explained/index.php/Energy_consumption_in_households, 2020-05-14.
- Fabra, N., Von Der Fehr, N.-H., and Harbord, D. (2002). Modeling Electricity Auctions. *The Electricity Journal*, 15(7):72–81.
- Faruqui, A., Sergici, S., and Sharif, A. (2010). The impact of informational feedback on energy consumption-A survey of the experimental evidence. *Energy*, 35(4):1598–1608.
- Fehr, E. and Gächter, S. (2000). Cooperation and punishment in public goods experiments. *American Economic Review*, 90(4):980–994.
- Fishbein, M. and Ajzen, I. (1975). *Belief, attitude, intention, and behavior: An introduction to theory and research*. Addison-Wesley, Reading.
- Flath, C. M., Ilg, J. P., Gottwalt, S., Schmeck, H., and Weinhardt, C. (2014). Improving Electric Vehicle Charging Coordination Through Area Pricing. *Transportation Science*, 48(4):619–634.
- Fleisch, E. (2010). What is the Internet of Things? An Economic Perspective What is the Internet of Things - An Economic Perspective. *Economics, Management, and Financial Markets*, 5(2):125–157.
- Fontoura, M., Ionescu, M., and Minsky, N. (2005). Decentralized peer-to-peer auctions. *Electronic Commerce Research*, 5(1):7–24.
- Fox, S. and Duggan, M. (2012). Mobile Health 2012. *Washington DC: Pew Research Center's Internet & American Life Project*, Report.

- Frederiks, E. R., Stenner, K., Hobman, E. V., and Fischle, M. (2016). Evaluating energy behavior change programs using randomized controlled trials: Best practice guidelines for policymakers. *Energy Research & Social Science*, 22:147–164.
- Frey, B. S. and Oberholzer-Gee, F. (1997). The Cost of Price Incentives: An Empirical Analysis of Motivation Crowding-Out. *The American Economic Review*, 87(4):746–755.
- Fridgen, G., Häfner, L., König, C., and Sachs, T. (2016). Providing Utility to Utilities: The Value of Information Systems Enabled Flexibility in Electricity Consumption. *Journal of the Association for Information Systems*, 17(8):537–563.
- Fridgen, G., Kahlen, M., Ketter, W., Rieger, A., and Thimmel, M. (2018). One rate does not fit all: An empirical analysis of electricity tariffs for residential microgrids. *Applied Energy*, 210:800–814.
- Froehlich, J., Dillahunt, T., Klasnja, P., Mankoff, J., Consolvo, S., Harrison, B., and Landay, J. A. (2009). UbiGreen: Investigating a Mobile Tool for Tracking and Supporting Green Transportation Habits. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.
- Froehlich, J., Findlater, L., and Landay, J. (2010). The Design of Eco-Feedback Technology. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*.
- Gahr, B., Ryder, B., Dahlinger, A., and Wortmann, F. (2018). Driver Identification via Brake Pedal Signals — A Replication and Advancement of Existing Techniques. In *IEEE International Conference on Intelligent Transportation Systems (ITSC)*, pages 1415–1420.
- Gatersleben, B., Steg, L., and Vlek, C. (2002). Measurement and Determinants of Environmentally Significant Consumer Behavior. *Environment and Behavior*, 34(3):335–362.
- German Federal Ministry for Economic Affairs and Energy (2016). Energy Efficiency Campaign: "Deutschland Macht's Effizient.". www.deutschland-machts-effizient.de/KAENEFF/Navigation/DE/Home/home.html, 2017-07-27.
- Gholami, R., Watson, R. T., Hasan, H., Molla, A., and Bjørn-Andersen, N. (2016). Information Systems Solutions for Environmental Sustainability: How Can We Do More? *Journal of the Association for Information Systems*, 17(8):521–536.
- Gigerenzer, G. and Brighton, H. (2009). Homo Heuristicus: Why Biased Minds Make Better Inferences. *Topics in Cognitive Science*, 1(1):107–143.

- Gillingham, K., Jenn, A., and Azevedo, I. M. L. (2015). Heterogeneity in the response to gasoline prices: Evidence from Pennsylvania and implications for the rebound effect. *Energy Economics*, 52:41–52.
- Gimpel, H., Nißen, M., and Görlitz, R. A. (2013). Quantifying the Quantified Self: A Study on the Motivation of Patients to Track Their Own Health. *Proceedings of the International Conference on Information Systems (ICIS)*.
- Glaser, F. (2017). Pervasive Decentralisation of Digital Infrastructures: A Framework for Blockchain enabled System and Use Case Analysis. *Hawaii International Conference on System Sciences (HICSS)*.
- Gneezy, U. and Rustichini, A. (2000). A Fine is a Price. *The Journal of Legal Studies*, 29(1):1–17.
- Gober, P., Kirkwood, C. W., and Macdonald, G. M. (2010). Vulnerability assessment of climate-induced water shortage in Phoenix. *Proceedings of the National Academy of Science (PNAS)*, 107(50):21295–21299.
- Gode, D. K. and Sunder, S. (1993). Allocative Efficiency of Markets with Zero-Intelligence Traders : Market as a Partial Substitute for Individual Rationality. *Journal of Political Economy*, 101(1):119–137.
- Goes, P. (2013). Information Systems Research and Behavioral Economics. *MIS Quarterly*, 37(3):iii–viii.
- Goes, P. B., Karuga, G. G., and Tripathi, A. K. (2010). Understanding Willingness-to-Pay Formation of Repeat Bidders in Sequential Online Auctions. *Information Systems Research*, 21(4):907–924.
- Goes, P. B., Karuga, G. G., and Tripathi, A. K. (2012). Bidding Behavior Evolution in Sequential Auctions: Characterization and Analysis. *MIS Quarterly*, 36(4):1021–1042.
- Goette, L., Stutzer, A., and Frey, B. M. (2010). Prosocial Motivation and Blood Donations: A Survey of the Empirical Literature. *Transfusion Medicine and Hemotherapy*, 37:149–154.
- Goldstein, N. J., Cialdini, R. B., and Griskevicius, V. (2008). A Room with a Viewpoint: Using Social Norms to Motivate Environmental Conservation in Hotels. *Journal of Consumer Research*, 35(3):472–482.
- Google Trends (2020). Google Trends. [trends.google.com/trends/explore?q=climate change&geo=US](https://trends.google.com/trends/explore?q=climate%20change&geo=US), 2020-05-15.

- Gottwalt, S., Ketter, W., Block, C., Collins, J., and Weinhardt, C. (2011). Demand side management-A simulation of household behavior under variable prices. *Energy Policy*, 39:8163–8174.
- Green, J. and Newman, P. (2017). Citizen utilities: The emerging power paradigm. *Energy Policy*, 105:283–293.
- Griego, D., Schopfer, S., Henze, G., Fleisch, E., and Tiefenbeck, V. (2019). Aggregation effects for microgrid communities at varying sizes and prosumer-consumer ratios. *Energy Procedia*, 159:346–351.
- Guo, Z., Koehler, G. J., and Whinston, A. B. (2012). A Computational Analysis of Bundle Trading Markets Design for Distributed Resource Allocation. *Information Systems Research*, 23(3-part-1):823–843.
- Gupta, A. (2017). Editorial thoughts: What and how ISR publishes. *Information Systems Research*, 28(1):1–4.
- Hahn, A., Singh, R., and Liu, C.-c. (2017). Smart Contract-based Campus Demonstration of Decentralized Transactive Energy Auctions. In *IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT)*.
- Hahnel, U. J., Herberz, M., Pena-Bello, A., Parra, D., and Brosch, T. (2019). Becoming prosumer: Revealing trading preferences and decision-making strategies in peer-to-peer energy communities. *Energy Policy*, 137(111098).
- Hales, S., de Wet, N., Maindonald, J., and Woodward, A. (2002). Potential effect of population and climate changes on global distribution of dengue fever: an empirical model. *The Lancet*, 360(9336):830–834.
- Halu, A., Scala, A., Khiyami, A., and González, M. C. (2016). Data-driven modeling of solar-powered urban microgrids. *Science Advances*, 2(1):e1500700.
- Hanelt, A., Busse, S., and Kolbe, L. M. (2017). Driving business transformation toward sustainability: exploring the impact of supporting IS on the performance contribution of eco-innovations. *Information Systems Journal*, 27(4):463–502.
- Hasse, F., von Perfall, A., Hillebrand, T., Smole, E., Lay, L., and Charlet, M. (2016). Blockchain - an opportunity for energy producers and consumers? *PwC Global Power & Utilities*, Report.
- Haynes, L., Service, O., Goldacre, B., and Torgerson, D. (2012). Test, Learn, Adapt: Developing Public Policy with Randomised Controlled Trials. *London: Cabinet Office Behavioural Insights Team*.

- Hentschel, M., Ketter, W., and Collins, J. (2018). Renewable energy cooperatives: Facilitating the energy transition at the Port of Rotterdam. *Energy Policy*, 121:61–69.
- Hermesen, S., Frost, J., Renes, R. J., and Kerkhof, P. (2016). Using feedback through digital technology to disrupt and change habitual behavior: A critical review of current literature. *Computers in Human Behavior*, 57:61–74.
- Herter, K., Wood, V., and Blozis, S. (2013). The effects of combining dynamic pricing, AC load control, and real-time energy feedback: SMUD’S 2011 Residential Summer Solutions Study. *Energy Efficiency*, 6(4):641–653.
- Hinsz, V. B., Kalnbach, L. R., and Lorentz, N. R. (1997). Using Judgmental Anchors to Establish Challenging Self-Set Goals Without Jeopardizing Commitment. *Organizational Behavior and Human Decision Processes*, 71(3):287–308.
- Hopf, K., Sodenkamp, M., Riechel, S., and Staake, T. (2017). Predictive Customer Data Analytics – The Value of Public Statistical Data and the Geographic Model Transferability. *Proceedings of the International Conference on Information Systems (ICIS)*.
- Hoppmann, J., Volland, J., Schmidt, T. S., and Hoffmann, V. H. (2014). The economic viability of battery storage for residential solar photovoltaic systems – A review and a simulation model. *Renewable and Sustainable Energy Reviews*, 39:1101–1118.
- International Energy Agency (2018). Market Report Series: Renewables 2018. webstore.iea.org/market-report-series-rene, 2019-10-10.
- International Energy Agency (2019). Energy Efficiency Indicators. www.iea.org/reports/energy-efficiency-indicators-2019, 2020-05-11.
- International Energy Agency (2020a). Data & Statistics. [www.iea.org/data-and-statistics?country=WORLD&fuel=Energy consumption&indicator=Total final consumption \(TFC\) by sector](http://www.iea.org/data-and-statistics?country=WORLD&fuel=Energy%20consumption&indicator=Total%20final%20consumption%20(TFC)%20by%20sector), 2020-04-14.
- International Energy Agency (2020b). Renewable power – Tracking Power – Analysis. www.iea.org/reports/tracking-power-2019/renewable-power, 2020-05-11.
- IPCC, Pachauri, R. K., Meyer, L., Hallegatte France, S., Bank, W., Hegerl, G., Brinkman, S., van Kesteren, L., Leprince-Ringuet, N., and van Boxmeer, F. (2014). Climate Change 2014 Synthesis Report. Technical report.
- Jain, R. K., Qin, J., and Rajagopal, R. (2017). Data-driven planning of distributed energy resources amidst socio-technical complexities. *Nature Energy*, 2(17112).

- Janze, C. (2017). Design of a decentralized peer-to-peer reviewing and publishing market. In *European Conference of Information Systems (ECIS)*.
- Kahneman, D. (1992). Reference points, anchors, norms, and mixed feelings. *Organizational Behavior and Human Decision Processes*, 51(2):296–312.
- Kahneman, D., Knetsch, J. L., and Thaler, R. H. (1990). Experimental Tests of the Endowment Effect and the Coase Theorem. *Journal of Political Economy*, 98(6):1325–1348.
- Kahneman, D. and Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2):263–292.
- Kang, J., Yu, R., Huang, X., Maharjan, S., Zhang, Y., and Hossain, E. (2017). Enabling Localized Peer-to-Peer Electricity Trading among Plug-in Hybrid Electric Vehicles Using Consortium Blockchains. *IEEE Transactions on Industrial Informatics*, 13(6):3154–3164.
- Karlan, D. and List, J. A. (2007). Does Price Matter in Charitable Giving? Evidence from a Large-Scale Natural Field Experiment. *The American Economic Review*, 97(5):1774–1793.
- Karlin, B., Zinger, J. F., and Ford, R. (2015). The Effects of Feedback on Energy Conservation: A Meta-Analysis. *Psychological Bulletin*, 141(6):1205–1227.
- Karneyeva, Y. and Wüstenhagen, R. (2017). Solar feed-in tariffs in a post-grid parity world: The role of risk, investor diversity and business models. *Energy Policy*, 106:445–456.
- Kastrati, G. and Weissbart, C. (2016). Kurz zum Klima: Blockchain - Potenziale und Herausforderungen für den Strommarkt. *ifo Schnelldienst*, 69(23):74–77.
- Kazan, E., Tan, C.-W., Lim, E. T. K., Sørensen, C., and Damsgaard, J. (2018). Disentangling Digital Platform Competition: The Case of UK Mobile Payment Platforms. *Journal of Management Information Systems*, 35(1):180–219.
- Kelly, J. and Knottenbelt, W. (2016). Does disaggregated electricity feedback reduce domestic electricity consumption? A systematic review of the literature. *arXiv preprint*, 1605:00962v2.
- Ketter, W. (2019). Are our smart meters smart enough? | World Economic Forum. www.weforum.org/agenda/2019/09/are-our-smart-meters-smart-enough/, 2019-09-16.

- Ketter, W., Collins, J., Gini, M., Gupta, A., and Schrater, P. (2012). Real-Time Tactical and Strategic Sales Management for Intelligent Agents Guided by Economic Regimes. *Information Systems Research*, 23(4):1263–1283.
- Ketter, W., Collins, J., and Reddy, P. (2013). Power TAC: A competitive economic simulation of the smart grid. *Energy Economics*, 39:262–270.
- Ketter, W., Collins, J., Saar-Tsechansky, M., and Marom, O. (2018). Information Systems for a Smart Electricity Grid. *ACM Transactions on Management Information Systems*, 9(3):1–22.
- Ketter, W., Peters, M., Collins, J., and Gupta, A. (2015). Competitive Benchmarking: An IS Research Approach to Address Wicked Problems with Big Data and Analytics. *MIS Quarterly*, 40(4).
- Khalilpour, R. and Vassallo, A. (2015). Leaving the grid: An ambition or a real choice? *Energy Policy*, 82:207–221.
- Kienscherf, P. A., Wörner, A., Ketter, W., and Tiefenbeck, V. (2020). Multi-Agent Reinforcement Learning for Electric Vehicle Charging Guided by Nodal Pricing. *Working Paper. To be submitted to the Journal of Artificial Intelligence in July 2020.*
- Kiesling, L., Munger, M., Theisen, A., Fitzgerald, T., Glachant, J.-M., Kaffine, D., Lazanski, D.-M., Leisten, M., Littlechild, S., Ohlers, A., and Zyontz, S. (2017). From Airbnb to Solar: Toward A Transaction Cost Model of a Retail Electricity Distribution Platform. *Working Paper, TILEC Workshop.*
- Klemperer, P. (2002). What Really Matters in Auction Design. *Journal of Economic Perspectives*, 16(1):169–189.
- Kloör, B., Monhof, M., Beverungen, D., and Braäer, S. (2018). Design and evaluation of a model-driven decision support system for repurposing electric vehicle batteries. *European Journal of Information Systems*, 27(2):171–188.
- Kluger, A. N. and DeNisi, A. (1996). The effects of feedback interventions on performance: A historical review, a meta-analysis, and a preliminary feedback intervention theory. *Psychological Bulletin*, 119(2):254–284.
- Komiak, S. Y. X. and Benbasat, I. (2006). The Effects of Personalization and Familiarity on Trust and Adoption of Recommendation. *MIS Quarterly*, 30(4):941–960.
- Koolen, D., Ketter, W., Qiu, L., and Gupta, A. (2017). The Sustainability Tipping Point in Electricity Markets. *Proceedings of the International Conference on Information Systems (ICIS).*

- Kranz, J., Nagel, E., and Yoo, Y. (2019). Blockchain Token Sale Economic and Technological Foundations. *Business & Information Systems Engineering*, 61.
- Lacity, M. C. (2018). Addressing Key Challenges to Manage Enterprise Blockchain Applications in Reality. *MIS Quarterly Executive*, 17(3).
- Lampinen, A. and Brown, B. (2017). Market Design for HCI: Successes and Failures of Peer-to-Peer Exchange Platforms. *Proceedings of the 2017 Conference on Human Factors in Computing Systems (CHI)*.
- Landry, C. E., Lange, A., List, J. A., Price, M. K., and Rupp, N. G. (2006). Toward an Understanding of the Economics of Charity: Evidence from a Field Experiment. *The Quarterly Journal of Economics*, 121(2):747–782.
- Lanier, J. (2010). *You are not a gadget: A manifesto*. New York: Knopf.
- Laszka, A., Dubey, A., Walker, M., and Schmidt, D. (2017). Providing Privacy, Safety, and Security in IoT-Based Transactive Energy Systems using Distributed Ledgers. *Proceedings of the International Conference on the Internet of Things (IoT)*.
- Latham, G. P. and Locke, E. A. (1991). Self-regulation through goal setting. *Organizational Behavior and Human Decision Processes*, 50(2):212–247.
- Lee, C.-S. (2001). An analytical framework for evaluating e-commerce business models and strategies. *Internet Research*, 11(4):349–359.
- Lee, K. and Ashton, M. C. (2004). Psychometric Properties of the HEXACO Personality Inventory. *Multivariate Behavioral Research*, 39(2):329–358.
- Lee, T. M., Markowitz, E. M., Howe, P. D., Ko, C.-Y., and Leiserowitz, A. A. (2015). Predictors of public climate change awareness and risk perception around the world. *Nature Climate Change*, 5:1014–1023.
- Lesic, V., Bruine de Bruin, W., Davis, M., Krishnamurti, T., and Azevedo, I. M. (2018). Consumers’ perceptions of energy use and energy savings: A literature review. *Environmental Research Letters*, 13(033004).
- Levitt, S. D. and List, J. A. (2007). What do Laboratory Experiments Measuring Social Preferences Reveal About the Real World? *Journal of Economic Perspectives*, 21(2):153–174.
- Levitt, S. D. and List, J. A. (2008). Homo economicus Evolves. *Science*, 319:909–910.
- Levy, P. E. and Baumgardner, A. H. (1991). Effects of Self-Esteem and Gender on Goal Choice. *Journal of Organizational Behavior*, 12:529–541.

- Li, I., Dey, A., and Forlizzi, J. (2010). A Stage-Based Model of Personal Informatics Systems. *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*.
- Li, R., Wu, Q., and Oren, S. S. (2014). Distribution locational marginal pricing for optimal electric vehicle charging management. *IEEE Transactions on Power Systems*, 29(1):203–211.
- Lindman, J., Rossi, M., and Tuunainen, V. K. (2017). Opportunities and risks of Blockchain Technologies in payments—a research agenda. *Proceedings of the 50th Hawaii International Conference on System Sciences (HICSS)*.
- List, J. A. (2004). Testing Neoclassical Competitive Theory in Multilateral Decentralized Markets. *Journal of Political Economy*, 112(5):1131–1156.
- List, J. A. (2011). Why Economists Should Conduct Field Experiments and 14 Tips for Pulling One Off. *The Journal of Economic Perspectives*, 25176110(3):3–15.
- List, J. A., Ashenfelter, O., Battalio, R., Benabou, R., Benjamin, D., Char-Ness, G., Falk, A., Glaeser, E., Gneezy, U., Harrison, G., Kahneman, D., Koch, L., Laibson, D., Levitt, S., Neal, D., Rabin, M., and Roth, A. (2006). The Behavioralist Meets the Market: Measuring Social Preferences and Reputation Effects in Actual Transactions. *Journal of Political Economy*, 114(1):1–37.
- Littman, M. L. (1994). Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings*, pages 157–163.
- LO3 Energy (2017). Exergy - Building a robust value mechanism to facilitate transactive energy. *Whitepaper*.
- Locke, E. A. (1996). Motivation through Conscious Goal Setting. *Applied & Preventive Psychology*, 5:117–124.
- Locke, E. A. and Latham, G. P. (2002). Building a practically useful theory of goal setting and task motivation: A 35-year odyssey. *The American Psychologist*, 57(9):705–717.
- Locke, E. A. and Latham, G. P. (2006). New Directions in Goal-Setting Theory. *Current Directions in Psychological Science*, 15(5):265–268.
- Loeser, F., Recker, J., Brocke, J. v., Molla, A., and Zarnekow, R. (2017). How IT executives create organizational benefits by translating environmental strategies into Green IS initiatives. *Information Systems Journal*, 27(4):503–553.

- Logg, J. M., Minson, J. A., and Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151:90–103.
- Loock, C.-M., Staake, T., and Thiesse, F. (2013). Motivating Energy-Efficient Behavior with Green IS: An Investigation of Goal Setting and the Role of Defaults. *MIS Quarterly*, 37(4):1313–1332.
- Lossin, F., Kozlovskiy, I., Sodenkamp, M., Staake, T., and Lossin, R. (2016). Incentives to go green: an empirical investigation of monetary and symbolic rewards to motivate energy savings. *Proceedings of the European Conference on Information Systems (ECIS)*.
- Lu, Y., Gupta, A., Ketter, W., and Heck, E. V. (2016). Exploring Bidder Heterogeneity in Multichannel Sequential B2B Auctions. *MIS Quarterly*, 40(3):1–18.
- Lupton, D. (2014). Self-tracking cultures: towards a sociology of personal informatics. *Proceedings of the 26th Australian Computer-Human Interaction Conference on Designing Futures: The Future of Design.*, ACM.
- Maki, A., Burns, R. J., Ha, L., and Rothman, A. J. (2016). Paying people to protect the environment: A meta-analysis of financial incentive interventions to promote proenvironmental behaviors. *Journal of Environmental Psychology*, 47:242–255.
- Malhotra, A., Melville, N. P., and Watson, R. T. (2013). Spurring Impactful Research on Information Systems for Environmental Sustainability. *MIS Quarterly*, 37(4):1265–1274.
- Malinova, K. and Park, A. (2016). Market Design for Trading with Blockchain Technology. *SSRN Electronic Journal*.
- Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press.
- Mattern, F., Staake, T., and Weiss, M. (2010). ICT for Green – How Computers Can Help Us to Conserve Energy. In *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking*.
- Mattila, J., Seppälä, T., Naucler, C., Stahl, R., Tikkanen, M., Bådenlid, A., and Seppälä, J. (2016). Industrial Blockchain Platforms: An Exercise in Use Case Development in the Energy Industry. *ETLA Working Papers*.
- McAfee, R. P. and Vincent, D. (1993). Declining Price Anomaly. *Journal of Economic Theory*, 60:191–212.

- Mccabe, K. A., Rassenti, S. J., and Smith, V. L. (1991). Smart Computer-Assisted Markets. *Science*, 254176110(5031):534–538.
- McCalley, L. T. and Midden, C. J. H. (2002). Energy conservation through product-integrated feedback: The roles of goal-setting and social orientation. *Journal of Economic Psychology*, 23:589–603.
- McDonald, R. I., Green, P., Balk, D., Fekete, B. M., Revenga, C., Todd, M., and Montgomery, M. (2011). Urban growth, climate change, and freshwater availability. *Proceedings of the National Academy of Sciences*, 108(15):6312–6317.
- McKerracher, C. and Torriti, J. (2013). Energy Consumption Feedback in Perspective: Integrating Australian Data to Meta-Analyses on In-Home Displays. *Energy Efficiency*, 6(2):387–405.
- Meeuw, A., Schopfer, S., Ryder, B., and Wortmann, F. (2018). LokalPower: Enabling Local Energy Markets with User-Driven Engagement Context. *CHI Extended Abstracts*.
- Meeuw, A., Schopfer, S., Wörner, A., Tiefenbeck, V., Ableitner, L., Fleisch, E., and Wortmann, F. (2020). Implementing a blockchain-based local energy market: Insights on communication and scalability. *Computer Communications*, 160:158–171.
- Melville, N. P. (2010). Information Systems Innovation for Environmental Sustainability. *MIS Quarterly*, 34(1):1–21.
- Mengelkamp, E., Gärttner, J., Rock, K., Kessler, S., Orsini, L., and Weinhardt, C. (2017a). Designing microgrid energy markets. A case study: The Brooklyn Microgrid. *Applied Energy*, 210:870–880.
- Mengelkamp, E., Gärttner, J., and Weinhardt, C. (2018). Decentralizing Energy Systems Through Local Energy Markets: The LAMP-Project. *Multikonferenz Wirtschaftsinformatik*.
- Mengelkamp, E., Staudt, P., Gärttner, J., and Weinhardt, C. (2017b). Trading on local energy markets: A comparison of market designs and bidding strategies. *International Conference on the European Energy Market (EEM)*.
- Mihaylov, M., Jurado, S., Avellana, N., Van Moffaert, K., de Abril, I. M., and Nowe, A. (2014a). NRGcoin: Virtual currency for trading of renewable energy in smart grids. In *International Conference on the European Energy Market (EEM14)*.
- Mihaylov, M., Jurado, S., Moffaert, K. V., Avellana, N., and Nowé, A. (2014b). NRG-X-Change: a Novel Mechanism for Trading of Renewable Energy in Smart Grids. *International Conference on Smart cities and Green ICT Systems (Smartgreens)*.

- Miller, A., Kreder, K., D'Agostino, M., Pop, C., Epstein, R., Chen, Y., Berarducci, P., and Walters, M. (2017). GridPlus. *Whitepaper*.
- Miscione, G., Klein, S., Schwabe, G., Goerke, T. M., and Ziolkowski, R. (2019). Hanseatic Governance: Understanding Blockchain as Organizational Technology. In *Proceedings of the International Conference on Information Systems (ICIS)*.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. (2013). Playing Atari with Deep Reinforcement Learning. *arXiv preprint*, 1312.5602v.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., and Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518:529–533.
- Montemayor, L. and Boersma, T. (2018). Comprehensive Guide to Companies involved in Blockchain and Energy. *Blockchain Business, Solar Plaza*.
- Moon, J. Y. and Sproull, L. S. (2008). The role of feedback in managing the internet-based volunteer work force. *Information Systems Research*, 19(4):494–515.
- Morstyn, T., Farrell, N., Darby, S. J., and McCulloch, M. D. (2018). Using peer-to-peer energy-trading platforms to incentivize prosumers to form federated power plants. *Nature Energy*, 3(2):94–101.
- Morstyn, T., Teytelboym, A., and McCulloch, M. D. (2019). Trading, Bilateral Contract Networks for Peer-to-Peer Energy Thomas. *IEEE Transactions on Smart Grid*, 10(2).
- Munsing, E., Mather, J., and Moura, S. (2017). Blockchains for decentralized optimization of energy resources in microgrid networks. In *IEEE Conference on Control Technology and Applications (CCTA)*. IEEE.
- Nakamoto, S. (2008). Bitcoin : A Peer-to-Peer Electronic Cash System. *Whitepaper*.
- Nicolaisen, J., Petrov, V., and Tesfatsion, L. (2001). Market Power and Efficiency in a Computational Electricity Market With Discriminatory Double- Auction Pricing. *ISU Economic Report Series*, 53.
- Notheisen, B., Cholewa, J., and Shanmugam, A. (2017). Trading Real-World Assets on Blockchain An Application of Trust-Free Transaction Systems in the Market for Lemons. *Business & Information Systems Engineering*, 59(6):425–440.

- Olmstead, S. M. and Stavins, R. N. (2009). Comparing price and nonprice approaches to urban water conservation. *Water Resources Research*, 45(4):1–10.
- Ordóñez, L. D., Schweitzer, M. E., Galinsky, A. D., and Bazerman, M. H. (2009). Goals Gone Wild: The Systematic Side Effects of Overprescribing Goal Setting. *Academy of Management Perspectives*, 23(1):6–16.
- O’Reilly, P. and Finnegan, P. (2010). Intermediaries in inter-organisational networks: building a theory of electronic marketplace performance. *European Journal of Information Systems*, 19(4):462–480.
- Orlov, A., Sillmann, J., and Vigo, I. (2020). Better seasonal forecasts for the renewable energy industry. *Nature Energy*, 5(2):108–110.
- Parag, Y. and Sovacool, B. K. (2016). Electricity market design for the prosumer era. *Nature Energy*, 1(4):16032.
- Parker, G. G. and Van Alstyne, M. W. (2005). Two-Sided Network Effects: A Theory of Information Product Design. *Source: Management Science*, 51(10):1494–1504.
- Pasaoglu, G., Fiorello, D., Martino, A., Scarcella, G., Alemanno, A., Zubaryeva, A., and Thiel, C. (2012). Driving and parking patterns of European car drivers - a mobility survey. Technical report, European Commission.
- Pavlou, P. A. and Gefen, D. (2004). Building Effective Online Marketplaces with Institution-Based Trust. *Information Systems Research*, 15(1):37–59.
- Peters, M., Ketter, W., Saar-Tsechansky, M., and Collins, J. (2013). A Reinforcement Learning Approach to Autonomous Decision-Making in Smart Electricity Markets. *Machine Learning*, 92:5–39.
- Piel, J.-H., Hamann, J. F., Koukal, A., and Breitner, M. H. (2018). Promoting the System Integration of Renewable Energies: Toward a Decision Support System for Incentivizing Spatially Diversified Deployment. *Journal of Management Information Systems*, 34(4):994–1022.
- PowerLedger (2017). PowerLedger. *Whitepaper*.
- Rafaeli, S., Noy, . A., and Noy, A. (2002). Online auctions, messaging, communication and social facilitation: a simulation and experimental evidence. *European Journal of Information Systems*, 11(3):196–207.
- Rahwan, I., Cebrian, M., Obradovich, N., Bongard, J., Bonnefon, J.-F., Breazeal, C., Crandall, J., Christakis, N., Couzin, I., Jackson, M., Jennings, N. R., Kamar, E.,

- Kloumann, I., Larochelle, H., Lazer, D., McElreath, R., Mislove, A., Parkes, D. C., Pentland, A. S., Roberts, M., Shariff, A., Tenenbaum, J. B., and Wellman, M. (2019). Machine behaviour. *Nature*, 548:477–480.
- Ramchurn, S. D., Vytelingum, P., Rogers, A., and Jennings, N. (2011). Agent-Based Control for Decentralised Demand Side Management in the Smart Grid. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 5–12.
- Ramchurn, S. D., Vytelingum, P., Rogers, A., and Jennings, N. R. (2012). Putting the ‘Smarts’ into the Smart Grid: A Grand Challenge for Artificial Intelligence. *Communications of the ACM*, 55(4):88–97.
- Rassenti, S. J., Smith, V. L., and Wilson, B. J. (2003). Discriminatory Price Auctions in Electricity Markets: Low Volatility at the Expense of High Price Levels. *Journal of regulatory Economics*, 23(2):109–123.
- Reddy, P. and Veloso, M. (2011a). Learned Behaviors of Multiple Autonomous Agents in Smart Grid Markets. *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, pages 1396–1401.
- Reddy, P. P. and Veloso, M. M. (2011b). Strategy Learning for Autonomous Agents in Smart Grid Markets. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1446–1451.
- Risius, M. and Spohrer, K. (2017). A Blockchain Research Framework What We (don’t) Know, Where We Go from Here, and How We Will Get There. *Business & Information Systems Engineering*, 59(6):385–409.
- Rivola, D., Medici, V., Nespoli, L., Corbellini, G., and Strepparava, D. (2018). Hive Power. *Whitepaper*.
- Robu, V., Gerding, E. H., Stein, S., Parkes, D. C., Rogers, A., and Jennings, N. R. (2013). An online mechanism for multi-unit demand and its application to plug-in hybrid electric vehicle charging. *Journal of Artificial Intelligence Research*, 48:175–230.
- Rocklöv, J. and Dubrow, R. (2020). Climate change: an enduring challenge for vector-borne disease prevention and control. *Nature Immunology*, 21(5):479–483.
- Roger Aitken (2017). What’s The Future Of Online Marketplaces & Blockchain’s Technology Impact? www.forbes.com/sites/rogeraitken/2017/10/24/whats-the-future-of-online-marketplaces-blockchains-technology-impact/#54c1bb4b63a0, 2019-03-11.

- Rogers, A., Ramchurn, S. D., and Jennings, N. R. (2012). Delivering the Smart Grid: Challenges for Autonomous Agents and Multi-Agent Systems Research. In *AAAI Conference on Artificial Intelligence*, pages 2166–2172.
- Rosen, C. and Madlener, R. (2013). An auction design for local reserve energy markets. *Decision Support Systems*, 56:168–179.
- Roth, A. E. (1991). Game Theory as a Part of Empirical Economics. *Source: The Economic Journal*, 101(404):107–114.
- Roth, A. E. (2000). Game theory as a tool for market design. In *Game Practice: Contributions from Applied Game Theory*, pages 7–18. Springer-Science+Business Media.
- Roth, A. E. (2008). What Have We Learned from Market Design? *The Economic Journal*, 118(527):285–310.
- Roth, A. E. (2018). Marketplaces, Markets, and Market Design. *American Economic Review*, 108(7):1609–1658.
- Saha, D. and Mukherjee, A. (2003). Pervasive computing: A paradigm for the 21st century. *Computer*, 36(3):25–31.
- Sandel, M. J. (2012). *What Money Can't Buy: The Moral Limits of Markets*. Macmillan.
- Sleich, J. ., Klobasa, M. ., Götz, S. ., and Brunner, M. (2013). Effects of feedback on residential electricity demand: Findings from a field trial in Austria) : Effects of feedback on residential electricity demand: Findings from a field trial in. *Energy Policy*, 61:1097–1106.
- Schmidt, J., Hildebrandt, B., Eisel, M., and Kolbe, L. (2015). Applying Demand Response Programs for Electric Vehicle Fleets. In *Proceedings of the International Conference on Information Systems (ICIS)*.
- Schopfer, S., Tiefenbeck, V., and Staake, T. (2018). Economic assessment of photovoltaic battery systems based on household load profiles. *Applied Energy*, 223:229–248.
- Schultz, P. W., Khazian, A. M., and Zaleski, A. C. (2008). Using normative social influence to promote conservation among hotel guests. *Social Influence*, 3(1):4–23.
- Schwartz, D., Bruin, W. B. D., Fischhoff, B., Lave, L., Schwartz, D., and Lave, L. (2015). Applied Advertising Energy Saving Programs : The Potential Environmental Cost of Emphasizing Monetary Savings Advertising Energy Saving Programs : The Potential Environmental Cost of Emphasizing Monetary Savings. *Journal of Experimental Psychology*, 21(2):158–166.

- Schwartz, D., Fischhoff, B., Krishnamurti, T., Sowell, F., Dietz, T., Mayo, E., and Roethlisberger, I. (2013). The Hawthorne effect and energy awareness. *Proceedings of the National Academy of Sciences*, 110(38):15242–15246.
- Schweizer, A., Schlatt, V., Urbach, N., and Fridgen, G. (2017). Unchaining Social Businesses – Blockchain as the Basic Technology of a Crowdlending Platform. *Proceedings of the International Conference on Information Systems (ICIS)*.
- Schweppe, F., Caramanis, M., Tabors, R., and Bohn, R. (1988). *Spot Pricing of Electricity*. The Kluwer International Series in Engineering and Computer Science: Power Electronics and Power Systems.
- Seidel, S., Fridgen, G., and Watson, R. T. (2017). The Sustainability Imperative in Information Systems Research The Sustainability Imperative in Information Systems Research. *Communications of the Association for Information Systems*, 40(3):40–52.
- Seidel, S., Kruse, L. C., Székely, N., Gau, M., and Stieger, D. (2018). Design principles for sensemaking support systems in environmental sustainability transformations. *European Journal of Information Systems*, 27(2):221–247.
- SFOE (2020). Faktenblatt 1 Änderung Stromversorgungsgesetz (StVG). www.news.admin.ch/news/message/attachments/60853.pdf, 2020-04-10.
- Sharon, G. and Stone, P. (2017). A Protocol for Mixed Autonomous and Human-Operated Vehicles at Intersections. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- Shogren, J. F., Cho, S., Koo, C., List, J., Park, C., Polo, P., and Wilhelmi, R. (2001). Auction mechanisms and the measurement of WTP and WTA. *Resource and Energy Economics*, 23:97– 109.
- Sikorski, J. J., Haughton, J., and Kraft, M. (2017). Blockchain technology in the chemical industry: Machine-to-machine electricity market Blockchain-based electricity market. *Applied Energy*, 195:234–246.
- Siler-Evans, K., Azevedo, I. L., Granger Morgan, M., and Apt, J. (2013). Regional variations in the health, environmental and climate benefits of wind and solar generation. *Proceedings of the National Academy of Sciences*, 110(29):11768–11773.
- Silva, P. G. D., Karnouskos, S., and Ilic, D. (2012). A Survey Towards Understanding Residential Prosumers in Smart Grid Neighbourhoods. In *IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe)*.

- Sjoeklint, M., Constantiou, I., Trier, M., and Constantiou, I. D. (2015). The Complexities of Self-Tracking -An Inquiry into User Reactions and Goal Attainment. *23rd European Conference on Information Systems (ECIS)*.
- Slavova, M. and Constantinides, P. (2017). Digital Infrastructures as Platforms: The Case of Smart Electricity Grids. *Proceedings of the European Conference on Information Systems (ECIS)*.
- Sovacool, B. K. (2014). Diversity: Energy Studies Need Social Science. *Nature*, 511(7511):529–530.
- Stanley, S. M. (2000). The past climate change heats up. *Proceedings of the National Academy of Sciences of the United States of America*, 97(4):1319.
- Strbac, G. (2008). Demand side management: Benefits and challenges. *Energy Policy*, 36(12):4419–4426.
- Strueker, J. and Dinther, C. (2012). Demand Response in Smart Grids: Research Opportunities for the IS Discipline. In *Proceedings of the American Conference on Information Systems (AMCIS)*.
- Subramanian, H. (2017). Decentralized blockchain-based electronic marketplaces. *Communications of the ACM*, 81(1):78–84.
- Sulyma, I., Tiedemann, K., Pedersen, M., Rebman, M., and Yu, M. (2008). Experimental Evidence: A Residential Time of Use Pilot. *ACEEE Summer Study on Energy Efficiency in Buildings*, pages 292–304.
- Sun Yin, H. H., Langenheldt, K., Harlev, M., Mukkamala, R. R., and Vatrappu, R. (2019). Regulating Cryptocurrencies: A Supervised Machine Learning Approach to De-Anonymizing the Bitcoin Blockchain. *Journal of Management Information Systems*, 36(1):37–73.
- Sutton, R. S. and Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. MIT Press.
- Sutton, R. S., Precup, D., and Singh, S. (1999). Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1-2):181–211.
- Swan, M. (2012). Sensor Mania! The Internet of Things, Wearable Computing, Objective Metrics, and the Quantified Self 2.0. *Journal of Sensor and Actuator Networks*, 1(3):217–253.

- Swan, M. (2013). The Quantified Self. *Big Data*, 1(2):85–99.
- Szabo, N. (1996). Smart Contracts: Building Blocks for Digital Markets. www.fon.hum.uva.nl/rob/Courses/InformationInSpeech/CDROM/Literature/LOTwinterschool2006/szabo.best.vwh.net/smart_contracts_2.html, 2018-01-11.
- Tabi, A., Hille, S. L., and Wüstenhagen, R. (2014). What makes people seal the green power deal? - Customer segmentation based on choice experiment in Germany. *Ecological Economics*, 107:206–215.
- The Economist (2015). The trust machine - How the technology behind bitcoin could change the world. [//www.economist.com/leaders/2015/10/31/the-trust-machine](http://www.economist.com/leaders/2015/10/31/the-trust-machine), 2018-03-15.
- The Economist (2019). The Economist explains - What are the school climate strikes? www.economist.com/the-economist-explains/2019/03/14/what-are-the-school-climate-strikes, 2020-05-11.
- Thøgersen, J. (1994). Monetary Incentives and Environmental Concern. Effects of a Differentiated Garbage Fee. *Journal of Consumer Policy*, 17:407–442.
- Tiefenbeck, V. (2016). On the Magnitude and Persistence of the Hawthorne Effect - Evidence from Four Field Studies. In *European Conference on Behaviour and Energy Efficiency*.
- Tiefenbeck, V. (2017). Bring behaviour into the digital transformation. *Nature Energy*, 2(17085):1–3.
- Tiefenbeck, V., Goette, L., Degen, K., Tasic, V., Fleisch, E., Lalive, R., and Staake, T. (2018a). Overcoming Salience Bias: How Real-Time Feedback Fosters Resource Conservation. *Management Science*, 64(3):983–1476.
- Tiefenbeck, V., Tasic, V., Schöb, S., and Staake, T. (2016). Long-lasting effects or short-term spark? On the persistence of behaviour change induced by real-time feedback on resource consumption. *Proceedings of the European Conference on Information Systems (ECIS)*, pages 1–17.
- Tiefenbeck, V., Wörner, A., Schöb, S., Fleisch, E., and Staake, T. (2018b). Shower data from six hotels. [//figshare.com/articles/Shower_data_from_six_hotels/6984323/1](http://figshare.com/articles/Shower_data_from_six_hotels/6984323/1), data set.
- Tiefenbeck, V., Wörner, A., Schöb, S., Fleisch, E., and Staake, T. (2019). Real-time feedback promotes energy conservation in the absence of volunteer selection bias and monetary incentives. *Nature Energy*, 4(1):35–41.

- Tim, Y., Pan, S. L., Bahri, S., and Fauzi, A. (2018). Digitally enabled affordances for community-driven environmental movement in rural Malaysia. *Information Systems Journal*, 28(1):48–75.
- Tinschert, P., Jakob, R., Barata, F., Kramer, J.-N., and Kowatsch, T. (2017). The Potential of Mobile Apps for Improving Asthma Self-Management: A Review of Publicly Available and Well-Adopted Asthma Apps. *JMIR mHealth and uHealth*, 5(8):e113.
- Tiwana, A. (2003). Affinity to Infinity in Peer-to-Peer Knowledge Platforms. *Communications of the ACM*, 46(5):77–80.
- Torbensen, A. C. G. and Ciriello, R. F. (2019). Tuning into Blockchain: Challenges and Opportunities of Blockchain-Based Music Platforms. In *Proceedings of the European Conference on Information Systems (ECIS)*.
- UK Department for Transport (2019). National Travel Survey 2018. www.gov.uk/government/statistics/national-travel-survey-2018.
- United Nations (2019). Sustainable Development Goals. www.un.org/sustainabledevelopment/sustainable-development-goals/, 2019-04-08.
- United Nations Conference of the Parties (COP) (2015). PARIS AGREEMENT. [//unfccc.int/sites/default/files/english_paris_agreement.pdf](http://unfccc.int/sites/default/files/english_paris_agreement.pdf), 2019-04-08.
- United States Environmental Protection Agency (2010). Customer Incentives for Energy Efficiency Through Program Offerings. *National Action Plan for Energy Efficiency*.
- US Energy Information Administration (2016). Average Monthly Bill - Residential. www.eia.gov/electricity/sales_revenue_price/pdf/table5_a.pdf, 2018-03-29.
- Valogianni, K. and Ketter, W. (2016). Effective demand response for smart grids: Evidence from a real-world pilot. *Decision Support Systems*, 91:48–66.
- Valogianni, K., Ketter, W., and Collins, J. (2013). Smart Charging of Electric Vehicles Using Reinforcement Learning. *Trading Agent Design and Analysis, Workshops at the Twenty-Seventh AAAI Conference on Artificial Intelligence*, pages 41–48.
- Valogianni, K., Ketter, W., and Collins, J. (2015). A Multiagent Approach to Variable-Rate Electric Vehicle Charging Coordination. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- Valogianni, K., Ketter, W., Collins, J., and Adomavicius, G. (2019). Heterogeneous Electric Vehicle Charging Coordination: A Variable Charging Speed Approach. In *Proceed-*

- ings of the Hawaii International Conference on System Sciences (HICSS)*, volume 6, pages 3679–3688.
- Valogianni, K., Ketter, W., Collins, J., and Zhdanov, D. (2014). Effective Management of Electric Vehicle Storage Using Smart Charging. In *AAAI Conference on Artificial Intelligence*.
- Valogianni, K., Ketter, W., Collins, J., Zhdanov, D., and Robinson, J. M. (2020). Sustainable Electric Vehicle Charging using Adaptive Pricing. *Production and Operations Management*, 29(6):1550–1572.
- Van Alstyne, M. W., Parker, G. G., and Choudary, S. P. (2016). Pipelines, Platforms, and the New Rules of Strategy - Scale now trumps differentiation. *Harvard Business Review*, 94(4):54–62.
- Vandelandotte, C., De Bourdeaudhuij, I., Sallis, J. F., Spittaels, H., and Brug, J. (2005). Efficacy of Sequential or Simultaneous Interactive Computer-Tailored Interventions for Increasing Physical Activity and Decreasing Fat Intake. *Annals of Behavioral Medicine*, 29(2):138–146.
- Vázquez-Canteli, J. R. and Nagy, Z. (2019). Reinforcement learning for demand response: A review of algorithms and modeling techniques. *Applied Energy*, 235:1072–1089.
- Venkatesh, V., Morris, M. G., and Ackerman, P. L. (2000a). A Longitudinal Field Investigation of Gender Differences in Individual Technology Adoption Decision-Making Processes. *Organizational Behavior and Human Decision Processes*, 83(1):33–60.
- Venkatesh, V., Morris, M. G., and Smith, R. H. (2000b). Why Don't Men Ever Stop to Ask for Directions? Gender, Social Influence, and Their Role in Technology Acceptance and Usage Behavior. *MIS Quarterly*, 24(1):115–139.
- Vine, E., Sullivan, M., Lutzenhiser, L., Blumstein, C., and Miller, B. (2014). Experimentation and the evaluation of energy efficiency programs. *Energy Efficiency*, 7(4):627–640.
- Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H., Powell, R., Ewalds, T., Georgiev, P., Oh, J., Horgan, D., Kroiss, M., Danihelka, I., Huang, A., Sifre, L., Cai, T., Agapiou, J. P., Jaderberg, M., Vezhnevets, A. S., Leblond, R., Pohlen, T., Dalibard, V., Budden, D., Sulsky, Y., Molloy, J., Paine, T. L., Gulcehre, C., Wang, Z., Pfaff, T., Wu, Y., Ring, R., Yogatama, D., Wünsch, D., McKinney, K., Smith, O., Schaul, T., Lillicrap, T., Kavukcuoglu, K., Hassabis, D., Apps, C., and Silver, D. (2019). Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, 575:350–354.

- vom Brocke, J., Watson, R. T., Dwyer, C., Elliot, S., Melville, N., and Ross, S. M. (2012). Green Information Systems: Directives for the IS Discipline - Panel Discussion. In *Proceedings of the International Conference on Information Systems (ICIS)*.
- Vytelingum, P., Ramchurn, S. D., Voice, T. D., Rogers, A., and Jennings, N. R. (2010). Trading Agents for the Smart Electricity Grid. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- Wang, J., Wang, Q., Zhou, N., and Chi, Y. (2017). A Novel Electricity Transaction Mode of Microgrids Based on Blockchain and Continuous Double Auction. *Energies*, 10(7):1–22.
- Wang, W. and Benbasat, I. (2005). Trust in and Adoption of Online Recommendation Agents *. *Journal of the Association for Information Systems*, 6(3):73.
- Warkentin, M., Goel, S., and Menard, P. (2017). Shared Benefits and Information Privacy: What Determines Smart Meter Technology Adoption? *Journal of the Association for Information Systems*, 18(11):758–786.
- Watson, R. T., Boudreau, M.-C., and Chen, A. J. (2010). Information Systems and Environmentally Sustainable Development: Energy Informatics and New Directions for the IS Community S. *MIS Quarterly*, 34(1):23–38.
- Weber, T. A. (2016). Product Pricing in a Peer-to-Peer Economy. *Journal of Management Information Systems*, 33(2):573–596.
- Weigert, A., Hopf, K., Weinig, N., and Thorsten, S. (2020). Detection of heat pumps from smart meter and open data. In *8th D-A-CH+ Energy Informatics Conference Proceedings*.
- Weinhardt, C., Zade, M., Mengelkamp, E., Cramer, W., Hambridge, S., Hobert, A., Kremers, E., Otter, W., Pinson, P., and Tiefenbeck, V. (2019). How far along are Local Energy Markets in the DACH+ Region? In *Proceedings of the Tenth ACM International Conference on Future Energy Systems*, pages 544–549, New York, New York, USA. ACM Press.
- WePower (2017). WePower. *Whitepaper*.
- Williams, J. H., Debenedictis, A., Ghanadan, R., Mahone, A., Moore, J., Morrow Iii, W. R., Price, S., and Torn, M. S. (2012). The Technology Path to Deep Greenhouse Gas Emissions Cuts by 2050: The Pivotal Role of Electricity. *Science*, 335:53–59.

- Williams, K. J., Donovan, J. J., and Dodge, T. L. (2000). Self-Regulation of Performance: Goal Establishment and Goal Revision Processes in Athletes. *Human Performance*, 13(2):159–180.
- Williamson, O. E. (1979). Transaction-Cost Economics: The Governance of Contractual Relations. *Journal of Law & Economics*, 22(2):233–261.
- Wooldridge, M. and Jennings, N. R. (2018). Intelligent agents: theory and practice. *The Knowledge Engineering Review*, 10(2):115–152.
- World Economic Forum (2020). Annual Meeting Davos 2020. www.weforum.org/events/world-economic-forum-annual-meeting-2020, 2020-05-11.
- Wörner, A., Ableitner, L., Meeuw, A., Wortmann, F., and Tiefenbeck, V. (2019a). Peer-to-peer energy trading in the real world: Market design and evaluation of the user value proposition. *Proceedings of the International Conference on Information Systems (ICIS)*.
- Wörner, A., Meeuw, A., Ableitner, L., Wortmann, F., Schopfer, S., and Tiefenbeck, V. (2019b). Trading solar energy within the neighborhood: field implementation of a blockchain-based electricity market. *Energy Informatics*, 2(1):1–12.
- Wörner, A. and Tiefenbeck, V. (2018). The role of self-set goals in IS-enabled behavior change. In *Proceedings of the European Conference on Information Systems (ECIS)*.
- Wörner, D., von Bomhard, T., Schreier, Y.-P., and Bilgeri, D. (2016). The Bitcoin Ecosystem: Disruption beyond Financial Services? *Proceedings of the European Conference on Information Systems (ECIS)*.
- Wu, J., Arturo, E., and Gaytán, A. (2013). The role of online seller reviews and product price on buyers’ willingness-to-pay: a risk perspective. *European Journal of Information Systems*, 22(4):416–433.
- Wu, Q., Nielsen, A. H., Ostergaard, J., Cha, S. T., Marra, F., Chen, Y., and Træholt, C. (2010). Driving Pattern Analysis for Electric Vehicle (EV) Grid Integration Study. In *IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT Europe)*.
- Wu, Y. (2019). ‘0% Success’: Why Blockchain Apps Just Aren’t Taking Off - CoinDesk. www.coindesk.com/0-success-why-blockchain-apps-just-arent-taking-off, 2019-03-07.
- Wüst, K. and Gervais, A. (2017). Do you need a Blockchain? *IACR Cryptology ePrint Archiv*, 375.

- Xia, M., Huang, Y., Duan, W., and Whinston, A. B. (2012). To continue sharing or not to continue sharing? An empirical analysis of user decision in peer-to-peer sharing networks. *Information Systems Research*, 21(1):247–259.
- Yang, Y., Hao, J., Wang, Z., Sun, M., and Goran, S. (2018). Recurrent Deep Multiagent Q-Learning for Autonomous Agents in Future Smart Grid. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 2136–2138.
- Yli-Huumo, J., Ko, D., Choi, S., Park, S., and Smolander, K. (2016). Where Is Current Research on Blockchain Technology?—A Systematic Review. *PLoS ONE*, 11(10).
- Yoo, Y., Henfridsson, O., and Lyytinen, K. (2010). Research Commentary: The New Organizing Logic of Digital Innovation: An Agenda for Information Systems Research. *Information Systems Research*, 21(4):724–735.
- Zhang, C., Wu, J., Zhou, Y., Cheng, M., and Long, C. (2018a). Peer-to-Peer energy trading in a Microgrid. *Applied Energy*, 220:1–12.
- Zhang, J., Zheng, Y., Qi, D., Li, R., Yi, X., and Li, T. (2018b). Predicting citywide crowd flows using deep spatio-temporal residual networks. *Artificial Intelligence*, 259:147–166.
- Zimmermann, S., Angerer, P., Provin, D., and Nault, B. R. (2018). Pricing in C2C Sharing Platforms. *Journal of the Association for Information Systems*, 19(8):672–688.

Appendix

A Literature Review on Blockchain-Based Energy Markets

Project/ Company	Basic idea	Market design	Status	Country	Blockchain protocol
Grid+ Miller et al. (2017)	P2P trading of electricity, and controlling activators	Different market models possible	Development phase	US	Public blockchain with payment channels (Ethereum, Raiden)
HivePower Rivola et al. (2018)	Platform for P2P trading of renewable energy	Different market models possible (e.g. central optimization)	Development phase, lab-scale prototype	Switzerland	Public blockchain + state channels (Ethereum, custom channels)
LO3 Energy LO3 Energy (2017), Mengelkamp et al. (2017a)	P2P trading of solar energy within local community	Auction mechanism: Iterative double auction	Field phase: local exchange within community in Brooklyn	US	Private blockchain (Tendermint)
PowerLedger PowerLedger (2017)	P2P trading of renewable energy	Different market models possible (e.g. central optimization, auction)	Development phase	Australia	Public-private hybrid blockchain (Ethereum, EcoChain)
WePower WePower (2017)	P2P trading of renewable energy tokens	Auctions mechanism: not further defined	Development phase	Gibraltar	Public blockchain (Ethereum)

Table A.1: Relevant industry articles/whitepapers discussing blockchain-based energy markets. Most projects are still in an early development stage.

A. Literature Review on Blockchain-Based Energy Markets

Study	Basic idea	Market design	Status	Country	Blockchain protocol
Mihaylov et al. (2014a); Mihaylov et al. (2014b)	P2P trading of NRG coins within micro-grid, obtained for locally produced energy	Externally fixed pricing functions	Proof of Concept (lab-scale prototype)	Spain, Belgium	Public blockchain (custom protocol, NRGCoin)
Mattila et al. (2016)	Trading electricity between solar panel, battery and apartments	Bilateral or mediated market between solar panel and devices within one building	Case study	Finland	n.a.
Aitzhan and Svetinovic (2016)	P2P trading of solar energy and storage capacity within microgrid	Auction mechanism for stored energy, bilateral market for ad-hoc transactions	Proof of concept (simulation)	United Arab Emirates	Private blockchain (Bitcoin)
Hahn et al. (2017)	Trading between distributed prosumers, prototype using simulated building loads and PV array	Buyers bid in one-sided Vickrey auction	Proof of concept (on campus demo with one solar panel)	US, China	Private blockchain (Ethereum)
Kang et al. (2017)	P2P trading of renewable energy for plug-in hybrid electric vehicles	Iterative double auction	Proof of concept (auction simulation, blockchain system only described)	China, Norway, Canada	Private blockchain (custom protocol)
Laszka et al. (2017)	P2P trading of renewable energy and storage within microgrid	n.a.	Conceptual case study	US	n.a. (PoS blockchain necessary)
Mengelkamp et al. (2017a)	P2P trading of solar energy within local microgrid	Iterative double auction with uniform pricing	Field Phase	Germany, US	Private blockchain (Tendermint)
Munsing et al. (2017)	Coordination and payment of DER in a microgrid, active control of batteries and flexible loads through smart contracts	Distributed optimal power flow algorithm	Proof of concept (simulation)	US	Private blockchain (Ethereum)
Sikorski et al. (2017)	M2M exchange of energy in chemical industry	Bilateral market	Proof of concept (prototype of 2 prosumers and 1 consumer)	Uk, Singapore	Private blockchain (MultiChain)
Wang et al. (2017)	P2P trading of renewable energy within microgrid	Continuous double auction with adaptive aggressiveness strategy	Case study	New Zealand, Singapore	n.a.
Mengelkamp et al. (2017a)	P2P trading of renewable energy within microgrid	Iterative double auction	Simulation Study	Germany	Private blockchain

Table A.2: Relevant scientific articles discussing blockchain-based energy markets. Most of the articles present case studies or proofs of concept.

B Quartierstrom Webapp

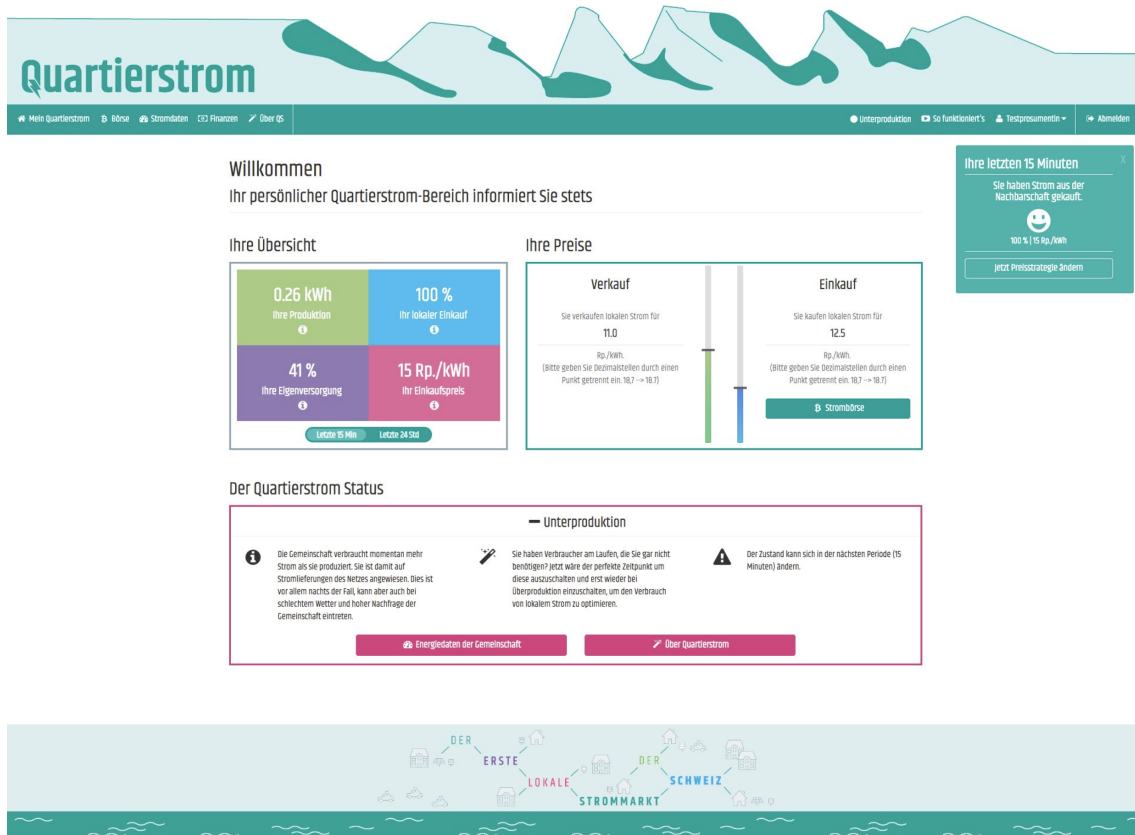


Figure B.1: Landing page of the Quartierstrom Webapp

Figure B.2: Bidding page of the Quartierstrom Webapp

B. Quartierstrom Webapp

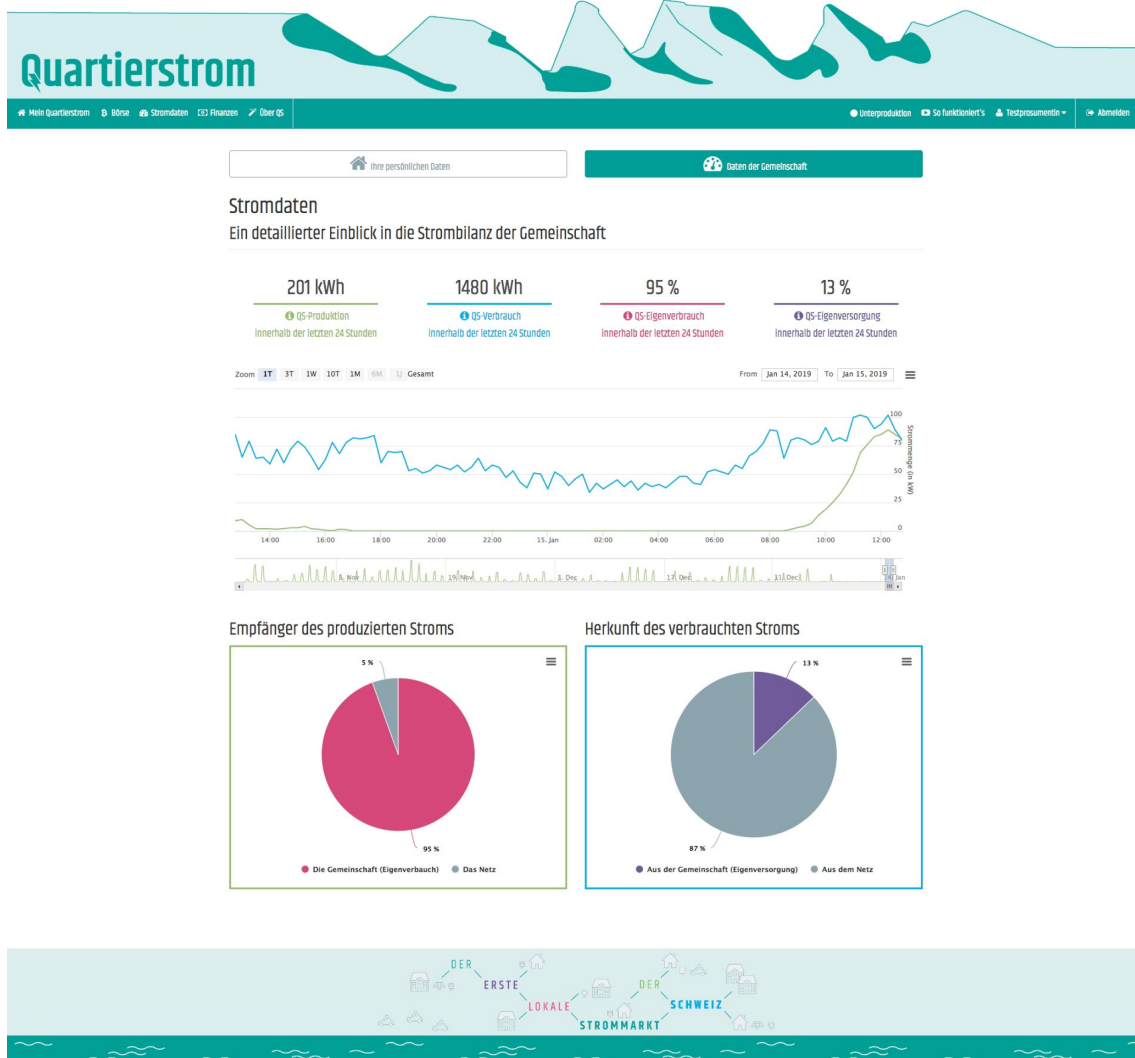
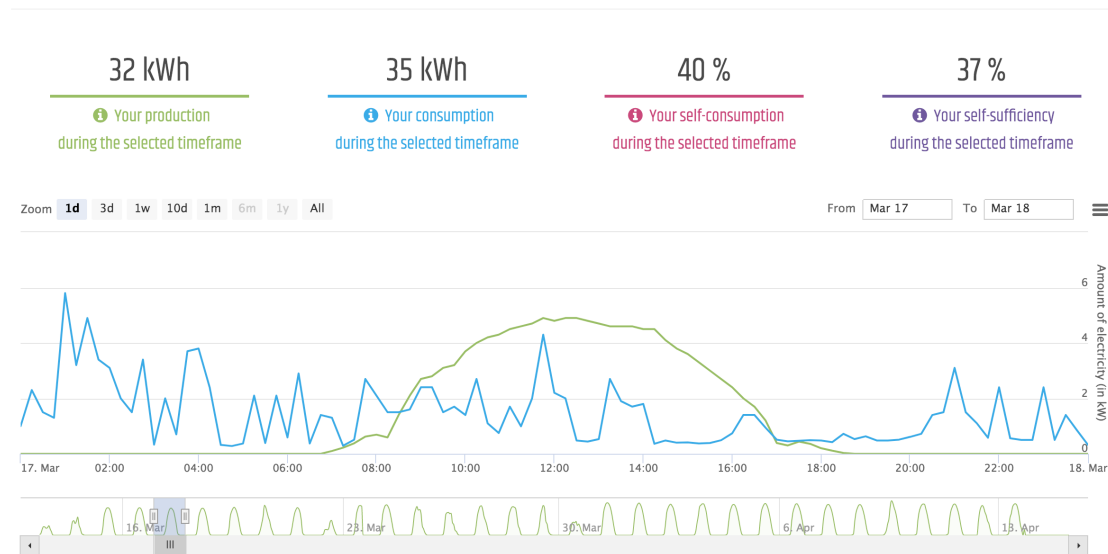


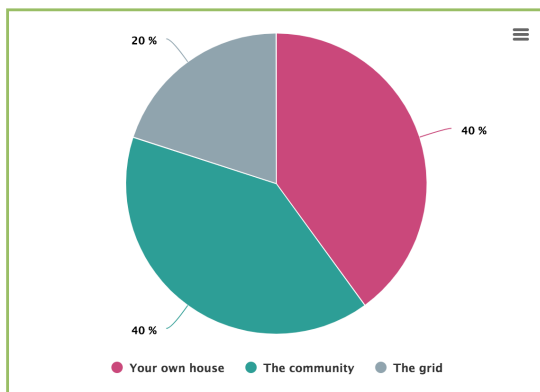
Figure B.3: Energy data of the community shown in the Quartierstrom Webapp

Electricity data

Your personal electricity production and consumption



Destination of your electricity



Origin of your electricity

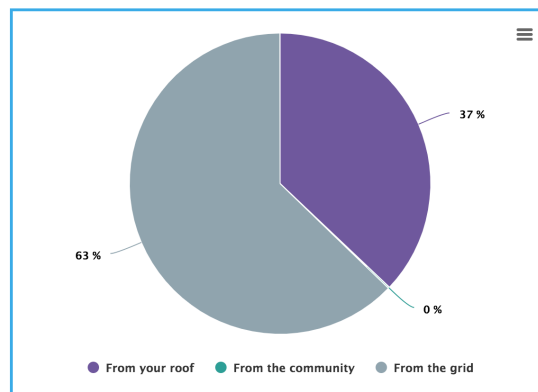


Figure B.4: Energy data of a sample household shown in the Quartierstrom Webapp (text was translated to English), see also Ableitner et al. (2020).

C Additional Analyses on Bidding Behavior

C.1 Simulation of Benchmark Strategies

The simulation is based on electricity demand profiles from an additional set of Swiss households from another rural area. The data constitutes a convenience sample of $n_{sim} = 223$ households from a partner utility provider which already had smart meters installed in apartment houses or single-family homes (see also Weigert et al. (2020)). The data is net-metered with a 15-minute resolution. The data covers the entire year of 2018, however only the month of September is used for the simulations. Energy consumption and production in that month are close to the year-round monthly averages and bidding behavior does not change as significantly anymore afterwards as it did for spring months. To put the bidding choices made in the field study into context and to understand the potential impact of extreme forms of bidding behavior, the simulation model serves as a ‘risk-free, cost-effective environment’ (Bapna et al., 2003): An agent-based model simulates a P2P market of the same size and setup, in which agents apply the a) Cost-minimizer and b) ‘Green’ bidding strategies, respectively, as defined in 7.5. The month of September for these scenarios is simulated using 10 different samples of 40 consumption profiles from the second dataset and 20 production profiles measured in the field study.

C.2 Explanatory Model for Bid Evolution

An explanatory data analysis supports the exploration of individuals’ bidding behavior over time. Caution is warranted in interpretation of the results, as sample size is small and activity in the auction varies strongly among participants. Monthly average price bids by participants are modeled using market information as independent variables:

$$y_{it} = \beta_0 + \mathbf{x}_{it}\boldsymbol{\beta} + u_{it} \quad (1)$$

with participants $i = 1, \dots, 37$ and months $t = 1, \dots, 12$ and y_{it} and u_{it} being scalars. \mathbf{x}_{it} includes five independent variables with coefficients $\boldsymbol{\beta}$. To explore the data observed, models with multiple different combinations of independent variables are computed. Below, the results for the model with independent variables $\mathbf{x}_{it} = \{month, total\ consumption, total\ production, Average\ buy/sell\ bid\ of\ all\ participants\ in\ prev.\ month, P2P\ trades\ in\ \%\ of\ consumption/production\}$ for each participant in each month. In this model, consumption and production reflect the seasonality of solar energy availability, share of P2P trades matched are a proxy for success in the auction. The average (sell/buy) bid by other participants reflects others bidding behavior which is communicated to individuals only in the monthly reports for the previous month.

C. Additional Analyses on Bidding Behavior

	Buy price bid [10 ⁻² CHF/kWh]	Sell price bid [10 ⁻² CHF/kWh]
Month	-0.002* (0.056)	-0.072 (0.057)
Total consumption [MWh]	0.017 (0.019)	0.060*** (0.019)
Total solar production [MWh]	0.016 (0.019)	-0.031* (0.014)
Average buy/sell bid of all participants in prev. month	0.184*** (0.064)	0.805*** (0.161)
% of own cons./prod. traded P2P in prev. month	0.001 (0.012)	-0.014*** (0.004)
Constant	13.406*** (2.28)	7.28*** (2.002)
R²	0.792	0.808
Participants	28	25
Observations	289	255

Table C.3: Analysis of average buy/sell price bid per participant per month (only for participants who registered on the web application). The model is a linear regression absorbing indicators for individual participants, standard errors are in parentheses; *, ** and *** indicate significance at the 5% 1% and 0.1% level respectively..

C. Additional Analyses on Bidding Behavior

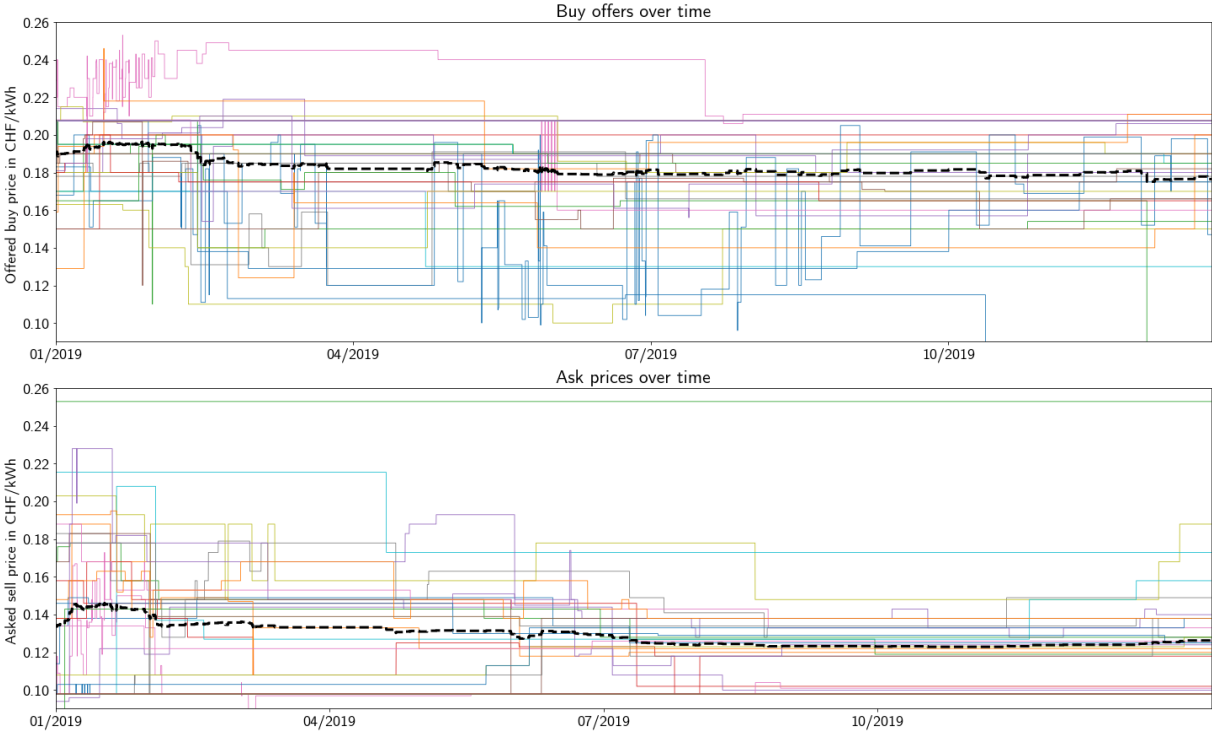


Figure C.5: Bidding curves over time. Overall trend for buy and sell bids is decreasing. Some individual bid curves display a seasonal pattern or slow, some seem erratic.

D Reinforcement Learning

D.1 Markov Decision Processes

A Markov Decision Process (MDP) is defined by a state space \mathcal{S} , an action space \mathcal{A} , a transition function $\delta(s, a)$, and a reward function $r(s, a)$ (Sutton and Barto, 1998):

$$M = \langle \mathcal{S}, \mathcal{A}, \delta, r \rangle \quad (9.1)$$

The action policy which is to be learned is defined by $\Pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ and assigns a probability of taking each possible action, based on the state of the environment. The state space \mathcal{S} includes information on the environment which is relevant for the agent to choose its action. The state S_t is a vector describing the environment at timeslot t .

Note that, in the standard formulation, state space \mathcal{S} and action space \mathcal{A} are constant finite states, such that the transition function δ can be modeled by one-step probabilities (Sutton et al., 1999). The state transition $\delta(s, a)$ depends, on the one hand, on the agent's action. On the other hand, the environment state is determined by other state variables which are not under the control of the agent. Transition probabilities are thus non-stationary.

$$\delta(s, a) = Pr\{s_{t+1} = s' | s_t = s, a_t = a\} \quad (9.2)$$

The reward function $r(s, a)$ for the agent can be formulated as:

$$r(s, a) = \mathbf{E}[r_{t+1} | s_t = s, a_t = a] \quad (9.3)$$

The action policy $\Pi(s, a)$ should maximize the expected future reward discounted by a factor γ which is defined by the action-value function, or 'Q-value' $Q(s, a)^\Pi$:

$$\begin{aligned} Q(s, a)^\Pi &= \mathbf{E}[r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \dots | s_t = s, a_t = a, \Pi] \\ &= \sum_{a \in \mathcal{A}} \Pi(s, a) [r(s, a) + \gamma v^*(\delta(s, a))] \end{aligned} \quad (9.4)$$

This means that the agent's objective is to learn the policy $\Pi(s, a)$ which chooses the action with maximal aggregate future rewards in each possible state of the environment:

$$Q(s, a)^* = \max_{\Pi} Q(s)^\Pi \quad (9.5)$$

As solving this optimization problem requires the expected value of discounted future rewards, the Q-value D.1, function approximators are usually used in practice. Initially,

linear functions were used to approximate the Q-function, but more recently, non-linear approximators such as neural networks have been used for this purpose, which will be described in Section D.3.

D.2 Multi-Agent Learning

In many real-world applications, intelligent agents interact with one another, rather than acting in an isolated single-player environment (Vinyals et al., 2019). The interaction of competing intelligent agents in a multi-agent system poses another level of complexity to the learning problem, as action spaces assume combinatorial structure in this setting.

Multi-agent learning can be modelled in the framework of markov games which extends the single-agent MDP to a matrix structure: Rewards and state changes depend on the actions of every one of the n agents, so they can be represented as an n -dimensional matrix (Littman, 1994).

In the framework defined above, an agent i 's Q-value is then defined as (Littman, 1994):

$$Q(s, \mathbf{a})_i^\Pi = Q(s, a_i, a_{-i})_i^\Pi = \sum_{a_i \in \mathcal{A}} \Pi(s, a_i) [r(s, a_i, a_{-i}) + \gamma v^*(\delta(s, a_i, a_{-i}))] \quad (9.6)$$

with a_{-i} representing a vector of the other $n - 1$ agents' actions. It is conceivable that maximizing the agent's reward by solving this problem analytically can become very complex already for only a few agents.

D.3 Deep Reinforcement Learning

The difficulty of applying reinforcement learning in practice lies in the perception and representation of real situations (Mnih et al., 2015). To be able to learn from past experience, agents need to estimate expected rewards and reduce the dimensionality of complex environments to a level which provides enough information to derive successful action policies, but is still computationally processable (Mnih et al., 2015).

In recent years, reinforcement learning techniques have improved significantly (Vinyals et al., 2019), in particular driven by the development of deep Q-learning (Mnih et al., 2015). Deep reinforcement learning involves the iterative training of a deep neural network based on the observed information about the environment and actions taken in the past. The neural net can process a wider variety of input than a mere Q-learning model can as it allows for non-linear function approximation (Mnih et al., 2015).

Deep reinforcement learning involves the iterative training of a neural network based on the observed information about the environment and actions taken in the past. The

neural net can process a wider variety of input than a simpler Q-learning model, as it allows for non-linear function approximation (Mnih et al., 2015). Mnih et al. (2013) and Mnih et al. (2015) propose the following algorithm for deep Q-learning with experience replay (see Algorithm 2). As it is difficult to derive the Q-value D.1 analytically, the optimal Q-value D.1 is approximated by a non-linear function approximator, in this case a neural network:

$$Q(s_t, a; \theta) \approx Q(s, a)^\Pi \tag{9.7}$$

with θ describing the weights of the neural network. The approach is model-free which means that the neural net does not explicitly estimate the transition function $\delta(\cdot)$, but learns how the environment behaves from the observed samples (Mnih et al., 2015).

The actions in each training step are chosen by an ϵ -greedy strategy, i.e. picking the greedy strategy $a_t = \Pi^*(s_t)$ which maximizes the expected cumulative reward $Q(s, a)^*$ with $1 - \epsilon$ probability, and picking a random action $a \in \mathcal{A}$ with probability ϵ . The rewards achieved in each step are stored in a dataset \mathcal{D} , also called replay memory. The agent then draws samples from this replay memory to update the weights θ of the neural network, i.e. the estimated Q-value iteratively, before making the next step. That way, she can potentially use each experience for multiple weight updates and use existing data more efficiently to learn (Mnih et al., 2013).

Algorithm 2 Deep Q-learning with Experience Replay (Mnih et al., 2013, p.5)

- 1: **initialize** replay memory \mathcal{D} to capacity N (memory size)
 - 2: **initialize** action-value function Q with random weights
 - 3: **for** episode = 1 to M (training episodes) **do**
 - 4: **for** $t = 1$ to T (timesteps per episode) **do**
 - 5: with probability ϵ select a random action a_t
 - 6: otherwise select $a_t = \max_a Q(s_t, a; \theta)$
 - 7: execute action a_t and observe reward r_t and state s_{t+1}
 - 8: store transition (s_t, a_t, r_t, s_{t+1}) in memory dataset \mathcal{D}
 - 9: sample random minibatch of transitions (s_j, a_j, r_j, s_{j+1}) from \mathcal{D}
 - 10: update $y_j = r_j + \gamma \max_{a'} Q(s_{j+1}, a'; \theta)$
 - 11: **end for**
 - 12: **end for**
-

List of Figures

2.1	Updated literature review in the Green IS framework	17
2.2	Overview of the studies presented in this thesis	19
3.1	Smart shower meter	29
3.2	Effect of consumption feedback	32
3.3	Effect of consumption feedback in different hotels	34
4.1	Self-set goals related to water consumption	49
4.2	Distribution of deviations from self-set goal	50
5.1	Mediated and decentralized P2P market	62
5.2	Analytical framework for P2P markets	64
5.3	Transition of the electricity market	74
6.1	Sample orderbook	89
6.2	Schematic representation of the main system components	91
6.3	Price evolution over time of day	92
6.4	Histogram of prices bid for local solar energy	95
6.5	Savings and additional revenue incurred in P2P market	96
7.1	Slider elements in the web application	111
7.2	Course of the field study	112
7.3	Energy sourcing during the field study	116
7.4	Bids and prices observed in the field	117
7.5	Average price bids	118
7.6	Evolution of bids over time	119
7.7	Number of price changes per month	120
7.8	Survey on price bidding vs. automated pricing	122
7.9	Simulation of different bidding strategies	124
8.1	Simplified scheme of a Markov Decision Process	131

- B.1 Landing page of the Quartierstrom Webapp 187
- B.2 Bidding page of the Quartierstrom Webapp 188
- B.3 Energy data of the community shown in the Quartierstrom Webapp 189
- B.4 Energy data of a sample household shown in the Quartierstrom Webapp . 190
- C.5 Bidding curves over time 193

List of Tables

3.1	Overview of the participating hotels	28
3.2	Main treatment effect	33
3.3	Main treatment effect with log transformation	34
3.4	Treatment effect for different hotels	35
4.1	Relationship between goal setting/goal difficulty and water consumption .	53
5.1	Analytical framework for P2P markets	67
5.2	Blockchain vs. centrally operated electricity exchange platform	71
5.3	Implementation of market tasks in blockchain-based energy market	77
7.1	Summary of the load profiles collected	114
7.2	Average price bids per month	120
A.1	Relevant industry articles/whitepapers on blockchain-based energy markets	185
A.2	Relevant scientific articles on blockchain-based energy markets	186
C.3	Analysis of average buy/sell price bid per participant per month	192