Essays on Market Design and Algorithms for Distributed and Sustainable Energy Systems

Inauguraldissertation

zur

Erlangung des Doktorgrades

der

Wirtschafts- und Sozialwissenschaftlichen Fakultät

der

Universität zu Köln

2025

vorgelegt

von

M. Sc. Philipp Artur Kienscherf

aus

Gütersloh

Referent:	Prof. Dr. Wolfgang Ketter
Korreferent:	Prof. Yixin Lu, Ph.D.
Tag der Promotion:	

Acknowledgements

Pursuing a PhD can be a long and bumpy ride. In my case it was certainly longer (and bumpier still) than most. Reflecting on this journey, it is shaped by the people you meet and the things you learn, in that order.

None of this work would have been possible without numerous people with whom I have crossed paths throughout my PhD. First and foremost, I would like to thank Prof. Wolfgang Ketter, who has been the catalyst to bring ideas to life, push forward a research agenda, and was a cornerstone in helping me to develop as a researcher. Wolf, you have introduced me to the realm of methods used throughout my research, to many colleagues and later collaborators, and, through your vision, to the IS3 staff.

Second, I would like to thank Yixin Lu, to whom I am extremely grateful for accepting to co-referee my thesis and who I have been lucky enough to have co-authored a paper with. Yixin, while we only collaborated intensely throughout the last phase of my PhD, you have helped me immensely in structuring and framing our research, providing invaluable lessons. I am further truly grateful to all of my other co-authors who helped bring this thesis to life. Verena Tiefenbeck and Anselma Wörner, Arthur Carvalho, and Long He have all been great contributors to shaping the chapters of this dissertation and have widened my perspective on research, both methodologically and personally.

I would like to thank Prof. Dr. Stefan Seidel, who kindly agreed to chair my PhD defense. It is an absolute honor to have a highly accomplished IS scholar like you chair my PhD committee. Thank you for dedicating your valuable time toward this.

Graduating from university, I was dead set on pursuing a PhD and stepping into academia. In hindsight, I was certainly oblivious to the implications and, frankly, what academia really was. But from my very first day at the University of Cologne, I was lucky enough to not only have a highly motivated supervisor but also a brilliant team to rely on. Nastaran, being the first one, and Karsten, co-joining with me at the beginning of a new semester, constituted the first members of Wolf's chair in Cologne. With none of us having the "home advantage" or having any idea what to do or how to do it, I felt we maneuvered our self-inflicted chaos superbly. Chris, without you we might have actually broken something beyond repair, so I really appreciate you having our backs and helping us "understand" university processes and policies.

From then on, our group only grew, and many inspiring personalities and researchers enriched my path. Christian, Muhammed, Ramin, Anna, Fred, Jannik, Lennard, Philipp Kai, Saber, and Julius, you all shaped our group, and it was a pleasure to spend time together, though in fact cut short by the pandemic.

I would like to express my gratitude towards the team at EWI, with which we co-located early on and whom I have been lucky enough to call my colleagues for the past few years. Even though this dissertation both conceptually and organizationally emerges from the field of Information Systems, the discussions with energy economists have greatly contributed to its practicality and have shaped my thinking immensely.

Lastly, I am eternally thankful to those closest to me. This journey has been supported by an ever-understanding family and incredible friends. Even when I had doubts, you encouraged me and helped me to complete this thesis. Annelie, you have been at my side for almost the entire time and at the forefront of all ups and downs. I could not have done this without you.

Philipp Artur Kienscherf — Cologne, September 2025

Summary

The dissertation “Essays on Market Design and Algorithms for Distributed and Sustainable Energy Systems” investigates how the large-scale electrification of transport can be coordinated through digital markets and intelligent algorithms. As electric vehicles become ubiquitous, millions of charging decisions intertwine with the real-time dynamics of power grids, creating a cyber-physical system. The dissertation finds that this challenge can be addressed by more means than engineering grid capacity: it proposes new market designs and machine learning methods that translate physical constraints into economic signals, allowing autonomous software agents to coordinate their behavior efficiently and sustainably.

The work is built on the insight that large-scale EV charging is a distributed decision problem. Each driver cares about having enough charge when needed, yet the collective charging pattern determines local grid stress, costs, and emissions. Instead of relying on centralized schedulers or detailed travel forecasts, the work develops decentralized mechanisms in which intelligent agents learn to act on real-time prices and limited local information. Through multi-agent reinforcement learning, household charging agents observe only their state of charge, time of day, and previous prices but still learn to bid effectively for charging power in repeated auctions. Simulations on realistic low-voltage grids show that simple linear learning agents can reach near-optimal outcomes within a few percentage points of a full-information benchmark while avoiding the instability and training burden of deep neural networks. This establishes a foundation for self-organizing, bottom-up control of electric loads.

Moving beyond individual grid situations, the thesis then explores how market architecture itself can be made trustless and transparent. It develops a blockchain-based bundle trading market that enables EV owners to buy and sell time-specific charging rights without a central auctioneer, using smart contracts that guarantee correctness and auditability. The system achieves the same efficiency as centralized clearing but adds resilience and privacy, illustrating how distributed ledger technology can support critical energy services. The concept of market-based coordination is further generalized to decentralized autonomous organizations (DAOs) that manage both physical and financial resources. Here the dissertation designs a mechanism in which agents not only trade physical bundles but also

issue contingent financial claims, allowing them to share risk while preserving autonomy and data privacy. The mechanism converges to the same equilibrium a risk-neutral central planner would choose, demonstrating that even complex, stochastic resource allocation can be governed through decentralized markets.

The final part of the dissertation widens the lens to the national scale. Using Germany's planned Deutschlandnetz of fast-charging stations as a natural laboratory, the dissertation develops a spatial competition model to understand how the location of stations, regional demand, and regulatory price caps shape prices and investment incentives. The analysis shows that uniform national price caps can unintentionally reduce service in sparsely populated areas and that lax ownership constraints can foster market power, harming consumer welfare. Policy options include regionally differentiated price caps, dynamic tendering processes, and real-time data-sharing requirements to foster both competition and equity.

What unites these strands is a sociotechnical vision of the future energy system. Rather than treating markets and technology separately, the dissertation demonstrates that market design, machine learning, and physical infrastructure must be co-designed. By coupling incentive-compatible mechanisms with adaptive algorithms, it shows how autonomous agents can collectively manage scarce resources and uncertainty, achieving reliable and cost-effective charging without heavy-handed central control. The work thus provides a conceptual and methodological blueprint for digital energy platforms that are at once economically efficient, technologically scalable, and sustainable.

Contents

List of Figures

List of Tables

1 Introduction

The increasing digitization of infrastructure systems, such as transportation, energy, and mobility, heralds a new era of tightly interwoven digital and physical processes. These systems, often referred to as *cyber-physical systems* (CPS), are characterized by dynamic interactions between computational intelligence, networked communication, and embedded physical components (E. Lee, 2008; Rajkumar et al., 2010). Their defining feature is the co-dependence of digital decision-making and real-world outcomes: software agents monitor, control, and optimize physical assets in real time, often autonomously and across decentralized networks. This melding of cyberspace and physical space, on the one hand, gives rise to unprecedented opportunities for efficiency, but, on the other hand, comes with novel challenges in ensuring safety, reliability, and resilience (Alur, 2023).

A prime example of such a transformation is the *electrification of transportation*. Electric vehicles (EVs) represent a shift in propulsion technology and act as mobile cyber-physical nodes that interact bidirectionally with the electric grid. Charging behavior depends on both individual mobility needs and aggregate grid conditions, creating a tight coupling between mobility patterns and electricity market dynamics. This interplay demands novel *market-mediated control* mechanisms: dynamic pricing, auctions, peer-to-peer trading, and real-time demand response (Robu et al., 2012). Such mechanisms translate physical constraints, like transformer capacity or line congestion, into economic signals that guide both human and algorithmic decision-makers.

In parallel, the *Information Systems* (IS) discipline has evolved to address these sociotechnical complexities by studying how information technologies enable and constrain organizational and societal processes (Markus et al., 2008; Orlikowski & Robey, 1991). Early MIS research focused on back-office automation and decision support, but contemporary IS scholarship emphasizes the design and governance of digital infrastructures that span multiple stakeholders and regulatory domains (Bygstad & Bendik, 2013; Yoo et al., 2012). Central to this evolution is the shift from monolithic, centrally controlled systems toward *platform ecosystems* and *decentralized architectures*. These digital platforms serve both as markets and coordination mechanisms, balancing conflicting objectives such as economic efficiency, equity, and sustainability across diverse user communities.

A prime example of increasingly decentralized information systems is the rise of *distributed ledger technologies* (DLTs) and *blockchain*. These innovations have introduced new possibilities for transparent, tamper-resistant coordination without centralized intermediaries (Andoni et al., 2019). In energy and mobility contexts, blockchain-based markets can record transactions, enforce smart contracts, and facilitate peer-to-peer exchanges of charging rights or grid services (Mengelkamp et al., 2018). Yet, integrating DLTs with existing regulatory frameworks and legacy infrastructure raises questions of scalability, governance, and interoperability.

Designing effective digital markets within CPS requires addressing multiple layers of complexity. At the **agent layer**, autonomous software or human actors must respond to local information and incentives while coping with uncertainty and strategic interactions (Myerson, 2013; Stone & Veloso, 2000). At the **smart market layer**, mechanism design must ensure incentive compatibility, budget balance, and computational tractability, often under regulatory constraints and privacy requirements (Bichler et al., 2010; Roth & Sotomayor, 1990). At the **infrastructure layer**, physical network constraints and reliability standards must be upheld despite fluctuating demand and decentralized control. Achieving coherence across these layers is a central challenge in both CPS and IS research.

Amid these technological advances, *learning-based* approaches, particularly *reinforcement learning* (RL), have emerged as a powerful tool for decentralized decision-making under uncertainty (Sutton & Barto, 2018). In CPS domains, RL agents can learn optimal control policies from real-time feedback, adapting to dynamic environments and co-evolving with market mechanisms. However, non-cooperative multi-agent settings introduce additional challenges of non-stationarity, credit assignment, and equilibrium selection. Understanding how to co-design market rules and learning algorithms is thus critical for deploying robust, scalable CPS solutions.

Before delving into the specific chapters, it is essential to recognize the unifying perspective: that *smart infrastructures* are inherently *sociotechnical ecosystems*, where digital platforms and market mechanisms serve as the glue binding individual preferences, strategic behavior, and physical realities into coherent, adaptive wholes.

Throughout the subsequent chapters, this cumulative dissertation develops and evaluates digital coordination mechanisms for EV charging systems across these dimensions. Anchored in the IS tradition of *design science* and sociotechnical system research, and grounded in the CPS tradition of embedded, real-time control, the work spans: (i) local agent learning and bidding strategies in micro-level, real-time markets, (ii) decentralized market architectures including blockchain-enabled platforms incorporating more expressive bidding languages in digital markets, (iii) integrated financial and operational decision-making in decentralized autonomous organizations, and (iv) macro-level spatial competition and regulatory analysis. Through simulation, analytical modeling, and case studies, this research seeks to answer: *How can information systems be designed to coordinate autonomous agents and markets, respecting physical constraints and societal objectives in cyber-physical electric mobility infrastructures?*

1. **Chapter 1: Parsimonious Multi-Agent Reinforcement Learning for Smart EV Charging.** We develop a *real-time, auction-based* charging mechanism in which lightweight reinforcement-learning agents learn to bid for charging slots under locational marginal pricing. This chapter demonstrates how MARL can achieve near-optimal coordination in a simulated residential grid without requiring full itinerary forecasts. Joint work with Anselma Wörner, Verena Tiefenbeck, and Wolfgang Ketter. I was responsible for the conceptualization of the research idea and implementation of the software framework. I was the lead author of the manuscript.

2. **Chapter 2: A Blockchain-Based Bundle Trading Market for EV Charging.** Building on chapter 1's auction framework, we introduce a *decentralized, blockchain-enabled* market for trading bundles of charging rights. We design smart-contract protocols that ensure transparency, auditability, and efficient resource allocation under regulatory constraints. Joint work with Arthur Carvalho and Wolfgang Ketter. I was responsible for the conceptualization of the research idea and implementation of the software framework. I was the lead author of the manuscript.

3. **Chapter 3: Market-Based Optimization in Decentralized Autonomous Organizations.** Extending the bundle-trading concept beyond EV charging, this chapter presents an *integrated physical-financial* auction mechanism for DAOs. We show how self-issued contingent claims and shared-resource auctions align incentives and manage risk across decentralized divisions. Joint work with Wolfgang Ketter. I was responsible for the

conceptualization of the research idea and implementation of the software framework. I was the lead author of the manuscript.

4. **Chapter 4: Spatial Competition in Fast-Charging Networks.** Returning to the EV charging domain at a macro scale, we develop a *computational framework* to analyze how locational and ownership structures shape pricing, investment incentives, and welfare in national fast-charging networks. Policy-relevant insights emerge on price caps, tender design, and regional equity. Joint work with Yixin Lu, Long He, and Wolfgang Ketter. I was responsible for the conceptualization of the research idea and implementation of the software framework. I was the lead author of the manuscript.

The remainder of the introduction is divided into three sections: Section 1.1 provides an outline of the content of the four essays. Section 1.2 discusses the methodological approaches of the essays, as well as limitations and potential areas of further research. Section 1.3 serves as a point of reference for notational choices throughout the thesis.

1.1 Outline

The four main chapters of this thesis are self-contained.

1.1.1 Parsimonious Multi-Agent Reinforcement Learning for Smart EV Charging

This chapter addresses the challenge of coordinating electric-vehicle (EV) charging in residential distribution grids without centralized scheduling by an aggregator or full knowledge of drivers' itineraries. We begin by framing EV charging as an *auction-based online resource-allocation* problem: each household agent bids for charging power at discrete time slots, facing locational marginal prices set by a benevolent system operator. Recognizing the combinatorial complexity of bidding over an entire time horizon, we propose a *real-time* mechanism in which lightweight reinforcement-learning agents use Q-learning with experience replay to learn bidding strategies based solely on current state-of-charge, time of day, and price trajectories.

After defining the state, action, and reward structure, we detail our simulation environment: a representative IEEE low-voltage feeder (Schneider et al., 2018) populated with realistic household load profiles and EV arrival/departure patterns. We describe how the system operator computes nodal prices via a linearized optimal power flow model, internalizing transformer and line congestions. In successive experiments, we compare linear function approximator agents against small deep Q-network (DQN) agents, evaluating convergence speed, policy stability, and welfare outcomes (aggregate charging cost and peak reduction).

Our results indicate that parsimonious linear learners rapidly converge to near-optimal schedules—within 5-10% of the omniscient and thus unrealistic optimum—while deep models require extensive tuning and may exhibit oscillations under non-stationarity. We close with insights on the trade-off between model complexity and robustness in multi-agent settings and discuss integration pathways for household-level charging agents in future smart-grid deployments.

1.1.2 A Blockchain-Based Bundle Trading Market for EV Charging

Building on the agent-level coordination of chapter 1, this chapter introduces a fully *decentralized market architecture* for trading EV charging rights using blockchain and smart-contract technology. We first review the regulatory context, in particular unbundling rules and the ability to integrate non-participating households. Building on these, we establish desiderata for a decentralized coordination system: transparency, auditability, privacy of driving itineraries, and resilience against single-point failures.

We then present the *Bundle Trading Market (BTM)* mechanism (Guo et al., 2007) adapted for EV charging: charging rights are partitioned into homogeneous time-slot bundles at each charging node; agents submit sealed bids for bundles; and a distributed dealer smart contract matches bids via a supply-demand clearing linear program. We detail the on-chain implementation for order submission, matching, and settlement, leveraging a strong-duality verification protocol to ensure correctness.

Through simulation on the same low-voltage grid model as in Chapter 1, we compare the blockchain-enabled BTM against the centralized auction and demonstrate that the on-chain approach attains equivalent allocative efficiency while preserving itinerary confidentiality and removing trust in a central operator.

We also examine gas-cost trade-offs and propose off-chain computations to reduce transaction overhead. The chapter concludes by outlining integration with Chapter 1's learning agents, enabling household software to bid autonomously via blockchain wallets.

1.1.3 Market-Based Optimization in Decentralized Autonomous Organizations

This chapter generalizes the BTM framework beyond EV charging to *integrated physical and financial decision-making* in fully decentralized autonomous organizations (DAOs). Motivated by scenarios in which divisions share scarce infrastructure and face idiosyncratic risk (e.g., grid-connected charging fleets), we develop a *smart-market mechanism* that auctions both physical resource bundles and self-issued contingent claims (Arrow-Debreu securities).

We formalize the central planner's stochastic cost-minimization problem under coherent risk measures and then derive a decomposition via decentralized bundle trading. Each agent solves a local optimization for physical activities and financial hedges, submitting improving-bundle bids to the DAO's market maker. Clearing solves a joint linear program whose dual prices reveal marginal values for both resources and states of nature.

Extensive experiments with synthetic cost distributions and risk-aversion parameters show that the DAO mechanism converges to the same *stochastic-endogenous equilibrium* as a fully centralized, risk-neutral planner, while requiring no agent to divulge private cost functions. We analyze scalability in the number of agents, resources, and securities and discuss governance considerations—on-chain versus off-chain—and potential applications in shared-mobility cooperatives and microgrid communities.

1.1.4 Spatial Competition in Fast-Charging Networks

The final chapter situates the micro-scale market mechanisms into a *macro-scale, spatial competition* setting. Focusing on Germany's national fast-charging rollout named the *Deutschlandnetz*, including 900 stations and multiple CPOs, we develop regional demand projections and a computational framework that combines a spatial competition model with congestion-dependent delay costs and locational attractiveness.

After calibrating demand densities and cost parameters for 2030 projections, we simulate equilibrium pricing and market shares under varying attributes of competition: (i) comparison of the tender's outcome to both ex-ante possible and impossible counterfactuals (ii) impact of price cap regulation on the competition. Key findings include: (a) prices inversely correlate with local station density but not with absolute demand; (b) uniform price caps can induce underutilization and deter investment in low-demand regions; and (c) ex-post market power emerges when ownership constraints are lax, harming consumer welfare.

We conclude by deriving policy recommendations: regionally differentiated price caps, dynamic tender designs that balance equity and competition, and real-time data sharing mandates to enable Chapters 1-3's digital market tools to function effectively at scale.

1.2 Methodology, Limitations, and Further Research

Each chapter of this dissertation addresses a distinct facet of digital coordination in cyber-physical mobility infrastructures. Accordingly, each employs a tailored methodological approach, alongside specific limitations and avenues for future work.

Chapter 2: Parsimonious Multi-Agent Reinforcement Learning for Smart EV Charging. This chapter formulates EV charging as a real-time, auction-based resource allocation problem and employs *multi-agent reinforcement learning* (MARL) to solve it. Each household agent uses Q-learning with experience replay to learn bidding strategies based solely on local state (state of charge, time to departure) and nodal price feedback (Mnih et al., 2013; Sutton & Barto, 2018). We simulate on a realistic IEEE distribution-feeder model, computing locational marginal prices via a linearized alternating current optimal power flow (AC-OPF) at each time step.

Limitations: Our agents assume perfect observability of nodal prices and ignore communication delays. The learning environment is stationary aside from agent interactions, whereas real grids may experience exogenous shocks (e.g., renewable variability). Moreover, we focus on single-day scheduling, leaving multi-day charging patterns for future research.

Further Research: Extending agents to handle partial observability (via partially observable Markov decision process (POMDP) formulations) and non-stationary price dynamics; integrating vehicle-to-grid services; and field experiments with hardware-in-the-loop to validate transferability beyond simulation.

Chapter 3: A Blockchain-Based Bundle Trading Market for EV Charging. Building on chapter 2's auction framework, this chapter designs a *decentralized on-chain market* using smart contracts. Charging rights are tokenized into time-slot bundles; agents submit sealed bids to a dealer contract that clears via a linear program and strong-duality verification protocol (Andoni et al., 2019; Christidis & Devetsikiotis, 2016). We implement core functions in Solidity pseudocode and benchmark gas costs and clearing latency in an Ethereum test net emulator.

Limitations: Gas-cost estimates depend on network conditions and may not generalize to production chains. Privacy is preserved only at the bid-level; on-chain clearing reveals aggregate demand, which could leak usage patterns. We abstract away settlement finality delays and potential blockchain reorganizations.

Further Research: Exploring Layer-2 scaling solutions (e.g., roll-ups) to reduce on-chain overhead; incorporating zero-knowledge proofs to enhance bid privacy; and prototyping a hybrid off-chain/on-chain matching engine for large-scale pilots.

Chapter 4: Market-Based Optimization in Decentralized Autonomous Organizations. This chapter generalizes the bundle-trading market to *integrated physical-financial* auctions within DAOs. We start with a stochastic, risk-averse cost-minimization problem under coherent risk measures, then decompose it into local bundle-determination problems and a centralized Market Matching Problem (MMP) whose duals yield resource and state-price signals (Bialas & Karwan, 1984; Gabriel et al., 2013). Agents iteratively submit improving-bundle bids for both physical assets and contingent claims (Arrow-Debreu securities).

Limitations: We assume complete markets for contingent claims and neglect transaction costs in issuing securities. The convergence proofs rely on convexity and risk-measure coherence, which may fail under non-convex operational constraints (e.g., integer decisions). Agent risk preferences are homogeneous coherent measures, ignoring behavioral heterogeneity.

Further Research: Extending to incomplete markets and non-convex resources (e.g., unit commitment); modeling heterogeneous and possibly non-coherent risk preferences; and implementing a proof-of-concept DAO on a permissioned ledger to study governance dynamics.

Chapter 5: Spatial Competition in Fast-Charging Networks. This chapter analyzes macro-scale competition among Charging Point Operators (CPOs) using a *spatial competition model* on a geographical graph. We incorporate congestion-dependent delay costs, locational attractiveness, and ownership constraints and solve for equilibrium prices and market shares via counterfactual simulations calibrated to Germany's national fast-charging rollout (AlixPartners, 2019; McKinsey, 2023). Tender-rule scenarios (unconstrained, zonal, price-cap) are compared for welfare and investment impacts.

Limitations: Demand projections for 2030 are uncertain and assume technology and subsidy trajectories that may not materialize. The model abstracts from dynamic entry and exit of CPOs and does not endogenize infrastructure investment decisions. We also assume static consumer preferences and perfect information on station locations.

Further Research: Incorporating dynamic investment models and endogenous network expansion; calibrating with real-world usage and pricing data as the German network matures; and extending the framework to multi-modal charging (e.g., home, workplace, en route) to capture interdependencies across charging ecosystems.

Beyond these chapter-specific considerations, an overarching limitation is the reliance on simulation and stylized models. Future empirical validation through pilot deployments, field data analysis, and stakeholder workshops will be crucial to refine these mechanisms and ensure their real-world applicability.

1.3 Notation

The different chapters of this thesis rely on mathematical notation to express the respective concepts. Even though the application domains are related, there are subtle differences between chapters to allow for self-containment alongside domain-specific expressiveness. Each of the chapters relying heavily on mathematical notation provides a table of notation as a point of reference. Some general choices are matrix-valued variables are bolded upper-case, e.g., $\mathbf{X}$. Scalar and vector-valued variables in general are lower-case. Sets or set cardinalities are upper-case, e.g., L for electric lines[1]. Prices are denoted by p, and costs are denoted by c (although c can also denote a coefficient, when there is no danger of

[1] Note that L_e denotes the length of an edge in chapter 5 as is a standard convention in the related literature for this chapter.

confusion). Electric power, even as a scalar variable, is denoted by an uppercase P to differentiate it from prices. Time is usually discretized in periods $t \in T$, where T can be either the number of periods such that $t \in \{1, ..., T\}$ or the set of periods itself.

2 Parsimonious Multi-Agent Reinforcement Learning for Smart EV Charging[2]

2.1 Introduction

The rapid adoption of electric vehicles (EVs) presents a formidable challenge for existing distribution grids. As EVs become more prevalent, their charging requirements threaten to impose significant stress on these systems, potentially leading to overloads and inefficiencies. EV drivers are generally more concerned with achieving the total required charge by a specific deadline than with the instantaneous charging power. There is thus inherent flexibility in EV charging, given that vehicles are typically idle for large portions of the day or night, which can be strategically managed to mitigate this stress. Recent events, like the Texas electricity market crisis, underscore the necessity for smart control of flexible electricity loads and pricing mechanisms to ensure reliable and affordable electricity supply (The Economist, 2021).

This scenario, characterized by the intertemporal complementarity of charging and constraints in available energy or grid infrastructure for multiple EVs, can be managed through either a combinatorial auction for a relevant time horizon or a simple real-time allocation scheme. The second approach, while presenting challenges such as the exposure problem, may benefit from the use of learning agents, as it reduces the dimensionality of the bidding space, potentially enhancing the learning speed of intelligent agents. Within a "smart grid", the coordination of charging does not only require electric vehicles to be coordinated, but also other distributed energy resources (Parag & Sovacool, 2016; Williams et al., 2012). As the power system has to be balanced at all times, this coordination is essential for the electrification of transportation and other sectors.

Recent advancements in market design research have aimed to address these allocation challenges through the development of various goods and services allocation mechanisms. Auction design, which entails setting allocation and

[2]Joint work with Anselma Wörner, Verena Tiefenbeck, and Wolfgang Ketter.

payment rules based on participants' bids, has garnered significant attention in applications ranging from electricity markets to wireless spectrum auctions. However, the increasing complexity of allocation tasks has led to more intricate auction designs, which may inadvertently offset efficiency gains due to higher transaction costs and the complexity of formulating effective bidding strategies. For human bidding in electric vehicle charging, such market designs may only exhibit limited usefulness, as the potential value of deviating charging schedules is outweighed by cognitive overload and transaction costs. Concurrently with advances in market design, artificial intelligence research, particularly in machine learning (ML), has made substantial strides, enabling intelligent software agents to tackle highly complex tasks with superhuman performance (Bichler et al., 2010; Ketter et al., 2016). This progress raises an intriguing question: can auction mechanisms and intelligent bidding agents be co-designed to optimize performance in resource allocation tasks, specifically for EV charging in distribution grids?

In addressing the complexity of EV charging, two primary approaches emerge: designing a complex auction and bidding protocol that captures all possible allocations, allowing agents to reveal their true valuations, or employing simpler auction protocols that complicate the bidding strategy inference. The former approach, while theoretically optimal, poses significant practical challenges due to its complexity. Conversely, the latter approach, though less precise, offers potential advantages by simplifying the learning space for intelligent agents. Our research explores the impact of simple, yet incomplete, resource allocation schemes on the learning of suitable bidding strategies by software agents. We employ multi-agent reinforcement learning (MARL) to address this question (Sutton & Barto, 2018). MARL, previously examined in the context of computer games to benchmark learning algorithms, demonstrates its power in multi-agent settings. For instance, Vinyals et al. (2019) demonstrated that a MARL algorithm could surpass human players in StarCraft, showcasing the method's power in multi-agent settings. By applying MARL to the practical, real-world application of smart load scheduling in a local energy market, we extend this research into a domain with significant societal implications.

Our study serves as a starting point for discussing the application of MARL in emerging market-based resource allocation and coordination problems, similarly seen in real-time road transportation markets (Cramton et al., 2018). AI's role is pivotal in managing the complexity and dynamism of distributed energy resources, where the electricity market must balance individual consumption patterns with dynamic energy supply, subject to physical grid constraints. The increasing demand for personal transportation and the growing adoption of EVs exacerbate the demand on the electricity grid. Intelligent control of charging loads is crucial to prevent grid failures. Online mechanisms for electricity allocation, as proposed in the literature, offer precise, dynamic incentives for "grid-serving" behavior by providing varying price signals based on grid usage, location, or marginal generation costs (Flath et al., 2014; S. Huang et al., 2015; Robu et al., 2012). Implementing these sophisticated online market mechanisms in the energy sector necessitates autonomous software agents that can act on market signals in real time. Our research explores the practical adoption of such mechanisms, using distributed, intelligent agents to manage individual charging loads and respond to smart market signals at the household level.

To mirror real-world conditions, we conducted a simulation experiment using residential load data in a representative IEEE distribution grid model based on a real UK setting. Our model incorporates physical grid restrictions and users' economic rationale to examine how load-shifting by learning agents can improve charging power allocation. In this scenario, reinforcement learning (RL) agents schedule EV charging based on nodal prices at their grid locations. These agents employ Q-learning with experience replay (Mnih et al., 2013) to make real-time charging and bidding decisions without forecasting future driving demand. We evaluated different learning models, including linear learners and small neural networks (DQN), for their performance in this simulation. We deliberately avoided very deep or large networks due to their higher sensitivity to the environment's instability caused by multi-agency and their impracticality in energy management due to training costs.

Our findings indicate that RL agents outperform current naive charging strategies. While well-parameterized deep learning models can surpass linear function approximation, they may perform substantially worse if not optimized correctly. Although our focus is on the smart grid, our findings are applicable to other dynamic market settings, such as Dutch flower auctions (Lu et al., 2019).

Our study bridges the gap between multi-agent reinforcement learning and mechanism design for a relevant and socially significant use case. This multidisciplinary research contributes to the understanding of intelligent agents in the smart grid by developing a novel techno-economic co-simulation of a residential smart grid, serving as a test bed for intelligent agent research. The insights gained can inform the development of smart load scheduling agents and aid policymakers in designing future energy markets that integrate distributed energy resources and support the electrification of transportation. By examining locational marginal pricing (LMP), we offer insights into how market-based MARL can achieve stable solutions in environments where prices coordinate agent behavior.

This article is organized as follows: First, we review related literature on nodal pricing and autonomous agents for load scheduling. Next, we outline the scheduling agents' decision problem, the determination of nodal prices, and our learning approach. Finally, we present the simulation results and conclude with remarks on the implications for future research on market design and intelligent agents for the smart grid.

2.2 Related Work

This paper is highly interdisciplinary, drawing on techniques from machine learning, transportation science, and (energy) economic theory. We consider intelligent agents as the connecting element between these disciplines, enabling a technically feasible and economically efficient solution for transportation applications.

2.2.1 Allocation Mechanisms for Transportation Resources

With the rise of shared mobility services, electric vehicles (EVs), and intelligent transportation systems, market mechanisms for optimizing transportation resources have become increasingly crucial (Ketter et al., 2023). These mechanisms aim to improve urban mobility, reduce congestion, and enhance the efficiency of transportation networks through dynamic pricing, auction-based allocations, and decentralized control systems.

A significant area of research and practice is the design of sharing systems, where real-time matching algorithms and pricing strategies optimize the allocation of rides to passengers. For instance, S. Ma et al. (2013) develop a dynamic ride-sharing model that uses real-time data to match drivers and passengers, reducing waiting times and increasing vehicle occupancy rates. Similarly, L. Wang et al. (2019) propose an auction-based approach for ride-sharing, where passengers bid for rides, and an optimal allocation is determined based on these bids. Schroer et al. (2022) investigate optimal management of a single fleet using a top-down approach.

Another critical area is congestion pricing, which aims to manage road traffic by charging users based on their usage of congested routes. This approach has been implemented in several major cities, such as London and Singapore, to reduce traffic congestion and improve air quality. Verhoef et al. (1996) discuss the effectiveness of second-best congestion pricing schemes, where tolls are set to reflect the marginal cost of congestion, thereby incentivizing users to choose less congested routes or alternative modes of transportation.

Parking management is also a vital aspect of transportation market mechanisms. Dynamic pricing for parking spaces, as studied by Geng and Cassandras (2013), allows for the efficient allocation of parking resources by adjusting prices based on demand. This approach helps reduce the time spent searching for parking and alleviates congestion in urban areas.

Decentralized control systems, facilitated by intelligent agents, are increasingly used to manage transportation resources. These systems allow individual users or vehicles to make autonomous decisions based on real-time information and market signals. For example, W. Zhang et al. (2011) propose a decentralized taxi dispatch system that uses real-time data and intelligent agents to match taxis with passengers efficiently. This system reduces idle times for taxis and waiting times for passengers, enhancing the overall efficiency of the taxi network.

A middle ground between centralized and decentralized approaches is the use of hybrid systems. These systems combine centralized coordination with decentralized decision-making to balance efficiency and flexibility. For instance, H. Yang and Lau (2015) introduce a hybrid model for taxi fleet management that integrates centralized route planning with decentralized execution by individual taxi drivers. This model optimizes overall fleet performance while allowing drivers to adapt to real-time conditions.

In the context of EV charging, market mechanisms are essential for managing the demand on the electricity grid and ensuring efficient use of charging infrastructure. Xu et al. (2020) explore a market-based approach to EV charging, where dynamic pricing signals incentivize users to charge their vehicles during off-peak hours, thereby reducing peak demand and avoiding grid overloads; similar work but for buses is presented by Abdelwahed et al. (2020).

These approaches to transportation problems illustrate the potential of dynamic pricing, auction mechanisms, and intelligent agents to enhance the efficiency and sustainability of urban mobility systems. By leveraging real-time data and advanced algorithms, these mechanisms can address the complex and dynamic nature of transportation networks, ultimately leading to improved resource allocation and user satisfaction. However, the mentioned allocation schemes typically rely on the use of centralized information, enabling top-down decision-making by a fleet or platform operator. Dispersed information in a distributed setting is seen much rarer in the literature, and if so, there is generally a reliance on voluntary information sharing.

2.2.2 Machine Learning in Transportation Resource Allocation

There is a wealth of research employing various machine learning approaches for transportation resource allocation tasks. Advances in distributed computing and intelligent devices enable bottom-up coordination of transportation resources (H. Yang & Lau, 2015; W. Zhang et al., 2011). AI methods and autonomous agents can leverage dynamic price incentives to optimize transportation demand based on real-time data and user preferences (Shann et al., 2017).

Numerous studies explore intelligent load scheduling by autonomous agents to operationalize "demand control" through AI (Ketter et al., 2018; Parag & Sovacool, 2016; Peters et al., 2018; Reddy & Veloso, 2011; Valogianni et al., 2015; Vázquez-Canteli & Nagy, 2019). However, real-time pricing can lead to herding behavior when rational agents receive the same price signal, causing network bottlenecks (Valogianni et al., 2015; Vytelingum et al., 2010).

In traditional tariff markets, incentives exist to avoid peak charging (Gottwalt et al., 2011) or to reduce charging rates (Valogianni et al., 2019), but effective incentive setting at the individual node level is challenging. The effect of tariff systems depends on the time-dependent price elasticity of demand, which is private information. Consequently, tariff schemes cannot ensure network restrictions are met at all times. With the increasing electrification of transport and decentralized energy supply, this poses a threat to grid stability and energy security (Ketter et al., 2018; Ramchurn et al., 2012), with significant welfare implications.

Despite the distributed nature of transportation networks, most multi-agent reinforcement learning (MARL) studies in the transportation domain focus on collaborative approaches to improve policy convergence. These techniques, similar to decomposition methods in mathematical programming, often distort agent incentives and lack economic grounding. However, the smart transportation domain, with its growing importance for sustainable mobility, is a critical application area for intelligent agents. Insights from this domain can facilitate the application of MARL in other distributed resource allocation areas, such as shared mobility and real-time traffic management.

While multi-agent settings based on bi-level optimization (cf. Gabriel et al., 2013) reflect autonomous decision-making, they often do not capture the complex uncertainties of individual behaviors. Deterministic planning requiring exact knowledge of future schedules is impractical. MARL addresses this by providing approximately optimal policies in stochastic, dynamic settings. We assume households have their software agents without requiring explicit knowledge of itineraries, enhancing convenience.

Approaches such as Nash Q-learning (cf. Casgrain et al., 2022; J. Hu & Wellman, 2003) have been proposed to solve similar problems but face challenges such as full observation of all actions and computational tractability. Consequently, we explore fictitious self-play in general-sum settings, despite known difficulties with environment stationarity and model convergence. By investigating MARL for autonomous EV charging in a transportation context, our research contributes to understanding how intelligent agents can enhance the efficiency and reliability of transportation systems through advanced resource allocation mechanisms.

2.3 The Model

Our market model reflects the interaction between self-interested agents, maximizing their individual utility, and a regulated market operator, maximizing social welfare. The interaction is primarily a coordination task, as grid resources are usually not scarce. In particular, it is not a zero-sum game, i.e., one agent's win is not necessarily another agent's loss. Still, it is not a cooperative game either, as agents in general have no incentive to co-optimize social welfare by themselves. We assume a well-regulated, i.e., benevolent, system operator who sets prices and allocates resources. An overview of notation is given in Table 2.1.

2.3.1 Agents' Decision Problem

Agents act in the market by bidding for the opportunity to charge their corresponding EVs. Note that we assume agents, EVs, and owners to be matched one-to-one , respectively; we therefore use the terms interchangeably. Agents interact in a real-time market mechanism. Such real-time pricing is currently

Variable	Domain	Explanation
J	$\mathbb{N}_+$	Number of agents
j	$\{1, ..., J\}$	Agent index
T	$\mathbb{N}_+$	Number of time steps
t	$\{1, ..., T\}$	Time index
Δt	$\mathbb{R}$ [h]	Time step length
$c_{i,t}$	$\mathbb{R}$ [MU]	i's cost in time step t
C_i	$\mathbb{R}$ [MU]	i's total cost
$SoC_{i,t}$	$[0, 1]$	i's state of charge at time t
$Battery_i$	$\mathbb{R}$ [kWh]	i's battery size
$P_{EV,i,t}$	$\mathbb{R}$ [kW]	i's charging power in time step t
$\rho(\cdot)$	$[0, 1] \to \mathbb{R}$ [MU]	i's penalty function
$T_{D,i}$	Set	i's departure time steps
$b_{EV,i}$	$\mathbb{R}$ [MU/kWh]	i's bid for power
$P_{EV,max,i}$	$\mathbb{R}$ [kW]	i's max. requested power
c_g	$\mathbb{R}$ [MU]	Grid energy price
P_g	$\mathbb{R}$ [kW]	Total power drawn from grid
$\mathbf{b_{EV}}$	$\mathbb{R}^N$ [MU/kWh]	Bid vector $(b_{EV,i})_i$
$\mathbf{P_{EV}}$	$\mathbb{R}^N$ [kW]	Charging power vector $(P_{EV,i})_i$
$\mathbf{P_{EV,max}}$	$\mathbb{R}^N$ [kW]	Max. charging power vector $(P_{EV,max,i})_i$
L	$\mathbb{N}_+$	Number of lines
l	$\{1, ..., L\}$	Line index
D_i	$\mathbb{R}$ [kW]	i's non-EV electricity demand
$\mathbf{D}$	$\mathbb{R}^N$	Non-EV demand vector
C_l	$\mathbb{R}$ [kW]	l's power limit
$\mathbf{C}$	$\mathbb{R}^L$ [kW]	Power limit vector
$\mathbf{PTDF}$	$\mathbb{R}^{L \times N}$	Power transfer distribution factor matrix

Table 2.1: Notation of the market model.

available in some retail electricity markets, though typically only containing a real-time component for energy. Our model is a real-time market to determine grid fees, whereas energy cost can be (in principle) due to any other tariff structure. Agents use the overall pricing signal to dynamically adapt their bidding strategy and to secure charging for their driving needs.

To ensure that agents' charging demands do not exceed the electricity grid's physical limits, agents in the simulation are not price-takers, opposed to what is usually the case in retail markets for electricity. Instead, each agent bids for the opportunity to charge, and infrastructure is allocated based on these bids using a repeated auctioning scheme, e.g., in 15-minute intervals. The auction determines both allocated charging power and (nodal congestion) prices for the agents. To participate in the auction, agent $j \in \{1, ..., J\}$ submits a tuple $a_j = (P_{EV,j,max}, b_{EV,j})$ of (maximal) demand, $P_{EV,j,max}$, and valuation, $b_{EV,j}$. As per standard nomenclature in machine learning, we call the choice submission of a_j agent j's *action*. For brevity, the indicated valuation $b_{EV,j}$ is referred to as agent j's *bid* for the remainder of this paper.

Demand and bid levels are discrete; that is, agents can decide between a *low*, a *medium*, and a *high* demand as well as placing a *low*, *medium*, *high* bid. Consequently, there are nine possible tuples for agents to submit to the auctioneer. This simplified bidding language compared to a continuous counterpart improves agents' learning ability and keeps complexity at bay. Further, in real-world applications, bidding would be limited to, e.g., €0.01 (in European markets) increments. Finding the optimal granularity, e.g., the number of bid levels, is in itself a complex trade-off between efficiency and simplicity of the mechanism. We find that, in our setting, three bid levels suffice to represent agents' preferences. It is reasonable to assume that most of the time, agents would choose a low bid, as there is sufficient flexibility in EV charging due to relatively large battery capacities and long idle times. The medium and high bids, respectively, can then be used to differentiate more urgent charging situations.

Compared to traditional power markets, devising a bidding strategy is more complex in the case of demand. The (instantaneous) valuation for charging is difficult to elicit, as there is an intertemporal coupling of charging decisions that is usually not accounted for in bidding languages. Even when block bids are used, as is usual in European electricity markets, they only allow for coupling of charging decisions within one day in the case of day-ahead markets. The key task for learning agents is thus to infer the valuation of charging. This valuation has to be derived from implicit signals. The agents that schedule EV charging therefore have to manage a dilemma: While they should minimize the cost incurred through EV charging, they should not risk securing too little energy to fulfill future driving needs. Opportunity costs of forgone charging can be substantial if a charged battery is necessary for a crucial task. Thus, an agent's goal is to individually optimize the corresponding household's charging costs by merely shifting the charging in time, not stopping to charge at all. As agents decide on their own, their policies constitute a bottom-up approach to the charging problem.

Given the willingness to pay a premium, their charging is prioritized in times of congestion. Again, the numeric levels themselves are not a crucial aspect of analysis; instead, they contribute to providing an intuition of the order of magnitude. As agents do not procure the energy for direct consumption, they cannot merely set their principal's current marginal value as bids. Instead, they have to infer the value of energy at each point based on the future stochastic use of the electric vehicle. Agents that bid too low may suffer from penalties due to forgone mobility needs or increased anxiety. Agents who are too eager to charge quickly may overpay. Striking a balance between patience and aggressiveness is hence the essence of any agent's successful policy.

The total cost of agent $j \in \{1, ..., J\}$ consists of the sum of costs over the entire planning horizon. The cumulative cost consists of instantaneous $c_{j,t}$ for agent j, at time t. Two terms contribute to the instantaneous cost: the monetary cost of electricity procurement, and the penalty costs. The charging power of the agent j' at time t is $P_{EV,j,t}$; multiplying by the length of time slots $\triangle t$ results in the charging energy. The price at the location of agent j' on the grid at time t is $p_{j,t}$ monetary units per kWh.

Whenever the driver wants to use their electric vehicle for a trip, i.e., time t is in agent j's set of departure times $T_{D,j}$, the driver will check the state of charge of the vehicle. Note that these departure times are not certainly known to the software agent. When the SoC at these departures deviates from a target SoC, then the software agent is penalized via a convex penalty function ρ. Convexity ensures that small deviations are tolerated, whereas larger deviations are not. Note that the penalty does not necessarily have to be a result of actual driving needs—and the corresponding energy demand—but can also comprise psychological aspects, for example, range anxiety (Valogianni et al., 2019). The higher the owner's anxiety, the more irritated she might be when large deviations from her set target occur.

The formal definition of the penalty term to apply only at departure times follows a rationale based on how vehicles would be used in real-world applications. Drivers employing a software agent to charge their vehicle would, and should, not constantly check how the charging is proceeding and micromanage the charging curve until the next use of the vehicle. Instead, they would plug in the vehicle when they return from their trips and not bother with any control until the next time they need their EV. Only then would they recognize whether the software agent fulfilled charging needs.

To summarize, the charging software should maximize long-term rewards

$$C_j = \sum_{t=1}^{T} c_{j,t} \tag{2.1}$$

$$= \sum_{t=1}^{T} \Big(\underbrace{p_{i,t} * P_{EV,j,t} * \triangle t}_{\text{Cost of energy}} \tag{2.2}$$

$$+ \underbrace{\rho \max\{SoC_{j,target} - SoC_{i,t}, 0\} * \mathbb{1}_{t \in T_{D,j}}}_{\text{Penalty}} \Big). \tag{2.3}$$

The agent's policy to maximize the reward connected to the cost of energy is straightforward—not charging at all. However, the development of the state of charge is uniquely determined by the charging:

$$SoC_{j,t+1} = \min\{SoC_{j,t} + \frac{P_{EV,j,t} * \triangle t}{Battery_j}, 1\}. \tag{2.4}$$

The optimal policy of software agents thus constitutes a trade-off between reducing the cost of charging and avoiding penalties. The agent's cost-reducing efforts then result in shifting the flexible EV load to periods with minimal nodal prices at their location. However, explicit information is not available to the charging agent to conduct optimal scheduling for a future time frame. Agents make their decision based on learned price signals and their current state of charge.

Beyond the state of charge that is easily accessible for agents, we assume their information to be sparse. Other observable environmental variables for charging agents are just the time of day and the time since the last return of the EV. This can be considered decision-making with minimal information, as this information is necessarily available on electric vehicle chargers. There are various proposals to enhance the capabilities of EV chargers and similar residential technologies in the smart home literature, e.g., using calendar information, monitoring the presence and movements of inhabitants, etc., to derive behavior and predict charging needs. While showing a growing presence in (new) homes, these techniques raise major questions of privacy and data security, making it quite likely that not everyone will be willing to use such technologies to automate their devices. Our decision to use limited information is thus warranted and should provide us with an understanding of intelligent agents' capabilities in tricky environments.

Reinforcement Learning

Stochastic charging demand results from driving behavior (and, possibly, range anxiety) outside the agent's strategy space. The price component of the bid, $b_{EV,j}$ on the other hand, influences when and at which price the trading mechanism allocates charging power to an agent. A software agent should thus adopt a bidding strategy to minimize costs, as defined in equation (2.3). However, in the present multi-agent setting, the costs incurred for one agent depend on the other agents' preferences (i.e., charging demand), which is private information. Given the imperfect information setting, the objective function is difficult to observe for the agent. Further, the lack of observability of other agents' behavior makes it difficult to model-based dynamic stochastic programming approaches to derive policies.

As a model-free approach to stochastic multi-stage decision-making, reinforcement learning generally is a well-tested tool. In stationary settings with only partially observable environments (Gupta et al., 2017), in which rewards depend on agents' actions—but not solely—and the exact reward function or its inputs are unknown, it leads to asymptotically optimal policies. Typically, an agent is trying to learn an action policy that maximizes the expected future reward discounted by a factor γ defined by the action-value function, or 'Q-value' $Q(s, a)$:

$$Q(s, a) = [r(s, a) + \gamma \cdot v^*(\delta(s, a))]. \tag{2.5}$$

In this standard formulation, $r(\cdot)$ is the reward for an agent's action a in the state s; γ is the discount factor for future rewards. $\delta(\cdot)$ represents the transition function, which defines the next state s' based on the agent's action in the current state. Note that γ need not be an (economically motivated) discount factor. As before, we assume that there is no actual discounting of future cost, as we consider short timeframes. Within the action-value function, γ contributes to the algorithmic convergence properties.

Even though our environment is non-stationary, we analyze agents employing reinforcement learning; the setting can therefore shed light on whether the proposed auction mechanism provides a sufficiently stable environment for agents to learn in. Specifically, the agents employ *Q-learning with experience replay* (DQN), a technique presented by Mnih et al. (2015). Here, the optimal Q-value (2.5), as defined by standard Q-learning, is approximated by a neural network (Mnih et al., 2015). Note that, for convenience, we are referring to DQN in this work when the function approximation uses at least one hidden layer, even though these architectures are not *deep* per se. Q-learning is a model-free learning approach, which means that the neural network does not explicitly estimate the transition function $\delta(\cdot)$, but learns from the observed samples of states, actions, and resulting rewards (Mnih et al., 2015).

(Negative) rewards in each time step, $r(\cdot)$, consist of the costs of charging derived from equation (2.3) depending on the nodal prices as per market clearing (2.8). The cost $c_{i,t}$ incurred at time step t depends on the bidding vectors $(\mathbf{b_{EV}}, \mathbf{P_{EV}})$, of which an agent only knows its bid $(b_{EV,j}, P_{EV,j})$. Each agent can observe further information that implicitly impacts the reward function. The state s contains the previous price at its location, the time of day, the EV's current state of charge, its ability to charge (i.e., its connection to the wall box). On that basis, the agents decide which action a to take, given their action space A. Recall that an action a is a submitted tuple to the auction mechanism described in 2.3.2. Agents act simultaneously in the environment, and electricity prices are dependent on the market-clearing, which, in turn, depends on all agents' actions.

During training, each step's actions are chosen by an ϵ-greedy strategy that picks the greedy strategy $a = max_a Q(s, a)$ with probability $(1 - \epsilon)$ and picks a random action $a \in \mathcal{A}$ with probability ϵ. During testing, agents are exclusively greedy. Q-values are updated using experience replay (Mnih et al., 2013), which means that all experiences the agent makes are stored. During training, samples of these experiences are replayed using Q-learning. This way, each action the agent takes can be learned from in future update steps. This technique also reduces the risk of diverging parameters, as the behavior distribution is averaged over many instances of previous experiences (Mnih et al., 2015).

2.3.2 Operator's Decision Problem

Historically, the retail electricity market was mostly a regulated industry with fixed prices. Recent decades have seen an increase in market liberalization, where price formation is subject to market forces, and various market mechanisms have been tested in practice and proposed in the scientific literature (Robu et al., 2013). We propose that there is a *system operator*, tasked with both the operation of the physical infrastructure and setting market-based grid fees. This deviates slightly from typical settings in practice, where customers can choose a (market-based) electricity tariff offered by market participants in liberalized retail electricity markets and additionally pay fixed grid fees to a network operator. This traditional setting of a tariff scheme has a major drawback in the envisioned scenario with large-scale EV adoption and corresponding electrical loads in residential networks: It is not easily possible to ensure that the electricity demand does not exceed the capacity of the grid. Integrating the physical aspects into a (second) market facilitates their reconciliation.

The distribution system operator is a natural monopolist, as there is only one distribution grid in place. Therefore, market forces (as in traditional retail competition) cannot be relied on to ensure the affordability of electricity. It is thus easy to conceive that the monopolist is regulated to act with the aim of maximizing the welfare. Our model concerns the short-term operation of an exogenously given electricity grid, i.e., investments to increase capacity are not within its scope. The system operator's (enforced) goal is thus to maximize consumer surplus by optimizing grid utilization without exceeding its capacity. The optimal grid fees can ensure that. Whenever grid utilization is at capacity, customers that have the highest valuation for electricity are served, while those who can do without (for any given instance) are not. In the following subsections, we assume the distribution system operator has received all relevant bids.

Auction Clearing — Allocation and Pricing

There are a multitude of proposed auction-based pricing mechanisms for electricity. One of the more elaborate and widely used schemes is *locational marginal pricing* (LMP), which is often used for electricity pricing at the transmission level—especially in the US power markets—while we use it at the distribution level. LMP sets the price at each location, that is, node, in an energy system to the cost of serving one (infinitesimal) more unit of energy at that node (Vytelingum et al., 2010). There are numerous formulations of LMP and corresponding bidding languages. Typically, bidding languages are inspired by a pricing scheme conceptually based on fossil energy generation (with non-zero marginal production costs). As a result, the price for a consumer, i.e., the marginal price at its location, is equal to the marginal cost of the generator that has to increase its production to serve the consumer's additional demand. Congestion surcharges occur if the cheapest generator in the system cannot serve any additional energy units to that consumer due to network constraints. The extra cost of ramping up a more expensive generator unit instead has to be carried by the consumer.

The auction scheme presented in this paper deviates from such purely (production) cost-driven pricing rules. Instead, it uses the marginal bidders' valuation of energy as the foundation for congestion pricing. The same intuition as for (production) marginal cost-based pricing applies. Whenever an agent's electricity consumption displaces another agent's consumption, the former has to pay a premium derived from the latter's valuation. In this sense, the proposed pricing scheme resembles common (second-price) auction mechanisms. The auction mechanism is formalized in the following.

We call the auctioneer the *system operator*, as it is not only in charge of the operation of the market—or auction—but also the physical grid. The system operator's task is to serve consumers within its territory through the existing physical infrastructure. In doing so, it has to adhere to two critical aspects of energy systems. First, the energy system must always be in balance, i.e., demand and supply must always match. Second, all grid components have physical limitations—they can carry maximum loads that must not be exceeded. In the investigated case, the components' limitations are the maximum capacity of electric lines. Each line has a maximal current it can carry, and, consequently, it can transfer a limited amount of energy. Note that low-voltage grids have characteristics exceeding their power transfer capabilities, e.g., phase symmetry requirements or voltage requirements. We do not model those here, as the resulting optimal power flow calculations are in general non-convex and therefore do not support the determination of competitive market prices based on the duals of the optimization procedure. Further, EV chargers are typically three-phase and may be equipped with a power factor correction, such that power limitations are most important.

Subject to these limitations, the system operator's activities are two-fold: first, it distributes electricity P_g from the upstream grid. The electricity flowing from the upstream grid must match the demand in the residential grid, i.e., the inflexible demand through household appliances and the EV charging demand that is allocated through the auction. Agents can be subject to any electricity tariff. As per typical European unbundling regulations, we assume that the distribution system operation and retail tariff are independent. In particular, even though an agent subscribes to an electricity tariff that is technically an option, the distribution system operator auctions the localized grid infrastructure as a distinct good. Non-allocation of infrastructure for an agent might prohibit the agent from exercising its (tariff) option.

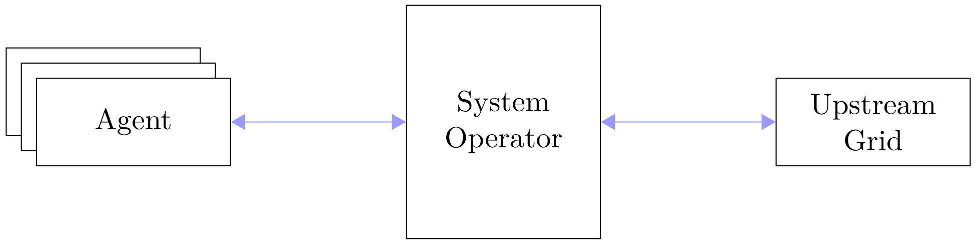

Figure 2.1: The *system operator* transmits energy from the upstream grid at a given price. It also repeatedly runs an auction market in which agents bid for energy to charge their vehicles.

Second, it distributes energy to customers within its service territory (while adhering to all physical constraints). As is typical in electricity markets, there is a strong interdependence between described powers and energies. Agents incur costs based on the energy they consume, but physical constraints are described in terms of maximal powers. In a discrete time setting, both are approximately linearly dependent with $E = P\Delta t$.). As outlined, agents provide (linear) bids $\mathbf{b_{EV}} = (b_{EV,1}, ..., b_{EV,J})^T \in \mathbb{R}^J$ per unit of capacity. Figure 2.1 schematically depicts the roles within the multi-agent setting.

Given that an operator would entail a natural monopoly, there would be institutions regulating it to act welfare-maximizing, as described in Equation (2.6). The objective here is described without a time index. The operator runs the described auction repeatedly to determine the allocation of charging electricity, $\mathbf{P_{EV}} = (P_{EV,1}, ..., P_{EV,J})^T$, close to real-time. Typical time steps in such markets are hourly periods (as in many day-ahead electricity markets around the world) or 15-minute intervals. The underlying notion is that repeated maximization of the instantaneous welfare leads to optimal long-term infrastructure utilization. Only calculating instantaneous prices shifts significant complexity from the system operator to the agents, as they have to predict how prices will develop when making charging decisions. However, the alternative of running an auction for more than one time step in advance is impractical. Predicting whether (or not) an EV will be available for charging at some point throughout the next hours or days is highly difficult even for an advanced learning algorithm. Having to procure electricity for such periods anyhow could lead to open positions that would need to be closed again at later times, requiring a much more complex trading strategy than just real-time procurement of electricity whenever the EV is actually available. Based on these considerations, we drop the time index in

(2.6):

$$\max_{\mathbf{P_{EV}}}. W = \mathbf{b_{EV}^T P_{EV}}. \tag{2.6}$$

As outlined, $\mathbf{b_{EV}} = (b_{EV,1}, ..., b_{EV,J})^T \in \mathbb{R}^J$ is the vector of the J agents' bids, with $b_{EV,j}$ being the bid of the scheduling agent $i \in \{1, ..., J\}$. Agents can only bid a constant price per unit of energy. Thus, the objective function is linear in the decision variables $\mathbf{P_{EV}} \in \mathbb{R}^J$ (charging electricity). The bid is the first element of a submitted tuple; the requested quantity is the second element. $P_{EV,max,j}$ denotes agent j's requested quantity, as it equals the maximum power its EV draws from the grid. The N agents' submitted tuples are thus $(b_{EV,i}, P_{EV,max,j}),\ i \in \{1, ..., N\}$. As with the bids, the requested quantities can be represented more concisely using the vector $\mathbf{P_{EV,max}} = (P_{EV,max,1}, ..., P_{EV,max,J})^T \in \mathbb{R}^J$. The welfare maximization in (2.6) so far does not account for infrastructure topology and limitations. Disregarding these would mean that the highest bids would win the auction, as it is intuitive, but would not necessarily be welfare-maximizing: depending on the grid's setup, there might be customers that can be served without displacing anyone else, while others are *behind* certain bottlenecks. Excluding the former "easy-to-serve" customers because the latter "hard-to-serve" customers bid more would not maximize welfare.

To provide a realistic case study of electric vehicle charging, network topologies in residential electricity grids must be understood. These are typically characterized by *radial* networks. A radial distribution network has a tree structure, i.e., each load point is connected by exactly one path to the upstream grid; see Figure 2.2i. The power flow to any load point is transferred fully through every line in this path and not through any other lines. The power transfer distribution factor (**PTDF**) matrix encodes the relationship between nodes and lines involved in transferring power. For J buses and L lines in the system (where a *bus* is a load point, i.e., a household in our case, and a *line* is an electrical connection), $\mathbf{PTDF} \in \{0,1\}^{L \times J}$. A small exemplary network and the corresponding **PTDF** matrix for four load buses (i.e., connection points) $\{1,2,3,4\}$ and five lines $\{a,b,c,d,e\}$ is given in Figure 2.2. Note that, generally, $\mathbf{PTDF} \in \mathbb{R}^{L \times J}$. Only for radial networks it is binary.

Lines, due to their maximal loading, are the limiting physical asset. The maximal loading of the line $l \in \{1, ..., L\}$ is denoted C_l. The vector $\mathbf{C} = (C_1, ..., C_L)^T \in \mathbb{R}^L$ then summarizes line limits. Whenever a line reaches its maximal loading, it is *congested*. The **PTDF** matrix attributes charging demand at a given bus i

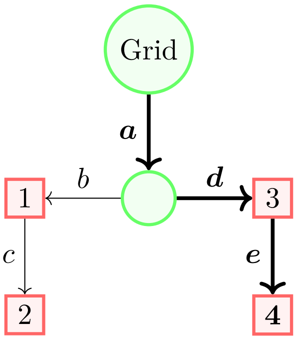

(i) Network topology. The load buses are marked as red squares. The green round bus does not serve a load, but is a distribution bus.

$$\mathbf{PTDF} = \begin{array}{c} \begin{array}{cccc} 1 & 2 & 3 & \mathbf{4} \end{array} \\ \begin{pmatrix} 1 & 1 & 1 & \mathbf{1} \\ 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & \mathbf{1} \\ 0 & 0 & 0 & \mathbf{1} \end{pmatrix} \end{array} \begin{array}{c} \\ \mathbf{a} \\ b \\ c \\ \boldsymbol{d} \\ \boldsymbol{e} \end{array} \in \{0,1\}^{L \times J}$$

(ii) **PTDF** matrix. Due to the tree topology, all entries are either 1 or 0. a is needed to serve every bus, whereas c and e are only needed to serve buses 2 and 4, respectively.

Figure 2.2: Exemplary network and corresponding **PTDF** matrix. An exemplary path connecting bus **4** to the grid using lines $\boldsymbol{a}$, $\boldsymbol{d}$, and $\boldsymbol{e}$ is bold.

to a (potentially congested) line l, assuming lossless conductors. The distribution network must not only serve EV charging but also other, here inelastic, demand. We denote these inflexible household loads by $\mathbf{D} = (D_1, ..., D_J)^T$. There is a body of research concerning the use of *smart appliances*, e.g., dishwashers, that react to price signals. In empirical studies, the behavioral change induced by using such appliances is small, as are the savings opportunities for households. Electric vehicles have much higher power draw and energy demand while providing a significant amount of flexibility due to long idle times. We therefore reckon that it is reasonable to assume that household flexibility would be limited to EV charging. With the convention for the bidding protocol, the operator's market-clearing task renders a linear program. Note that the auction clears repeatedly; the subscript t is dropped for brevity.

$$\max_{\mathbf{P_{EV}}}. \quad \mathbf{b_{EV}^T P_{EV}} \tag{2.7a}$$

$$\text{s.t.} \quad \mathbf{PTDF} \cdot (\mathbf{P_{EV}} + \mathbf{D}) \preccurlyeq \mathbf{C} \quad (\boldsymbol{\mu}) \tag{2.7b}$$

$$\mathbf{0} \preccurlyeq \mathbf{P_{EV}} \preccurlyeq \mathbf{P_{EV,max}} \tag{2.7c}$$

As the inflexible household demand $\mathbf{D}$ is exogenous, only the EV charging powers $\mathbf{P_{EV}}$ (and consequently the power drawn from the grid P_g) can be subject to the auction's allocation. The operator $\preccurlyeq$ denotes element-wise (in)equality. The flow constraints (2.7b) have Lagrangian multipliers $\boldsymbol{\mu} = (\mu_1, ..., \mu_L)^T \in \mathbb{R}^L$. Existence and well-behavedness of the Lagrangian multipliers and existence of a globally optimal solution to the maximization procedure (2.7) trivially follow from LP theory for all relevant settings in our case study.

The electricity prices at the nodes now comprise two parts. First, the electricity cost that may be individual for households, depending on their tariff. For simplicity and to focus on the congestion management of the market, we assume a constant grid cost $c_g \forall i \in \{1, ..., N\}$. Second, there is a cost from line congestion; all nodes that use a congested line should pay a premium based on that line's marginal user's valuation, $\boldsymbol{\mu} = (\mu_1, ..., \mu_L)^T$. Using the PTDF matrix to account for line costs to nodes, the optimal price vector $\mathbf{p}^* \in \mathbb{R}^N$ that summarizes prices at all nodes is

$$\mathbf{p}^* = c_g + \mathbf{PTDF^T} \boldsymbol{\mu}^*. \tag{2.8}$$

Equation (2.8) emphasizes that the bidding protocol calls for a willingness to pay premiums beyond the electricity, rather than absolute willingness to pay. Agents effectively bid for the *transmission right*, not the energy (see Hogan, 1992). Pricing implicitly accounts for the balance constraint. This requires that the supply to the grid be sufficient, meaning that it is at least as large as the lines permit, which can be safely assumed for any residential distribution grid.

Incentives

To be welfare-maximizing, the optimization procedure in 2.3.2 has to rely on truthful bidding of the agents. The truthfulness of LMP mechanisms is a matter of extensive studies (e.g. Karaca & Kamgarpour, 2017, 2018; F. Wu et al., 1996). In the given applications, agents have an incentive to shade their bids, as they are setting the price at their node with positive probability; the mechanism is thus not strategy-proof: assume that the network was only serving a single agent that demands more energy than can be supplied by the single line from the network connection to its location. Its bid will always be the shadow price of the line capacity; hence, it has an incentive to bid zero. Furthermore, any coalition of agents can choose to manipulate the prices in their favor.

There are reasons to choose LMP over other pricing mechanisms, despite opportunities for manipulation. First, LMP is the most well-studied and understood nodal pricing mechanism in electricity markets. Hence, this work embeds itself into existing research. Second, in the specific case presented, the opportunity for manipulation is slimmer than in some extreme examples presented in the literature. The bidding protocol only accepts positive bids, discouraging perverse results. Unlike generators that can in principle bid infinitely high, there is thus a natural limitation to exercising market power for demand. For these reasons, it is safe to assume that LMP is a suitable candidate mechanism. It is further easier to calculate than alternative pricing rules based on the well-known Vickrey-Clark-Groves (VCG) (Clarke, 1971; Groves, 1973; Vickrey, 1961) mechanisms. These, among other widely discussed weaknesses of VCG mechanisms (see Ausubel and Milgrom (2006) for an extensive discussion), make LMP preferable.

An important issue against the backdrop of learning agents is the non-stationarity of the environment in the presence of multiple interdependent agents. In these cases, canonical results from single agent theory do not hold, and convergence of the mechanism is difficult to attain and even more so to prove generally. The aforementioned VCG mechanism features dominant-strategy incentive-compatibility (DSIC), at least for a single auction. This quality can be expected to improve convergence in a multi-agent reinforcement learning setting, as each agent has a (learnable) optimal strategy, independent of other agents' behavior. Although the stability emerging from DSIC is desirable, particularly in a MARL setting, we argue that the mechanisms' disadvantages still outweigh this benefit, keeping in mind our foundational motivation to study a real-world application, where employing such complex mechanisms is not seriously discussed. Our research therefore also sheds light on the question of to what degree a market mechanism without a strong equilibrium concept, such as DSIC, can provide sufficient stationarity for the use of MARL.

2.4 Experimental Evaluation

To examine the applicability of this model and the efficacy and information requirements of reinforcement learning in this multi-agent setting, we conduct an experimental evaluation using a representative distribution grid model provided by IEEE (Schneider et al., 2018) and realistic price ranges and energy loads. The maximum charging power for the electric vehicles is set to 22 kW, corresponding to values seen in the industry. While this charging power is fairly high and not all models currently support it, the evolution towards higher charging powers at home warrants using this value when discussing a scenario with high EV penetration in a residential area.

The bids are discrete, taking realistic values. The typical price for one kilowatt-hour (the standard unit in residential applications) of energy in Europe is roughly 0.20 − 0.30 €. To abstract from currencies, we are using 0.2, 0.4, and 0.8 monetary units/kWh (MU/kWh) as *low*, *medium*, and *high* bids, respectively. We use slightly higher values in the simulation, i.e., with a grid price of 0.2, the minimum bid is set to $0.2 + \epsilon$. This is necessary, as otherwise the market-clearing optimization (2.7) could reject *low* bids.

2.4.1 Data

Most car trips are made to and from work (YouGov, 2018). We are thus splitting the agent population into three archetypes, accounting for heterogeneity in the population. We propose that 60% are early workers, with working times (the working times include commuting, i.e., they represent the mean times of arrival and departure) from 7:30 AM to 4:30 PM. A standard deviation of one hour is attached to both times to determine the mean departure and arrival times of these agents. Similarly, there are late workers with working times from 9:30 AM to 6:30 PM, again with a one-hour standard deviation. The remaining agents are gone between 12 PM and 3 PM, with two hours of standard deviation attached to arrival and departure times. These choices are in line with travel surveys (Department for Transport, 2019; YouGov, 2018); however, precise times are not as relevant as the general notion of heterogeneity between agents: too similar driving profiles would overstate the value of smart charging, as peak demands in the evening hours would be exaggerated. Synthetic data is necessary, as driving profiles with high granularity are not publicly available. This synthetic data is specifically designed to reflect an everyday schedule. Because of this, the policies derived by the software agents are generally not ideal for days with fundamentally different driving patterns (e.g., long-distance drives toward travel destinations). Assuming that these days are rare, manual overrides by drivers are a suitable solution to this issue.

After agents' initialization with their mean and departure times, a driving itinerary for the working week is created, based on the mean times and random normal noise. Departure times in the morning are distorted with a 15-minute standard deviation, and arrival times in the afternoon or evening with a 30-minute

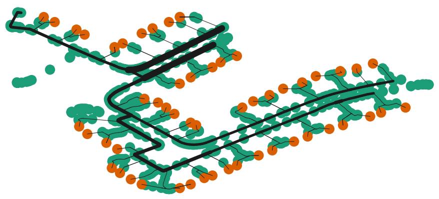

(i) Original IEEE Test Feeder with 55 load buses (orange) and 905 lines (black). Green buses are connection points without loads.

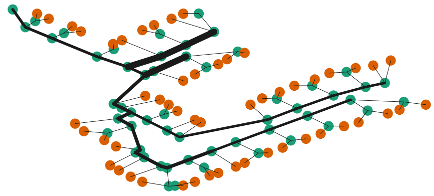

(ii) Reduced IEEE Test Feeder with load 55 buses (orange) and 110 lines (black). Green buses are connection points without loads.

Figure 2.3: Original and reduced test grid. Both topologies are equivalent for the presented use case.

standard deviation. Both are truncated to ensure realism in the simulation. Not using a fully deterministic driving itinerary per agent reflects that there are always some variations due to traffic conditions or irregularities in work patterns. Software agents tasked with charging control should be able to handle this distortion.

The energy demand for driving is drawn from a scaled $Beta(2,5)$ distribution. Using this right-skewed distribution accounts for the fact that many commutes are short, while some are quite long. The distribution is scaled so that the energy demand per day (i.e., for a return trip) corresponds to a driving distance of overall 40 km for a typical EV (cf. van Velzen et al. (2019)). We are using an adaptation of the IEEE European Low Voltage Test Feeder (LVTF) to model the residential distribution grid (Schneider et al., 2018). The LVTF was explicitly designed to analyze the impact of an increased amount of distributed resources on distribution grids. However, note that it was also designed to support investigations regarding phase symmetry and is thus a three-phase system. Because we are investigating only power flows, we are using a single-phase version instead. This approach is reasonable as long as chargers themselves are assumed to be three-phase, which is undoubtedly the case for higher charging powers. Figures 2.3i and 2.3ii show the original test grid and a reduced version, respectively. The reduced version maintains all electrical characteristics while removing buses that are not necessary for modeling.

2.4.2 Benchmark strategies

To evaluate the DQN agents' performance, they are compared to two different benchmark strategies, the *naïve* and the *linear* policy. In the naïve case, agents' behavior is modeled on patterns frequently observed in current real-world scenarios. Whenever the respective EV connects to a charger, its managing agent requests charging at full speed. As the pricing mechanism based on displaced agents' bids requires a valuation for those bids, we assume that agents always bid the highest valuation when employing the naïve strategy. Albeit a crude benchmark, it still enables a comparison to an expensive allocation.

To provide a more reasonable benchmark, we are also studying agents that approximate the Q-function using linear function approximation. As the auction mechanism already encodes much of the problem's actual complexity, linear function approximation is a potent benchmark for policies using neural networks. Even using only linear function approximation, employing learning agents in this realistic multi-agent setting is much more elaborate than existing solutions utilized in practice. As to scalability and computational efficiency, such simpler learning strategies can well be preferable in real-world settings.

2.4.3 Results

Multi-agent reinforcement learning is, in real-world applications, still largely in its infancy. Due to the multitude of parameters to set in a DQN application, we thus investigate numerous possible implementations, based on typical values seen in comparable studies. Table 2.2 displays the varied parameters together with investigated values. Whereas more parameters could be altered, we reckon

Hyperparameter	Values
Hidden Layers	$0, 1, 2, 3$
Hidden Units/Layer	$8, 12, 16$
Activation Function	Sigmoid, ReLu
Discount Factor γ	$0.95, 0.99, 1.0$
Min. Expl. Rate ϵ_{min}	$0.1, 0.05, 0.01$
Batch Size	$32, 128, 512$

Table 2.2: Varied hyperparameters and corresponding values for the RL implementation. Note that not all combinations are sensible. Shaded parameters are only varied for non-zero hidden layers.

Parameter	Values
$\rho_j(SoC_j, t)$	$10 * (SoC_{j,target} - SoC_{j,t})^2 \; \forall j$
$SoC_{j,target}$	$1.0 \; \forall j$
$Battery_j$	30kWh $\forall j$
Δt	15 min
Episode	5 Days

Table 2.3: The simulation parameters capture important aspects of real-world settings with electric vehicles. Batteries are relatively small; because of this, the target SoC is set high.

that these constitute a solid choice for variation. Note that not all parameter combinations are sensible, as a linear approximation (without hidden layers) does not implement activation functions or hidden units. Important simulation parameters are given in Table 2.3. The simulation parameters are designed to resemble a typical residential (e.g., suburban) commuter setup.

We investigated a set of over 500 learning configurations ($3 \times 3 \times 3 = 27$ linear learners and $3 \times 3 \times 2 \times 3 \times 3 \times 3 = 486$ DQN learners). Even though each learner's problem instance is small and can be learned quickly, no matter the architecture's complexity, the number of models to evaluate puts a limit on potential combinations. Figure 2.4 shows the agents' average relative performance of different learning algorithms, i.e., their relative cost compared to the aforementioned constant-price scenario. There is little difference in the performance of deep learners compared to linear learners, with both architectures on average outperforming the naïve strategy after training for 1,000 episodes. While it is possible that deep learners could ultimately outperform linear learners given sufficient training time, the added benefit of increased model complexity must outweigh the (computational) cost in this type of practical application in which distributed computational power is limited. The preliminary results indicate that linear function approximation, while parsimonious, is promising for a smart grid application. In the following, neural network learning models are thus disregarded, and different learning strategies for linear models are focused on.

Figure 2.5 shows the training progress distribution of the learners relative to the naïve strategy in their corresponding scenarios. Both linear and deep learning strategies lead to a substantial reduction in relative cost quickly. The dispersion between different linear learners is much lower, shown by the gray band around

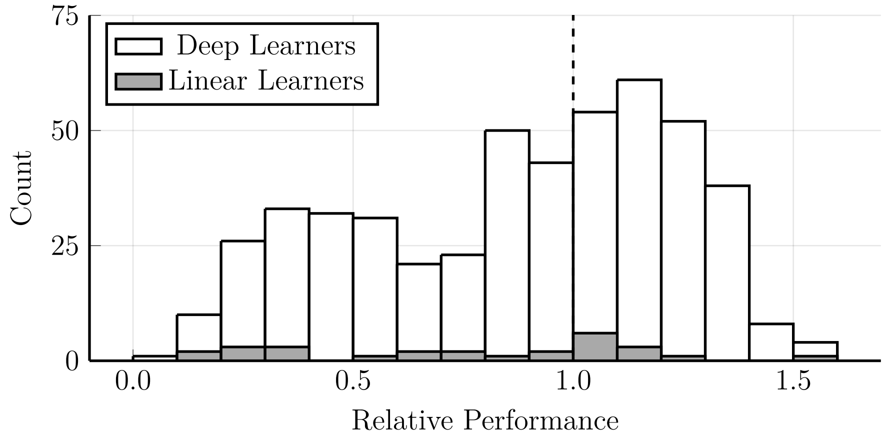

Figure 2.4: Relative performance of different learning algorithms after 1,000 episodes of training. Both deep and linear learners are on average better than the benchmark. With longer training, DQN learners might become more competitive. However, very long training periods are impractical from both an application and energy intensity perspective.

the black line representing the standard deviation and mean of relative cost, respectively.

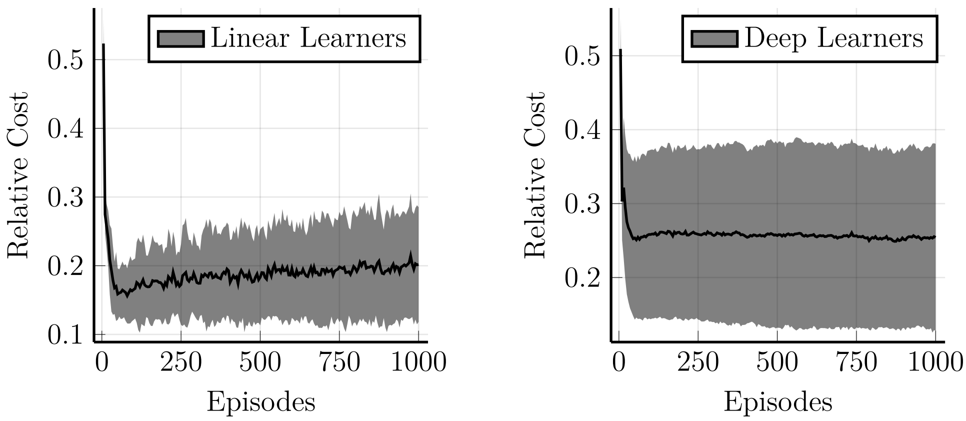

Figure 2.5: Both linear and deep learners learn quickly. Dispersion between different linear learner configurations is lower than for deep learners.

2.5 Discussion

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, volatile, and sustainable energy resources is “one of the greatest engineering challenges of our day", p. 2166—as unfortunately illustrated by the market failures of the California and Texas energy markets in 2000/2001 and 2021, respectively (Borenstein et al., 2002; The Economist, 2021). While there are various approaches to alleviate congestion in the grid, including infrastructural expansions or power caps, many of those require high hardware investments or come with hard restrictions on individual usage.

We develop a novel techno-economical co-simulation of a residential smart grid, with the structure of a partially observable stochastic game (or a partially observable Markov decision process (POMDP) from an agent’s perspective), in which owners of electric vehicles (EVs) pay for their use of the grid infrastructure when charging their EVs. In a repeated auction setting, we evaluate the performance of intelligent agents using real energy load profiles and simulated EV charging demands in a representative IEEE distribution grid to schedule loads. Herein, we go beyond the theoretical proposals and often very crude and simplified test cases in earlier literature and provide a test case for an online market mechanism for EV scheduling. In particular, we incorporate physical grid limitations as well as preferences of self-interested agents. The feasibility and efficiency of learning models is crucial for leveraging online market mechanisms in practice to realize the vision of a smart grid integrating distributed energy resources.

Our experiments provide insights into the trade-off between simple, lighter learning models and more powerful but computationally expensive deep learning models, which is relevant for the practical implementation of such machine learning models in the real world in a distributed network. Training complex machine learning models is by itself not only computationally expensive but thereby also energy-intensive. In applications that seek improvements regarding environmental sustainability, the much-touted application of artificial intelligence techniques must provide not only operational efficiency but also reasonable upfront investments in the models. The burden of training complex models to competitiveness seems overly high. The parsimony of simple linear models, on the other hand, provides convincing performance combined with low computational effort. The trade-off between model complexity and model performance is visualized in Figure

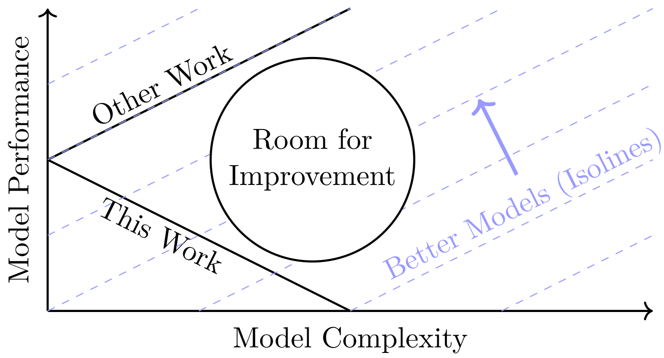

Figure 2.6: Trade-off between complexity and performance

2.6. While in this work, more complex models even performed worse than simpler ones, it is easily comprehensible that other complex models might outperform the presented linear learners. If, however, increased complexity must be met by substantial performance improvements, then the simplest models might still be better. Acknowledging this trade-off instead of over-engineering machine learning models as an end in itself is crucial for artificial intelligence to actually impact many real-world use cases positively.

The presented results are thus not only informing research on multi-agent settings in the domain of smart grid applications but also any domain that ultimately seeks widespread adoption in decentralized environments. Highly elaborate algorithms and learning techniques are worthwhile in their own right, but parsimonious models should not be underestimated as means to improved coordination.

On the theoretical side, our experiments revealed that the definition of the reward function in this setting is vital for the performance of RL agents. Defining a penalty for not driving due to insufficient charging of an electric vehicle is straightforward theoretically; however, eliciting an actual value of this penalty is highly complex and depends on individual preferences in practice (Shann et al., 2017). It is not constant but depends on the specific trip that should have been made, individuals' range anxiety, and may be evolving. The penalty should thus be seen not only as a hyperparameter of the learning environment but also as an actual economic variable.

2.6 Conclusions & Future Work

We contribute to the theoretical research on multi-agent reinforcement learning, as well as to the practical work on autonomous agents for the smart grid. Our work addresses a societally relevant issue by leveraging computational tools to coordinate distributed energy resources efficiently. Herein, we can show that nodal pricing mechanisms, which provide desirable economic properties but had previously been impractical to implement in the energy grid, could be utilized in combination with intelligent agents. In combination, these tools can alleviate pressure from distribution grids and respect individual preferences without requiring central control over individuals' consumption schedules.

As the complexity of learning models influences the requirements on distributed computational power, and hence their applicability in practice, a careful and realistic evaluation of the learning models implemented by intelligent agents is crucial. Our research addresses this issue by examining the trade-off between model complexity and performance improvements in a simulation of a physical distribution grid. Herein, we can show that, in this setting and with a reasonably (low) amount of model training, relatively simple linear model-free reinforcement learning algorithms converge to a more stable equilibrium compared to a DQN approach.

Further research could expand the simulations presented above further by using empirical data on various flexible loads, such as heat pumps or home battery storage, and various grid models. Moreover, the market mechanism itself could be developed further: the presented model uses a linear optimal power flow that is known to not perfectly describe all aspects of low-voltage systems. With complex market mechanisms including reactive power or phase symmetry, the trade-off between complexity and performance might need to be revisited. In addition, the presented market mechanism closely follows design choices first made for making the mechanism explainable for human decision-makers. This criterion is not in the same way relevant to markets that are designed purely for algorithmic trading agents. Approaches such as stochastic prices could be an interesting avenue for future research in this context of autonomous intelligent agents interacting.

Within our model, architecture and hyperparameters are homogeneous and exogenous for all agents, as otherwise the search space would become intractably large. However, when discussing the value of parsimony in the context of a smart grid application, it is also important to realize that the adoption of an

AI technology in this context, i.e., implementing different smart charging agents varying in complexity, is in itself a distributed decision by human decision makers. If more complex algorithms such as DQN could outperform simpler ones, the efficient *equilibrium* of simple models is unstable. Instead, a *prisoner's dilemma* of AI adoption could cause widespread inefficiencies, countering all operational gains through smart technology. Steering such equilibria in beneficial directions is a fascinating research question that would be highly informative for policymakers.

3 A Blockchain-Based Bundle Trading Market for EV Charging[3]

3.1 Introduction

The global push towards decarbonization has placed electric vehicles (EVs) at the forefront of sustainable transportation solutions. According to the International Energy Agency (IEA), the number of EVs on the road is projected to reach well over 400 million by 2030, significantly contributing to the reduction of greenhouse gas emissions and reliance on fossil fuels (International Energy Agency, 2024). However, this rapid adoption of EVs introduces new challenges for electricity grid management, particularly related to the high and variable demands of EV charging.

One of the most pressing issues is the potential for grid congestion caused by simultaneous charging of multiple EVs, especially during peak hours. The traditional grid infrastructure, designed for relatively stable and predictable energy consumption patterns, may struggle to accommodate these new, highly variable loads. This can lead to increased operational expenses, the need for substantial infrastructure upgrades, and, in the worst-case scenario, grid failures (Flath, Ilg, & Weinhardt, 2012).

To address these challenges, it is crucial to develop innovative market mechanisms that can dynamically allocate limited grid resources in an efficient and fair manner. Existing approaches, such as time-of-use pricing and centralized control algorithms, have shown limitations in scalability and flexibility. These methods often fail to account for the decentralized nature of information and control in modern electricity grids (Bichler et al., 2010; Valogianni et al., 2020). Parker et al. (2019) provides an extensive review of some related operational challenges associated with the transformation of power systems.

[3] Joint work with Arthur Carvalho and Wolfgang Ketter.

This paper proposes a novel solution to these challenges: a distributed bundle trading market (BTM) for EV charging. The BTM framework leverages blockchain technology to create a transparent, auditable, and resilient market for allocating grid resources. By treating the rights to use grid infrastructure as tradable commodities, the BTM ensures that these resources are allocated to where they are most valued, based on real-time market dynamics. The use of blockchain technology in the proposed BTM offers several key advantages. Firstly, it enhances transparency and trust among participants by providing an immutable record of all transactions. This reduces the risk of manipulation and increases the market's credibility (Andoni et al., 2019). Secondly, the decentralized nature of blockchain removes the need for a central authority, thereby reducing operational expenses and potential bottlenecks. Finally, smart contracts on the blockchain automate the execution of market rules, ensuring that all trades are conducted fairly and efficiently. Overall, this work lies at the intersection of smart market design, smart grid research, and distributed systems. The relationship and interfaces of these streams can be seen in Figure 3.1.

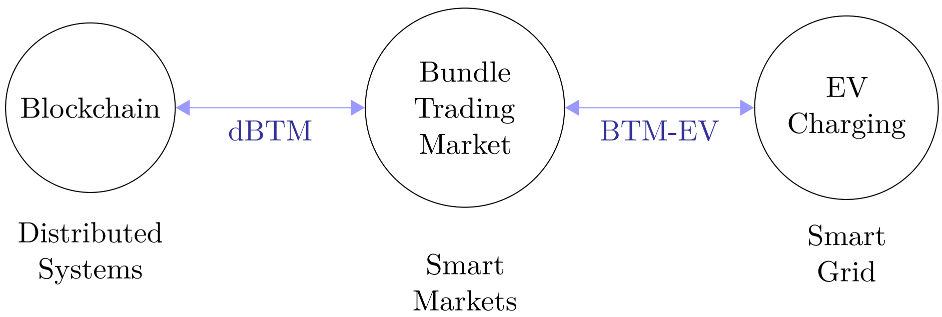

Figure 3.1: Placement of this work in the relevant literature streams.

The proposed BTM for EV charging significantly extends previous (smart) market designs for charging services. It provides an auction mechanism to allocate rights of use for infrastructure assets. Consequently, it is a real-time mechanism for market-based grid or charger fees rather than a market for energy. It thus adheres to existing unbundling regulation (cf. Pollitt, 2008). By design, it can allow for non-participation of (some) grid users, making it more robust against asymmetric adoption and adverse network effects. The auction process is highly transparent, yet it does not reveal actual charging (and thereby driving) itineraries. To our knowledge, no existing market designs provide these features. Using a novel trade verification procedure, the *Strong Duality Verification* (SDV), the presented design enables an efficient smart contract-based implementation of the BTM. To summarize, in this paper, we make the following key contributions:

1. We introduce a novel market mechanism for the efficient allocation of EV charging rights, addressing the unique challenges posed by the high and variable demand of EV charging.
2. We leverage blockchain technology to enhance the transparency, auditability, and resilience of the proposed market mechanism.
3. We conduct extensive simulations to evaluate the performance of the BTM, demonstrating its efficacy in real-world scenarios.
4. We provide practical insights into the implementation of the BTM using blockchain technology, offering guidance for future research and development.

The remainder of this paper is structured as follows. Section 2 reviews related work on distributed resource allocation, dynamic control of EV charging, and blockchain applications in the energy sector. Section 3 outlines the BTM framework, detailing the mathematical formulation and bundle determination process. Section 4 presents a case study evaluating the BTM's performance using the IEEE European Low Voltage Test Feeder. Section 5 discusses how the BTM model can be implemented in a distributed fashion using smart-contract-enabled blockchains. Section 6 discusses practical implementation insights and challenges. Section 7 concludes the paper with a summary of findings and suggestions for future research.

3.2 Related Work

Our research touches upon different streams of literature, distributed resource allocation problems and distributed systems among them. As we see electric vehicle charging as an instance of a resource allocation problem, we also visit the more general stream of auction-based solutions to resource allocation problems in this literature review. It is a broad and mature field of research; for brevity, we present work that is either seminal in the area or has a significant impact on the research presented in this paper. We also review the technical literature on smart charging, particularly but not exclusively from the engineering field. Lastly, we refer to research on distributed ledger technologies (DLTs), such as blockchain, in the energy sector. Lately, this has been a vital area of research, connecting DLTs to numerous exciting use cases.

3.2.1 Distributed Resource Allocation

The interpretation of dual variables in optimization problems as shadow prices provides a fundamental understanding of mathematical programs for resource allocation as models of economic activity. Decomposition techniques allow different actors to emerge from the program's structure, enabling market-based control of distributed resources (Bertsekas, 1992). In traditional centralized systems, a central planner has complete information and controls all resource allocations. However, in distributed systems, information and control are decentralized, requiring mechanisms that can coordinate and optimize resource allocation among self-interested agents. Market-based mechanisms, such as auctions, have been effectively used to allocate resources in distributed settings. These mechanisms align individual incentives with global objectives, ensuring efficient resource utilization (Clearwater, 1996; Fan et al., 2003). Our work contributes to this literature stream by adopting distributed decision-making but also decentralizing the information system.

3.2.2 Dynamic Control of Electric Vehicle Charging

Research on dynamic control of EV charging includes top-down algorithms, pricing mechanisms, and auction-based approaches. Top-down methods involve a central entity that schedules charging based on available information, but this approach is often infeasible due to the high information requirements and lack of scalability (Flath, Ilg, & Weinha, 2012). Pricing mechanisms, such as time-of-use pricing, encourage EV owners to charge during off-peak times. However, these mechanisms can lead to herding effects, where too many users charge simultaneously, causing demand peaks (Valogianni et al., 2015). Non-linear or capacity-based pricing can alleviate some of these issues but often fails to account for the physical constraints of the grid.

Auction-based mechanisms offer a promising alternative by allowing EV owners to bid for charging slots based on their preferences and willingness to pay. These mechanisms can dynamically adjust to changing demand and supply conditions, ensuring efficient allocation of charging infrastructure. The proposed BTM builds on this approach, providing a transparent and resilient market for EV charging rights. By incorporating the physical constraints of the grid into the market design, the BTM ensures that the allocated charging rights do not exceed the available infrastructure capacity. In general, a balance has to be found between energy system-optimal electric vehicle charging and the mobility value of the vehicle (cf. Kahlen et al., 2018). In our work, direct disclosure of charging preferences manages this trade-off.

3.2.3 Blockchain in Supply Chain Management and the Energy Sector

There is a large body of work discussing the use of blockchain technology within supply chains. Much of this literature focuses on issues such as traceability and auditability, i.e., ensuring the provenance of goods or authenticity of identities (Hastig & Sodhi, 2020; Shen et al., 2022). Andoni et al. (2019) provide a comprehensive review focusing on the energy sector applications, including peer-to-peer energy trading and transaction logging. The Blockchain's decentralized and immutable nature makes it suitable for applications requiring transparency and trust among participants. For example, in peer-to-peer energy trading, blockchain can facilitate direct transactions between energy producers and consumers without intermediaries. This can reduce transaction costs and enhance market efficiency (Mengelkamp et al., 2017). Similarly, blockchain can improve the transparency and auditability of energy markets by providing a tamper-proof record of transactions, similar to the general supply chain applications.

The proposed BTM leverages blockchain not only for logging transactions but also as an active component in the market mechanism. Smart contracts on the blockchain can automate market operations, ensuring that trades are executed according to predefined rules. This enhances the transparency and trustworthiness of the market, making it more robust against manipulation and fraud (Wood, 2014).

3.2.4 Review of the Bundle Trading Market Framework

The Bundle Trading Market (BTM) framework was initially designed to manage resource allocation in supply chains through a market-based mechanism. It involves decentralized decision-making, where agents trade bundles of resources to achieve optimal allocations (Guo et al., 2007). The BTM framework ensures incentive compatibility and convergence to optimal solutions through iterative auctions. The BTM framework can be particularly effective in settings where resources are scarce and must be allocated efficiently among competing users. By allowing users to trade resource bundles, the BTM aligns individual incentives with global objectives, ensuring that resources are used where they are most valuable. This approach can be applied to various domains, including supply chain management, energy markets, and telecommunication networks (Bichler et al., 2010). However, the original BTM framework assumes a central entity with complete information, which is not feasible in real-world applications like EV charging. Therefore, adapting the BTM for a distributed setting using blockchain technology addresses this limitation, enabling decentralized control and enhancing transparency.

Though the BTM is a well-known framework, we briefly review it in this work to be able to discuss relevant design choices in the distributed implementation. The framework assumes that a supply chain organization's goal is a linear optimization problem (LP), implying infinitely divisible resources and linear production functions. To correspond to standard notation, the LP is formulated as a minimization problem, as Equation (3.1) shows. The index j represents the J individual agents; the corresponding activity vector is x_j, whereas d_j denotes the cost vector (or negative gross margin) of the agents' activities. The different agents have local constraints $N_j x_j \leq n_j$, where n_j is a non-tradable resource endowment and N_j is the associated resource consumption matrix. In the setting of electric vehicle charging, the most relevant local constraints would be the power rating of the EV charger, or charging restrictions determined by the private driving itineraries. Further, there are shared resources $c \in \mathbb{R}^m$ that prevent the problem's full decomposition. The joint consumption of these resources is constrained by $\sum_{j=1}^{J} C_j x_j \leq c$, where $C_j x_j$ is Agent j's consumption of shared resources at the activity level x_j.

$$
\begin{aligned}
Z(c) = \min_{x_j \geq 0} \quad & \sum_{j=1}^{J} d_j' x_j \\
\text{s.t.} \quad & N_j x_j \leq n_j, \qquad j = 1, ..., J \\
& \sum_{j=1}^{J} C_j x_j \leq c,
\end{aligned}
\tag{3.1}
$$

The organization's goal is thus minimizing the operating expenses of the system, as defined by the aggregation of agents' costs. If a central entity, e.g., a holding or a social planner, had complete information, then this optimization procedure would be viable and ensure optimality. Within the BTM, however, it is assumed that *the optimal solution to the central problem cannot be directly calculated due to the lack of access to relevant information.* As a remedy, the BTM framework uses a market-based resource allocation mechanism to coordinate decentralized decision-making in which independent, self-interested agents trade resource bundles in a double-auction market run by a dealer. Trading bundles instead of individual resources accounts for complementarities and avoids agents suffering from the exposure problem. The central problem is decomposed into sub-problems and distributed among the agents, as Equation (3.2) shows. The level of the resource the agent can use, c_j, is the result of bundle trading and the initial endowment. Given the endowment of resources, agents locally optimize their activities.

$$\begin{aligned} z_j(c_j) = \min_{x_j \geq 0} \quad & d_j' x_j \\ \text{s.t.} \quad & N_j x_j \leq n_j, \\ & C_j x_j \leq c_j \end{aligned} \tag{3.2}$$

To improve their objective, agents can trade for bundles of the shared resources. Bundles that improve the agent's positions are called *improving bundles.* The formal definition of both terms is as follows.

Definition 1 (Bundle)**.** *A bundle $w \in R^m$ is an m-dimensional vector of shared resources. Each of the m elements corresponds to an amount of one specific shared resource. A negative sign of an element component signifies a sell amount; a positive sign signifies a buy amount.*

Definition 2 (Improving Bundle Set)**.** *Assuming current prices $p \in R^m$ of shared resources, the improving bundle set is the set of bundles that lower an agent's operating expenses (at the given price levels):*

$$\begin{aligned} W_j(c_j|p) = \quad & \{w : \exists x_j \geq 0 \ni d_j' x_j + p' w \leq z_j(c_j), \\ & N_j x_j \leq n_j, C_j x_j \leq c_j + w\} \end{aligned} \tag{3.3}$$

A rational agent will choose an improving bundle set for trading that weakly increases its wealth level. Agent j's selection of its most improving bundle w results in the *Bundle Determination Problem*:

$$\begin{aligned} \min_{x_j \geq 0, w} \quad & d_j' x_j + p' w \\ \text{s.t.} \quad & N_j x_j \leq n_j, \\ & C_j x_j \leq c_j + w. \end{aligned} \tag{3.4}$$

To trade in the BTM, agents submit sealed bids to the dealer. The dealer maintains an individual order book for each agent. Each order book contains two different order sets, the set for limited orders $\{w_j^i : i \in I_j\}$, and the set for unlimited orders $\{u_j^h : h \in H_j\}$. The corresponding bids are $l_j(\cdot)$. If an agent submits orders, but there is no trade execution, the orders are collected and accumulated for the respective agent. Otherwise, any agent's trade will clear the individual order book, as an *old* bundle does not safely improve the agent's objective value, given its new resource endowment. To determine the maximal trade surplus, the dealer solves the *Market Matching Problem* (MMP) by finding the optimal degrees of fulfillment y_j^i for limited orders and t_j^h for unlimited orders:

$$\begin{aligned} \max_{y_j^i \geq 0, t_j^h \geq 0} \quad & \sum_{j=1}^{J} \left(\sum_{i \in I_j} l_j(w_j^i) y_j^i + \sum_{h \in H_j} l_j(u_j^h) t_j^h \right) \\ \text{s.t.} \quad & \sum_{j=1}^{J} \left(\sum_{i \in I_j} w_j^i y_j^i + \sum_{h \in H_j} u_j^h t_j^h \right) \leq c_0 \\ & \sum_{i \in I_j} y_j^i \leq 1, \quad j = 1, ..., J. \end{aligned} \tag{3.5}$$

Considering the first set of constraints, if $c_0 = 0$, it is required that buy amounts meet sell amounts. If $c_0 > 0$, the dealer may supply additional resources from inventory to meet buys. The second set of constraints indicates that the market handles limited and unlimited bundles differently. The trades of unlimited bundles in H_j are unrestricted, whereas limited bundles in I_j are restricted. Limited trades are restricted to be convex combinations of bundles in I_j. The dual values (or *shadow prices)* of the first set of constraints in the MMP determine the market-clearing prices $p \in R^m$. The market closes when no trade takes place, and the market prices remain unchanged. If the MMP has a nonzero solution for y_j^{i*}, for $i \in I_j$, or t_j^{h*}, for $h \in H_j$, then agent j will have traded the following:

$$w_j^* = \sum_{i \in I_j} w_j^i y_j^{i*} + \sum_{h \in H_j} u_j^h t_j^{h*}$$

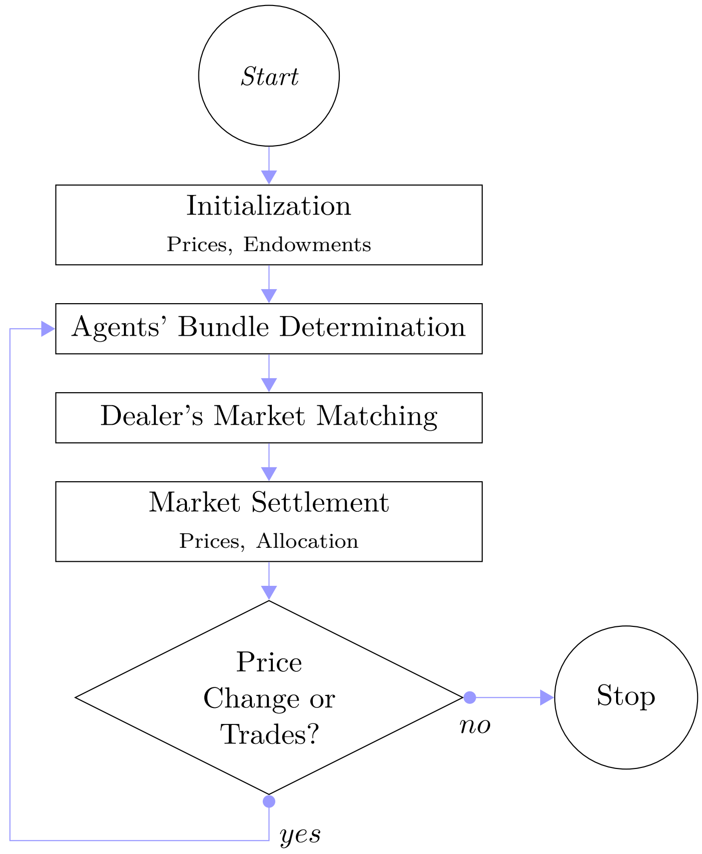

Figure 3.2: Flowchart of the centralized iterative bundle trading market (BTM).

For any initially given resource endowments c_j, for $j = 1, ..., J$, the mechanism enables convergence to the centrally optimal solution through a direct, iterative, and wealth-improving bidding process among the self-interested agents (Guo et al., 2007). Based on the bidding protocol and clearing prices, agents have no incentives to bid untruthfully, as underbidding will lead to underprovision of resources, and persistent overbidding will eventually lead to overpaying. Truthful bidding in this context refers to the fact that agents bid the difference of the optimal objective function value of problems (3.2) and (3.4), i.e., their gain from bundle acquisition. Figure 3.2 shows the operational process of the BTM framework.

3.3 A Bundle Trading Market for EV Charging

The increasing penetration of EVs creates a critical need for efficient and fair allocation of charging infrastructure. Current grid systems are not designed to handle the simultaneous charging demands of a growing EV population, leading to potential congestion and inefficiencies. A novel market mechanism is required to manage these demands dynamically. EV charging presents unique challenges due to the variability in charging needs and the impact of charging on grid stability. Fast charging, in particular, can create significant spikes in demand, leading to potential grid overloads. Traditional approaches to managing grid stability, such as expanding grid capacity or implementing static pricing models, are often costly and inefficient. A dynamic, market-based approach can provide a more flexible and efficient solution.

The primary objective of an ideal solution is to design a market mechanism that maximizes social welfare by efficiently allocating charging rights to EVs. The mechanism should be transparent, resilient, and compliant with existing regulatory frameworks. It should also support non-participation, allowing for a fair initial allocation of resources. The proposed BTM framework treats infrastructure usage rights as the traded commodity. Agents (EVs) bid for these rights, ensuring efficient allocation based on their valuations. The market operates in discrete periods, and the allocation problem is modeled as a linear program, respecting both individual and global constraints. As we discuss later, the power transfer distribution (PTDF) matrix links EV charging to line utilization, and the dual variables represent the shadow prices of line capacities.

The BTM framework is implemented using blockchain technology to enhance transparency and trust. Smart contracts on the blockchain automate the market operations, ensuring that trades are executed according to predefined rules. This eliminates the need for a central authority, reducing the risk of manipulation and increasing the resilience of the market. A peculiarity of the BTM framework is that bundles of resources are traded instead of individual goods. At first sight, this aspect does not extend to EV charging in a straightforward manner, as the charging electricity is typically considered a homogeneous product. We argue that in a potentially congested grid infrastructure, the energy itself is not the sole good, but rather agents have to pay for the rights of use for infrastructure assets (e.g., usage of lines) to enable welfare-optimal charging. The BTM for EV charging does not account for the electricity cost *per se*, but only the (opportunity) cost

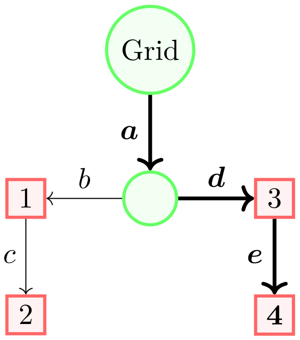

(i) Network topology. The agents are marked as red squares. The green round node does not serve a load, but is a distribution node.

$$\textbf{PTDF} = \begin{array}{c} \begin{array}{ccccc} \boldsymbol{a} & b & c & \boldsymbol{d} & \boldsymbol{e} \end{array} \\ \begin{pmatrix} 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ \mathbf{1} & 0 & 0 & \mathbf{1} & \mathbf{1} \end{pmatrix} \end{array} \begin{array}{c} \\ 1 \\ 2 \\ 3 \\ \mathbf{4} \end{array} \in \{0,1\}^{J \times L}$$

(ii) Power Transfer Distribution (**PTDF**) matrix. Due to the tree topology, all entries are either 1 or 0. a is needed to serve every bus, whereas c and e are only needed to serve agents 2 and 4, respectively.

Figure 3.3: Exemplary network and corresponding Power Transfer Distribution (**PTDF**) matrix. Bold typeset marks an exemplary path connecting node **4** to the grid using lines $\boldsymbol{a}$, $\boldsymbol{d}$, and $\boldsymbol{e}$.

of its transmission. We assume the grid capacity remains unchanged for the time horizon we are considering; agents thus have to bid for a specific set of lines for electricity to be transmitted to them. This intuition is formalized in what follows, using a small grid as a running example.

3.3.1 Mathematical Formulation

Suppose there are four electric vehicles, or agents ($J = 4$). The agents are connected to an upstream electricity grid, as depicted in Figure 3.3i. A total number of five lines ($L = 5$) connects the four agents to the (upstream) grid. The power transfer distribution factor (PTDF) matrix indicates which of the lines transmit electricity from the grid connection point to the individual agents. Due to the radial, or tree, configuration of the network, there is a unique path from the grid to each agent, i.e., the PTDF matrix is binary. This example can represent several interesting settings. First, if only one line a has limited capacity, then only the total charging is limited, and the agents' location within the system is not relevant. If only d and b are limiting, the setup represents two chargers with two cables each, which are restricted in their total delivered power.

If a system operator had full knowledge of agents' preferences for charging, its task of resource allocation would entail a mathematical program, similar to the original bundle trading market. Unlike the original BTM, we are using a multi-period formulation, following approaches in the literature, and also corresponding better to a real-world application. Briefly, it means that agents bid for and purchase bundles of rights of use to charge their electric vehicles now and in the future. For convenience, we are defining one period $t \in \{1, 2, ..., T\}$ as a 15-minute interval, corresponding to intraday trading in the European electricity market, and we are considering a bidding horizon of one hour, i.e., $T = 4$. In practice, the horizon could be substantially larger, e.g., one day or even a week, but the shorter horizon enables us to provide a more concise numerical example. Given these modeling choices, the operator's mathematical program under perfect information is:

$$\min_{\mathbf{X} \geq \mathbf{0}} \quad \sum_j d_j' x_j \tag{3.6a}$$

$$\text{s.t.} \quad \mathbf{X} \leq \mathbf{X_{max}} \tag{3.6b}$$

$$\sum_t \delta t x_{j,t} \leq \Delta_j \quad \forall j \tag{3.6c}$$

$$\mathbf{X} * \mathbf{PTDF} \leq \mathbf{L} \quad (\boldsymbol{\mu}). \tag{3.6d}$$

All inequalities are element-wise operators. The decision variables $\mathbf{X} \in \mathbb{R}^{T \times J}$ ($x_j \in \mathbb{R}^T, x_{j,t} \in \mathbb{R}$) is a matrix of charging powers allocated to agents for given time slots. $\mathbf{X_{max}} \in \mathbb{R}^{T \times J}$ is a matrix of maximum power ratings of agents' chargers per time. If the rating is constant, $\mathbf{X_{max}}$ is row-wise constant. The maximum charge requested per agent j is Δ_j. Assuming knowledge of the true time-specific (negative) valuations $d_j \in \mathbb{R}^T \; \forall j$, this procedure is hence welfare-maximizing. To correspond to the notation of the original BTM, we are posing the welfare maximization as a minimization program, hence the negated valuation vector. The allocation procedure further respects both the individual agents' maximum power and maximum energy constraints (3.6b)-(3.6c) and the line limits (3.6d), where $\mathbf{L} \in \mathbb{R}^{T \times L}$ indicates the lines' limits at given points in the future (cf. Staudt et al., 2018). The time slot length is accounted for by the factor δt. There could also be minimum charging constraints similar to (3.6c), indicating the necessity to charge *at any cost*. We are refraining from defining them here for brevity, as d_j encodes the value of charging. Such constraints would not change the structure of the problem.

As there would be other appliances using the same network, we are imposing that the line limits are augmented to correspond to the residual capacity for EV charging at any given time. Further, the PTDF matrix links the vehicles' charging to the line utilization. The dual variables $\boldsymbol{\mu}$ can now be interpreted as the shadow price matrix of lines' capacity limitations at given times. Without knowing the (negative) valuations d_j, these prices are indeterminate. As such, *the BTM enables price discovery for the networks' lines.*

Note that the BTM for EV charging (3.6) and the original BTM (3.1) are equivalent; the simpler linkage through the PTDF matrix enables the concise notation in (3.6). The BTM for EV charging consists of a central objective to maximize welfare, individual constraints of each agent, and a linkage between agents' activities (or charging) to the limited resources. Agent j's bundle determination problem for EV charging is then:

$$\min_{x_j \geq 0, w_j}., \quad d_j' x_j + PTDF_j * \mathbf{P}' * w_j \tag{3.7a}$$

$$\text{s.t.} \quad x_j \leq x_{max,j} \tag{3.7b}$$

$$\sum_t \delta t x_{j,t} \leq \Delta_j \tag{3.7c}$$

$$x_j \leq C_j + w_j. \tag{3.7d}$$

Note that $PTDF_j \in \mathbb{R}^{1 \times L}$ is the j-th row of the PTDF matrix, which acts as a filter for the line price matrix $\mathbf{P} \in \mathbb{R}^{T \times L}$. By including $PTDF_j$ in the objective, agents do not have to bid directly for individual line capacities. Instead, their requested bundles $w \in \mathbb{R}^T$ correspond directly to the time-specific charging rights they ask for. $C_j \in \mathbb{R}^T$ denotes agent j's current charging rights, either obtained by an initial endowment or through trading. Assuming non-negative prices, agents will never request unlimited bundles (as their charging power is limited). The dealer's simplified market matching problem is:

$$\begin{aligned} \max_{y_j^i \geq 0} \quad & \sum_{j=1}^{J} \sum_{i \in I_j} l_j(w_j^i) y_j^i \\ \text{s.t.} \quad & \sum_{j=1}^{J} \sum_{i \in I_j} w_j^i PTDF_j y_j^i \leq \mathbf{C_0} \\ & \sum_{i \in I_j} y_j^i \leq 1, \quad j = 1, ..., J, \end{aligned} \tag{3.8}$$

$$\mathbf{P} = \begin{array}{c} \begin{matrix} a & b & c & d & e \end{matrix} \\ \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} \end{array} \in \mathbb{R}^{T\times L}$$

(i) Initial price matrix. Price information is available for all lines and open time slots.

$$-\mathbf{D} = \begin{array}{c} \begin{matrix} j & 1 & 2 & 3 & 4 \end{matrix} \\ \begin{pmatrix} 5 & 5 & 10 & 0 \\ 4 & 4 & 0 & 0 \\ 3 & 0 & 0 & 1 \\ 2 & 0 & 0 & 0 \end{pmatrix} \end{array} \in \mathbb{R}^{T\times J}$$

(ii) Valuation matrix. Agents have preferences to charge at given points. Note the negative sign.

Figure 3.4: Price and valuation matrices during an ongoing auction.

where again $l_j(w_j^i)$ are the truthful limit orders of agents j for bundles i. Similar to (3.7a), (3.8) has been adjusted to account for the fact that agents do not bid for line capacities at given times but rather directly for charging rights. Past market matching, agent j's resources C_j are adjusted depending on fulfilled trades y_j^{i*}.

3.3.2 Numerical Example

Assume the BTM's auction was just starting and instantiated with the price matrix shown in Figure 3.4i, indicating the lines' prices for the next four time steps. In this example, initial rates are zero, but other approaches are conceivable (Guo et al., 2012). Figure 3.4ii indicates the agents' valuations for the same periods. An *equal share rationale* dictates the initial endowment with rights of use for lines. Each line's capacity is divided by the number of agents that use it to calculate the fair share. Then, each agent's endowment is the minimum fair share of lines in its path. To simplify the example, assume that the line limits $C_{lt} = 12\ \forall l, t$, meaning that all lines' residual capacity (after other demands) equals 12 for all considered time steps. We are deliberately dropping units in this example for brevity. Then, line a's fair share is the smallest, as all 4 agents have to use it. The endowment is thus $\frac{12}{4} = 3$ units of rights of use for every agent in all time steps, $C_j = (3, 3, 3, 3)'\ \forall j$. Also simplifying, assume that EVs have a maximal charging power of $x_{max,j,t} = 10\ \forall j, t$ and that all agents require at most $\Delta_j = 5\ \forall j$ units of energy.

Corresponding to (3.4), agents now determine their most improving bundle. Consider Agent 1. Its current charging rights limit its power draw to $x_{1,t} = 3\ \forall t$.

Neither maximum charging nor maximum energy constraints are binding in this case. The optimal objective function value without bundle trading is thus $d_1'x_1^* = -(5,4,3,2)(3,3,3,3)' = -42$. Given the current prices, which are all zero, the agent would now wish to adjust its charging schedule by acquiring a bundle. Solving (3.7), the agent increases $x_{1,1}$ and $X_{1,2}$, as these charging times are most valuable to it. Its requested bundle is thus $w' = (7,7,0,0)$. Its truthful limit order for that bundle is $-42 + (5,4,3,2)(10,10,0,0) = 48$. Similarly, the other agents determine their limit orders and bid in the market. In the first step, due to the zero prices, all agents wish to acquire rights of use, but none wish to sell. Hence, the first iteration of the auction market does not execute any trades. However, as the updated prices after the first round are nonzero, agents also start to bid for sales. After 21 iterations, the trading converges (i.e., agents place no more orders, and the prices do not change anymore), and the final schedule from the decentralized scheduling in (3.7) is:

$$X^* = \begin{array}{c} \begin{array}{cccc} j & 1 & 2 & 3 & 4 \end{array} \\ \begin{pmatrix} 0 & 2 & 10 & 0 \\ 2 & 10 & 0 & 0 \\ 10 & 0 & 0 & 2 \\ 8 & 0 & 0 & 0 \end{pmatrix} \end{array} \in \mathbb{R}^{T \times J}.$$

In the first three periods, line capacities are fully exhausted. In the last period, there are 4 remaining units of rights of use, as Agent 1, the only agent with a positive utility coefficient in the last period, already reaches its energy limit when drawing 8 units of power. Note that the charging schedules are not public information after trading; neither are they necessary for market clearing. Public information is only the final charging right allocation, i.e.,

$$\mathbf{C} \approx \begin{array}{c} \begin{array}{cccc} j & 1 & 2 & 3 & 4 \end{array} \\ \begin{pmatrix} 0 & 2 & 10 & 0 \\ 2 & 10 & 0 & 0 \\ 10 & 10 & 0 & 2 \\ 10 & 0.15 & 0 & 0 \end{pmatrix} \end{array} \in \mathbb{R}^{T \times J}, \quad \mathbf{C} * \mathbf{PTDF} \leq \mathbf{L}.$$

The system or network operator would thus only have to have sufficient control to limit the EV chargers' fuses but would not obtain direct control of charging power. Further, agents may decide not to use their secured charging rights so that their schedule remains private.

3.4 Case Study

The efficacy of the proposed BTM, in particular when implemented on a distributed ledger using smart contracts, depends on the speed of convergence of the market. Slower convergence times or an increased number of trades, respectively, result in difficulties in clearing the market close to real-time and, therefore, satisfyingly for a power system application. Whereas the previously provided stylized example serves to build intuition for the proposed mechanism, it is not a realistic setting. To provide such a one, we build on a well-known and widely used distribution grid representation, i.e., we adapt the IEEE European Low Voltage Test Feeder (LVTF) to model the residential distribution grid (cf. Schneider et al., 2018). The LVTF was explicitly designed to analyze the impact of an increased amount of distributed resources on distribution grids. However, note that it was also designed to support investigations regarding phase symmetry and is thus a three-phase system. Because we are investigating only power flows, we are using a single-phase version instead. This approach is reasonable as long as chargers themselves are assumed to be three-phase, which is undoubtedly the case for higher charging powers. Figures 3.5i and 3.5ii show the original test grid and a reduced version, respectively. The reduced version maintains all electrical characteristics while removing nodes that are not necessary for modeling.

For the case study, we assume initial prices of zero for all lines in the network. A better initialization based on experience would generally provide faster convergence (cf. Guo et al., 2012), in particular in this relatively repetitive environment. At the same time, e.g., time-of-use grid fees as initial price vectors would not suffer from herding, as the market mechanism ensures allocation of only available grid capacity. The emergence of strictly positive prices depends highly on the grid's topology. If the line (or transformer) connecting to the upstream grid is most limiting, it will be relatively often price-setting, inducing equal grid fees in the entire system. If line capacities and user numbers are more balanced, prices might be highly localized. In the LVTF, the upstream line's capacity is small

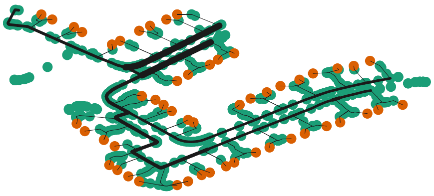

(i) Original IEEE Test Feeder with 55 load nodes (orange) and 905 lines (black). Green nodes are connection points without loads.

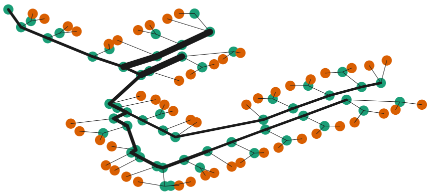

(ii) Reduced IEEE Test Feeder with load 55 nodes (orange) and 110 lines (black). Green nodes are connection points without loads.

Figure 3.5: Original and reduced test grid. Both topologies are equivalent for the presented use case.

compared to the other lines, reducing most congestion occurring right behind the transformer.

The valuation of charging is difficult to obtain in our case study, as most current tariff schemes do not provide sufficient pricing differentiation to allow for consumers' self-selection. We therefore rely on synthetic data to investigate the efficacy of the proposed mechanism in this setting. We vary valuation distribution, the (average) number of connected vehicles, charging power, and market horizon as follows:

- *Valuation*: we assume willingness-to-pay for grid usage in the same order of magnitude as grid fees in European markets.We assume that the valuation to pay for grid usage (e.g., on top of a constant component) is distributed between 0 and 10.0 monetary units (MU). We use $Beta(1, 1)$, $Beta(2, 5)$, and $Beta(0.5, 0.5)$ distributions to allow for a wide range of potential demand characteristics.
- *Average number of connected vehicles*: we assume a random presence of vehicles. Connection probabilities of vehicles are set to 25%, 50%, and 100%, respectively.
- *Charging power*: we set the charging power to either 3 kW, 11 kW, or 22 kW.
- *Market horizon*: we set the market horizon to one of three options: 2 hours, 4 hours, or 8 hours.

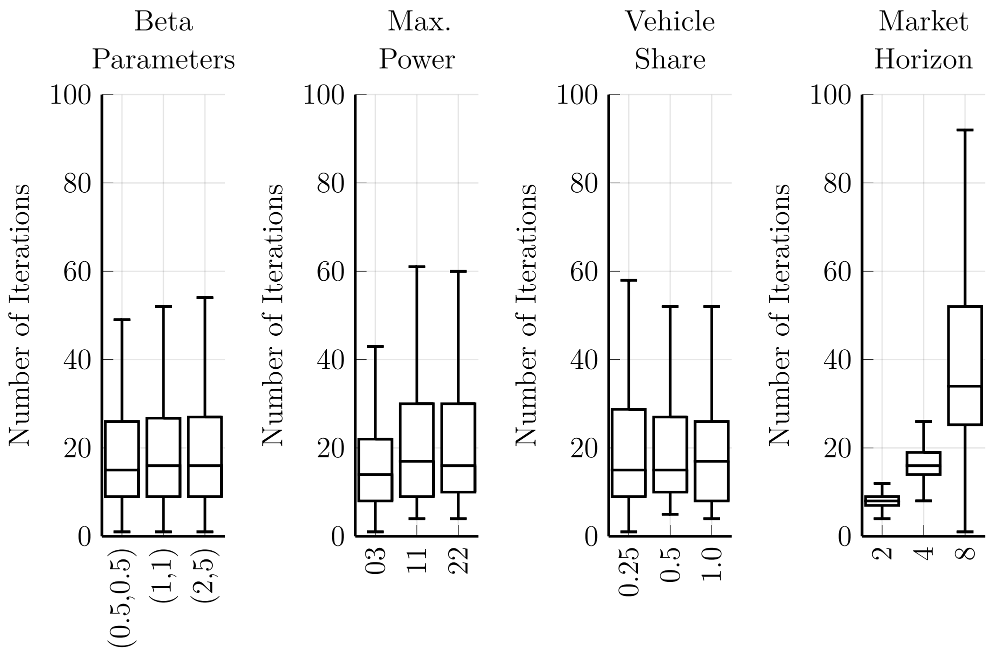

Figure 3.6: Iteration analysis of the proposed algorithm. The clearest dependence can be seen regarding the market horizon. All other input parameters are mostly uncorrelated with the resulting iteration count.

We investigate all combinations of parameters, yielding $3^4 = 81$ simulation scenarios. To allow sufficient variation in the random parameters, we ran 100 simulations of each scenario, resulting in 8,100 total simulation runs. Of the 8,100 simulation runs, 97.5% converged in 100 iterations. All other cases did not converge in that time but achieved virtually the same social welfare as the corresponding central solution, indicating that a (randomized) early stopping of the mechanism is a feasible approach to ensure both swift market clearing and high efficiency. The distribution of iteration duration for the simulation runs with up to 100 iterations is given in Figure 3.6.

The convergence is practically independent of the concrete distribution of charging valuations, indicated by the β-tuples representing both parameters of the associated Beta-distributions. The average number of vehicles is, compared to other varied parameters, a less significant driver of convergence time. While numerous vehicles increase the average steps needed for market clearing, it is much less pronounced compared to the increase for larger market horizons or larger charging powers. Both a longer horizon and larger charging powers make the

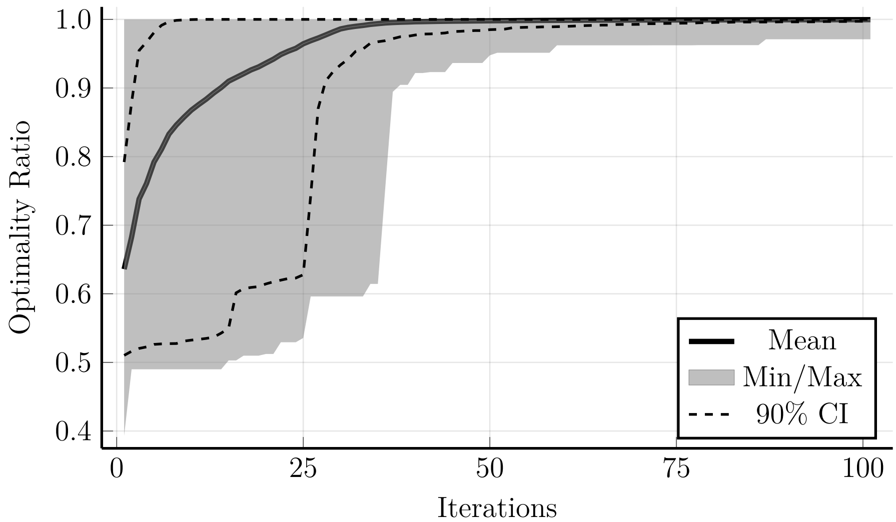

Figure 3.7: Convergence characteristics of the proposed mechanism in the case study for Beta$(1, 1)$ valuation and a long market horizon. In most cases, good performance can be reached within 30–40 iterations.

coordination effort more challenging, as the best solution will depend more heavily on non-constant charging schedules to be coordinated with one another. Moreover, low charging powers in particular reduce the occurrence of grid congestion and thus the cases in which there is actually a need for (scarcity)pricing of lines.

To analyze the development of social welfare during the market mechanism, we focus on the most complex market setting with the longest convergence times, i.e., a long market horizon. As the concrete distribution of charging valuations did not indicate a significant impact, we chose the Beta$(1, 1)$ distribution case for this analysis. Figure 3.7 shows the evolution of optimality ratios, i.e., the ratio of welfare obtained in the distributed mechanism and the centralized optimization throughout iterations. The dashed lines representing the 90% confidence interval of optimality ratios indicate that most settings achieve high optimality after roughly 30 iterations. After 40 iterations, all settings achieve optimality ratios over 90%.

To improve computational performance, a premature stopping criterion can be developed. To avoid gaming of market agents, this stopping criterion could be randomized and thus be used to determine market clearing, as could other heuristics regarding the cleared trades. Accounting for the fact that convergence can potentially be sped up using a non-zero initialization of the price vector and

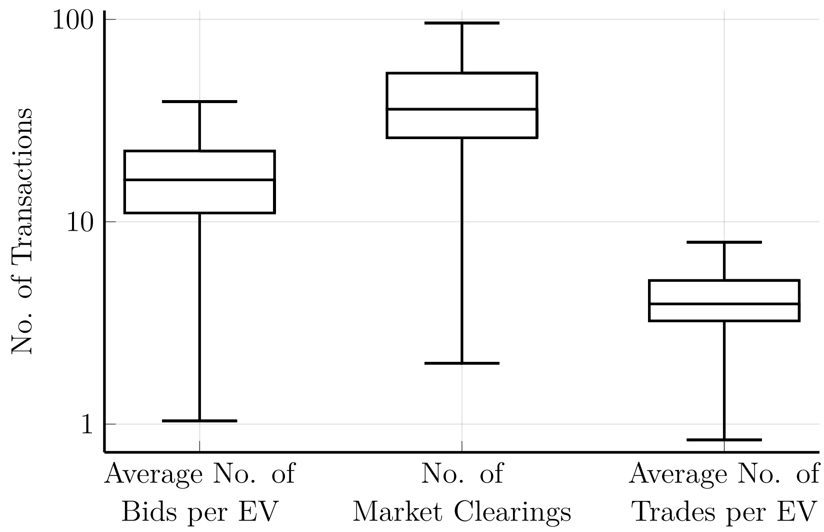

Figure 3.8: Necessary transactions of the proposed mechanism in the case study for Beta$(1, 1)$ valuation and a long market horizon shown on a logarithmic scale. In most cases, only a few bids and trades are necessary per electric vehicle to achieve convergence.

accounting for the fact that pre-existing allocations from earlier market runs (recall the rolling horizon) will further improve computational characteristics, timely convergence of the proposed market mechanism can be asserted.

For a more profound understanding of the market mechanism, we further analyze the transactions necessary for market clearing in the case of a long market horizon, i.e., the most computationally taxing clearing. Figure 3.8 shows the distribution of the numbers of bids, trades, and market clearings across those settings.

3.5 Distributed Implementation

Whereas the bundle determination and bidding of the agents in the original BTM happened locally, the market matching and settlement are tasks of a central auctioneer. Blockchain is a technology candidate to remove this intermediary (Andoni et al., 2019) and, thus, create a distributed BTM (dBTM). To discuss whether a Blockchain-based implementation of the BTM framework could be worthwhile, we follow the ten-step procedure developed by Pedersen et al. (2019).

Figure 3.9 depicts the decision process. In summary, the dBTM design fulfills most of the steps outlined for assessing the suitability of blockchain implementation. It effectively uses blockchain's strengths to address the needs of a distributed market for EV charging, with considerations for shared databases, trust, intermediaries, governance, immutability, public access, and privacy. Scalability is addressed through off-chain processing, though it remains a partial fulfillment due to the inherent limitations of some blockchain models.

The removal of the intermediary can potentially reduce transaction costs, particularly in the area of network operation that is prevalent in EV charging. The network operator owns a natural monopoly and must thus be regulated to avoid that welfare-reducing monopolistic rents are obtained. The use of blockchain can hence be interpreted as a digitally ensured regulation, using affordances such as smart contracts to codify regulatory needs. In fact, smart contracts are the key functionality we use to implement the aforementioned iterative market, effectively resulting in real-time grid fees. To ensure the efficacy of the proposed market mechanism, we emphasize the design of a suitable consensus mechanism for the dBTM.

The first blockchains (e.g., Bitcoin) were designed to enable digital currencies. In digital—as well as *regular* – currencies, the system needs to validate that units of value are spent only once, i.e., there is no double-spending. Conventionally, this is done by banks or other centralized and trusted financial institutions that act as intermediaries. In some cases, however, central management is either infeasible or unwanted by stakeholders. Reasons for this can be intermediary costs or a lack of trust of the users in any intermediary to operate the system; central management further has the disadvantage that it constitutes a single point of failure. Depending on the application, the centralized system is fragile to technical problems and to external malicious attacks (Andoni et al., 2019). Consequently, the main rationale behind blockchain technologies is the development of technological features that enable the removal of such intermediaries. This feat is achieved through a distributed network of various computational nodes, which are incentivized to cooperatively verify transactions and protect the validity of the ledger.

There are numerous blockchain models available, all having different idiosyncrasies and applicability levels, depending on the respective application. We do not set out to discuss the advantages and disadvantages of particular implementations in detail. Instead, we argue how functionalities of modern DLTs can be leveraged to automatically run an iterative market, i.e., to use the blockchain

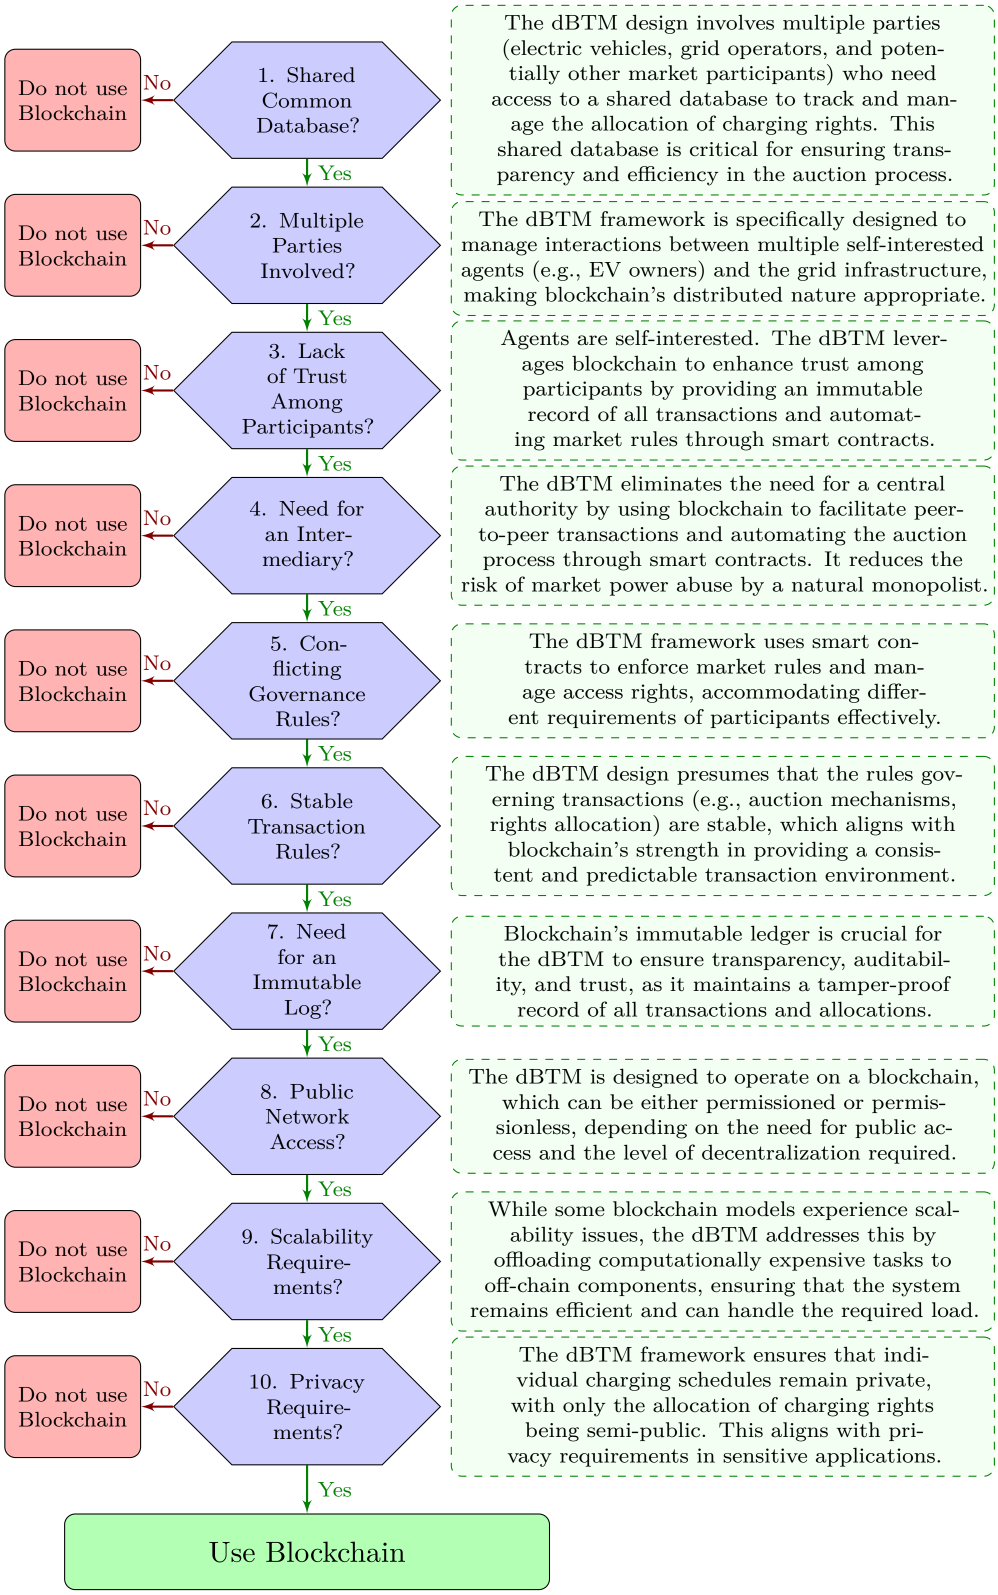

Figure 3.9: Decision process for the use of blockchain (Pedersen et al., 2019). Hexagonal nodes show the relevant decisions for a blockchain application, and dashed boxes on the right explain the fulfillment of the criteria.

as an auction engine. Due to its technological maturity and widespread use when compared to other blockchain variants with smart contract functionality, we loosely base our following discussion on *Ethereum* (Wood, 2014). Nonetheless, we reckon that our approach can easily be adapted to accommodate other smart contract-enabling blockchains.

3.5.1 Ethereum

Ethereum is a self-proclaimed transaction-based state machine (Wood, 2014). The state machine starts with a so-called *genesis state*, with no executed transactions. It then executes transactions incrementally, and the state of the *Ethereum world* changes into new states. Through this, it morphs into some final state. Consequently, the final state has universal acceptance as the canonical version of the world of Ethereum and is agreed upon by all agents interacting with the state machine.

Transactions are the method for agents (or *accounts*) to interact with each other. There are two different types of transactions: those that result in a so-called *message call*, similar to a traditional transfer, and those that result in *contract creation*. The blockchain does not process transactions individually but accumulates them in blocks (hence the name). An Ethereum block contains a list of transactions, the state root representing the final state after processing these transactions, and a header with meta information linking it to the blockchain, including hashes of the state, transactions, and receipts. Including representations of the preceding state chains blocks together. The sequence of blocks that make up the blockchain holds a complete list of transaction records (Zheng et al., 2017). In the dBTM framework, a transaction can, for instance, be a bid on a specific bundle as determined by the bundle determination problem. Due to the distributed nature of the blockchain, transactions must be broadcast through the network. We will hence outline the transaction lifecycle (with a message call as an example) on Ethereum.

A message call (e.g., a transfer) has a destination (implied by a *to* attribute). The transaction with the stated attributes will be first created and then submitted to an Ethereum node. After the receiving node validates the transaction, it sends it to its peer nodes, distributing it. Figure 3.10 outlines this process. It also indicates that there are both mining and non-mining nodes in the Ethereum network. The mining nodes take transactions from the pool and process them into a pending block. There are no specific rules on how the mining nodes sort the transactions. In public implementations, posted transactions have an associated *gas* price, i.e., a transaction fee. Therefore, miners typically sort transactions by the gas price. In private blockchains, this does not have to be the case.

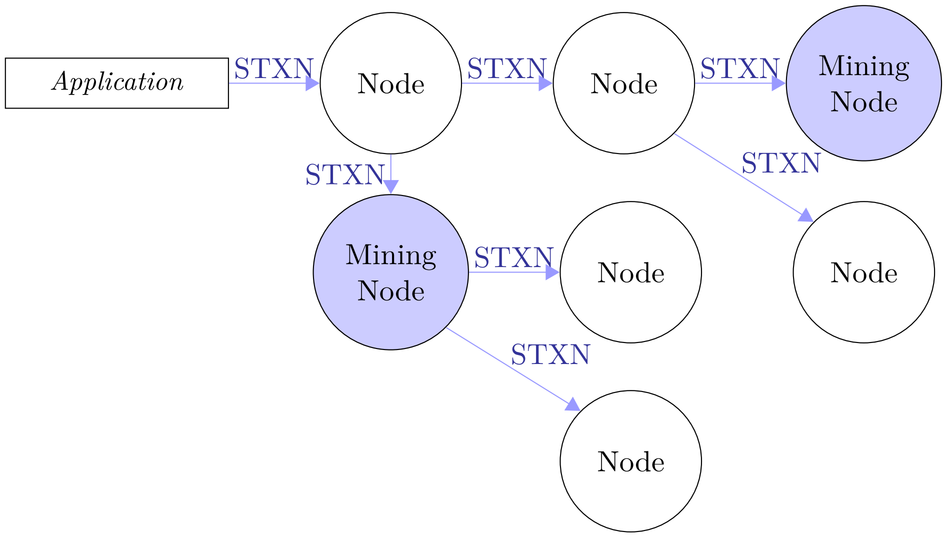

Figure 3.10: Node network demonstrating a transaction broadcast. STXN denotes a signed transaction.

The process that generates and validates a new block is called a *consensus mechanism*. Various consensus algorithms exist, and each of them provides different advantages and disadvantages. To a large extent, the key performance drivers of a blockchain model, like transaction speed, scalability, and security, depend on the embedded consensus algorithm (Andoni et al., 2019). At the time of writing, the Ethereum blockchain, which hosts the cryptocurrency *Ether*, uses the Proof of Stake (PoS) mechanism, having shifted from Proof of Work (PoW) in September 2022. PoW enables high participation in terms of the number of miners (Vukolić, 2016), although it results in a reduced transaction speed and high-power consumption due to the solving of a cryptographic puzzle, which requires significant computational effort (Andoni et al., 2019). Due to these drawbacks, other consensus mechanisms like PoS, or Proof of Authority (PoA), have gained importance in (corporate) blockchain implementations, depending on the underlying application and trust relationship between parties. For a more detailed description of consensus mechanisms, see, e.g., Zheng et al. (2017). Note that widespread criticism of blockchain technologies concerning their sustainability is directly linked to the use of PoW as their consensus mechanism, while PoS alleviates such concerns.

Smart contracts are the enabling element of Ethereum to make it a utilizable infrastructure for the BTM framework. A smart contract is an immutable computer program that is executed deterministically by the Ethereum Virtual Machine (EVM) as part of the network protocol. As the smart contract is immutable, its source code cannot be changed after its deployment to the blockchain. Determinism means that the output of the smart contract must depend solely on its source code and the world state, which is mutually agreed upon by all parties. Each smart contract is identified and reachable via a known Ethereum address, which can be used in a transaction as the recipient to send funds to the contract or to call one of the contract functions. In contrast to an externally owned account, there are no private keys associated with an account of a smart contract. The creator of a smart contract does not have any exclusive rights on the protocol level. However, it is possible to implement these exclusive rights in the source code, for instance, by limiting who is allowed to invoke the contract. Smart contracts are only active and run if a transaction calls them.

3.5.2 dBTM Architecture

The proposed distributed BTM uses blockchain technology as the underlying IT infrastructure. The agents participating in the BTM submit and receive all relevant data, information, and payments via the blockchain. Moreover, a dealer implemented as a smart contract, with a conventional software client, operates the double-auction market that enables agents to trade their bundles. The *smart contract dealer* encodes all relevant information regarding the market mechanism, like submitted orders, trades, dealer's inventory, and market prices. Agents exclusively communicate with the smart contract dealer. This ensures that the mechanism to clear the market and allocate resources is transparent to all agents.

As Figure 3.11 shows, the agents submit their orders to the smart contract dealer. Afterward, the dealer's conventional software client, labeled *Off-Chain Dealer*, fetches all existing orders contained in the smart contract and solves the market matching problem of the BTM. Then, the off-chain dealer updates all relevant information in the smart contract, such as the settled orders, the trades, the new market prices, and the resource inventory. Finally, the dealer's smart contract informs the agents regarding their respective trades.

The off-chain dealer is necessary due to the potentially expensive calculation of the MMP. Even though solving an LP has limited computational complexity, it can be substantially costly for sufficiently large problems. As a consequence, implementation and execution of the MMP in the smart contract itself would be impractical, as running code on a public blockchain is expensive. A well-designed smart contract allows one to move computational complexity away from the blockchain and focus more on providing information.

A drawback of moving any functionality off the blockchain is that the off-chain source code is not visible to the stakeholders anymore. To still ensure that all agents can validate the correctness of results from off-chain functions, it must be ensured that the returned values of the off-chain functions are authenticatable given the inputs. We address this issue by introducing a novel trade verification mechanism based on strong duality theory, *Strong Duality Verification* (SDV). The details of SDV are described later in this section. The process of the proposed dBTM then has three parts, representing the thematically summarized actions.

1. *Bidding*: Agents submit bids to the smart contract dealer.
2. *Market Matching Mechanism*: The market dealer solves the MMP.
3. *Trade Verification* Agents verify whether the solution is valid.

The behavior of a peer-to-peer system such as blockchain differs from that of conventional client-server architecture. For the upcoming discussion, it is important to note that two different method invocations are available in the Ethereum blockchain environment: the *transact* method and the *call* method. The difference lies in the fact that the *transact* method requires the submitting of verified transactions that potentially change the blockchain's state, whereas the *call* method is free of charge as it only reads the state of the blockchain. A *transact* method requires network verification, which causes a delay due to the

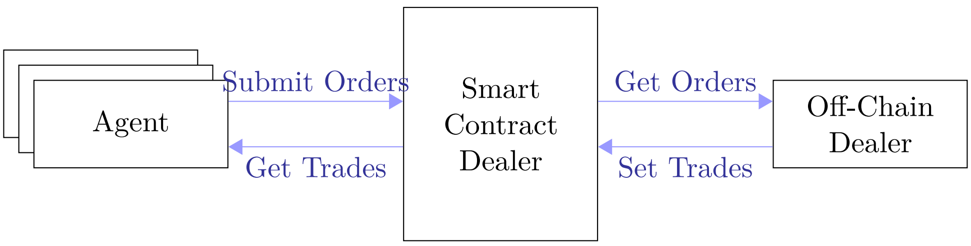

Figure 3.11: Communication flow between agents, on-chain, and off-chain dealer.

mining procedure. Consequently, before the values written into the blockchain by a *transact* method are accessible, some time passes. It is essential to take this into account when trying to read specific values from the blockchain that are set in previous steps.

Bidding

The agents get the current market prices from the smart contract dealer to determine the improving bundle set. In the first round of the auction, these rates can be either initialized with zero or with other arbitrary positive values. While the right choice of initial prices can influence the speed of convergence, convergence itself is unaffected. More sophisticated mechanisms could use market intelligence to forecast prices to increase convergence speed (Guo et al., 2012).

Receiving the market prices from the dealer's smart contract is a *call* method and hence proceeds rapidly, as there is no state-changing transaction. Subject to the price vector, the agents solve the bundle determination problem and submit the orders in the dealer's smart contract. In contrast to receiving market prices, the placement of orders requires the submission of verified transactions that potentially change the blockchain's state. As mentioned earlier, this causes a delay. Since the market matching mechanism needs to access the set bundle orders, it ensures the finality of all newly submitted orders before market clearing. Figure 3.12 outlines the complete process of agent bidding.

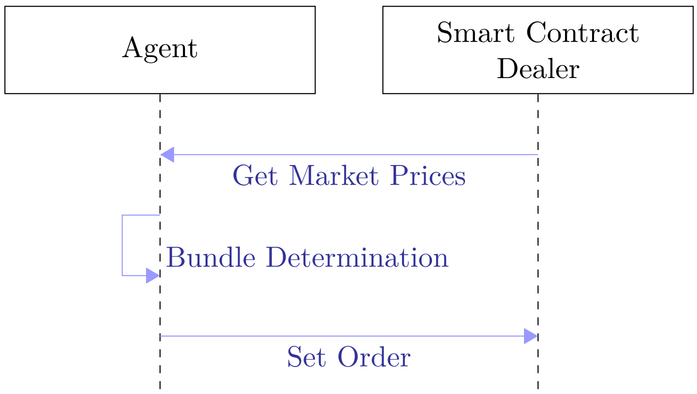

Figure 3.12: Agent bidding.

Market Matching Mechanism

The smart contract dealer waits for the finality of the agents' bundle orders and then fetches them. The off-chain dealer then obtains the orders from the smart contract dealer, which is also implemented by a *call* method. The transfer from on- to off-chain dealer takes place virtually instantly. Based on the received orders, the off-chain dealer creates and solves the MMP. As the MMP is solved off-chain, the concrete implementation to solve the problem instance is not restricted, which means that commercial solvers enable large problem instances. However, this also means that the solution approach is not auditable by agents in a fully distributed setting. The lack of auditability makes verification enhancements necessary but dramatically increases the flexibility of the design we present.

After solving the MMP and calculating the respective trades, the off-chain dealer sets the smart contract dealer's trades and removes the settled orders, as well as other orders of agents that participated in trade in this step, as their remaining bundles are not necessarily improving anymore. The newly calculated market prices, the MMP's solution, and its duals, which are necessary for SDV, are set in the smart contract dealer. The setting of values in the smart contract requires verified transactions that use the *transact* method. As the values are needed to confirm the trades, the agents have to wait for the finality of the transaction. Figure 3.13 shows the dealer's market matching mechanism.

Trade Verification

Moving calculations off the Blockchain comes along with the drawback of opacity. To address this issue and to enhance trust in the solution of the MMP, we developed the trade verification procedure SDV, which builds on the strong duality theorem. Agents retrieve the values of the optimal solution of the MMP and the corresponding duals from the smart contract dealer. To verify if the off-chain dealer calculated the respective values correctly, the developed SDV is applied. For brevity, we adhere to the standard notation of linear programming. If a primal LP has an optimal solution x^*, then the dual also has an optimal solution y^*, such that $c'x^* = b'y^*$. As a result, it is possible to verify that the solution of the dealer is optimal without solving the corresponding MMP. The agents calculate whether strong duality holds and decide based on the result whether they accept the trade. After SDV, the agents notify the smart contract dealer whether they accept or reject their trades. In case of acceptance, the

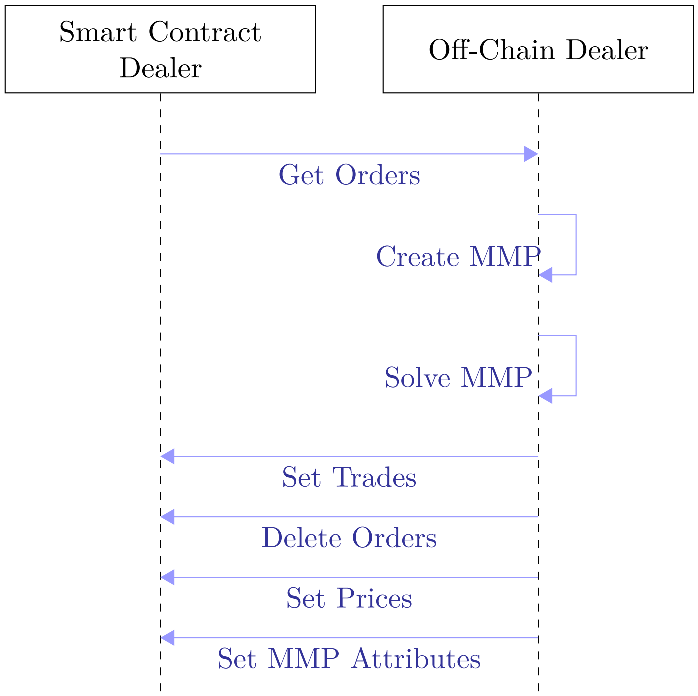

Figure 3.13: Market Matching Mechanism.

smart contract dealer returns the trade. Otherwise, in case of rejection, the smart contract dealer deletes the respective trade. In the case of acceptance and subsequent return of the trade, the agent adds the respective trade to their shared resources. Figure 3.14 shows the process of agents verifying the trades that have been posted by the smart contract dealer.

The described blockchain solution allows for almost all aspects of the BTM to be transparent and auditable by all actors. Solving the MMP that might be computationally expensive is moved to an off-chain software client to enhance flexibility, but SDV alleviates concerns regarding the subsequent opacity. Even though in practical implementations a market entity (e.g., a system operator) would typically calculate the MMP, this does not have to be the case. Every participant in the market may solve it and propose its solution to be verified by the other participants.

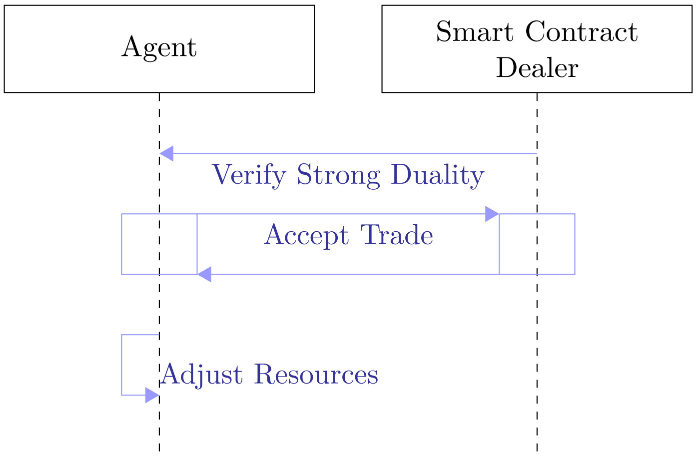

Figure 3.14: Trade verification.

3.6 Discussion

Several merits of the proposed distributed BTM for EV charging significantly extend the current state-of-the-art. First, the fact that agents bid for rights of use instead of units of energy or power is more compatible with unbundling regulations in European electricity markets. After obtaining time-specific rights of use via the proposed mechanism, they have contractual freedom to buy energy from arbitrary suppliers. The role of the network operator and utility can thus be separated. Second, agents do not have to participate in the market. The initial *fair share* endowment ensures that agents are endowed at least with an equal share of the existing infrastructure. They *may* either sell excess capacity in high-price times or buy additional charging rights when desiring more than their fair share, but they do not have to do either. This rationale enables the introduction of the BTM for EV charging, even if not all potential participants own intelligent software agents. The opportunity costs of non-participation may drive the investment incentive for smart technology.

The proposed BTM for EV charging is also more transparent than most mechanisms in the literature. Being a fully distributed market, no one central player is setting charging schedules, as would be the case in top-down control. Further, by bidding for rights of use, there is no indication of that right's execution. Agents with a high valuation for privacy might decide not to participate in the BTM; others might overbuy or undersell to mask their actual charging power. The only semi-public (i.e., known by the participants) information is the allocation of charging rights, not charging itself.

The distributed implementation allows for an auditable and resilient design. Unlike other application areas of blockchain technology, such as finance, energy systems necessarily have known institutional intermediaries like network operators. Further, full anonymity is infeasible because the agent's location has to be known to enable welfare-maximizing allocation. However, we argue that the increase in transparency makes the distributed approach more widely applicable; for instance, in the developing world, where both trusted institutions and infrastructure may be scarce. Moreover, given recent developments in blockchain technology, the notion that a distributed ledger implementation must be less efficient than a centralized design is not guaranteed (cf. Androulaki et al., 2018).

It is still important to acknowledge that some limitations persist. The linear and deterministic formulation does not capture all aspects of distribution system management, i.e., reactive power and voltage control. Moreover, the randomness of available capacity due to other forms of usage is not accounted for. While the separation between EV and non-EV load seems a practical decision, it might open up the opportunity for participants' strategic behavior. However, due to the ongoing adoption of smart meter technology, it is possible to include non-EV loads easily. Lastly, price formation necessitates knowledge of preferences and valuations for charging of EV users. This information is mostly tacit, as for an individual EV owner to derive its valuation for energy at a given moment might be cumbersome. Developments in intelligent software agents and the uses of big behavioral data might help overcome this gap by inferring (numerical) valuations without EV owners having to set them explicitly.

3.7 Conclusion and Future Work

We propose a variant of the bundle trading market to auction off rights of use for electricity infrastructure between self-interested agents. The distributed BTM for EV charging is a welfare-maximizing, incentive-compatible, iterative auction mechanism running distributed on a blockchain. The use of blockchain improves auditability and transparency and enhances trust; it further increases system resilience. The proposed mechanism is transparent and allows for limited participation.

Based on the current limitations, some extensions seem highly worthwhile. First, redesigning the original BTM to allow for a convex structure of the underlying mathematical program allows for more complex—and more accurate—mapping of physical limitations onto the market clearing. A stochastic extension might be valuable (cf. Dvorkin, 2019).

Lastly, as with all mechanisms designed for real-world applications, field research should be conducted. At the same time, this would enable learnings on charging preferences and valuations that are necessary for the proposed market design to work. Throughout the process, the impact of other technologies, such as heat pumps or behind-the-meter battery storage, could also be investigated. This exploration into a distributed, blockchain-based bundle trading market not only highlights the practical feasibility of decentralized EV charging infrastructure but also sets the stage for more innovative, scalable, and sustainable solutions in energy management. As technologies evolve and adoption grows, the proposed market could become a keystone in achieving a resilient and efficient energy ecosystem, aligning economic incentives with sustainable progress and enhancing both accessibility and operational integrity in smart grids.

4 Market-Based Optimization in Decentralized Autonomous Organizations[4]

4.1 Introduction

Organizational decisions are often subject to risk. Reserving computing resources to serve a load, for instance, might lead to uncertain future costs or payoffs. Although high earnings (in expectation) are generally favorable, some organizational units may want to avoid high costs in adverse scenarios, depending on their risk aversion. Decision makers have different methods at their disposal to incorporate their risk attitudes into their actions. On the one hand, decisions concerning physical assets (such as computing resources) can be altered. [5] A physical decision with moderate expected earnings could be preferred as long as it is never ruinous. On the other hand, financial instruments might be available to manage exposure. If an operational decision performs terribly in only specific scenarios and securities exist with a payout in precisely these situations, then, depending on the securities' prices, the combined physical and financial decision might be sensible, even for a highly risk-averse decision-maker.

In decentralized autonomous organizations (DAOs), the challenge of managing risk is emphasized. Different members of the organization, which we call agents, may have a different perception of risk while at the same time exhibiting individual goals that might not align perfectly with the organization's goals. Thus, these DAOs must not only find a way to allocate physical resources and find the optimal *production plan*[6], but also allocate financial resources to account for the heterogeneous risk aversion of the agents. This setup is notably different from standard distributed optimization approaches, in which it is assumed that there is a central objective and all distributed resources are expected to work towards it. The distribution is then merely a tool for managing computational tasks

[4] Joint work with Wolfgang Ketter.

[5] In this paper, physical assets do not only refer to tangible commodities but any goods related to the physical world, such as utilization of computing hardware.

[6] We use *production plan* as an umbrella term for decision-making regarding physical assets.

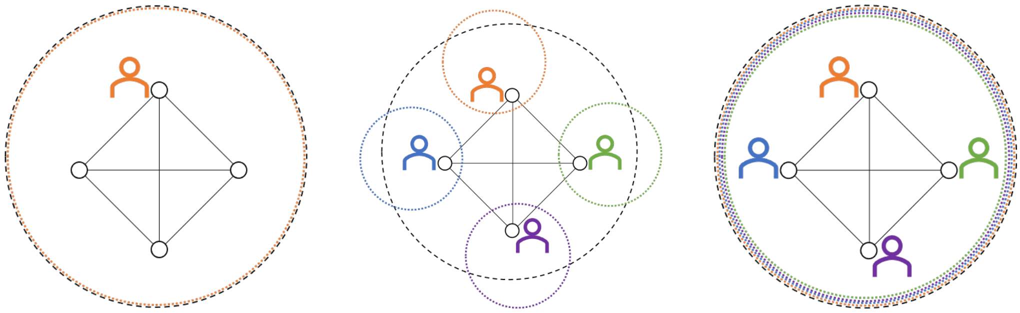

(i) Distributed Optimization. A single decision-maker uses distributed computational resources.

(ii) Decentralized Decision-Making. Individual agents' incentives do not have to be aligned with the common goal.

(iii) Incentive-aligned Decentralized Decision-making. Every agent has incentives to optimize the common objective.

Figure 4.1: Different forms of distributed and decentralized decision-making. Dotted lines show agents' incentives, and the dashed line shows the common objective.

or reducing communication overhead. In a truly decentralized setting, decision makers must be motivated to work towards the common objective (cf. Wörner et al., 2022); their incentives have to be aligned. Figure 4.1 visualizes the distinction. Our work considers the following scenario: different agents form a DAO, e.g., a cooperative or a supply chain network, with pronounced responsibilities of agents. The agents decide on their *division's* activities subject to a local objective value. They optimize their individual results. The decentralized organization has shared resources that must be allocated to these agents. To acquire them, they bid on an internal auction market. We build on existing research (originally posed by Ba et al. (2001a); Fan et al. (2003); Guo et al. (2007, 2012) and further developed subsequently), but extend it by assuming that the agents' utility is not uniquely determined by their activities but is stochastic. For any choice of activities, for example, any production plan, the actual outcome depends on random factors. To account for the risk associated with this uncertainty, agents not only internally trade to acquire shared physical resources but also issue and trade financial securities among themselves. By doing so, they can trade risk to find an efficient risk-aware allocation within the organization. Beyond important theoretical insights, our work is motivated by recent technological advances in the area of finance, naturally abbreviated *fintech*. Heterogeneous advances in this area share the commonality that they reduce transaction costs and entry barriers in trading activities, improving the general accessibility of financial products. One technology that particularly spurred the work on this paper is Blockchain, especially when enabling the so-called *smart contracts* (Wood, 2014). Smart contracts enable encoding of (almost) any contractual logic into software, guaranteeing fulfillment of the contract subject to the defined terms. This concept corresponds naturally to contingent claims—or Arrow-Debreu securities—that prompt a payout if a specific *state* of the world occurs Arrow and Debreu (1954). These securities can be used to internally finance distributed organizations and allow for integrated risk management that combines physical and financial decision-making. Furthermore, smart contracts can be used to implement business logic such as the mentioned auction protocol. In principle, all parts of the proposed mechanism are, thus, suitable for decentralized implementation.

Our study contributes to a growing body of research related to the use of auctions and market mechanisms for intra-organizational resource allocation (Ba et al., 2001a; Fan et al., 2003; Guo et al., 2007). Our focus lies on the joint allocation of physical and financial assets. Previous research has explored various methods to allocate physical resources within organizations, and some have also included externally supplied financial instruments to optimize resource allocation under uncertainty. To our knowledge, no existing research provides an intra-organizational market mechanism that optimizes the use of physical resources while using self-issued financial instruments to manage risk within the firm. Building upon previous work on incentive-compatible auctions, we find that our market design provides an efficient way to issue securities and allocate them jointly with shared resources. The method scales well in the number of agents, shared and private resources, and activities. Lower scalability concerning the number of securities can be met by offering just a subset of all future outcomes as contingent claims, still achieving good allocative efficiency. Given the well-described task of the market operator, the mechanism can be implemented using distributed ledger technologies, as has been proposed for prediction markets, deepening decentralization (Carvalho, 2020). Within such an approach, the mentioned mechanism would constitute a novel integrated risk management approach within a decentralized autonomous organization (DAO), allowing disintermediation in a complex allocation setting (cf. Zhao et al., 2022).

The remainder of this work is structured as follows. We first provide an overview of related work in the domains of prediction markets, as these are conceptually similar to the securities trading we describe and the distributed implementation of market mechanisms for privacy concerns. We then discuss related research on *stochastic-endogenous* equilibria that inspired our internal financing approach. We also discuss the general body of work on auction mechanisms for intra-organizational resource allocation. In particular, we discuss the bundle trading market, as it is the second pillar that strongly inspired our work.

After discussing related work, we proceed to describe our market model that is purposefully abstract to accommodate different organizational setups. We proceed to review coherent risk measures briefly, as these are a requirement to find stochastic-endogenous equilibria. After pointing out the central optimization of allocating resources under perfect information, we discuss bundle trading as a decomposition method for the central problem. We evaluate the proposed mechanism using extensive simulations. As sensitivities, we also include two types of friction, limited participation, and incomplete financial markets to assess their impact on allocative efficiency. We conclude with a discussion of our results and contributions, as well as managerial implications and future work.

4.2 Related Work

Our research touches on several streams of research. In particular, our model is closely related to work in the area of contingent claims, prediction markets, and securities in the operational field. Additionally, auction or market-based mechanisms for resource allocation within firms heavily inspired our work. Lastly, research on integrated optimization of physical and financial flows in supply chains and other organizations relates to our work.

4.2.1 Prediction markets and securities trading

Research on prediction markets has been vivid recently, with numerous works contributing to efficient mechanisms to clear such markets (Agrawal et al., 2011; M. K. Chen et al., 2008; Healy et al., 2010). Carvalho (2020) designs a decentralized prediction market using a Blockchain as the underlying information system (IS). Here, securities are issued using smart contracts. Although often discussed with different applications in mind, prediction markets are conceptually the same as the (Arrow-Debreu) securities market that we discuss in this work, as agents manage their risk exposure by *betting* on future outcomes. Carvalho (2020) also addresses privacy concerns as a reason to implement a decentralized prediction market within a firm. In a similar spirit, Clifton et al. (2008) discuss the use of cryptographic techniques to avoid the need for brokerage in interorganizational resource sharing. They find an incentive-compatible mechanism for doing so in a load swap application for trucking companies. However, the mechanism does not account for uncertain outcomes, and in particular does not discuss the hedging of such outcomes using financial instruments.

Ralph and Smeers (2012, 2015) discuss agents' simultaneous trading of physical and financial assets for the case of perfectly competitive and complete markets and convex or coherent risk measures, respectively. They assume central optimization by a *social planner* with perfect information, as is common practice in economic analysis under the assumption of sufficient competitiveness. Strikingly, they find that a market risk measure, or most risk-neutral agent, emerges from the population's risk measures. This market risk measures the price risk for all agents, constituting a *stochastic-endogenous equilibrium*.

4.2.2 Market-based resource allocation within organizations

The use of market mechanisms, auctions in particular, to allocate resources is a traditional economic topic. In recent decades, auctions have also received increasing attention from managerial and operational research, as technical advances made the use of auctions as practical tools possible (Bichler et al., 2010). Ba et al. (2001a) discuss the use of an incentive-compatible auction mechanism to invest in knowledge within a firm, i.e., they provide an intraorganizational use of an auction.

Combining ideas for centralized decision-making under uncertainty and deterministic decision-making in decentralized autonomous organizations, our model enables integrated risk management in such settings by extending work on the bundle trading market framework (BTM) (Fan et al., 2003; Guo et al., 2007, 2012). Our extension of the BTM enables risk-averse agents' trading in physical and financial assets, providing an iterative mechanism that results in the same stochastic-endogenous equilibrium as in the case of perfect competition presented by Ralph and Smeers (2015), but reduces information disclosure requirements. The mechanism constitutes a smart market, as it involves "interactive market design, respective decision support, and computational tool" (Bichler et al., 2010), and simultaneously an integrated risk management tool for distributed organizations.

Nadiminti et al. (2002) discuss intra-organizational resource allocation with a focus on information asymmetries and negative externalities. Especially the aspect of information asymmetry is mirrored in our work, as we do not only expect "top management" not to know about resources' benefits for managers or divisions, but rather that such top management might not even exist. Beyond this distinction, Nadiminti et al. unlike us do not account for uncertainty. There is a large body of research on deterministic market mechanisms in decentralized networked businesses, for example, carrier alliances (Agarwal & Ergun, 2010; Houghtalen et al., 2011).

4.2.3 Integrated risk management

Risk management in general is a widespread issue in supply chains (cf. Narasimhan & Talluri, 2009), particularly with recent disruptions. Beyond purely extra-organizational risk management that mostly focuses on sourcing or sales, research on the joint optimization of risk exposure and operational decisions within organizations is very active in the operational sciences. Gaur and Seshadri (2005) describe an approach to reduce the risk of short-lived item inventory using financial instruments that are correlated with inventory demand. A similar approach is presented by Kouvelis and Li (2019). However, in their case, the financial instruments are supplied exogenously, contrary to our work. Qian and Olsen (2020) provide a method for the integrated operational and financial decision-making of an agricultural cooperative. Their work mirrors an organizational structure that is similar to the one we conceive, i.e., agents with individual and common objectives. However, they only use simple financial products that are again supplied externally. Internal financing is discussed in Ning and Sobel (2018), but without the existence of the variety of securities that we propose.

4.2.4 Our contribution

To our knowledge, no existing research covers all three aspects that are distinctive for our paper: 1) a decentralized organizational structure with (potentially) divergent incentives, 2) joint operational and financial decision-making, and 3) internal financial management using self-issued securities. Using a market-based decomposition approach, we provide a mechanism that achieves optimal allocation of physical and financial resources within a distributed organization.

4.3 Market Model

Think of the decentralized autonomous organization and its agents as a (virtual) parent company in which division managers maximize local performance. To allocate its joint resources, the parent company may choose an auction; that is, it orchestrates a market mechanism in which managers trade for shared resources. Managers do not share private information (their local cost functions or constraints) with the parent company but simply acquire resources through trading. Depending on the auction's design, their trading reveals preferences for the resources. However, in our proposed mechanism, beyond just trading physical assets, managers may also trade in securities to manage their risk. Given that divisions are smaller units, and worst-case results would be extremely detrimental even if unlikely, such risk aversion is easily conceivable. With the opportunity to trade financial assets, division managers can share or reallocate the risk. Consider the following example. There are two divisions in a company, of which one would be extremely successful in a specific scenario. It could sell a security for this scenario to the other division, leading to higher average profits (owing to the earnings from selling the security) and costs in the payout scenario that are tolerable due to the success in the *physical* business. Under certain assumptions of risk preferences, this financial trading can improve overall organizational utility.

Whereas the subdivision company can serve as an easy-to-grasp example, the mechanism we propose is not limited to this case. The method describes any DAO in which individual agents are risk-averse and trade physical and financial assets, while retaining private information. Aligning the incentives of individual agents with some global goal has been a long-discussed issue in research (Ba et al., 2001b). The iterative characteristic of our proposed mechanism initially increases complexity for the participating agents, as they repeatedly have to evaluate their possible actions subject to changing prices for physical and financial assets. Given the aforementioned technological innovations, however, we argue that the increased complexity can be managed using software agency, e.g., algorithmic trading, and that the benefits of the proposed mechanism regarding privacy and strategy-proofness may outweigh concerns of computational complexity.

To allow for generalizability of the proposed mechanism, we are formulating it as generically as possible. For brevity and clarity, we restrict ourselves to cases we consider most applicable in real-world scenarios, omitting some mathematical intricacies of a more extensive discussion. When applicable, we hint at extensions that might be worthwhile for future research or could be discussed analogously to the presented method, at the expense of comprehensiveness.

The considered decentralized autonomous organization consists of J agents, each endowed with a coherent (not necessarily equal) risk measure (see §4.4). Each agent $j \in \{1, ..., J\}$ chooses an activity vector x_j, subject to local restrictions. Such activities could be the utilization of computing resources or something similar. Depending on their activities, agents have a stochastic cost. Subject to their risk measure, agents now engage in trading of physical and financial assets to minimize their risk-aware cost (or, equivalently, maximize their risk-aware profits). Note that prices for both physical and financial assets are endogenous to the agents' trading behavior. The global availability of some resources and the necessity of balanced trades of financial assets couple the agents' optimization tasks, making it a nontrivial system. The complete optimization problem of this distributed system for an omniscient central planner is given in §4.5, and the respective decomposition in §4.6. A complete overview of the notation is given in Table 4.1.

4.4 Risky Costs and Coherent Risk Measures

For brevity, we restrict ourselves to the case where organizationally relevant aspects of the future can be described by a discrete and finite set of scenarios $\{\omega\}_{\omega=1,...,\Omega}$. This restriction is not only useful for the tractability of the upcoming optimization schemes but is also proper given our underlying reasoning, as we consider trading of state-price securities. Each scenario ω is associated with a cost $z_\omega \in \mathbb{R}$, where $\mathbb{R}$ are the real numbers. The uncertain cost vector $z = (z_1, ..., z_\Omega)' \in \mathcal{Z}, \mathcal{Z} = \mathbb{R}^\Omega$ then represents the cost in all scenarios. The corresponding set of probability densities is $\mathcal{P} := \{\Pi \in \mathbb{R}^\Omega : \Pi \succcurlyeq 0, \mathbf{1}'\Pi = 1\}$. Here, $\succcurlyeq$ denotes the element-wise (in-)equality, and $\mathbf{1}$ is a vector of ones of appropriate dimension. For any probability density $\Pi \in \mathcal{P}$, the expected cost is $\mathbb{E}_\Pi[z] := \Pi' z$. As an example, for $\Omega = 5$, the risky costs for an agent j may be $z_j = (-7, -3, -5, -1, -4)'$, which means that in scenario $\omega = 1$, the agent's cost is -7 (or, equivalently, her profit is 7). Assuming the uniform probability

Variable	Domain	Explanation
J	$\mathbb{N}_+$	Number of agents
j	$\{1, ..., J\}$	Agent index
ν_j, M	$\mathbb{N}_+$	Number of local and shared resources, respectively
A_j	$\mathbb{N}_+$	Number of activities of agent j
x_j	$\mathbb{R}^{A_j}$	Activity decisions of agent j
N_j, n_j	$\mathbb{R}^{\nu_j \times A_j}, \mathbb{R}^{\nu_j}$	Local resource coefficient matrix and local resources of agent j
C_j, c_j	$\mathbb{R}^{M \times A_j}, \mathbb{R}^{M}$	Shared resource coefficient matrix and shared resources of agent j
Ω	$\mathbb{N}_+$	Number of scenarios
ω	$\{1, ..., \Omega\}$	Scenario index
Ξ_j	$\mathbb{R}^{A_j \times \Omega}$	Risky cost coefficient matrix of agent j
z, z_j	$\mathbb{R}^{\Omega}$	Risky cost (of agent j)
$r_j(\cdot)$	$\mathbb{R}^{\Omega} \to \mathbb{R}$	Coherent risk measure of agent j
s	$\mathbb{R}^{\Omega}$	Security with payout s_ω in scenario ω $(\forall \omega)$
c_{j0}, s_{j0}	$\mathbb{R}^{M}, \mathbb{R}^{\Omega}$	Initial endowment of agent j with physical and financial resources
c_0, s_0	$\mathbb{R}^{M}, \mathbb{R}^{\Omega}$	Initial endowment of dealer with physical and financial resources.
p_C	$\mathbb{R}^{M}$	(Current) prices of physical resources
p_S	$\mathbb{R}^{\Omega}$	(Current) prices of securities

Table 4.1: Notation of the market model.

of all scenarios (i.e., $\Pi_{uni} = \frac{1}{5}\mathbf{1}$), the expected cost subject to this measure is $E_{\Pi_{uni}}[z_j] = -4$.

In their seminal paper, Artzner et al. (1999) describe risk measures $r : \mathcal{Z} \to \mathbb{R}$ (given an uncertain outcome space $\mathcal{Z}$) as coherent if they exhibit "four desirable properties for risk measures". Following the cost-oriented notation of Ralph and Smeers (2015), these properties are the following:

SA **Subadditivity:** $r(z_1 + z_2) \leq r(z_1) + r(z_2) \quad \forall z_1, z_2 \in \mathcal{Z}$.

MO **Monotonicity:** $r(z_1) \leq r(z_2) \quad \forall z_1, z_2 \in \mathcal{Z} : z_1 \preccurlyeq z_2$.

TI **Translation invariance:** $r(z + \lambda \mathbf{1}) = r(z) + \lambda, \quad z \in \mathcal{Z}, \lambda \in \mathbb{R}$.

PH **Positive homogeneity:** $r(\lambda z) = \lambda r(z), \quad z \in \mathcal{Z}, \lambda \in \mathbb{R}_+$.

Fulfilling requirements 1. and 4. ensures the convexity of the risk measure. A notable example of a coherent risk measure (CRM) is conditional value-at-risk

(CVaR)[7]. CVaR at a level $100\epsilon\%$ (CVaR_ϵ) denotes the expected value of a random outcome under the condition that one of the $100\epsilon\%$ worst scenarios occurs. In the above example, CVaR at the level 40%, i.e., $\mathrm{CVaR}_{0.4}(z_j)$ (with the underlying probability distribution Π_{uni}), is -2, the conditional expectation of the two worst outcomes. Note that z_j is a cost vector, meaning that lower results are better, and negative results represent profits.

In this particular example, the two worst results are z_{j2} and z_{j4}, which means that $\mathrm{CVaR}_{0.4}(z_j) = \frac{1}{2}(z_{j2} + z_{j4})$. Obviously, this need not be the case, as other scenarios than 2 and 4 might just as well be worse. Let $\mathcal{D}^{\mathrm{CVaR}}_{\Pi_{uni},0.4} \subseteq \mathcal{P}$ be a *risk set*, defined as follows (conv$\{\cdot\}$ denoting the convex hull).

$$\mathcal{D}^{\mathrm{CVaR}}_{\Pi_{uni},0.4} = \mathrm{conv}\{\frac{1}{2}(1,1,0,0,0)', \frac{1}{2}(1,0,1,0,0)', ..., \frac{1}{2}(0,0,0,1,1)'\} \tag{4.1}$$

Then, $\mathrm{CVaR}_{0.5}(z_j)$ is defined as

$$\mathrm{CVaR}_{0.4}(z) = \max_{\Pi \in \mathcal{D}^{\mathrm{CVaR}}_{\Pi_{uni},0.4}} \mathbb{E}_\Pi[z]. \tag{4.2}$$

Consequently, the risk set for the general case of any ϵ and any underlying probability distribution Π_0 is

$$\mathcal{D}^{\mathrm{CVaR}}_{\Pi_0,\epsilon} = \{\Pi \in \mathcal{P} : \Pi \preccurlyeq \frac{1}{1-\epsilon}\Pi_0\}. \tag{4.3}$$

All CRMs can be described by a convex risk set, which means that a CRM is the expected worst-case value (i.e., maximum cost) regarding a convex set of probability distributions (Ralph & Smeers, 2015). In the remainder of this paper, we use CVaR as a running example of a CRM.

4.5 Central Optimization

Suppose that every agent j has an activity vector $x_j \in \mathbb{R}^{A_j}$. The affine risky cost function $\Xi_j : \mathbb{R}^{A_j}_+ \to \mathbb{R}^\Omega$ maps her activity level to a risky cost vector[8]. Note that agents might have different activity spaces but share a common understanding of future scenarios Ω, meaning that the world can take specific states upon

[7] Just as notable, a prominent *in*coherent risk measure is value at risk (VaR) (Artzner et al., 1999).

[8] Any deterministic cost components can be included in Ξ, as a direct result of the translation invariance of CRMs.

which agents agree. This restriction is necessary to issue tradable securities, in particular, when thinking of them as smart contracts.

The agents have local restrictions $N_j x_j \preccurlyeq n_j$ consider these to be limitations due to nontradable (physical) assets) and restrictions $C_j x_j \preccurlyeq c_j$ that limit the use of shared resources c (these are tradable physical assets of which only a finite amount exists so that $\sum_j c_j \preccurlyeq c$). Agents now evaluate their risky cost $\Xi_j(x_j)$ using a CRM r_j with an associated risk set $\mathcal{D}_j$, minimizing their exposure. They do so not only by choosing a particular activity level, but also by trading in securities $s \in \mathcal{Z}$. As an example, a security $s_a = (1, 0)'$ gives a payout of 1 unit in scenario $\omega = 1$ and none otherwise.

Assume that agent j receives an exogenous asset endowment s_j and an exogenous resource endowment c_j, compactly represented by the endowment vector. She then determines her optimal activity level by solving

$$\begin{aligned} z(s_j, c_j) \equiv \min_{x_j} \quad & r_j\left(\Xi_j(x_j) - s_j\right) \\ \text{s.t.} \quad & C_j x_j \preccurlyeq c_j,\ N_j x_j \preccurlyeq n_j,\ x_j \succcurlyeq 0. \end{aligned} \tag{4.4}$$

If there was an omniscient planner able to control each agent, that is, their endowments and activity levels, the optimization of the distributed system could be achieved centrally. This central planner's problem is

$$\begin{aligned} \min_{x_1,\dots,x_J,s_1,\dots,s_J} \quad & \sum_j r_j\left(\Xi_j(x_j) - s_j\right) \\ \text{s.t.} \quad & \sum_j C_j x_j \preccurlyeq c_0 & \text{(4.5a)} \\ & \sum_j s_j = 0 & \text{(4.5b)} \\ & x_j \succcurlyeq 0,\ N_j x_j \preccurlyeq n_j \quad \forall j. \end{aligned}$$

Equation (4.5b) denotes that the securities must be balanced between the trading agents, i.e., there is no exogenous source of such securities. In a competitive equilibrium, the dual variables (or shadow prices) to the constraints (4.5a) and (4.5b) determine the prices for physical and financial assets. This equilibrium is the benchmark for any decentralized solution approach.

Let r_0 be the CRM with the associated risk set $\mathcal{D}_0 = \cap_j D_j$. As a CRM uses the worst-case expectation regarding any probability distribution in the risk set, this means that r_0 leads to weakly lower risk aversion compared to any individual agent. As Ralph and Smeers (2015) shows, the equilibrium in (4.5) also solves the optimization problem

$$
\begin{aligned}
\min_{x_1,\ldots,x_j} \quad & r_0\left(\sum_j \Xi_j(x_j)\right) \\
\text{s.t.} \quad & \sum_j C_j x_j \preccurlyeq c_0 \\
& x_j \succcurlyeq 0,\ N_j x_j \preccurlyeq n_j \quad \forall j.
\end{aligned}
\tag{4.6}
$$

This means that the compound risky cost of all agents is evaluated with a least risk-averse CRM r_0. It also directly indicates the benefit of diversification within the organization, since risky costs are added before being evaluated with the risk measure.

Solving the optimization problem (4.5) or equivalently (4.6) requires in-depth knowledge of information that is typically regarded as private to the agents. Both local constraints $N_j x_j \preccurlyeq n_j$ and individual cost functions Ξ_j would typically not be known by a central planner, as well as resource usage C_j. Therefore, we propose an iterative mechanism for simultaneous bundle trading of physical and financial assets that allows convergence to equilibrium with reduced information disclosure requirements.

4.6 Decentralized Algorithm

The central solution to the problem (4.6) is only possible if the *central planner* has all the necessary information to do so. In particular, it needs detailed information on the local constraints and coefficient matrices of the agents. Depending on the organizational context, this information will be considered private to agents. Therefore, we assume that agents find a resource allocation by bidding for and trading in physical and financial assets with each other, each maximizing local performance. Suppose, therefore, that agents are not exogenously endowed with bundles of securities s and physical asset,c, but may trade for them at price levels p_s and p_c, respectively. Assume further that agent j currently holds securities s_{j0} and physical assets c_{j0}. To determine her optimal activity and trading levels, she then solves

$$\begin{aligned}\min_{x_j, s_j, c_j} \quad & p_c' c_j + p_s' s_j + r_j \left(\Xi_j(x_j) - s_{j0} - s_j \right) \\ \text{s.t.} \quad & C_j x_j \preccurlyeq c_{j0} + c_j, \; N_j x_j \preccurlyeq n_j, \; x_j \succcurlyeq 0. \end{aligned} \tag{4.7}$$

Two aspects are notable. First, neither s_j nor c_j have to be element-wise positive; that is, sales of resources are possible. Short-selling physical resources is implicitly prohibited by the resource constraint. Second, the trading of resources and securities might exhibit large complementarities. Purchasing physical resources to change the activity level might only be profitable if, at the same time, traded securities could offset large risks in some scenarios, given the increased activity. Trading for only resources or securities individually could, just as trading for single resources instead of bundles, cause an exposure problem for the agent (cf. Krishna, 2010).

4.6.1 Agent Bidding

As alluded to before, we impose that prices are endogenous to agents' trading behavior, and not exogenously given. At any time, current price levels are merely an informative aid to help agents bid efficiently for traded assets. Let $z(s_{j0} + s_j, c_{j0} + c_j)$ be the objective value of (4.4), which is endowed with $s_{j0} + s_j$ and $c_{j0} + c_j$, while $z(s_j, c_j)$ is the objective value without trading. The value of acquiring the bundle (s_j, c_j) is then $\bar{v}_j(s_j, c_j | s_{j0}, c_{j0}) = z(s_j, c_j) - z(s_{j0} + s_j, c_{j0} + c_j)$ (note the sign convention). If (4.7) is bounded with $\bar{s}_j$, $\bar{c}_j$ denoting optimal values for s_j and c_j, then $\bar{v}_j(\bar{s}_j, \bar{c}_j | s_{j0}, c_{j0})$ is finite, and we call the limited bundle $(\bar{s}_j, \bar{c}_j)$ the most improving bundle. Due to the convexity of z_j, $\forall \lambda \in [0, 1], \bar{v}_j(\lambda s_j, \lambda c_j) \geq \lambda \bar{v}_j(s_j, c_j)$, that is, the partial fulfillment of any desired limited bundle to the degree of λ will bring at least λ of the bundle value. Following Guo et al. (2007), we assume that agents truthfully bid $\bar{v}_j(\bar{s}_j, \bar{c}_j | s_{j0}, c_{j0})$ for limited bundles.

The solution to (4.7) may be unbounded. This situation occurs if, at given price levels, an agent wants to trade infinite amounts of securities or resources. In this case, the requested trade cannot be properly described by a (limited) bundle. As (4.7) is convex, it can only be unbounded if there exists an extreme ray $(\hat{x}_j, \hat{s}_j, \hat{c}_j)$ along which the value of the objective function is weakly decreasing. We call $(\hat{s}_j, \hat{c}_j)$ an unlimited bundle. Assuming reasonable values of the other parameters, two cases are notable:

1. The agent trades in physical assets $\hat{c}_j$) (and financial assets $\hat{s}_j$), adjusting its activity levels by $\hat{x}_j$. This situation may only occur if there are no binding local constraints along the ray.
2. The agent *speculates* with securities, i.e., $\hat{c}_j, \hat{x}_j = 0$.

Resource Trading

Starting from an initial endowment c_{j0}, s_{j0}, the agent must determine the per-unit value $\hat{v}_j(\hat{s}_j, \hat{c}_j)$ of the unlimited bundle $(\hat{s}_j, \hat{c}_j)$. Assume that the agent increases her activity levels starting from some values x_0. Given the properties of a CRM, a lower bound to $\hat{v}(\cdot)$ is described by

$$\hat{v}_j(\hat{s}_j, \hat{c}_j) = \frac{r_j(\Xi_j(x_0) - s_{j0}) - r_j\left(\Xi_j(x_0 + \alpha\hat{x}_j) - \alpha\hat{s}_j - s_{j0}\right)}{\alpha} \tag{4.8}$$

$$\overset{\text{SA}}{\geq} \frac{r_j(\Xi_j(x_0) - s_{j0}) - r_j\left(\Xi_j(x_j) - s_{j0}\right) - r_j(\Xi_j(\alpha\hat{x}_j) - \alpha\hat{s}_j)}{\alpha} \tag{4.9}$$

$$\overset{\text{lin.}}{=} \frac{-r_j(\alpha\Xi_j(\hat{x}_j) - \alpha\hat{s}_j)}{\alpha} \tag{4.10}$$

$$\overset{\text{PH}}{=} -r_j(\Xi_j(\hat{x}_j) - \hat{s}_j), \tag{4.11}$$

where α is the amount of the unlimited bundle. For large quantities of α, that is, large quantities of the requested bundle, the impact of the initial activity levels and the security endowment vanishes. Then, $\hat{v}_j$ is a tight bound on the value of the unlimited bundle. For smaller quantities, the actual value might actually be higher. We assume that the agents bid $\hat{v}_j(\hat{s}_j, \hat{c}_j)$ for unlimited bundles $(\hat{s}_j, \hat{c}_j)$, erring on the side of caution.

Security Speculation

An agent arbitrages securities if she is certain that scenarios occur differently than priced by the market. Assume the agent is risk-neutral, using expectation as CRM. Her risk set is then a singleton containing just Π. Let $\Pi = (1, 0)'$, i.e., the agent is certain that scenario $\omega = 1$ occurs, whereas scenario $\omega = 2$ does not. If the current market price for the security $s_1 = (1, 0)$, i.e., a payout in scenario $\omega = 1$, is below 1, then purchasing any positive amount of such a security would increase the agent's utility.

Again, due to the sub-additivity property of CRM, the issue of additional securities $\hat{s}_j$ will decrease the total risk by at least the risk-aware value of these securities, $-r(-\hat{s}_j)$. This is especially true when agents seek to purchase an unlimited number of securities, as the impact of initial endowment then vanishes. The truthful per-unit valuation of an unlimited security bundle is thus

$$\hat{v}_j(\hat{s}_j) \geq -r(-\hat{s}_j), \tag{4.12}$$

where the right-hand side might be an underestimator for smaller amounts.[9] To revisit the example from above, $\hat{v}_j(\hat{s}_j) = -r_j(-\hat{s}_j) = -\Pi'(-s_1) = -(1,0)'(-(1,0)) = 1$.

Bid Submission

We assume truthful bidding of the agents, following the discussion of Guo et al. (2007). Given the additional complexity due to trading in both resources and securities, as well as the potential nonlinearity of the cost function, we deem this a reasonable assumption. Agents therefore submit their (un)limited bundle orders as quadruples, either $l_j^k = \left(\bar{s}_j^k, \bar{c}_j^k, \bar{v}_j^k, 1\right)$ or $u_j^h = \left(\hat{s}_j^h, \hat{c}_j^h, \hat{v}_j^h, \infty\right)$, k and h indexing individual orders. The market operator maintains two order books per agent j, an order book $\mathcal{L}_j = \{l_j^k\}_k$ for limited orders and an order book $\mathcal{U}_j = \{u_j^h\}_h$ for unlimited orders.

[9] The equation follows directly from (4.11) for $\hat{x}_j = 0$.

4.6.2 Market Clearing

The market program's task is to maximize the gains from trade.[10] It does so by solving the market matching problem (MMP)

$$\max_{\sigma_j^k, \rho_j^k} \quad \sum_j \left(\sum_{l_j^k \in \mathcal{L}_j} \sigma_j^k \bar{v}_j^k + \sum_{u_j^h \in \mathcal{U}_j} \rho_j^h \hat{v}_j^h \right) \tag{4.13}$$

$$\text{s.t.} \quad \sum_j \left(\sum_{l_j^k \in \mathcal{L}_j} \sigma_j^k \bar{c}_j^k + \sum_{u_j^h \in \mathcal{U}_j} \rho_j^h \hat{c}_j^h \right) \preccurlyeq c_0 \tag{4.14}$$

$$\sum_j \left(\sum_{l_j^k \in \mathcal{L}_j} \sigma_j^k \bar{s}_j^k + \sum_{u_j^h \in \mathcal{U}_j} \rho_j^h \hat{s}_j^h \right) \preccurlyeq s_0 \tag{4.15}$$

$$\sigma_j^k, \rho_j^h \geq 0, \quad \sum_{l_j^k \in \mathcal{L}_j} \sigma_j^k \leq 1, \ \forall j. \tag{4.16}$$

The formulation of the MMP largely follows the aforementioned bundle trading market, with the addition of the constraint (4.15) that accounts for the trading of securities. It requires a weak supply-demand balance between traded securities. The reason to not use a strict balance here is the increased liquidity when the market maker may (temporarily) hold securities. Allowing this flexibility avoids premature halting of the market mechanism that might otherwise occur. Note that, upon convergence of the mechanism, the securities inventory of the market maker is zero, unless there are states that have zero probability under the risk-neutral measure. In that case, the market maker may hold securities after convergence of the mechanism; however, they are priced at zero by the market.

Initializing the market with $s_0 = 0$ captures that there is no external source of securities, they are solely issued by market participants. Because of this, agents are just trading risk among peers. The prices of the resources are now calculated based on the dual variables of Equations (4.14) and (4.15), respectively.

The market now iterates as follows. During initialization, agents are endowed with resources, and prices for financial and physical assets are set. We do not focus on initialization strategies in this work; see Guo et al. (2012) for an extensive discussion of these. In our simulations, all agents are endowed with no resources

[10] The market logic itself, that is, the auction protocol, may be implemented using smart contracts as part of the decentralized autonomous organization. Therefore, the disintermediation is maintained.

at all and the prices of the securities are set uniformly to $p_s = (\frac{1}{\Omega})_\omega$. Then, agents determine their (un)limited bundle orders, and trades might occur, changing the price levels. Order books of all agents involved in trades are cleared, and endowments are updated. When no trades are executed and prices do not change anymore, the market operation ceases.

4.7 Evaluation

For the evaluation of the proposed market mechanism, two aspects are paramount. First, the computational aspects, especially the necessary number of iterations to reach equilibrium, are important. This is particularly true in light of the potential application of the proposed mechanism in a DAO setting, providing an operational tool to allocate resources within the DAO. Second, the added benefit of trading securities beyond physical assets must be evaluated to infer if the added complexity is worthwhile. To evaluate the market mechanism regarding these questions, we run three experiments:

Vanilla Mechanism The proposed mechanism with full participation and complete markets

Limited Participation (LP) Only some agents trade in securities.

Incomplete Market (IM) Not all future states can be traded as securities.

Whereas the *vanilla mechanism* serves as a benchmark for the alternative cases, the LP and IM cases represent frictions in the market.[11]

We investigate these settings by performing extensive simulations. To do so, we randomly create environments with different numbers of agents, resources, and state space sizes. To limit the overall number of necessary simulations, we impose an important restriction. That is, all agents share a common physical measure Π_0 and use CVaR as CRM. The shared understanding of the underlying probability of future scenarios ensures that there is an equilibrium in financial transactions. Recall that when there is no intersection of the risk sets of the agents, there cannot be an equilibrium in the market (Ralph & Smeers, 2015). Using a common physical measure and CVaR ensures that the intersection of the risk sets at least contains the physical measure Π_0, CVaR as CRM for all agents is also particularly interesting, as it is widely used in practice.

[11]While we report speed of convergence, acceleration of the market mechanism in itself is not the focus of this work.

4.7.1 Vanilla Mechanism

Before the detailed experiments on specific modifications are evaluated, the proposed mechanism is tested as is under various parameter setups; see Table 4.2.

Parameter	Explanation	Values
J	No. of agents	$\{10, 20, 50\}$
A	No. of activities	$\{2, 5, 10, 20\}$
M	No. of shared resources	$\{2, 5, 10, 20\}$
ν	No. of local resources	$\{2, 5, 10, 20\}$
Ω	No. of scenarios	$\{2, 5, 10, 20\}$

Table 4.2: Parameter variation in testing of the *vanilla mechanism*. All 768 combinations are tested 50 times, totaling 38,400 simulations.

The dependence of the number of rounds required in the auction market on the input parameters is given in Table 4.3a. The most striking impact on the required rounds for market clearing is due to the number of scenarios, Ω. Table 4.3a shows that for all parameters except Ω, the average number of rounds is around $1,000$, with high standard deviations. For Ω, a much clearer relationship is shown. For a smaller number of traded securities, the number of auction rounds decreases drastically. As the relationship between the parameters, particularly Ω, and the required number of rounds seems nonlinear, we analyzed the convergence using a log-linear model; see Table 4.3b. With good model fit, it can be seen that convergence times increase approximately quadratically in the number of traded securities and sub-linearly in all other parameters. The different impact of the number of physical and financial assets is particularly striking, as the impact of M is much less substantial than that of Ω.

4.7.2 Limited Participation

Due to the complexity of financial asset trading, it seems conceivable that not all agents would participate in securities trading. To evaluate the performance of the proposed mechanism in such a setting, we investigate another simulation setting. As the effect of large values for Ω is detrimental to convergence and the choice of the remaining parameters does not severely impact convergence, we fix the parameters for this setting. In particular, J is set to 20; all other parameters

J	10	1165.6545 (1766.8576)	ν	2	1041.9598 (1651.3948)
	20	1052.6808 (1663.8291)		5	1061.2382 (1674.1345)
	50	878.6697 (1481.3815)		10	1028.0426 (1640.71)
				20	998.0994 (1616.049
A	2	912.341 (1522.428)	Ω	2	67.316 (414.6372)
	5	997.4838 (1613.8605)		5	114.3293 (287.4245)
	10	1083.8653 (1677.5273)		10	491.5486 (519.2815)
	20	1135.6499 (1752.175)		20	3456.146 (1539.2108)
M	2	917.7161 (1561.0706)			
	5	1069.1688 (1687.3862)			
	10	1084.3358 (1680.4289)			
	20	1058.1193 (1646.0782)			

(a) Number of rounds subject to parameter variation. The average number of rounds and standard deviation of number of rounds (given in parentheses) subject to parameter variation. Note that simulations were cut off at $5,000$ rounds, i.e., the values for $\Omega = 20$ are inconclusive.

	log(Rounds)
(Intercept)	1.431*** (0.029)
$\log(J)$	-0.184*** (0.007)
$\log(A)$	0.137*** (0.005)
$\log(M)$	0.159*** (0.005)
$\log(\nu)$	-0.017** (0.005)
$\log(\Omega)$	2.061*** (0.005)
Estimator	OLS
N	38,400
R^2	0.803

(b) Log-linear regression model of auction rounds. OLS regression results of a log-linear model of the number of auction rounds. The number of rounds grows sublinearly in all parameters except Ω. Ceteris paribus, the number of rounds has a quadratic relationship to Ω. The goodness of fit and the significance of the coefficients are high. This indicates a good scalability in all parameters but Ω.

Table 4.3: Simulation analysis for a variation of the respective parameters.

are set to 10. Then, the number of agents participating in securities trading is varied between 2 and 19, i.e., from only bilateral trading of securities to almost full participation. Each experiment is repeated 100 times. Figures 4.2i and 4.2ii show the results of the simulation.

The auction runs faster when only very few agents participate; this effect is negligible with even a moderate number of participants. The effect concerning the performance, compared to full participation, of the mechanism is substantially more pronounced. The relative performance starts around 70% for only two

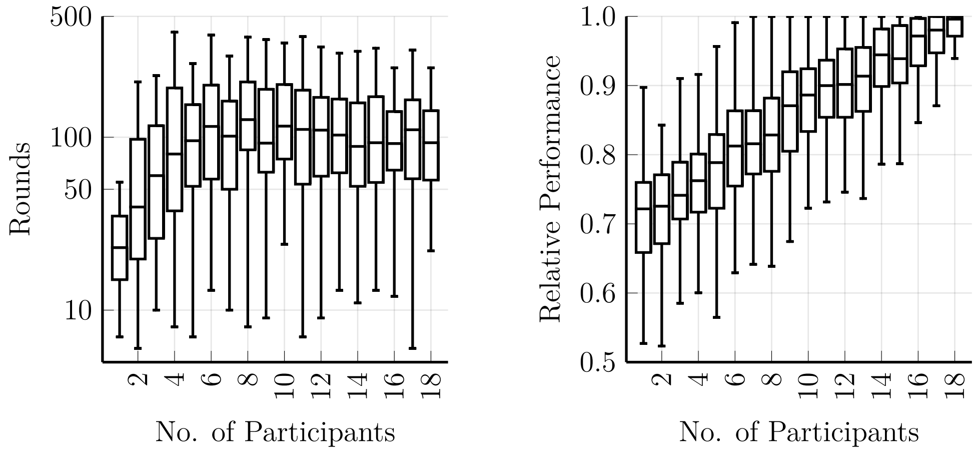

(i) Auction rounds for LP case. Number of auction rounds necessary to clear the market for a varying number of agents participating in securities trading (out of 20 agents). The number of auction rounds is almost constant for a wide range of participant numbers.

(ii) Relative performance of LP case. Relative allocative efficiency of the auction for a varying number of agents participating in securities trading (out of 20 agents overall). The relative performance of the mechanism increases approximately linearly in the number of participants.

Figure 4.2: Computational analysis of the limited participation case.

participants. Note, however, that it would not be zero even with no participation in financial trading, as physical trading alone yields some welfare benefits. The performance then grows approximately linearly with the number of participants, with slightly smaller added benefit for the very last participants. Agents not participating in financial trading are chosen ex ante, i.e., disregarding their individual benefit of trading in securities. Because of this, the addition of even single agents can benefit the system greatly if they have a favorable risk assessment and can thus provide liquidity to the market.

4.7.3 Incomplete Securities Market

We next discuss the case where not all future states are tradable as securities. This evaluation is of paramount importance, as the convergence of the mechanism

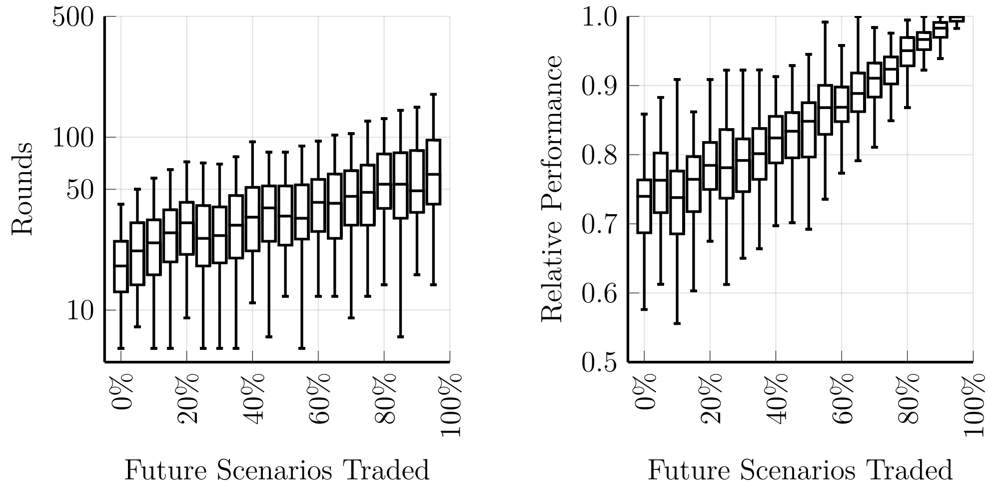

(i) Number of auction rounds necessary to clear the market with a varying share of future outcomes traded as securities. The number of rounds increases sharply with an increased share of traded securities

(ii) Relative allocative efficiency of the auction for a varying share of future outcomes traded as securities. Efficiency increases continuously with the share of traded securities; good performance in terms of allocative efficiency is already obtained at moderate shares.

Figure 4.3: Computational analysis of the incomplete markets case.

deteriorates quickly with the number of traded financial assets; see §4.7.1. Unlike the case of limited participation, §4.7.2, here we are not selecting tradable states randomly. Instead, we select them to match a predetermined total probability of all tradable states. Thus, we account for the fact that, in cases of incomplete markets, the most probable states would be those traded at first, while markets for more obscure outcomes would be thinner. Figures 4.3i and 4.3ii show the results for the incomplete securities markets.

The number of auction rounds necessary to clear the market increases significantly in the share of scenarios traded; note the logarithmic scale. This behavior mirrors the analysis of the *vanilla* mechanism, in which the absolute number of securities was the main driver of the auction runtime. The increase in relative performance of the mechanism is again approximately linear. Based on these results, it is conceivable that the proposed mechanism could not only be used if all states of the future can be conclusively mapped to securities, but also if only some percentage of potential outcomes should be traded. By doing so, runtime is reduced significantly, whereas good performance results are still obtained.

4.8 Discussion

We extend the bundle trading market framework to account for risk, developing an enhanced auction mechanism that allows market-based optimization of decentralized autonomous organizations, even under uncertainty. By doing so, we add an important dimension to the self-optimization capabilities of DAOs, given that risk management is paramount, e.g., in supply chains. In our model, individual agents manage their risk, evaluated using coherent measures, by simultaneously trading in markets for physical assets and securities, accounting for all potential complementarities. As there is no external source of securities, agents have to issue them themselves. Given the continuous developments in the area of financial technologies and comparative ease in setting up online markets, this is progressively conceivable as a viable solution for real-world DAOs.

4.8.1 Contribution

Auction or market-based mechanisms for intra-organizational resource allocation are typically limited to the deterministic case. While there exist approaches for distributed decision-making for integrated risk management, none known to the authors consider self-issued securities, i.e., risk management without externally supplied financial instruments. We significantly extend existing work on the bundle trading market framework by introducing uncertainty as an additional dimension of complexity. We provide an iterative auction mechanism in which agents simultaneously trade physical and financial assets, converging to an efficient stochastic endogenous equilibrium.

The decentralized mechanism requires much less information disclosure from participating agents, as they only have to reveal their valuation of bundles at current market prices. Given the pricing of resources, they have strong incentives to bid truthfully, thus resulting in an overall favorable allocation. We further consciously limit the role of the market program, as strict allocation and pricing rules can be easier implemented as software artifacts themselves. Therefore, the proposed auction mechanism can be automated and allow deep decentralization of the organizational structure.

4.8.2 Managerial Implications

In many areas of industry, decentralization of decision-making is prevalent. In the sharing economy, resources are not owned by a central party, but by many smaller agents. In many supply chains, different firms along the value chain have some common, but also pronounced, individual interests. Finding mechanisms to coordinate resource allocation is crucial to the efficient operation of such organizations. The introduction of the risk dimension in our work allows for a more granular risk sharing between firms or agents. Through these possibilities, some players who otherwise might not have been able to participate in these value chains due to their risk aversion can overcome this entry barrier.

As an example, consider an energy cooperative. The cooperative owns a shared asset, such as battery storage, that may be used by the agents. In addition to using the battery, agents can purchase additional energy on an external market for an uncertain price. Depending on the usage of the battery, their flexibility, and other internal constraints, they have thus an uncertain cost vector. By jointly trading the battery usage rights and securities that are used to hedge energy procurement costs, the agents can manage their risk and simultaneously use the cooperative's battery optimally. As the cooperative issues the securities internally, agents are not limited to trading standard financial products that are commonly available but might not suit its needs (e.g., because of participation constraints or unit sizes) but can design them freely.

The proposed mechanism allows for a high degree of automation, as the market program's activities are limited. This is in line with the emergence of distributed ledger technologies. There has been an increased interest in finding suitable applications for the technology, both in academia and in practice, as the tamperproof ledger and automation opportunities through smart contracts help to encode business logic in software. The proposed mechanism interacts nicely with these approaches, and we can easily envision the auction being run using smart contracts for auction clearing and security issuance.

4.8.3 Future Work

The simulation results indicate that convergence speed becomes an issue primarily when the number of traded securities is high. However, even when only a subset of all possible future scenarios is traded as securities, the relative performance of the mechanism is solid, and the auction is significantly accelerated. The mechanism therefore seems particularly applicable in cases where relevant aspects of the future can be described by a relatively limited set of scenarios.

As the availability of a mechanism for joint physical and financial optimization in distributed systems might also greatly benefit settings with a large number of future scenarios, speeding up the auction seems to be a very fruitful path for future research. In general, two approaches are worthwhile in doing so. First, the operation of the market itself could be increased, similar to the approaches presented by Guo et al. (2012). However, this would increase the complexity in envisioning the mechanism as fully decentralized, i.e., the market program being a software artifact. Thus, it might take away some merits presented concerning the developments in financial technologies, online markets, and distributed systems as a whole. The second approach is to analyze the incentive structures in more efficient decomposition algorithms and to analyze whether they can be made fully incentive compatible, without requiring agents to be intrinsically motivated to help convergence (cf. Boyd, 2010).

Lastly, all reported evaluations are based on randomly created distributed optimization problems to enable a wide variety of settings. In reality, problems could be assumed to have much more pronounced sparsity patterns, that is, agents being only interested in subsets of physical and financial assets. Designing hierarchical sub-auctions for that case to help market clearing could provide a fruitful avenue for future research.

5 Spatial Competition in Fast-Charging Networks[12]

5.1 Introduction

Many countries have enacted legislation to promote the shift from internal combustion engine vehicles (ICEVs) to electric vehicles (EVs) in response to climate change. Despite the steady increase in EV adoption rates, lack of public charging infrastructure and long charging time remain to be major hurdles to the transition to electric mobility (AlixPartners, 2019). A necessary path to overcoming these hurdles is to increase the accessibility of public fast-charging infrastructure (McKinsey, 2021): Unlike the most common chargers that are used at home, workplace, or public parking lot, fast charging allows EVs to be fully recharged in 20-40 minutes rather than hours, which not only provides greater convenience and flexibility (BCG, 2023) but also makes EV ownership attainable for those without access to charging at home[13] or workplace.

While EV fast-charging and traditional gas stations both provide essential infrastructure for motorists, there are critical differences that pose unique challenges for EV charging networks. First, unlike typical gas stations where vehicles can fully refuel in 2-3 minutes, fast-charging stations require significantly longer service times, creating risks of more frequent queuing and higher delay costs for consumers. Second, fast-charging equipment and its installation requires greater capital investment: Excluding grid upgrading and integration, which can cost up to millions, the hardware and installation costs for a fast-charging unit range between \$85,000 and \$250,000 (McKinsey, 2023). For any charging point operator (CPO), without government subsidies, both the upfront capital costs and operational expenses have to be covered by the operating income. Without a high utilization rate, the operating income can hardly achieve a break-even point or a positive net profit (H. Lee & Clark, 2018). This motivates operators

[12] Joint work with Yixin Lu, Long He, and Wolfgang Ketter.

[13] In the United States, the share of home charging is expected to fall to 50 percent by 2030 (McKinsey, 2023).

to target lucrative areas and implement dynamic pricing mechanisms (e.g., peak or congestion-based pricing) that help improve station utilization and recoup investments. These differences compound the existing hurdles in creating equitable access to charging, which is critical to ensure broad EV adoption.

To boost investment in fast-charging infrastructure, policymakers around the world have created various subsidy programs. Take the United States, for example. The Bipartisan Infrastructure Law has dedicated $5 billion through the NEVI program to building a robust, publicly accessible fast-charging network (The White House, 2023). In the rest of the world, governments have also accelerated the rollout of fast-charging infrastructure. Specifically, the federal government of Germany launched a public tender scheme to invest €2 billion for the construction of a national fast-charging network with 900 stations, each of which is equipped with 4 to 16 charging points (German National Centre for Charging Infrastructure, 2023).

Despite the consensus on the importance and urgency of investing in fast-charging infrastructure, there is a lack of understanding about the underlying economics of the emerging fast-charging market. From the government's perspective, it is important to ensure a wide dispersion of investments, such that fast-charging stations materialize not only in densely populated metropolitan and economically striving regions but also in rural and disadvantaged regions. From private investors' perspective, government subsidies, especially those awarded through public tenders, can be a double-edged sword: while they can significantly reduce initial capital requirements and financial risk, they often come with various restrictions, among which, geolocation requirements (German National Centre for Charging Infrastructure, 2023) and price caps (LaMonaca & Ryan, 2022) are the most common ones. These restrictions may complicate CPOs' operation decisions (e.g., pricing), which in turn affect their profitability and consumer welfare. In particular, given EVs' limited range and the travel costs associated with driving detours to charging stations, the competition in the fast-charging market is inherently regional: Any CPO's operating income will depend on its ability to capture consumers in its immediate vicinity. A case in point is Tesla's recent strategic move: the forefront runner in the fast-charging market decided to open its charging stations to competitors' vehicles to increase the utilization rate of its superchargers (Wired, 2024). This move, however, raises concerns that the economies of scale enjoyed by large enterprises like Tesla may lead to regional monopolies (The New York Times, 2023).

Motivated by these practical challenges, in this paper, we study the competitive dynamics in fast-charging networks. In particular, we focus on the following research questions: *How do spatial and ownership structures of charging points affect CPOs' pricing strategies and consumer welfare? Whether and under what circumstances can CPOs' competition lead to regional monopoly?*

To address these questions, we adopt the sociotechnical lens (Sarker et al., 2019) and propose a novel computational framework to capture the complex interplay of strategic and informational factors in a fast-charging network. At the core of our framework is a spatial competition model where (1) consumers are price-takers and choose between different charging points based on the total cost incurred in travel, waiting, and charging, together with their locational attractiveness, and (2) CPOs within the charging network would consider consumers' choices when making their pricing decisions. Since there are multiple CPOs in the charging network, each CPO also needs to consider its competitors' pricing strategies. This may lead to the curse of dimensionality, making CPOs' optimal pricing decisions in real-world charging networks intractable. To tackle this challenge, we develop a computationally efficient method to optimize CPOs' pricing decisions.

We instantiate our framework in the recently inaugurated national fast-charging network in Germany (German National Centre for Charging Infrastructure, 2023). Here, a public tender (i.e., a procurement auction) has been held to solicit investments in fast-charging stations in predefined regions. Using demand projections for 2025 and 2030, we conduct a series of counterfactual experiments to examine the welfare implications of different competitive scenarios that are jointly shaped by the tender design and CPOs' decisions. Our analysis leads to four sets of findings. First, price exhibits strong negative correlation with local competitiveness (measured by the density of charging stations), but shows no clear dependence on absolute demand levels. Second, CPOs' profits are shaped by both local competitiveness and absolute demand. In particular, some high-demand regions yield strong profitability albeit lower prices because utilization rates are high, whereas some low-demand regions are less profitable albeit higher prices. There are a few high-demand regions with low projected charging station densities, and hence yield the highest profits. Interestingly, the same holds true for regions with low demand and very low projected charging station densities. Third, a uniform, stringent price cap may create inefficiencies in network utilization and discourage infrastructure investment, particularly in less lucrative regions. Finally, spatial and ownership constraints in tender rules may fail to ensure competitive outcomes in large-scale charging networks, allowing CPOs to strategically select locations that strengthen their market power ex-post. This can harm consumer welfare by limiting charger availability and inflating prices. Taken together, these findings highlight the necessity of regulatory interventions to ensure regionally equitable access to fast charging while ensuring efficient network utilization.

Our paper contributes to the growing literature on smart sustainable mobility (Ketter et al., 2023) by identifying and investigating the key economic forces underlying the competitive landscape of the fast-charging market. Despite their instrumental role in facilitating the transition to electric mobility, the research on design and operation of EV charging networks is still in its infancy. To the best of our knowledge, this is the first paper that studies welfare implications of spatial and ownership structures in fast-charging networks. Our paper provides timely insights for practitioners and policymakers. For practitioners, our computational framework can serve as an effective decision support tool: Regardless of funding schemes, understanding how the addition of a new charging station affects the competitive dynamics of a fast-charging network—especially the prices and profitability of the incumbent and new charging stations—is crucial for CPOs' investment decisions. For policymakers looking to boost investments in fast-charging infrastructure, our findings provide the much-needed guidance. In particular, our case study of Germany's national fast-charging network offers guidance for policymakers in other countries that are seeking to solicit investment in fast-charging infrastructure through public tenders.

5.2 Literature Review

The EV charging dilemma is reminiscent of the classic chicken-and-egg problem, which is known for being a major obstacle to large-scale adoption of infrastructure-dependent technologies (Constantinides et al., 2018; G. Ma et al., 2019): On the one hand, consumers are hesitant to purchase EVs without convenient access to charging infrastructure. On the other hand, firms are reluctant to invest in the infrastructure when there are too few adopters to make it profitable. Thus, to understand the emerging market for fast charging, its network characteristics, especially the spatial and ownership structures, are paramount. Successful participation in the market requires sufficient penetration and associated platforms to build profitability (Bhargava et al., 2021).

To address the chicken-and-egg dilemma in charging infrastructure investment, previous studies have explored various policy interventions (e.g., Shi et al., 2022; Yu et al., 2022; W. Zhang & Dou, 2022). Additionally, there is a substantial body of literature focused on optimizing strategic and operational decisions for charging services, including location planning for charging networks (e.g., Avci et al., 2015; Guo et al., 2016; Lamontagne et al., 2023; Mak et al., 2013), dynamic pricing and scheduling for demand management (e.g., L. Chen et al., 2024; O. Q. Wu et al., 2022). Few papers have studied the spatial competition in the charging market. One exception is Yuan et al. (2015), where the authors analyzed the strategic interactions between two charging stations and consumers through a Stackelberg game and characterized the equilibrium conditions. Our paper contributes to this stream of work by introducing a spatial competition model that incorporates various complexities of real-world EV charging markets, such as price caps, locational heterogeneity of demand, and the ownership of multiple charging stations by the same CPO.

The spatial competition dynamics of EV charging networks share similarities with traditional gas station networks, where firms compete over location and pricing (Chan et al., 2007; Houde, 2012; Kim & Park, 2024). However, a key distinction lies in the substantially higher delay cost associated with EV charging relative to gas refueling, which alters competitive strategies (Y. Huang & Kockelman, 2020; D. Yang et al., 2021). Unlike gas refueling, which is largely uniform across stations, EV charging varies significantly in duration and availability (Hardman et al., 2018). Moreover, EV users often have access to and utilize real-time charging station availability and delay cost information before deciding (Asensio et al., 2025), introducing an anticipatory element that is essential in station choice and thus the analysis of EV charging. These factors necessitate theoretical and practical insights that specifically account for dynamic pricing, congestion effects, and locational differentiation in ways that extend beyond prior research situated on gas station networks.

In the transportation sector, technological advances (e.g., smart cards, smartphones, and the Internet of Things) unlock the potential to characterize urban mobility patterns and optimize infrastructure and services accordingly (e.g., Munizaga & Palma, 2012; Schroer et al., 2022; Y. Wang et al., 2022). Furthermore, the urgent need to meet environmental targets, such as reducing greenhouse gas emissions and achieving net-zero goals, has motivated researchers to investigate various sustainable mobility solutions. In particular, a set of papers explored design and operational issues of shared mobility systems (e.g., He et al., 2020, 2021; Kabra et al., 2020; Kahlen et al., 2024), as well as the social impact of these systems (e.g., Babar & Burtch, 2020; Y. Zhang et al., 2023). Our paper complements the existing literature on sustainable (urban) transportation by examining the market dynamics and their impact on the development and deployment of EV charging infrastructure. In this regard, our paper is also related to Crönert and Minner (2021), although the two papers have different focuses: Crönert and Minner (2021) focused on competitive facility location optimization in the context of hydrogen refueling stations, whereas we focus on spatial competition between CPOs in the context of EV charging.

Our paper is also related to the growing stream of IS literature that studies the necessity and effectiveness of regulations in emerging markets (e.g., W. Chen et al., 2023; Li & Wang, 2025). For example, W. Chen et al. (2023) studied the effects of limiting the number of properties a host can manage in the peer-to-peer home-sharing market. In a similar vein, Li and Wang (2025) analyzed the impact of a commission fee cap for on-demand delivery platforms on revenue generation of independent and chain restaurants. These papers studied the regulatory impacts on players with different market powers. We, on the other hand, study welfare implications of market structures that are subject to regulations. One paper that also considers the impact of market structures on the strategic interactions of different players is Zhu et al. (2021). However, the two papers have different focuses: Zhu et al. (2021) focused on the impact of network connectivity on competition between incumbent and new entrant platforms, whereas we focus on spatial competition of service providers in the emerging EV charging market, in which competitors jointly enter.

Finally, our study is also related to the literature on product differentiation and competition. Since the seminal work of Hotelling (1929), researchers have extended the canonical model (*à la Hotelling*) to study competitions in various markets. We refer readers to Drezner and Eiselt (2024) for a comprehensive review. One paper in this stream that is closest to ours is Bimpikis et al. (2019), where the authors studied competition between firms that offer a homogeneous product in networked markets, which in general also captures the essence of EV charging. The major difference between Bimpikis et al. (2019) and our paper is that they focused on a *Cournot*-style quantity competition, whereas we are interested in a *Bertrand*-style price competition subject to government regulations. By incorporating elements such as anticipatory delay costs and regulatory constraints, our study aims to provide a more nuanced understanding of spatial competition in EV charging markets. This enables policymakers to better understand the interplay between spatial competition and network characteristics in their consideration, particularly against the backdrop of large-scale service networks offered by incumbents.

5.3 Model Description

Consider a set of stations in a fast-charging network. Each CPO may operate several charging stations and set different prices. Consumers can, in principle, choose to charge their vehicles at any charging station within the network. We consider a specific share of consumption, i.e., regional fast-charging demand. This demand typically originates from consumers that a) lack access to private charging infrastructure (e.g., residents of existing apartment buildings, particularly in densely populated areas) and b) do not charge en route to a destination. Given the relatively low value of a single charging transaction, consumers are typically unwilling to incur the cost to travel to a distant charging location (Hardman et al., 2018). As such, we assume that consumers would only consider charging stations in their direct vicinity (e.g., close to home or workplace).

We capture the interactions between consumers and CPOs in the network through a graph $G(V, E)$ with the set of vertices V and the set of edges E. Each vertex $v \in V$ denotes a charging station with $\kappa_v \geq 1$ charging points and a charging price p_v^t at time t; each edge $e \in E$ represents a continuum of consumers with a demand density f_e^t at time t. Figure 5.1 provides an illustrative example of a charging network. Here, vertices 1 to 6 represent charging stations. In this example, CPO 1 controls vertices 1, 4, and 6; CPO 2 controls vertices 2 and 5; and CPO 3 controls vertex 3. The edges represent consumers, i.e., the edge between vertices 3 and 4 are consumers who may choose between those vertices to fulfill their charging demand. The graph representation captures the spatial competition between charging stations, e.g., vertex 4 is in direct competition with vertices 2, 3, and 5.

We assume consumers act as price-takers and do not strategically account for their impact on the CPOs' pricing decisions. Such a price-taking assumption is in line with the intuition of consumer decision-making in similar situations, e.g., refilling at gas stations, where drivers do not ponder whether their decisions to (not) refill impact the prices. In our case, there is a multitude of (potential) consumers distributed along the edges of the graph. As such, demand can be

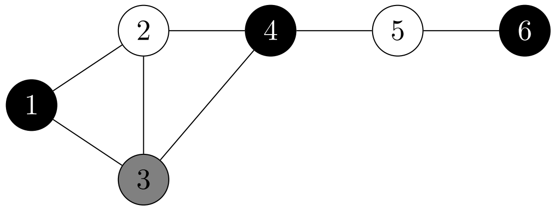

Figure 5.1: A simple EV charging network. Vertices 1, 4, and 6 (black) are controlled by CPO 1, vertices 2 and 5 (white) by CPO 2, and vertex 3 (gray) by CPO 3.

viewed as a continuum along the graph's edges by way of simplification, rendering each consumer infinitesimal regarding its impact on the total demand on any edge. In that sense, the interaction between a CPO and consumers resembles a leader-follower game (i. a. Fudenberg & Tirole, 1991). Since there are more than one CPO competing in the network, and price is the sole decision variable for CPOs (i.e., CPOs compete *à la Bertrand*), combining the aforementioned characteristics, the competition between CPOs constitutes a sequential Bertrand game in a networked environment.

In particular, given competitors' pricing, an individual CPO's decision problem is a bi-level optimization problem (i.e., the CPO's pricing decision in the upper level and consumers' responses in the lower level); it can be formulated as a mathematical program with equilibrium constraints (MPEC) (Bialas & Karwan, 1984; Colson et al., 2007): the CPO optimizes its profit, considering that consumers respond optimally to the prices charged by other CPOs. When CPOs account for both their consumers' responses and their competitors' pricing, the competition is in equilibrium, i.e., the game constituted by the MPECs of all CPOs is an equilibrium program with equilibrium constraints (EPEC) (Gabriel et al., 2013). Here, in addition to the equilibrium of a single CPO's decision and the affected consumers, we have to consider the equilibrium between CPOs.

To establish the sequential Bertrand game formally, we first introduce the CPO model in Section 5.3.1. For easier reference, we provide an overview of the key notations in Table 5.1.

Symbol	Description
	Network data and CPOs
$i \in N$	Set of CPOs, indexed by i
$t \in T$	Set of periods, indexed by t
$G(V,E)$	Graph with edges E and vertices V
$v \in V$	Vertex v in the set of vertices V
$e \in E$	Edge e in the set of edges E
f_e^t	Demand density on edge e at time t
L_e	Length of edge e
c_v^t	(Marginal) Cost at vertex v at time t
κ_v	Number of charging points at vertex v
a_v	Attractiveness of vertex v
Δa_e	Attractiveness difference $a_{\mathrm{src}(e)} - a_{\mathrm{dst}(e)}$ along edge e
$\mathrm{src}(e)$	Source vertex of edge e
$\mathrm{dst}(e)$	Destination vertex of edge e
$\mathcal{V}_i$	Set of vertices controlled by CPO i
$\mathcal{E}_v^- = \{e \in E : \mathrm{src}(e) = v\}$	Set of edges with v as source vertex
$\mathcal{E}_v^+ = \{e \in E : \mathrm{dst}(e) = v\}$	Set of edges with v as destination vertex
$\mathcal{E}_v = \mathcal{E}_v^- \cup \mathcal{E}_v^+$	Set of edges with v as source or destination vertex
N_v	Set of neighbors of vertex v
	Scalar parameters
τ	Travel cost parameter
δ	Delay cost parameter
q	Charging amount
	Competitive decision variables
x_e^t	Indifferent consumer on edge e at time t
p_v^t	Price at vertex v at time t
$\mathbf{p}^t = (p_v^t)_{v \in V}$	Vector of all prices at time t
$\mathbf{p}_i^t = (p_v^t)_{v \in \mathcal{V}_i}$	Vector of all prices at vertices controlled by CPO i at time t
Δp_e^t	Price difference $p_{\mathrm{src}(e)}^t - p_{\mathrm{dst}(e)}^t$ along edge e at time t
$D_{v,e}^t(\Delta p_e^t)$	Demand of edge e at vertex v at time t
$D_v^t(p^t) = \sum_{e \in \mathcal{E}_v} D_{v,e}^t(\Delta p_e^t)$	Demand at vertex v at time t
	General notation
$*$	Optimal choice
$\sim$	Augmented term accounting for delay cost
$'$	Approximate term used in the case study presented in section 5.4

Table 5.1: Overview of notation.

5.3.1 CPO Pricing Problem

Each CPO $i \in N$ is assumed to be risk-neutral and maximizes its profit Π_i. We first consider the cost structure of CPOs. In general, there are two types of costs: fixed costs and variable (or marginal) costs. Fixed costs (e.g., construction costs) are independent of the charging points' usage. However, as fixed costs are sunk, only the marginal cost c_v is relevant to the (short-run) best-response strategy of any CPO. Marginal costs may have two components: (1) CPOs' individual costs, of which electricity costs account for a significant portion, and (2) costs due to regulation, e.g., taxes, levies, or remuneration payments for funding support. We assume that the second type of costs is homogeneous for all CPOs. CPOs' electricity costs depend on many factors, e.g., time of use (as electricity prices are volatile) and grid fees at the respective connection points. While individual CPOs may be able to reduce electricity procurement costs to some degree through better management of electricity procurement, most of the factors are beyond CPOs' control, and are thus exogenous. We further assume that CPOs price-discriminate in different locations. This allows us to derive more nuanced insights versus those under a single-price strategy or price tiers. There are two notable scenarios that can materialize for a given network:

- A vertex $v \in \mathcal{V}_i$ only has edges to vertices controlled by CPOs $j, j \neq i$. In this case, the pricing at v is an isolated decision, and is independent of the pricing at i's other vertices $v' \in \mathcal{V}_i \setminus \{v\}$.
- An edge e connects two vertices $v, w \in \mathcal{V}_i$. In this case, i is a monopolistic CPO serving all the demand from e, and we call e a *monopolistic edge*. Recall that in our demand model, the total demand on an edge is constant; only the demand shared at the source and destination vertices are functions of the price difference. The monopolist's best strategy would thus be to (potentially) forgo demand on all other connecting edges and raise prices p_v, p_w to an arbitrarily high level, yielding infinite profits. The regulator can prevent this situation explicitly with a price cap or implicitly by surrounding charging stations of one lot with charging stations of other lots, i.e., by avoiding monopolistic edges.

For CPO $i \in N$, to maximize the total revenue over all controlled charging stations, it solves the following optimization problem:

$$\begin{aligned} \max_{\mathbf{p}_i^t} \quad & \Pi_i = \sum_t \Pi_i^t = \sum_t \sum_{v \in \mathcal{V}_i} \sum_{e \in \mathcal{E}_v} D_{v,e}^{t*}(\mathbf{p}) \Big(p_v^t - c_v^t \Big) \\ \text{s.t.} \quad & \mathbf{p}_i^t \preccurlyeq p_{cap}. \end{aligned} \tag{UL(i)}$$

where p_{cap} denotes a price cap set by the regulator and c_v^t is the variable cost at vertex v at time t. The optimization problem (UL(i)) includes the consumers' decision-making via the demand term $D_{v,e}^{t*}(\mathbf{p})$ and may be non-convex. Intuitively, such non-convexity arises from the network structure: When lowering prices, an operator may capture all demand on an edge without a further increase in demand for even lower prices. At the same time, increasing prices only affects demand until all consumers on that edge switch to the competitor. In other words, "saturation" would emerge at both pricing extremes, leading to a sigmoidal and hence non-convex demand response.

5.3.2 Consumer Demand

To understand the impact of CPOs' pricing decisions on the realized demand at the vertices, we need to model consumers' demand. Consider a single edge e with length L_e and demand density f_e^t, as illustrated by Figure 5.2. Time dependency of demand density can stem from consumer preferences to charge at certain times of the day or could be associated with more convenient charging opportunities arising from a driver's itinerary. Each consumer between 0 and L_e is characterized by her position x on the edge. For expositional convenience, for each edge e, we specify a *source* and a *destination* vertex, denoted by src(e) and dst(e), respectively. We measure position x from the source vertex. To solve for the consumers' response to pricing strategies, consider edge e as Hotelling's "linear city" market model (Hotelling, 1929). Then, there exists a point of division x_e so that all consumers located at $x < x_e$ choose the source src(e), and those located at $x > x_e$ choose the destination dst(e). We call the consumer located at x_e the *indifferent consumer*. Thus, from the edge e, $f_e^t x_e$ consumers visit the source vertex and $f_e^t(L_e - x_e)$ consumers visit the destination vertex, respectively.

A consumer's utility is determined by the difference between her valuation of charging and the associated cost. We assume that the valuation of charging is sufficiently high so that consumers always charge *somewhere*, i.e., there is no so-called self-elasticity. Furthermore, we abstract from temporal cross-elasticity (Bao et al., 2021) and concentrate on spatial elasticity, as regional characteristics of the charging network as well as their impact on the competitive outcomes constitute our main research focus. We propose that the associated cost breaks down into three components: energy cost, travel cost, and delay cost. The energy cost for a charging amount of q at a vertex v equals qp_v^t (p_v^t denotes the price at the vertex v at time t). Let τ denote the unit travel cost parameter. Depending on whether she chooses edge e's source or destination vertex for charging, the travel cost for the consumer at location x is either $x\tau$ or $(L_e - x)\tau$.

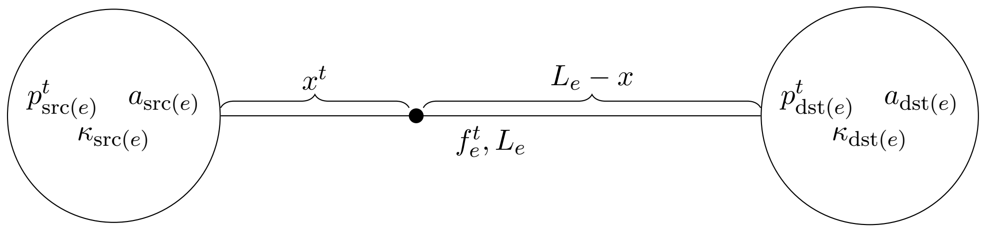

Figure 5.2: Demand on a single edge.

Both product and travel costs are part of standard models in spatial product differentiation (e.g., T. Wang & Wang, 2018). Beyond that, we assume that each vertex v has an attractiveness a_v, which can be influenced by nearby points of interest such as restaurants, banks, hospitals, etc. A more attractive location will offset charging and travel cost, effectively increasing the overall utility for consumers. We further introduce an additional cost term accounting for congestion at charging stations: delay cost. Delay costs may be incurred as highly popular charging stations capture more demand and might not be able to provide sufficient capacity to serve all consumers without waiting times. Unlike the energy or travel cost, the delay cost incurred by a consumer directly depends on other consumers' charging decisions. The consumers' waiting time at vertex v is thus increasing in the number of consumers visiting vertex v and decreasing in the number of charge points κ_v at that location. Note that κ_v is not a function of time, i.e., there is no (short-term) option to increase the offering of chargers at a given location. Rational consumers maximize their total utility.

We consider the waiting time at a vertex to be linearly increasing in the demand and inversely decreasing in the number of chargers at vertex v. In our base model, we assume consumers form a "locally-aware" estimation of the waiting time by accounting for the arrivals from the *same* edge, i.e., consumers perceive the arrivals at the source and destination vertices to be $f_e^t x_e^t$ and $f_e^t(L_e - x_e^t)$ without accounting for arrivals from other edges. Therefore, their estimated delay cost is expressed as $\delta \frac{f_e^t x_e^t}{\kappa_{\mathrm{src}(e)}}$ and $\delta \frac{f_e^t(L_e - x_e^t)}{\kappa_{\mathrm{dst}(e)}}$ at the source and destination vertex, respectively, where δ is the unit delay cost parameter. In general, during peak demand times with higher f_e^t, delay cost is increasing. We relax this simplification and consider globally aware consumers (i.e., consumers that account for arrivals from all edges in the network) in an extension (see section 5.6.2).

If the price difference $\Delta p_e^t = p_{\mathrm{src}(e)}^t - p_{\mathrm{dst}(e)}^t$ between the source and destination vertices is small, the indifferent consumer can be determined by equating all cost terms at both vertices. However, such a closed-form expression may not exist for larger price differentials because in that case, all demand will converge to one of the vertices. The following lemma characterizes the (saturated) indifferent consumer. All proofs of lemmas and propositions are presented in section 5.6.1

Lemma 1. *The indifferent consumer on an edge e is characterized by*

$$x_e^{t*} = \min\left(L_e, \max\left(\frac{-q\Delta p_e^t + \Delta a_e + L_e(\tau + f_e^t \delta \kappa_{\mathrm{dst}(e)}^{-1})}{2t + f_e^t \delta\left(\kappa_{\mathrm{src}(e)}^{-1} + \kappa_{\mathrm{dst}(e)}^{-1}\right)}, 0\right)\right), \qquad (xe)$$

where $\Delta p_e^t = p_{src(e)}^t - p_{dst(e)}^t$.

Lemma 1 indicates that an interior solution of the indifferent consumer has a linear dependence on Δp_e^t. However, as the unit travel cost parameter (t) increases, the relative importance of the price differential decreases compared to the preference of traveling to the nearest station. In particular, when $\tau \to \infty$, the indifferent consumer moves toward the center of the edge, i.e., $x_e^{t*} \to \frac{L_e}{2}$. Moreover, with larger δ (unit delay cost parameter) or f_e^t (demand density), consumers would experience longer waiting times and hence higher delay costs. When $\delta \to \infty$, $x_e^{t*} \to \frac{L_e}{2} \frac{\kappa_{\mathrm{dst}(e)}^{-1}}{\kappa_{\mathrm{src}(e)}^{-1} + \kappa_{\mathrm{dst}(e)}^{-1}}$. That is, demand will split at a weighted average of inverse station capacities.

Using Lemma 1, the demand at the source vertex and destination vertex of e is given by $D_{\mathrm{src}(e),e}^{t*} = q f_e^t x_e^{t*}$ and $D_{\mathrm{dst}(e),e}^{t*} = q f_e^t (L_e - x_e^{t*})$, respectively. We sometimes formulate the demand as a function of p^t, i.e., $D_{v,e}^{t*}(p^t)$, to highlight its dependence on p^t. As CPOs' pricing decisions affect the demand on their associated edges, the impact of the price difference on the optimal demand level must be evaluated. When making pricing decisions, CPOs have to assess the impact of prices on demand. As established before, the indifferent consumer is a piecewise linear function of the price differential. Where defined, the price sensitivity of demand $D_{v,e}$ from edge e at vertex v is therefore a piecewise constant function, as stated in the following proposition.

Price Sensitivity of Demand

CPOs' pricing decisions have to account for consumers' responses. As such, the price sensitivity of demand is crucial to CPOs' pricing.

Proposition 1. *$\forall v, e$, the derivative of demand $D^t_{v,e}$ from edge e at vertex v with respect to p^t_v is uniquely defined by*

$$\frac{dD^{t*}_{v,e}}{dp^t_v} = \begin{cases} -\frac{q^2 f^t_e}{2\tau + f^t_e \delta\left(\kappa^{-1}_{\mathrm{src}(e)} + \kappa^{-1}_{\mathrm{dst}(e)}\right)}, \quad \text{for } \Delta p^t_e \in \left(\frac{\Delta a_e - L_e\left(\tau + f^t_e \delta \kappa^{-1}_{\mathrm{src}(e)}\right)}{q}, \frac{\Delta a_e + L_e\left(\tau + f^t_e \kappa^{-1}_{\mathrm{dst}(e)}\right)}{q}\right) \\ 0, \quad \text{for } \Delta p^t_e \notin \left[\frac{\Delta a_e - L_e\left(\tau + f^t_e \delta \kappa^{-1}_{\mathrm{src}(e)}\right)}{q}, \frac{\Delta a_e + L_e\left(\tau + f^t_e \kappa^{-1}_{\mathrm{dst}(e)}\right)}{q}\right] \end{cases} \tag{dDdp}$$

Proposition 1 allows us to study settings that are not fully competitive because, as in standard monopoly models, CPOs with market power use information on consumers' sensitivity to their prices to determine profit-maximizing prices. The proposition further shows the impact of the different model parameters on the price sensitivity. If the price sensitivity of demand is low, there is a stronger incentive for CPOs to increase prices.

5.3.3 Equilibrium Analysis

When CPOs make their pricing decisions, they have to account for both consumers' behavior, and their competitors' responses. Given competitors' prices $\mathbf{p}^t_{-i} = (p^t_v : v \notin \mathcal{V}_i)$, CPO i's optimal prices can be found by reformulating (UL(i)) as a mixed-integer program, which can be solved via commercial solvers (Gabriel et al., 2013). We refer to the solution $\mathbf{p}^{t*}_i = (p^{t*}_v)_{v \in \mathcal{V}_i}$ as player i's *best response* to prices $\mathbf{p}^t_{-i}$. The corresponding EPEC is to find a vector of prices $\mathbf{p}^{t*}$ that is mutually consistent, i.e., prices that constitute a fixed point of the best response vector. This leads to our following definition:

Definition 3 (Global Nash Equilibrium)**.** *Let $G(V, E)$ be a graph with vertices V and edges E, and let N be the set of CPOs competing on G. A vector*

$\mathbf{p}^{t*} = (\mathbf{p}_i^{t*} : i \in N)$ *that simultaneously solves (UL(i))* $\forall i \in N$ *and (xe)* $\forall e \in E$ *is a global Nash equilibrium.*

If a global Nash equilibrium exists, it can be found by the *diagonalization* method: CPOs calculate their optimal pricing decisions while taking competitors prices as given iteratively. Once this process converges, i.e., no competitor has incentives to deviate, an equilibrium emerges (Gabriel et al., 2013; X. Hu & Ralph, 2007). Below, we will analyze CPOs' equilibrium prices as best responses to competitors' decisions.

CPOs' equilibrium pricing decisions are best understood using small-scale examples. Figure 5.3 provides an illustration of two stylized networks: one with two vertices, and one with three vertices. For expositional convenience, we assume

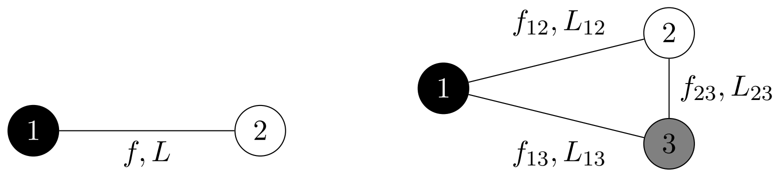

(i) Two-vertex network.

(ii) Competitive Three-vertex network.

Figure 5.3: Exemplary networks. For easier visualization, vertices (i.e., CPOs) are indexed and colored.

symmetric demand density, i.e., $f_e = f \;\forall e$, as well as equal attractiveness a (i.e., same number of amenities nearby the charging station) and capacities κ (i.e., the same number of charging points per charging station). We further assume consumers have a unit demand (i.e., $q = 1$), and drop the time index t in the following discussion of small-scale examples. With these simplifications, we can rewrite equation (xe) as

$$x_e^* = \min\left(L_e, \max\left(\frac{-q\Delta p_e}{2\tilde{\tau}} + \frac{L_e}{2}, 0\right)\right).$$

where $\tilde{\tau} := \tau + f\delta\kappa^{-1}$ is an augmented travel cost term accounting for waiting time under the presented symmetry assumptions. We assume all CPOs' marginal costs are below the price cap p_{cap}, as otherwise CPOs could not make an operational profit regardless of the competitive setting.

Two-vertex network

The first example, shown in Figure 5.3i, is a single edge example that is widely used in the spatial economics literature (e.g., Hotelling, 1929). Extending the results from T. Wang and Wang (2018), the following two conditions guarantee the existence of a global Nash equilibrium:

(C2) $|c_1 - c_2| \leq 3L\tilde{\tau}$

(T2) $L\tilde{\tau} + \frac{1}{3}(\sum_j c_j + \max_i\{c_i\}) \leq p_{cap}$,

where Condition (C2) can be viewed as a special case of Proposition 2 in T. Wang and Wang (2018). Condition (T2) ensures that the price cap is not binding in equilibrium. Condition (C2) ensures that the **c**ost difference between CPOs is not too large; the "2" indicates that this condition is only directly valid for the two-node case. With a relatively small cost difference between vertices, resulting equilibrium prices are also small in difference, thus leading to an interior solution of the indifferent consumer. With a higher cost differential, the cheaper CPO would capture all demand, thus leading to a binding restriction on the position of the indifferent consumer and an inability to find a closed-form solution. Condition (T2) requires that **t**ravel cost $\tilde{\tau}$ cannot be too high: if getting to a distant charging point is very costly for consumers, a nearby CPO can increase its price without losing *too much* demand; ultimately, the price will reach the cap set by the regulator. When both (C2) and (T2) hold, $p_i^* = \tilde{\tau} + \frac{1}{3}\sum_j c_j + c_i$ is an inner solution to both the upper- and lower-level problems.

Since $c_i < p_{cap}$ $\forall i$, we can show that $\neg$(C2) $\implies$ (T2), and therefore, $\neg$(T2) $\implies$ (C2). Without loss of generality, assume $c_1 \leq c_2$, that is, the cost of CPO 1 is no greater than that of CPO 2. Table 5.2 summarizes the different settings.

(C2)	(T2)	Price	Regime
yes	yes	$p_i^* = L\tilde{\tau} + \frac{1}{3}(\sum_j c_j + c_i)\ \forall i$	Compete
no	yes	$p_1^* = c_2 - L\tilde{\tau}$, $p_2^* = c_2$	Concede
yes	no	$p_i^* = p_{cap}\ \forall i$	Split

Table 5.2: Resulting prices in two-vertex network.

If $\neg$(C2), the cheaper CPO is better off by reaping all demand. In this setting, there exists only a weak global Nash equilibrium, as for $p_1 = c_2 - \tilde{\tau}$, CPO 2 is indifferent between any prices $p_2 \in [c_2, p_{cap}]$. If $\neg$(T2), it is beneficial for the

CPOs to split the demand and both raise their prices to the cap. As the settings specified by Table 5.2 are exhaustive, we have the following result.

Theorem 1. *Let $G(V, E)$ be a graph with $|V| = 2$ and $|E| = 1$. Let $N = \{1, 2\}$ be the set of CPOs competing on G. There always exists a (weak) global Nash equilibrium on G.*

To this point, we note that the (non-)satisfaction of conditions (C2) and (T2) spawns different competitive *regimes*, as outlined in Table 5.2. Despite the simplicity of the two-vertex network, these regimes serve as the building blocks of competition in any larger network.

Three-vertex network

Conditions (C2) and (T2) can be augmented to incorporate the fully connected three-vertex network shown in Figure 5.3ii:

(C3) $|c_{\text{src}(e)} - c_{\text{dst}(e)}| \leq \frac{5}{2} L_e \tilde{t} \quad \forall e \in E$

(T3) $L\tilde{t} + \frac{1}{5}(\sum_j c_j + 2\max_i\{c_i\}) \leq p_{cap}$

Note that these conditions do not translate directly to three-vertex networks that are not fully connected. Again, when both (C3) and (T3) are satisfied, all edges are in the *competitive* regime, and equilibrium prices are analytically tractable (corresponding to an inner solution of the EPEC). Since there are three edges, there are now three conditions corresponding to the requirement of small cost differences. Further, (C3) and (T3) are not mutually exclusive anymore. For example, consider $\epsilon > 0$ $f_e = f, L_e = L$ $\forall e$ and $c_1 = p_{cap} - \frac{5}{2}L\tilde{\tau} - 2\epsilon$; $c_2, c_3 = p_{cap} - \epsilon$; both (C3) and (T3) are violated when $L\tilde{t} > \frac{12}{5}\epsilon$. In general, $\neg$(T3) implies that (C3) holds for (at least) one edge.

Proposition 2. *Let $G(V, E)$ be a graph with $|V| > 2$ and $|E| > 1$. Let N be the set of CPOs competing on G. A global Nash equilibrium may not exist.*

Proposition 2 states that a global Nash equilibrium may not always exist for a network with more than two vertices. Given this result, we adopt an alternative solution concept, i.e., *local equilibrium.* The concept was first brought forward by X. Hu and Ralph (2007) and can be used to characterize CPOs' optimal pricing decisions: While the non-existence result offers important theoretical insights to the spatial competition model, CPOs need practical tools to assess profitability.

Intuitively, local Nash equilibria capture *locally* best responses of competitors to one another. To formalize this idea, we start by deriving the KKT conditions of (UL(i)), which are necessary conditions for local optimality. Here, we consider all aspects of charging station heterogeneity again, e.g., varying capacities (κ_v), time-varying charging prices (p_v^t), and different demand densities (f_e^t). We first determine the gradient $\nabla_{\mathbf{p}_i^t}\Pi_i^t$. For a vertex $v \in \mathcal{V}_i$ that is only connected to vertices controlled by other CPOs, i.e., $w \in V_j, j \neq i$, this is straightforward, as the derivative of the profit depends solely on the price at vertex v. However, for a vertex v with adjacent monopolistic edges, i.e., there are vertices w such that $w \in N_v \cap \mathcal{V}_i$, the impact of a change in p_v^t on the demand (and subsequently, profit) at those vertices w also need to be considered:

$$
\begin{aligned}
\nabla_{\mathbf{p}_i^t}\Pi_i^t &= \left(\frac{\partial \Pi_i^t}{\partial p_v^t}\right)_{v\in\mathcal{V}_i} \\
\frac{\partial \Pi_i^t}{\partial p_v^t} &= \sum_{e\in\mathcal{E}_v}\left(D_{v,e}^t + \frac{\mathrm{d}D_{v,e}^t}{\mathrm{d}p_v^t}(p_v^t - c_v^t)\right) + \underbrace{\sum_{w\in N_v\cap\mathcal{V}_i}\frac{\mathrm{d}D_{w,e}^t}{\mathrm{d}p_v^t}(p_w^t - c_w^t)}_{\text{monopoly correction}} \\
&= \sum_{e\in\mathcal{E}_v}\left(D_{v,e}^t + \frac{\mathrm{d}D_{v,e}^t}{\mathrm{d}p_v^t}(p_v^t - c_v^t)\right) - \sum_{w\in N_v\cap\mathcal{V}_i}\frac{\mathrm{d}D_{v,e}^t}{\mathrm{d}p_v^t}(p_w^t - c_w^t).
\end{aligned}
$$

The third line of the above derivation follows from that any demand increase at the source vertex corresponds to a decrease at the destination vertex.

Let $\overline{\boldsymbol{\mu}_i^t} = (\overline{\mu_v^t})_{v\in\mathcal{V}_i}$ be the vector of dual multipliers corresponding to the upper bound (cap) on prices $\mathbf{p}_i^t$. We can obtain the KKT conditions for (UL(i)) as

$$
\begin{aligned}
\frac{\partial \Pi_i^t}{\partial p_v^t} p_v^t = \hat{p}_v^t + \overline{\mu_v^t} &= 0 \quad \forall v \in \mathcal{V}_i,\ \forall t \\
\overline{\mu_v}(p_{cap} - \hat{p}_v^t) &= 0 \quad \forall v \in \mathcal{V}_i,\ \forall t \\
\hat{\mathbf{p}}_i^t &\preccurlyeq p_{cap} \quad \forall t \\
\overline{\boldsymbol{\mu}_i^t} &\succcurlyeq 0 \quad \forall t,
\end{aligned}
\tag{KKT(i)}
$$

where $\hat{\mathbf{p}}_i^t := (\hat{p}_v^t)_{v\in\mathcal{V}_i}$ is a vector of prices that satisfies (KKT(i)). With (KKT(i)), we can find a point that is locally stationary for (UL(i)). This leads to our formal definition of local Nash equilibrium.

Definition 4 (Local Nash Equilibrium)**.** *Let $G(V, E)$ be a graph with vertices V and edges E, and let N be the set of CPOs competing on G. A vector of prices*

$\hat{\mathbf{p}}^t = (\hat{\mathbf{p}}_i^t : i \in N)$ *that simultaneously solves (KKT(i))* $\forall i$ *and (xe)* $\forall e$ *is a local Nash equilibrium.*

Here, $\hat{\mathbf{p}}^t$ captures the *local* best response of a CPO to its competitors. Such a local response is a much weaker concept compared to a globally best response. However, in an applied setting, competitors might be hesitant to abruptly change their prices, as price jumps could affect their reputation with consumers and harm business eventually. In this regard, a gradual adjustment to prices seems more appropriate.

5.3.4 Equilibrium Calculation

Even though the non-linear local Nash equilibrium is a simpler problem to solve than its global counterpart, it is still of significant difficulty (Gabriel et al., 2013). For any large-scale EV charging market, the associated price vectors are high-dimensional. This is because (1) real-world instances of charging networks would consist of numerous charging locations; (2) although the market exhibits regionality, the pricing decision at one location in the network may impact pricing at a distant location via various dependencies. As a result, finding the exact solution of such local equilibria in general cases is a significant computational challenge and does not lend itself to scenario-based policy or competitive analysis in need of multiple simulation runs. In light of this, we propose an approximation method to speed up the search for local equilibria and to ensure that our findings can be used by practitioners and policymakers seeking to analyze real-world settings.

To achieve computational efficiency, we introduce a continuously differentiable approximation of the demand function, which, combined with sequential least squares quadratic programming (Kraft, 1988, 1994), significantly speeds up the search for local equilibria. Let

$$
\begin{aligned}
x_e^{t'} &= \frac{L_e}{2} - \frac{1}{2}\sqrt{\frac{q}{2\tau_e'}\left(4s + \frac{q}{2\tau_e'}\left(\Delta p_e^t - \frac{\Delta a_e}{q} + \frac{L_e'\tau_e'}{q}\right)^2\right)} \\
&\quad + \frac{1}{2}\sqrt{\frac{q}{2\tau_e'}\left(4s + \frac{q}{2\tau_e'}\left(\Delta p_e^t - \frac{\Delta a_e}{q} - \frac{L_e'\tau_e'}{q}\right)^2\right)} \\
\frac{\mathrm{d}x_e^{t'}}{\mathrm{d}\Delta p_e^t} &= \frac{\left(\frac{q}{2\tau_e'}\right)^2\left(\Delta p_e^t - \frac{\Delta a_e}{q} - \frac{L_e'\tau_e'}{q}\right)}{2\sqrt{\frac{q}{2\tau_e'}\left(4s + \frac{q}{2\tau_e'}\left(\Delta p_e^t - \frac{\Delta a_e}{q} - \frac{L_e'\tau_e'}{q}\right)^2\right)}} \\
&\quad - \frac{\left(\frac{q}{2\tau_e'}\right)^2\left(\Delta p_e^t - \frac{\Delta a_e}{q} + \frac{L_e'\tau_e'}{q}\right)}{2\sqrt{\frac{q}{2\tau_e'}\left(4s + \frac{q}{2\tau_e'}\left(\Delta p_e^t - \frac{\Delta a_e}{q} + \frac{L_e'\tau_e'}{q}\right)^2\right)}},
\end{aligned}
$$

with the smoothing parameter $s > 0$, and $\tau_e' := \tau + \frac{f_e^t\delta}{2}\left(\kappa_{\mathrm{src}(e)}^{-1} + \kappa_{\mathrm{dst}(e)}^{-1}\right)$, $L_e' := L_e\frac{\tau + f_e^t\delta\kappa_{\mathrm{dst}(e)}^{-1}}{\tau_e'}$.

Note that $\lim_{s\to 0} x_e^{t'} = x_e^{t*}$, and wherever the non-smoothed demand is differentiable, we have $\frac{\mathrm{d}x_e^{t'}}{\mathrm{d}\Delta p_e^t} = \frac{\mathrm{d}x_e^{t*}}{\mathrm{d}\Delta p_e^t}$.

The approximation provides a good fit to the piecewise-linear demand, and enhances regularity for gradient-based solution techniques (see Figure 5.15 in section 5.6.2 for an illustration of the asymptotic properties of our approximation method). In addition to the advantage of continuous differentiability, using $x_e^{t'}$ instead of x_e^{t*} reduces the number of variables and complementarity constraints from $\mathcal{O}(|V| + |E|)$ to $\mathcal{O}(|V|)$, and thereby greatly improves the computational efficiency. The smooth approximation with $s = 0.001$, which we use for the case study, reduces the computation time of local equilibria from 22 hours to roughly three minutes on a high-power computer.

5.4 Case Study: Spatial Competition in the Deutschlandnetz

In this section, we apply the spatial competition model and the computational method presented in Section 3 to study the competition between CPOs over the recently inaugurated national fast-charging network in Germany, i.e., the *Deutschlandnetz*.

5.4.1 Deutschlandnetz

In 2021, the federal government of Germany initiated a public tender (i.e., a procurement auction) for the *Deutschlandnetz*. This initiative aimed to create a comprehensive public fast-charging infrastructure that would serve not only densely populated urban centers but also rural areas throughout Germany. To achieve a balanced coverage, the tender divides the country into six regions: North-West, North-East, Central Germany, South-East, South-West, and West. As shown in Table 5.3, each region consists of several lots. These lots are *scattered* over their respective regions, avoiding local monopolies. Each lot contains multiple search spaces, i.e., designated areas where the winner of the lot has to set up charging stations. These search spaces vary in capacity, with between 4 and 16 charging points each. In total, the tender design includes 7,821 charging points at 900 charging stations.

Region	Lots	Number of Charging Stations in Each Lot
North-West	4	20, 47, 48, 48
North-East	3	20, 41, 42
Center	3	20, 44, 45
South-East	3	20, 49, 50
South-West	5	20, 43, 43, 44, 44
West	5	20, 48, 48, 48, 48
SUM	23	900

Table 5.3: Overview of charging stations across different regions.

Network Structure

The vertices of the charging network correspond to the charging stations, with locations determined through the tender. To determine the edges in the network, we use Delaunay triangulation (i. a. De Berg, 2000), where we remove edges unless 99% of their length falls within 5 km of Germany's borders; see Figure 5.4i for an illustration. The resulting network has 900 vertices, each of which represents a charging location, and 2,616 edges. On average, each vertex is adjacent to approximately six edges. Figure 5.4ii provides an overview of the charging network.

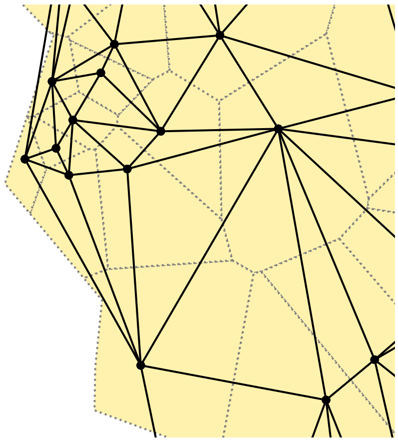

(i) Charging stations in the west border of Germany, next to Belgium and the Netherlands. Black lines represent edges of the network, gray dotted lines are borders of the Voronoi regions.

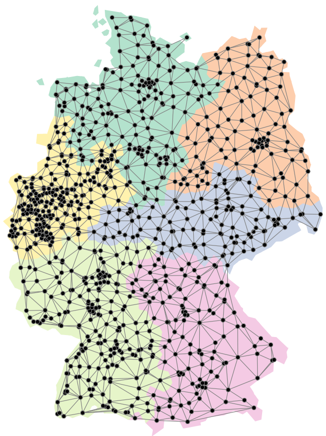

(ii) The *Deutschlandnetz* derived using Delaunay triangulation. The density of vertices is substantially higher in metropolitan regions. Colors represent the six regions of the tender.

Figure 5.4: Network structure derived by Delaunay triangulation.

Pricing Rules

Since the tender covers all fixed-cost components according to CPOs' bids, the utilization risk for CPOs is minimized. Upon utilization, a compensation payment has to be made to the government, which discourages CPOs from offering charging at the marginal cost of electricity and thereby protects CPOs of existing fast-charging stations. Such compensation payment, upon a high utilization rate, also makes the *Deutschlandnetz* tender significantly cheaper for the federal budget, retroactively.

The retail price of charging electricity at vertex v thus consists of the price of electricity, the compensation payment, and the margin:

$$p_v = \underbrace{c_{grid,v} + c_{comp,v}}_{c_v} + p_{margin,v}.$$

At the time of the initial tender design, the electricity price (including fees, applicable taxes, and surcharges) is 0.2023€/kWh, and the compensation payment is 0.1785€/kWh. Accounting for other minor marginal cost components reducing the margin, we assume homogeneous marginal costs of 0.39€/kWh for all CPOs. Originally, a price cap of 0.44€/kWh was discussed. In the final tender design, such a strict price cap was not enforced, yet measures to ensure competitive pricing were announced. In our base case, we assume that CPOs price below 0.44€/kWh to avoid interventions (we later relax this assumption and consider alternative price caps and study their impact on pricing and consumer welfare; see Section 4.4). This leads to a margin of up to 0.05€/kWh, making EVs cost-competitive with efficient diesel vehicles.

Taken together, the tender rules make product differentiation—apart from the location—difficult for CPOs by strictly regulating hardware qualities, payment modes, and reliability. Because search spaces restrict the locations, CPOs are participating in an almost pure price competition.

Tender Outcome

In September 2023, the German National Center for Charging Infrastructure announced the ten winners of the Deutschlandnetz tender. Notably, while some large incumbent CPOs won lots in the tender, German market leaders Tesla and EnBW did not. Table 5.4 shows the winners and the number of charging stations they will operate, respectively.

CPO	**Lots**	**Locations**	**Chargers**
Eviny Elektrifisering AS	3	142	1232
E.ON Drive Infrastructure GmbH	3	139	1216
Total Energies Marketing Deutschland GmbH	3	134	1092
Fastned Deutschland GmbH & Co. KG	2	92	868
EWE Go HOCHTIEF Ladepartner GmbH & Co. KG	2	96	852
Via Deutschlandnetz (VINCI Concessions Deutschland GmbH)	3	106	828
Mer Germany	3	83	700
Allego GmbH	1	48	432
Pfalzwerke AG	2	40	400
BayWa Mobility Solutions GmbH	1	20	192

Table 5.4: Winners of the Deutschlandnetz tender.

5.4.2 Demand

We calculate future EV stock based on the predicted development of EV share and the current overall vehicle stock, assuming a 100% market share of EVs in the passenger segment by 2050, in line with the European climate targets. While the total vehicle ownership is subject to change due to various socioeconomic and technological factors, it should not affect our analysis for the years 2025 and 2030 at large. We further assume a mileage of 13,700 km per vehicle per year—the current German average—with a demand of 20 kWh/100 km and a share of 10% of demand that is consumed at fast-charging points (cf. Funke et al., 2019). Both mileage and consumption are representative of the EVs in the market. Per charging operation, the estimated demand is $q = 50$ kWh, which is roughly equivalent to charging a 70 kWh battery from 10% state of charge (SoC) to 80% SoC.[14] These assumptions translate to an average of roughly one fast-charging transaction per vehicle per week, which provides a solid baseline for our analysis, although we acknowledge that absolute demand levels are subject to uncertainty.

We assume the unit travel cost $t = 0.04$€/km and the unit delay cost $\delta = 2$€/kWh. Here, the unit of €/kWh does not refer to the charging amount of an individual consumer but relates to the overall hourly demand at the chosen charging station, which affects the waiting time and thus delay cost. We present a sensitivity analysis of travel and delay costs in section 5.6.3.

Spatial Structure

To model the demand side, we need spatially granular EV charging data. Because the competition we set out to investigate is only realized in the future (i.e., the charging stations are still under construction), we use projections of the associated statistics. To estimate the regional distribution of electric vehicles for relevant future years—2025 and 2030—within the operation of the *Deutschlandnetz*, we extrapolate past vehicle registration data from the German Federal Motor Transport Authority (*Kraftfahrt-Bundesamt*) (Kraftfahrt-Bundesamt, 2024). We assume a sigmoidal development of the EV stock following Bass diffusion (Bass, 1969). In

[14] Funke et al. (2019) provided lower battery capacities, though through cost decline in battery technologies, larger batteries are slowly becoming more prevalent]. Because of the decline in charging speeds at high SoCs, charging at a fast-charging point is typically stopped around 80% SoC.

particular, we fit the following function

$$\mathrm{share}_i(y) = \frac{1}{1 + \exp\left(-\beta_r (y - y_0)\right)}, \tag{5.1}$$

using the least squares approach with the *dogbox* optimization algorithm (e.g. Powell, 1970). Here, y_0 is the first year of consideration in our study, i.e., 2022. The parameter β_r governs the diffusion speed per region r, accounting for the spatial heterogeneity of EV adoption.

This spatially heterogeneous adoption of electric vehicles translates to a total EV stock of roughly 11 million by 2030, or approximately 25% of all passenger vehicles (see Figure 5.5), which is much higher than today's stock of 1 million EVs.

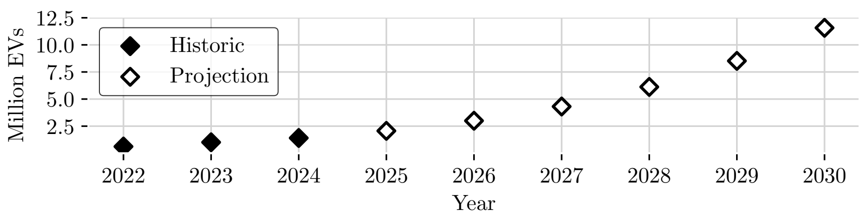

Figure 5.5: Germany-wide projection of EV stock

Temporal Pattern

Beyond the spatial structure, the temporal pattern of charging operations is also crucial for CPOs' pricing decisions. Figure 5.6i shows the distribution of starting times of charging operations at publicly funded high-performance chargers in Germany in 2022 (German National Centre for Charging Infrastructure, 2024). It is evident that the bulk (and virtually the single mode) of transactions occurs midday, translating to higher hourly demand densities f_e^t in peak times (10 AM to 7 PM) compared to shoulder (7 AM to 10 AM and 7 PM to 9 PM) or off-peak times (9 PM to 7 AM).

In peak times, one can expect that larger charging stations have more leverage to obtain monopoly rents, as the expectation of delay costs at smaller stations may deter consumers. In principle, CPOs could adopt dynamic pricing and set different prices at different times of a day. To keep the analysis tractable, we differentiate between three price levels for all operators: peak, shoulder, and off-peak pricing. Specifically, based on practical observations, we assume that CPOs' profits consist of 3,285 hours of peak, 1,825 hours of shoulder, and 3,650 hours of off-peak operation in a year. Together, these price tiers can be used to calculate a *weighted average price*, i.e., the average price per unit of energy charged at a respective charging station.

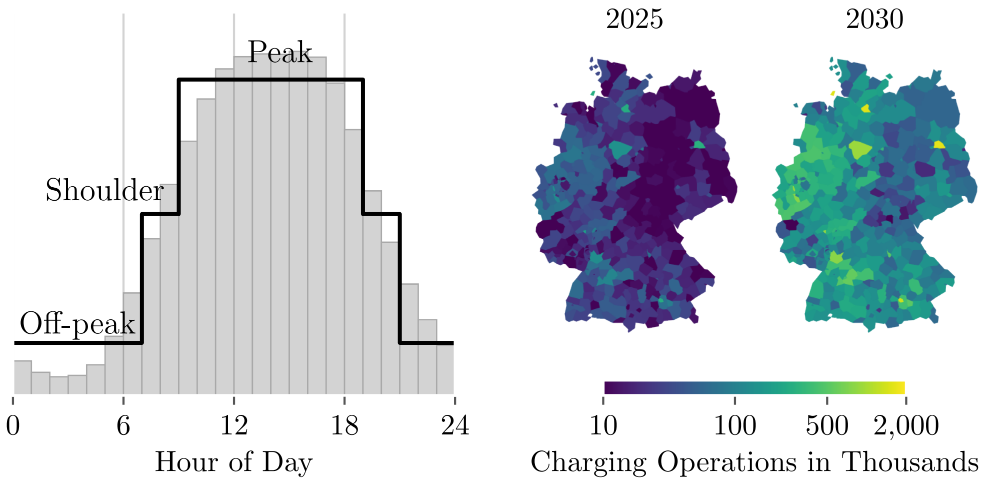

(i) Histogram of starting times of fast-charging operation in Germany in 2022. The bulk of charging operations begins at midday.

(ii) Projected number of charging operations in 2025 and 2030. By 2030, there will be beyond one million charging operations yearly in high-adoption regions, e.g., in metropolitan areas.

Figure 5.6: Starting times of charging operations (a) and projected numbers of charging operations (b).

We map the regional demand (per zip code level) to edges on the charging network in the following manner: whenever only one edge intersects a region, we attribute all the region's demand to that edge. When several edges intersect the region, we attribute demand proportionally to the intersections' lengths. The demand projections for the years 2025 and 2030 are shown in Figure 5.6ii. As expected, most charging operations will be in densely populated and metropolitan areas, such as North Rhine-Westphalia in the West, areas around Frankfurt and Stuttgart in the South, and Hamburg in the North. These high-demand areas correspond to the densely meshed areas in the associated graph.

Attractiveness

Given that charging operations, even in high-performance charging infrastructure, take longer than fueling conventional vehicles, consumers may prefer charging locations with amenities close to them. Depending on individual preferences, such amenities may include food and beverage offerings, synergetic services such as banks, gyms, and salons.

To approximate attractiveness values for the nodes in the charging network, we follow prior research (Babar & Burtch, 2024) and count relevant amenities within 150 meters of each charging location, using OpenStreetMap data (OpenStreetMap, 2024). Recognizing that additional amenities provide diminishing marginal benefits, we measure the attractiveness of a charging location using a logarithmic utility function of the amenity count. Figure 5.7 shows the distribution of the attractiveness scores that result from charging stations within the network. Under this assumption, most locations score an attractiveness between 0€ and 1€ (translating to less than 10 nearby amenities), that is, a minor increase in willingness to pay, all else being equal.[15] There are, however, 10 locations without any amenities in their direct proximity, thus lacking attractiveness compared to others.

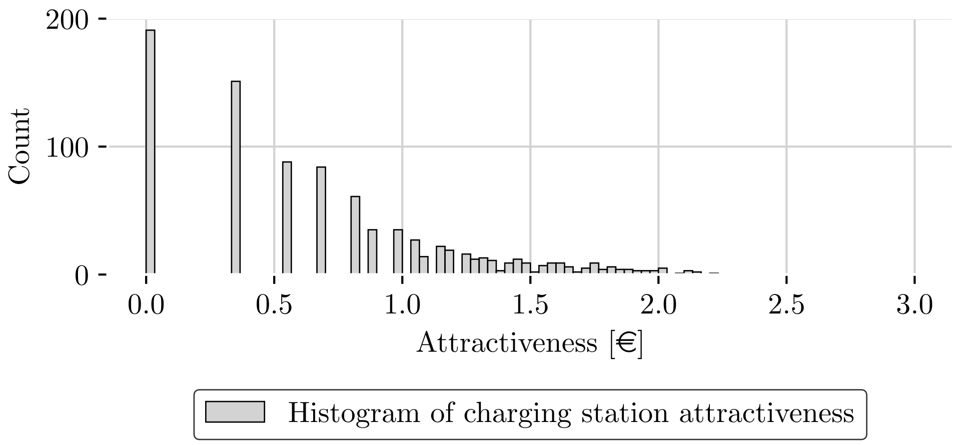

Figure 5.7: Distribution of attractiveness scores within the network based on an assumed log-utility of amenities.

[15] Note that individual willingness to pay to charge near a specific amenity might be higher. On the other hand, there are consumers who may not be interested in any amenities and thus look for the cheapest option.

5.4.3 Pricing over *Deutschlandnetz*

Based on the network structure and demand projections, we calculate the equilibrium prices charged by winners of the *Deutschlandnetz* tender. Figure 5.8 shows the resulting average prices of peak, shoulder, and off-peak operation from our competition model for the year 2030. We first examine the spatial structure of prices without accounting for the attractiveness of charging stations (Figure 5.8i). Interestingly, prices are slightly lower in high-demand areas, particularly in Western Germany, which, at first sight, appears counterintuitive. Contrary to basic economic principles, where higher demand typically drives up prices, we observe that densely populated metropolitan areas have lower prices than the rural areas (e.g., Eastern and South-Eastern Germany). This pattern emerges primarily because the tender designates a higher density of charging stations in populous areas, thereby encouraging competition and lowering prices.

Comparing Figures 5.4ii and 5.8i, we observe that equilibrium prices are largely determined by the network structure, with lower prices emerging in densely populated areas. However, since these areas also have higher demand, profits are not compromised. In general, operational profits are higher in the aforementioned high-demand areas, although some very low-demand areas in the South-East generate among the highest profits due to lack of competition. When comparing Figures 5.8i and 5.8ii, we note that accounting for the attractiveness of charging stations dilutes this otherwise clear spatial pattern. This is particularly pronounced in high-demand areas such as in Western Germany, where the attractiveness of charging stations creates a sharp contrast, resulting in the coexistence of extremely low and extremely high prices close to each other, i.e., a mixture of light and dark spots.

To identify the specific drivers of regional pricing patterns, we perform a more in-depth analysis of the relationship between network characteristics and associated prices through a set of metrics. The first is *local competitiveness* of a vertex v, LC_v, which is defined as the proximity of competitors of v. i.e., the sum of reciprocal lengths of adjacent edges that connect to a competitor. Standard measures for competitiveness, e.g., the Herfindahl–Hirschman Index, HHI, are not directly applicable to our case, as they typically do not account for spatial characteristics. In the *supercompetitive* setting, each edge connects to a competitor, as no operator has control over more than one charging station. Let $ME = \{e \in E \mid \exists\, i \in N : \mathrm{src}(e), \mathrm{dst}(e) \in \mathcal{V}_i\}$ the set of monopolistic edges, $\mathrm{LC}_v := \sum_{e \in \{O_v \cup I_v\} \setminus ME} \frac{1}{L_e}$.

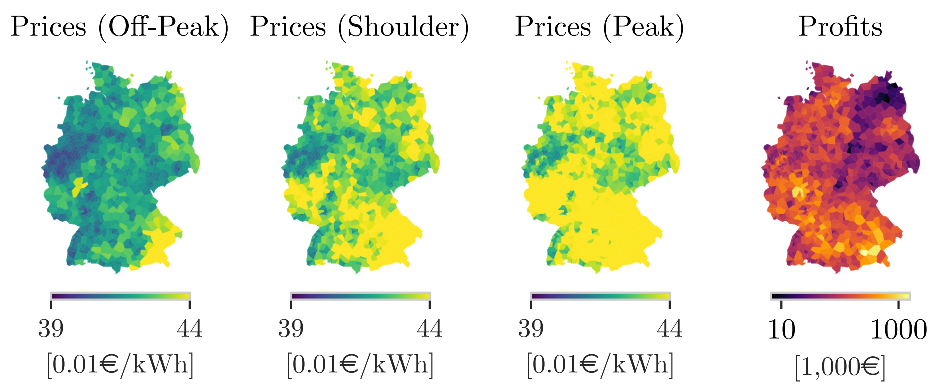

(i) Prices and profits without accounting for attractiveness.

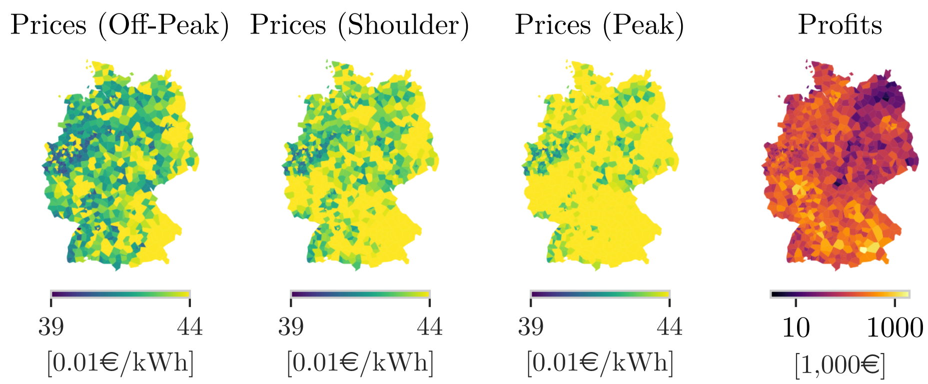

(ii) Prices and profits, accounting for attractiveness.

Figure 5.8: Results for 2030. Prices and profits in the Deutschlandnetz setting shown in the Voronoi regions derived from the Delaunay triangulation. Areas with a high demand potential generally show lower prices and slightly higher profits. Accounting for attractiveness dilutes the spatial characteristics within the network.

Beyond the competitive pressure created by geographic proximity, the capacity of competing charging stations also affects a CPO's optimal pricing decision: when consumers anticipate shorter waiting times at competitors' facilities due to larger capacity, those stations become inherently more attractive. Therefore, CPOs facing such competition must lower their prices to remain competitive. To quantify this effect, we use a simplified metric to capture competitor size: for each vertex in the network, we calculate the arithmetic mean of the capacity of all neighboring vertices, which we refer to as "mean competitor size".

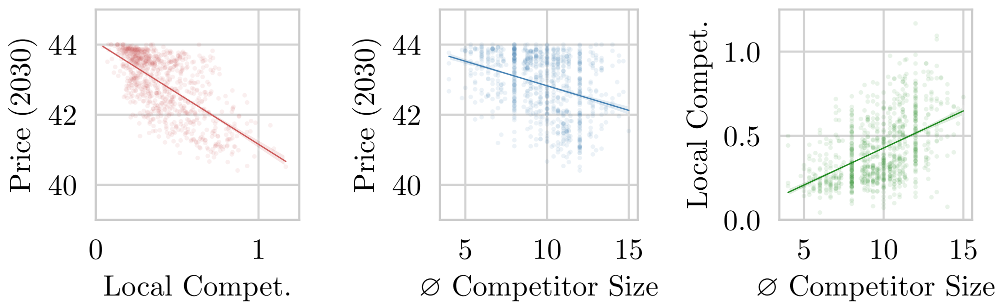

Figure 5.9: Correlation analysis between local competitiveness, mean competitor size, and (weighted average) price for 2030.

According to Figure 5.9, both (local) competitiveness and mean competitor size have a pronounced relationship with prices. The correlation coefficients of competitiveness and prices are $\rho_{\text{Pearson}}(p_v, LC_v) = -0.47$ ($p-$value < 0.001), $\rho_{\text{Spearman}}(p_v, LC_v) = -0.40$ ($p-$value < 0.001). Similarly, being surrounded by larger charging stations leads to lower prices at the respective vertex. With Pearson and Spearman correlation coefficients of -0.22 ($p-$value < 0.001) and -0.19 ($p-$value < 0.001), respectively, the relationship is even slightly more pronounced as for the local competitiveness. Moreover, we observe a strong correlation between the mean competitor size and the competitiveness of a charging station. According to the tender design, areas with higher charging station densities also feature larger average station sizes. This correlation between local competitiveness and station size (capacity) is not surprising: intuitively, higher-demand areas would have a denser network of charging stations with larger sizes. Therefore, simple regression analysis offers only limited insight into the driving factors behind pricing patterns.

To understand the dominant factors influencing pricing decisions, we conduct a multivariate linear regression analysis. Although prices in our case study are deterministic, making the standard interpretation of regression results not entirely applicable, OLS regression remains a suitable tool for identifying the best linear model to characterize pricing patterns. It also allows us to distinguish between the specific impact factors under consideration and pseudorandom disturbances not captured in our model. Table 5.5 presents the regression results. Here, *Demand Potential* refers to the sum of all demand on edges connected to a vertex, i.e., the maximal demand a charging station could theoretically capture. We can observe that both local competitiveness and mean competitor size have distinctly negative coefficients.The same holds true for the number of competitors. Conversely, demand potential exhibits a positive relationship with optimal prices: When serving a larger consumer base, CPOs can set higher prices (and hence achieve a higher profit margin), as losing a portion of consumers has less impact on their overall profits. Among all factors examined, local competitiveness demonstrates the strongest influence on equilibrium pricing strategies.

	coef	**std err**	**t**	**P> |t|**	**[0.025**	**0.975]**
Const	44.5348	0.131	339.726	0.000	44.278	44.792
Demand Potential	0.0002	1.89e-05	9.745	0.000	0.000	0.000
Number of Competitors	-0.0628	0.015	-4.067	0.000	-0.093	-0.033
Local Competitiveness	-2.2407	0.129	-17.324	0.000	-2.494	-1.987
Mean Competitor Size	-0.0788	0.013	-6.300	0.000	-0.103	-0.054
Attractiveness	0.0062	0.000	16.611	0.000	0.005	0.007

Table 5.5: OLS regression results for the Deutschlandnetz setting (2030).

5.4.4 Impact of Price Caps on Profit Margins and Consumer Welfare

We further examine the spatial distribution of CPO profit margins under different price caps and analyze its implication for consumer welfare. To focus on the direct effects of price regulation (i.e., price caps), we disregard the effect of charging station attractiveness here. Our analysis draws on equilibrium model outputs under uniform marginal costs of 0.39€/kWh and three price cap levels: 0.44, 0.50, and 0.56€/kWh. Recall that the price cap of 0.44€/kWh was proposed in the initial tender design but not enforced.

Spatial Distribution of Profit Margins

The spatial distribution of CPO profit margins varies substantially across the three cap levels (see Figure 5.10). Under the least restrictive cap of 0.56€/kWh, the profit margin distribution exhibits significant dispersion, with several regions showing excess margins approaching 0.17€/kWh. These high profit margins concentrate primarily in rural and infrastructure-poor areas, particularly in South-East Germany. In these regions, limited charger availability and low competitive intensity enable CPOs to extract local monopoly profits. By contrast, metropolitan areas such as Berlin, Hamburg, and the Rhine-Ruhr region exhibit lower and more uniform profit margins. These areas benefit from denser charging networks and more active CPO competition, which exert downward pressure on prices.

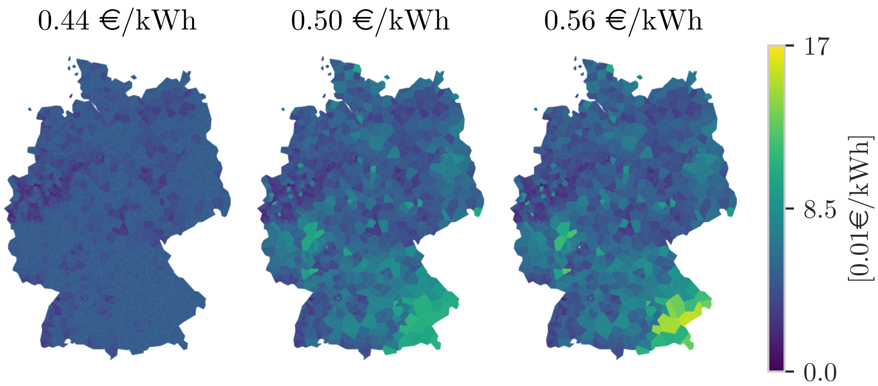

Figure 5.10: Profit margin distributions under different price caps in the Deutschlandnetz setting (2030).

As the cap tightens to 0.50€/kWh, the upper tail of the distribution contracts, significantly reducing profit margins in previously high-priced areas. The most pronounced effects occur in peripheral regions. Nevertheless, spatial heterogeneity in profit margins persists.

When the strictest cap of 0.44€/kWh applies, the profit margin distribution becomes sharply compressed. Most regions converge to margins of 0.02–0.05€/kWh, effectively flattening spatial differences. While this convergence promotes greater uniformity in pricing, it also limits CPOs' ability to adjust prices according to local demand and cost conditions.

Consumer Welfare

To analyze the impact of price caps on consumer welfare, we decompose the total consumer cost into three components: markup payments (i.e., CPOs' profit margins), travel costs, and delay costs. As shown in Figure 5.11, markup costs decline monotonically as caps tighten, reflecting reduced consumer exposure to pricing power. However, these direct savings are partially offset by rising indirect costs: under the low price cap of 0.44€/kWh, delay costs slightly increase, while travel costs show a marginal decrease. Overall, the effect on travel and delay costs is much smaller compared to the effect on markup cost (CPOs' profit margin).

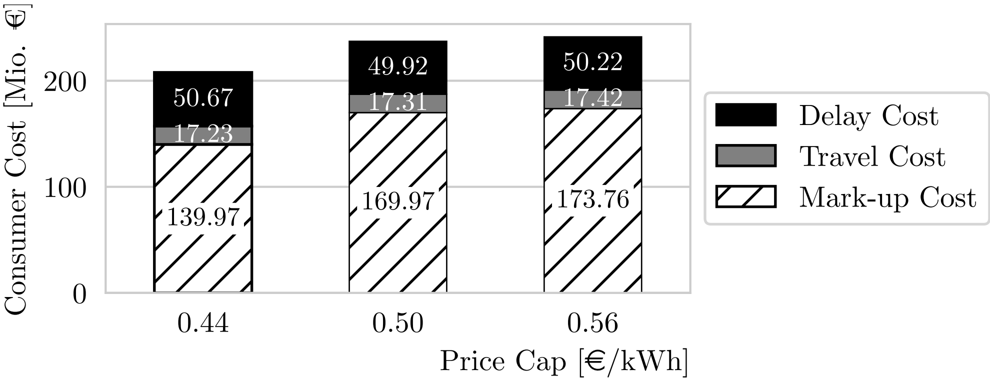

Figure 5.11: Components of consumer cost under different price caps in the Deutschlandnetz setting (2030).

These findings underscore the importance of accounting for spatial market structure when designing regulatory interventions. While uniform price caps promote equity and shield consumers from market power, they may inadvertently distort usage patterns and undermine incentives for infrastructure investment—particularly in less profitable regions. Policymakers should consider differentiated or dynamic pricing mechanisms or regionally calibrated caps to better align welfare objectives with spatial economic realities.

5.4.5 Comparative Analysis of Alternative Market Structures

Given the tender rules, numerous outcomes could have materialized. Since a complete enumeration of all possible outcomes is not tractable, we select four representative market structures (counterfactual scenarios) to benchmark the *Deutschlandnetz* outcome, namely, supercompetitive, competitive, oligopolistic, and monopolistic settings.

- *Supercompetitive Setting.* This configuration features a market where every charging station is operated by a different CPO (or at least managed independently). By design of the tender, this configuration cannot materialize, as stations are tendered in lots. However, it provides a benchmark for maximal competitiveness given the established network structure and charging station distribution.

- *Competitive Setting.* This configuration features a market where no CPO wins more than one lot, resulting in 23 distinct CPOs. While somewhat unrealistic—as the market is likely to show some consolidation—this scenario serves as a good benchmark for competitive pricing. Moreover, many configurations where CPOs own geographically distant lots are equivalent to this setting, as CPOs cannot extract monopoly rents.
- *Oligopolistic Setting.* This configuration represents a market with the minimal number of CPOs per tender rules. In particular, there are eight CPOs: five operate in the western regions (North-West, West, and South-West), and the remaining three operate in the eastern regions (i.e., North-East, Central, and South-East).
- *Monopolistic Setting.* This configuration represents a market where all charging stations are operated by a single CPO. By design of the tender, this configuration cannot materialize, as charging stations are tendered in lots, and one bidder cannot win all lots. Nevertheless, it provides a benchmark for maximal market power given the established network structure and charging station distribution. When investigating price distributions, we omit this scenario, as prices are set to the cap everywhere. However, we include this scenario in our comparative analysis of consumer welfare.

While these scenarios feature extreme cases, a solid understanding of these scenarios serves as the foundation for analyzing any intermediate competitive configuration—such as settings where some regions exhibit higher competition whereas others resemble oligopolistic structures.

Comparison of Equilibrium Prices

Figure 5.12 shows the complementary empirical cumulative distribution function of weighted average prices in the *Deutschlandnetz* setting and three counterfactual scenarios. We make two observations. First, there is a distinct deviation between the supercompetitive setting and the other settings. In particular, when comparing the *Deutschlandnetz*, competitive, and oligopolistic settings, we observe only marginal differences in their price distributions. While approximately 80% of vertices have prices exceeding €0.425/kWh in these less competitive settings, this proportion falls below 60% in the supercompetitive scenario. Second, prices

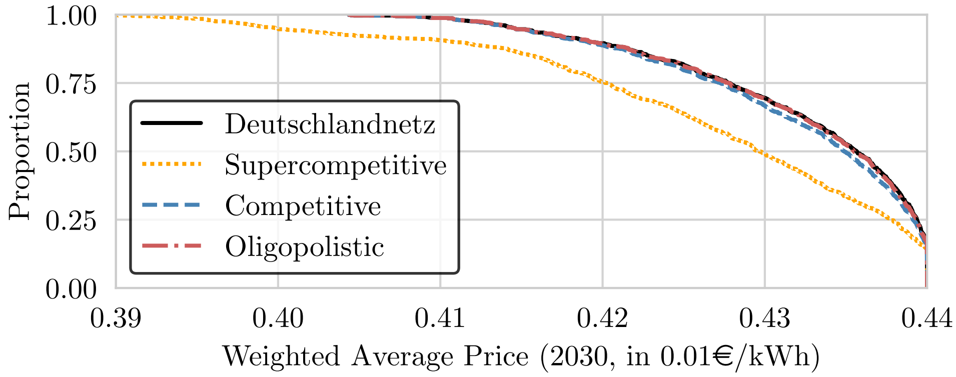

Figure 5.12: Empirical complementary cumulative distribution function of weighted average prices (with respect to price tiers) under different competitive settings, calculated with $s = 0.001$. Prices are highest in the Deutschlandnetz setting in the sense of first-order stochastic dominance.

are higher in the *Deutschlandnetz* setting compared to the three counterfactuals. Though counterintuitive, we propose that the higher average prices in the *Deutschlandnetz* setting stem from its increased number of monopoly edges compared to the other settings: there are 394 (competitive), 426 (oligopolistic), and 432 (*Deutschlandnetz*) monopoly edges, respectively. That is, despite the *Deutschlandnetz* setting having more total suppliers than our oligopolistic counterfactual, it actually contains more monopoly edges, resulting in higher prices. Nevertheless, the most significant difference remains between the supercompetitive setting and all other scenarios.

The differences between the supercompetitive scenario and the other settings that could have materialized raise questions about the effectiveness of the tender rules designed to promote competition. We find that the similarities among the three possible settings primarily result from the tender design, which prevents CPOs from winning multiple lots within the same region. Thus, even though both the *Deutschlandnetz* and oligopolistic settings contain more monopoly edges than the competitive setting, the dispersion is fairly low, particularly compared to the network's total of 2,616 edges in the network. It is still notable, however, that (aspiring) CPOs strategically targeted locations in a way that maximizes their market power ex-post, given the constraints imposed by the tender rules.

Comparison of Consumer Welfare

To compare consumer welfare under different settings, we apply a price cap of 0.44 €/kWh (i.e., the originally proposed price cap in the tender design) everywhere and simulate consumer charging behavior accordingly. Figure 5.13 shows the consumer cost under different competitive settings.

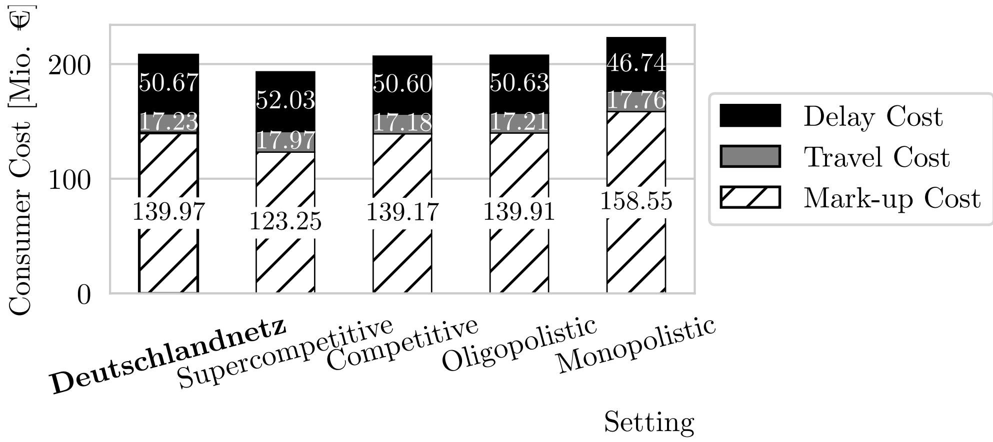

Figure 5.13: Comparison of consumer cost under different competitive settings (2030).

As expected, the three feasible settings (i.e., *Deutschlandnetz*, competitive, and oligopolistic) yield similar levels of consumer costs. In the supercompetitive setting, markup costs are notably lower, whereas in the monopolistic setting, they are substantially higher. Consumer cost is approximately 7% lower in the feasible settings compared to the monopolistic setting. The supercompetitive setting reduces cost by an additional 7% compared to the feasible settings. Overall, from a consumer perspective, a more competitive regime could potentially deliver up to 14% savings. Consistent with our findings on price distribution, the *Deutschlandnetz* setting actually results in the highest consumer cost among the feasible scenarios, albeit by a small margin compared to the oligopolistic setting. This is primarily driven by increased markup costs, reflecting CPOs' enhanced ability to set higher prices.

Taken together, our case study indicates that under the current tender rules, local monopolies may emerge and influence prices compared to a theoretical benchmark with 900 different CPOs. The exact set of winners within these tender rules, however, only weakly impacts equilibrium pricing. Moreover, all investigated market configurations that could materialize under the tender rules are substantially more competitive than a monopolistic setting. Nevertheless, the actual set of winners of the tender features the least competitive outcome among the possible scenarios that could have materialized. Thus, while the *Deutschland-netz* tender rules have largely achieved their competitiveness objectives, CPOs appear to have strategically bid to maximize their surplus within the given rules.

5.5 Conclusion

In this paper, we study the competitive dynamics in fast-charging networks through a novel computational framework, where we capture both the key strategic factors associated with decision-makers (i.e., CPOs and consumers) in these networks and the idiosyncrasies that may arise from government interventions (e.g., price caps, geolocation constraints). Our analysis reveals several insights. From the consumer's perspective, we find that prices are negatively correlated with the local density of charging stations but largely independent of absolute demand levels. For both current and future CPOs, the results from our case study indicate that profitability hinges on both competitiveness and demand: high-demand areas with lower prices can be highly profitable due to strong utilization, while low-demand areas with higher prices often yield weaker returns. From the regulator's perspective, our findings suggest that imposing a uniform, stringent price cap may deter infrastructure investment, particularly in less profitable areas. Furthermore, we find that spatial and ownership restrictions in tender rules may fall short of fostering fully competitive markets in large-scale charging networks.

Our study contributes to the understanding of the emerging EV charging market. While the design and operational challenges associated with fast-charging networks have attracted increasing attention from academics across multiple disciplines, existing studies have mostly treated the network and ownership structures and operational decisions (e.g., pricing) in isolation, which may miss nuanced aspects of the strategic interactions between different stakeholders. Our computational framework accounts for the interplay between network characteristics and economic forces, and hence enables novel insights into the competitive dynamics of fast-charging networks.

From the practical perspective, the contribution of our study is two-fold. First, our findings provide actionable insights into CPO's optimal pricing strategies. Specifically, we show that optimal prices are largely dependent on the regional structure of demand rather than the absolute demand levels. Such findings help to simplify CPOs' pricing decisions, since compared to the absolute demand levels, the relative demand levels (between regions) are easier to predict. Second, our study provides much-needed guidance for policymakers involved in infrastructure planning. In particular, the results from our counterfactual experiments indicate that (1) ensuring local competitiveness is critical to the transition to electric mobility at large, and (2) price caps should be carefully gauged to balance consumer protection against exploitative pricing while preserving sufficient incentives for network expansion and operational efficiency.

It is worth noting that although the current study is specifically geared towards the emerging fast-charging market, our framework can be applied to study other competition with a networked structure of product differentiation. For example, consider restaurants collaborating with on-demand food delivery platforms such as Uber Eats or DoorDash. In this context, these establishments must account for consumers' inconvenience related to travel for pickup or waiting for delivery when setting prices, as they compete with neighboring restaurants for customers. Take cloud computing infrastructure as another example. A consumer may choose between different data centers at different locations. Given the varying distance, there is heterogeneous travel (latency) cost between data centers. Data centers, having different marginal costs due to, e.g., varying electricity prices, are anticipating consumers' decision-making and setting prices for their services accordingly.

In closing, we note that our current work bears several limitations. To start with, while the proposed spatial competition model already accounts for several complications arising from real-world settings, it is still a relatively rough approximation of a true two-dimensional competitive environment. For example, we focus on regional charging networks rather than en-route charging during long trips. Extending our model to capture both types of demand would be a fruitful direction for future work. In addition, we (implicitly) assume a constant travel speed across the network such that travel cost can be characterized by travel distance. To capture more nuanced characteristics of a real-world EV charging market, a *planar* analysis might be helpful, although it would require high-fidelity data (e.g., travel speed data) and more extensive use of numerical techniques, as the resulting demand and price would most likely be analytically intractable. Next, we choose to focus on pure-strategy equilibria in the theoretical analysis based on practical considerations—it is unclear how exactly CPOs would employ mixed strategies (e.g., stochastic pricing) and how consumers would react to such a *lottery*. Still, analyzing characteristics of mixed-strategy equilibria (which exist under much more lenient requirements) would be an interesting direction to explore in the future. Finally, in the case study of *Deutschlandnetz*, we assume homogeneous cost for all CPOs and thus abstract from economies of scale in charging point operation. As cost mostly depends on the remuneration payment and energy cost, this approximation is valid. However, it is conceivable that eventually, larger CPOs can profit from economies of scale regarding various aspects of operation. Incorporating the heterogeneity of CPOs' operational expenses into the analysis of competitive pricing will lead to more fruitful results.

5.6 Appendix

5.6.1 Proofs

Proof of Lemma 1.

Total cost at source vertex $\mathrm{src}(e)$ and destination vertex $\mathrm{dst}(e)$ of edge e, respectively, for an infinitesimal consumer located at x are given by

$$C^t_{\mathrm{src}(e)}(x) = x\tau + x^t_e f^t_e \delta \kappa^{-1}_{\mathrm{src}(e)} + q\, p^t_{\mathrm{src}(e)} - a_{\mathrm{src}(e)}$$
$$C^t_{\mathrm{dst}(e)}(x) = (L_e - x)\tau + (L_e - x^t_e) f^t_e \delta \kappa^{-1}_{\mathrm{dst}(e)} + q\, p^t_{\mathrm{dst}(e)} - a_{\mathrm{dst}(e)},$$

assuming that x^t_e is the indifferent infinitesimal consumer. Equating cost to find the consistent indifferent consumer results in

$$x^t_e = \frac{-q\Delta p^t_e + \Delta a_e + L_e(\tau + f^t_e \delta \kappa^{-1}_{\mathrm{dst}(e)})}{2\tau + f^t_e \delta \left(\kappa^{-1}_{\mathrm{src}(e)} + \kappa^{-1}_{\mathrm{dst}(e)}\right)}.$$

Here, for $\Delta p^t_e < \frac{\Delta a_e - L_e\left(\tau + f^t_e \delta \kappa^{-1}_{\mathrm{src}(e)}\right)}{q}$ or $\Delta p^t_e > \frac{\Delta a_e + L_e\left(\tau + f^t_e \kappa^{-1}_{\mathrm{dst}(e)}\right)}{q}$, respectively, $x^t_e \notin [0, L_e]$. As the indifferent consumer must not lie outside edge e, we account for this saturation by defining the indifferent consumer as

$$x^{t*}_e = \min\left(L_e, \max\left(\frac{-q\Delta p^t_e + \Delta a_e + L_e(\tau + f^t_e \delta \kappa^{-1}_{\mathrm{dst}(e)})}{2\tau + f^t_e \delta \left(\kappa^{-1}_{\mathrm{src}(e)} + \kappa^{-1}_{\mathrm{dst}(e)}\right)}, 0\right)\right),$$

which is a piece-wise linear function of Δp^t_e.

Proof of Proposition 1.

As stated in the proof of lemma 1, the indifferent consumer is a piecewise linear function of Δp^t_e. Its derivative regarding Δp^t_e is therefore piecewise constant outside the kinks, where the derivative is not defined. The kinks of x^{t*}_e are located at $\Delta p^t_e = \frac{\Delta a_e - L_e\left(\tau + f^t_e \delta \kappa^{-1}_{\mathrm{src}(e)}\right)}{q}$ and $\Delta p^t_e = \frac{\Delta a_e + L_e\left(\tau + f^t_e \kappa^{-1}_{\mathrm{dst}(e)}\right)}{q}$, respectively. Therefore, we obtain the derivative of the position of the indifferent consumer on edge e, x_e^{t*}, regarding Δp_e^t as

$$\frac{\mathrm{d}x_e^{t*}}{\mathrm{d}\Delta p_e^t} = \begin{cases} -\frac{q}{2\tau + f_e^t\delta\left(\kappa_{\mathrm{src}(e)}^{-1} + \kappa_{\mathrm{dst}(e)}^{-1}\right)}, & \text{for } \Delta p_e^t \in \left(\frac{\Delta a_e - L_e\left(\tau + f_e^t\delta\kappa_{\mathrm{src}(e)}^{-1}\right)}{q}, \frac{\Delta a_e + L_e\left(\tau + f_e^t\kappa_{\mathrm{dst}(e)}^{-1}\right)}{q}\right) \\ 0, & \text{for } \Delta p_e^t \notin \left[\frac{\Delta a_e - L_e\left(\tau + f_e^t\delta\kappa_{\mathrm{src}(e)}^{-1}\right)}{q}, \frac{\Delta a_e + L_e\left(\tau + f_e^t\kappa_{\mathrm{dst}(e)}^{-1}\right)}{q}\right] \end{cases} \quad \forall v, e.$$

Demand at source and destination vertex of edge e are defined as $D_{\mathrm{src}(e),e}^{t*} = qf_e^t x_e^{t*}$ and $D_{\mathrm{dst}(e),e}^{t*} = qf_e^t(L_e - x_e^{t*})$, respectively. The derivatives of $D_{\mathrm{src}(e),e}^{t*}$ and $D_{\mathrm{dst}(e),e}^{t*}$ are

$$\begin{aligned} \frac{\mathrm{d}D_{\mathrm{src}(e),e}^{t*}}{\mathrm{d}p_{\mathrm{src}(e)}^t} &= \frac{\mathrm{d}D_{\mathrm{src}(e),e}^{t*}}{\mathrm{d}x_e^{t*}} \frac{\mathrm{d}x_e^{t*}}{\mathrm{d}\Delta p_e^t} \frac{\mathrm{d}\Delta p_e^t}{\mathrm{d}p_{\mathrm{src}(e)}^t} = qf_e^t \frac{\mathrm{d}x_e^{t*}}{\mathrm{d}\Delta p_e^t} \\ \frac{\mathrm{d}D_{\mathrm{dst}(e),e}^{t*}}{\mathrm{d}p_{\mathrm{dst}(e)}^t} &= \frac{\mathrm{d}D_{\mathrm{dst}(e),e}^{t*}}{\mathrm{d}x_e^{t*}} \frac{\mathrm{d}x_e^{t*}}{\mathrm{d}\Delta p_e^t} \frac{\mathrm{d}\Delta p_e^t}{\mathrm{d}p_{\mathrm{dst}(e)}^t} = -qf_e^t \frac{\mathrm{d}x_e^{t*}}{\mathrm{d}\Delta p_e^t}(-1). \end{aligned}$$

As both derivatives yield the same expression, we only have to determine $\frac{\mathrm{d}D_{v,e}^{t*}}{\mathrm{d}p_v} = qf_e^t \frac{\mathrm{d}x_e^{t*}}{\mathrm{d}\Delta p_e^t}$:

$$\frac{\mathrm{d}D_{v,e}^{t*}}{\mathrm{d}p_v^t} = \begin{cases} -\frac{q^2 f_e^t}{\tau + f_e^t\delta\left(\kappa_{\mathrm{src}(e)}^{-1} + \kappa_{\mathrm{dst}(e)}^{-1}\right)}, & \text{for } \Delta p_e^t \in \left(\frac{\Delta a_e - L_e\left(\tau + f_e^t\delta\kappa_{\mathrm{src}(e)}^{-1}\right)}{q}, \frac{\Delta a_e + L_e\left(\tau + f_e^t\kappa_{\mathrm{dst}(e)}^{-1}\right)}{q}\right) \\ 0, & \text{for } \Delta p_e^t \notin \left[\frac{\Delta a_e - L_e\left(\tau + f_e^t\delta\kappa_{\mathrm{src}(e)}^{-1}\right)}{q}, \frac{\Delta a_e + L_e\left(\tau + f_e^t\kappa_{\mathrm{dst}(e)}^{-1}\right)}{q}\right] \end{cases} \quad \forall v, e.$$

This completes the proof.

Proof of Proposition 2

To prove Proposition 2, we only need a counterexample. Consider two three-vertex networks, A and B. Let $f_e, L_e = 1\ \forall e,\ p_{cap} = 100, \tilde{\tau} = 10,\ c_A = (50, 60, 70)^T$ and $c_B = (10, 60, 70)^T$, respectively. Figure 5.14 shows equilibrium plots.

Here, the horizontal axis shows p_1, the price of the first vertex. On the vertical axis, the black lines represent the best-response equilibrium of CPO 2 (dashed) and CPO 3 (dash-dotted), i.e., the prices p_2^* and p_3^* that are in equilibrium when p_1 is given exogenously. The solid red line shows CPO 1's best response to the corresponding (p_2^*, p_3^*) pair. The dotted red line is the identity function. An intersection of the two red lines (i.e., the best response and identity) is a fixed point, and thus a global Nash equilibrium. As an example for the competitive

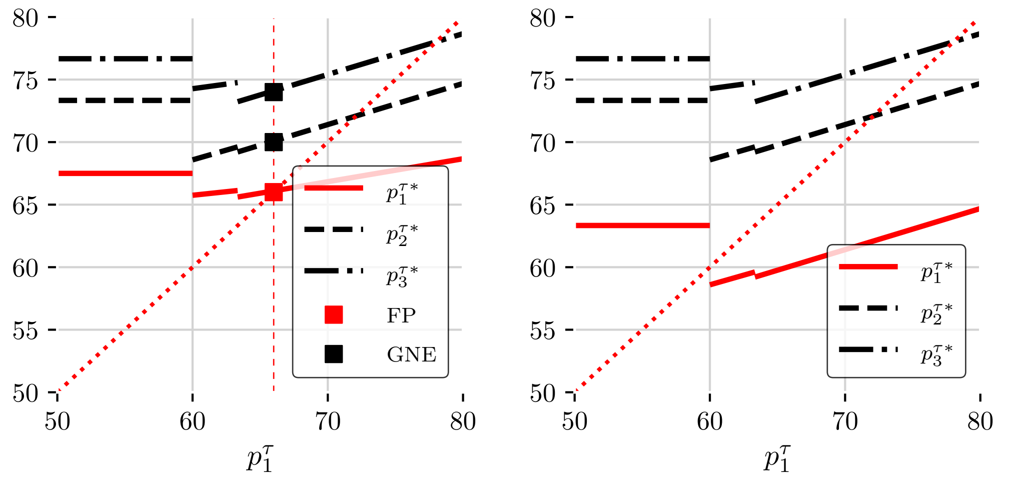

(i) Conditions (C3) and (T3) are satisfied. p_1^* intersects the identity function at the global Nash equilibrium $(p_1^*, p_2^*, p_3^*)^T = (66, 70, 74)^T$.

(ii) Condition (C3) is not satisfied, condition (T3) is. p_1^* and the identity function do not intersect; there is thus no global Nash equilibrium.

Figure 5.14: Equilibrium plots of two three-vertex networks, a (left) and b (right).

regimes derived for two-vertex networks in Table 5.2, see $p_1 < 60$. Here, CPOs 2 and 3 concede on edges connecting to CPO 1, as they cannot fruitfully compete on these; the partial equilibrium p_2^*, p_3^* is thus independent of p_1. Clearly, we can see that while network A has a global Nash equilibrium, network B does not.

5.6.2 Supplementary Results to Section 3

Globally Aware Consumers

In our base model, we assume that consumers are locally aware in estimating the waiting time without the cognitive overhead of a complete assessment of arrivals from all edges. Departing from this assumption, here, we consider the consumers' globally aware delay estimation of waiting time. Indifferent consumers with a globally aware delay estimation are characterized as follows:

$$
\begin{aligned}
&x_e^t \tau + \delta \kappa_{\mathrm{src}(e)}^{-1} \sum_{e' \in \mathcal{E}_{\mathrm{src}(e)}} \frac{D^t_{\mathrm{src}(e),e'}(\mathbf{p}^t)}{q} + q\, p^t_{\mathrm{src}(e)} \\
=&(L_e - x_e^t)\tau + \delta \kappa_{\mathrm{dst}(e)}^{-1} \sum_{e' \in \mathcal{E}_{\mathrm{dst}(e)}} \frac{D^t_{\mathrm{dst}(e),e'}(\mathbf{p}^t)}{q} + q\, p^t_{\mathrm{dst}(e)} \qquad \forall e \in E.
\end{aligned}
\tag{5.2}
$$

In Eq. (5.2), we highlight that demand $D_{v,e}^t$ from edge e at vertex v does not only depend on the price differential Δp_e^t on that edge, but on all prices $\mathbf{p}^t$ in the network. When there is an interior solution to Eq. (5.2), we can rewrite Eq. (5.2) as a linear system of equations for $(\mathbf{p}^t, \mathbf{x})$:

$$\begin{aligned} & x_e^t \tau + \delta \kappa_{\mathrm{src}(e)}^{-1} \left(\sum_{e' \in \mathcal{E}^-_{\mathrm{src}(e)}} f_{e'}^t x_{e'}^t + \sum_{e' \in \mathcal{E}^+_{\mathrm{src}(e)}} f_{e'}^t (L_{e'} - x_{e'}^t) \right) + q\, p_{\mathrm{src}(e)}^t - a_{\mathrm{src}(e)} \\ =& (L_e - x_e^t)\tau + \delta \kappa_{\mathrm{dst}(e)}^{-1} \left(\sum_{e' \in \mathcal{E}^-_{\mathrm{dst}(e)}} f_{e'}^t x_{e'}^t + \sum_{e' \in \mathcal{E}^+_{\mathrm{dst}(e)}} f_{e'}^t (L_{e'} - x_{e'}^t) \right) + q\, p_{\mathrm{dst}(e)}^t - a_{\mathrm{dst}(e)} \\ & \forall e \in E. \end{aligned} \tag{5.3}$$

Let $\mathbf{x}^{t*}(\mathbf{p}^t) = (x_e^{t*})$ be the solution to Eq. (5.3). We have $D_{\mathrm{src}(e),e}^{t*}(\mathbf{p}^t) = q f_e x_e^{t*}(\mathbf{p}^t)$ and $D_{\mathrm{dst}(e),e}^{t*}(\mathbf{p}^t) = q f_e^t \left(L_e - x_e^{t*}(\mathbf{p}^{\mathbf{t}})\right)$. Each CPO $i \in N$ must now make pricing decisions considering not only the direct impact on adjacent edges of its controlled vertices, but also indirect effects that propagate through the network, which leads to the following decision problem:

$$\begin{aligned} \max_{\mathbf{p}_i^t} \quad & \Pi_i = \sum_t \Pi_i^t = \sum_t \sum_{v \in \mathcal{V}_i} \sum_{e \in \mathcal{E}_v} D_{v,e}^{t*}(\mathbf{p}^t) \Big(p_v^t - c_v^t \Big) \\ \text{s.t.} \quad & \mathbf{p}_i^t \preccurlyeq p_{cap}. \end{aligned} \tag{5.4}$$

The optimization problems defined in Eq. (5.4) $\forall i \in N$ constitute an equilibrium problem (Gabriel et al., 2013), which can be solved with standard techniques using the corresponding linear system arising from the KKT conditions. If $\mathbf{p}^{t*} = (p_v^{t*})_v$ solves Eq. (5.4) $\forall v$ and $\mathbf{x}^{\mathbf{t}*}(\mathbf{p}^t)$ is an interior solution, $\mathbf{p}^{t*}$ is a (global) equilibrium price vector. Nevertheless, both consumer and operator decisions are significantly more complex under the assumption that consumers estimate delays by accounting for price information on all edges. Further, for general prices $\mathbf{p}^t$, $\mathbf{x}^{\mathbf{t}*}(\mathbf{p}^t)$ might be a boundary solution, giving rise to intractable non-linear demand patterns in Eq. (5.2). For the sake of tractability, we adopt the locally aware delay estimation for consumers in the large-scale case study presented in Section 4.

Illustration of the Equilibrium Approximation

To speed up the calculation of local Nash equilibrium, we propose an approximation method that smooths the demand function (see Section 3.4). Figure 5.15 provides an illustration of the asymptotic properties of our method, assuming no delay costs. We can observe that our proposed approximation method provides a close fit to the piecewise-linear demand when s (the smoothing parameter) is sufficiently small.

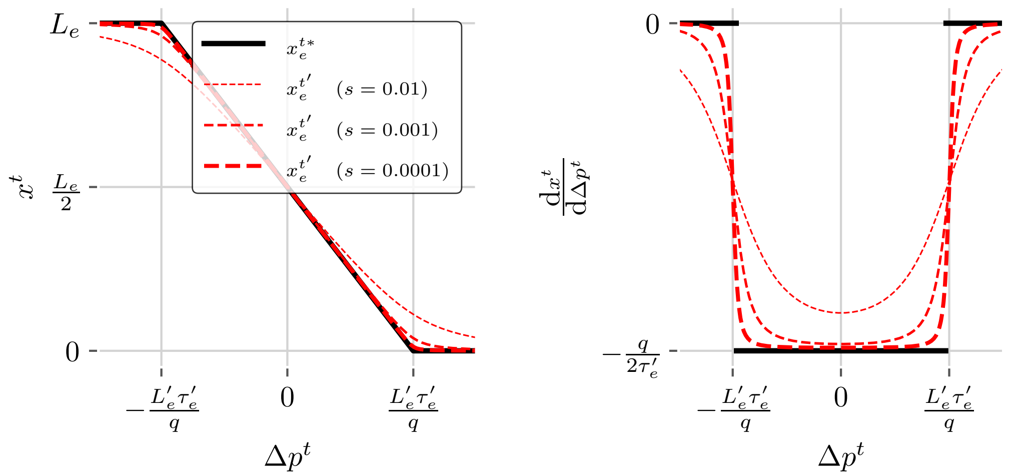

(i) Smoothing of the indifferent consumer x_e on edge e. The largest approximation errors are around the kinks of the exact x_e^*.

(ii) Smoothing of $\frac{\mathrm{d}x_e^t}{\mathrm{d}\Delta p^t}$. The exact x^* is not differentiable at $\Delta p^t = \pm\frac{L'_e\tau'_e}{q}$, elsewhere the approximation is exact for $s \to 0$.

Figure 5.15: Plot of the continuously differentiable demand approximation and its derivative.

5.6.3 Supplementary Results to the Case Study

Sensitivity Analysis

An important driver of equilibrium prices is travel cost. These costs encompass not only the time or fuel expenses incurred for additional travel but also reflect consumers' willingness to visit more distant locations—sometimes beyond what would be strictly economically rational—in order to reduce their charging expenditures. In light of this, we conduct a sensitivity analysis by varying the travel cost parameter. Figure 5.16 presents the equilibrium prices under the competitive setting in 2030 with the delay cost parameter $\delta = 1$ €/kWh.

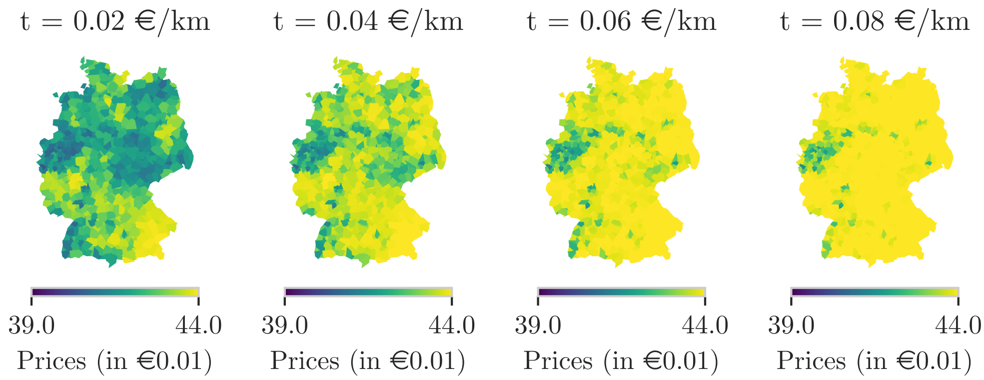

Figure 5.16: Sensitivity analysis of travel costs (Deutschlandnetz setting in 2030).

For higher travel costs, the price cap set by the regulator becomes binding across almost all regions; only in metropolitan areas with the highest charging station densities—where competitive pressure is strongest—can CPOs set prices below the cap. Eastern German regions are among the first where competition breaks down as consumers' willingness to travel vanishes; here, even the increased supply through the tender—compared to free market investments—proves insufficient to overcome the lack of competitiveness.

We conduct a similar set of sensitivity analyses regarding delay cost, see Figure 5.17 for the results. As delay costs increase, prices converge toward the cap. This occurs because extended waiting times at popular, lower-priced charging stations effectively reduce their competitive advantage, reducing the overall competitive pressure on other CPOs in the network.

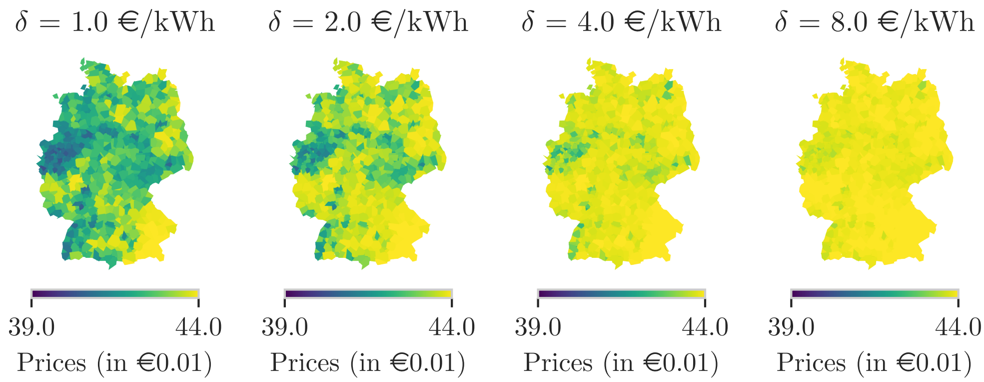

Figure 5.17: Sensitivity analysis of delay costs (Deutschlandnetz setting in 2030).

Bibliography

Abdelwahed, A., van den Berg, P. L., Brandt, T., Collins, J., & Ketter, W. (2020). Evaluating and optimizing opportunity fast-charging schedules in transit battery electric bus networks. *Transportation Science*, *54*(6), 1601–1615.

Agarwal, R., & Ergun, Ö. (2010, 11). Network design and allocation mechanisms for carrier alliances in liner shipping. *Operations Research*, *58*(6), 1726–1742. doi: 10.1287/opre.1100.0848

Agrawal, S., Delage, E., Peters, M., Wang, Z., & Ye, Y. (2011, 6). A Unified Framework for Dynamic Prediction Market Design. *Operations Research*, *59*(3), 550–568. Retrieved from `http://pubsonline.informs.org/doi/abs/10.1287/opre.1110.0922` doi: 10.1287/opre.1110.0922

AlixPartners. (2019). *International electric-vehicle consumer survey.* (Accessed: March 20, 2025)

Alur, R. (2023). *Principles of Cyber-Physical Sysems* (2nd ed.). MIT PRESS.

Andoni, M., Robu, V., Flynn, D., Abram, S., Geach, D., Jenkins, D., ... Peacock, A. (2019, 2). Blockchain technology in the energy sector: A systematic review of challenges and opportunities. *Renewable and Sustainable Energy Reviews*, *100*, 143–174. Retrieved from `https://linkinghub.elsevier.com/retrieve/pii/S1364032118307184` doi: 10.1016/j.rser.2018.10.014

Androulaki, E., Barger, A., Bortnikov, V., Muralidharan, S., Cachin, C., Christidis, K.-n., ... Yellick, J. (2018, 4). Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains. In *Proceedings of the 13th eurosys conference, eurosys 2018* (Vol. 2018-Janua, pp. 1–15). ACM. Retrieved from `https://doi.org/10.1145/3190508.3190538` doi: 10.1145/3190508.3190538

Arrow, K. J., & Debreu, G. (1954, 7). Existence of an Equilibrium for a Competitive Economy. *Econometrica*, *22*(3), 265. Retrieved from `https://www.jstor.org/stable/1907353?origin=crossref` doi: 10.2307/1907353

Artzner, P., Delbaen, F., Eber, J. M., & Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, *9*(3), 203–228. doi: 10.1111/1467-9965.00068

Asensio, O. I., Buckberg, E., Cole, C., Heeney, L., Knittel, C. R., & Stock, J. H. (2025). *Charging uncertainty: Real-time charging data and electric vehicle adoption.*

Ausubel, L. M., & Milgrom, P. (2006). The lovely but lonely Vickrey auction. In P. Cramton, Y. Shoham, & R. Steinberg (Eds.), *Combinatorial auctions* (pp. 17–40). London: The MIT Press.

Avci, B., Girotra, K., & Netessine, S. (2015). Electric vehicles with a battery switching station: Adoption and environmental impact. *Management Science*, *61*(4), 772–794.

Ba, S., Stallaert, J., & Whinston, A. B. (2001a, 9). Optimal Investment in Knowledge Within a Firm Using a Market Mechanism. *Management Science*, *47*(9), 1203–1219. Retrieved from `http://pubsonline.informs.org/doi/abs/10.1287/mnsc.47.9.1203.9781` doi: 10.1287/mnsc.47.9.1203.9781

Ba, S., Stallaert, J., & Whinston, A. B. (2001b, 9). Research Commentary: Introducing a Third Dimension in Information Systems Design - The Case for Incentive Alignment. *Information Systems Research*, *12*(3), 225–239. Retrieved from `http://pubsonline.informs.org/doi/abs/10.1287/isre.12.3.225.9712` doi: 10.1287/isre.12.3.225.9712

Babar, Y., & Burtch, G. (2020). Examining the heterogeneous impact of ride-hailing services on public transit use. *Information Systems Research*, *31*(3), 820–834.

Babar, Y., & Burtch, G. (2024). Recharging retail: Estimating consumer demand spillovers from electric vehicle charging stations. *Manufacturing & Service Operations Management*, *26*(3), 797–813.

Bao, Z., Hu, Z., Kammen, D. M., & Su, Y. (2021). Data-driven approach for analyzing spatiotemporal price elasticities of ev public charging demands based on conditional random fields. *IEEE Transactions on Smart Grid*, *12*(5), 4363–4376.

Bass, F. M. (1969, 1). A New Product Growth for Model Consumer Durables. *Management Science*, *15*(5), 215–227. doi: 10.1287/mnsc.15.5.215

BCG. (2023). *What electric vehicle owners really want from charging networks.* `https://www.bcg.com/publications/2023/what-ev-drivers-expect-from-charging-stations-for-electric-cars`. (Accessed: March 20, 2025)

Bertsekas, D. P. (1992). Auction algorithms for network flow problems: A tutorial introduction. *Computational optimization and applications*, *1*, 7–66.

Bhargava, H., Boehm, J., & Parker, G. G. (2021). How Tesla's charging stations left other manufacturers in the dust. *Harvard Business Review*, *27*.

Bialas, W. F., & Karwan, M. H. (1984). *Two-Level Linear Programming* (Vol. 30; Tech. Rep. No. 8). Retrieved from `https://www.jstor.org/stable/2631591`

Bichler, M., Gupta, A., & Ketter, W. (2010, 12). Designing smart markets. *Information Systems Research*, *21*(4), 688–699. Retrieved from `http://pubsonline.informs.org/doi/abs/10.1287/isre.1100.0316`

Bimpikis, K., Ehsani, S., & İlkılıç, R. (2019). Cournot competition in networked markets. *Management Science*, *65*(6), 2467–2481.

Borenstein, S., Bushnell, J. B., & Wolak, F. A. (2002). Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market. *The American Economic Review*, *92*(5), 1376–1405. Retrieved from `http://www.jstor.org/stable/3083255http://about.jstor.org/terms`

Boyd, S. (2010). Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers. *Foundations and Trends® in Ma-*

chine Learning, *3*(1), 1–122. Retrieved from http://www.nowpublishers.com/article/Details/MAL-016 doi: 10.1561/2200000016

Bygstad, O. H., & Bendik. (2013, 9). The Generative Mechanisms of Digital Infrastructure Evolution. *MIS Quarterly*, *37*(3), 907–931. Retrieved from https://misq.umn.edu/the-generative-mechanisms-of-digital-infrastructure-evolution.html

Carvalho, A. (2020, 3). A permissioned blockchain-based implementation of LMSR prediction markets. *Decision Support Systems*, *130*, 113228. Retrieved from https://linkinghub.elsevier.com/retrieve/pii/S016792361930257X doi: 10.1016/j.dss.2019.113228

Casgrain, P., Ning, B., & Jaimungal, S. (2022, 1). Deep Q-Learning for Nash Equilibria: Nash-DQN. *Applied Mathematical Finance*, *29*(1), 62–78. doi: 10.1080/1350486X.2022.2136727

Chan, T. Y., Padmanabhan, V., & Seetharaman, P. (2007). An econometric model of location and pricing in the gasoline market. *Journal of Marketing Research*, *44*(4), 622–635.

Chen, L., He, L., & Zhou, Y. (2024). An exponential cone programming approach for managing electric vehicle charging. *Operations Research*, *72*(5), 2215–2240.

Chen, M. K., Ingersoll, J. E., & Kaplan, E. H. (2008, 8). Modeling a Presidential Prediction Market. *Management Science*, *54*(8), 1381–1394. Retrieved from http://pubsonline.informs.org/doi/abs/10.1287/mnsc.1080.0872 doi: 10.1287/mnsc.1080.0872

Chen, W., Wei, Z., & Xie, K. (2023). Regulating professional players in peer-to-peer markets: evidence from airbnb. *Management Science*, *69*(5), 2893–2918.

Christidis, K., & Devetsikiotis, M. (2016). Blockchains and smart contracts for the internet of things. *IEEE access*, *4*, 2292–2303.

Clarke, E. H. (1971, 9). Multipart pricing of public goods. *Public Choice*, *11*(1), 17–33. Retrieved from http://link.springer.com/10.1007/BF01726210 doi: 10.1007/BF01726210

Clearwater, S. H. (1996). *Market-Based Control*. WORLD SCIENTIFIC. Retrieved from https://www.worldscientific.com/worldscibooks/10.1142/2741 doi: 10.1142/2741

Clifton, C., Iyer, A., Cho, R., Jiang, W., Kantarcıoğlu, M., & Vaidya, J. (2008, 1). An Approach to Securely Identifying Beneficial Collaboration in Decentralized Logistics Systems. *Manufacturing & Service Operations Management*, *10*(1), 108–125. Retrieved from http://pubsonline.informs.org/doi/abs/10.1287/msom.1070.0167 doi: 10.1287/msom.1070.0167

Colson, B., Marcotte, P., & Savard, G. (2007, 6). An overview of bilevel optimization. *Annals of Operations Research*, *153*(1), 235–256. doi: 10.1007/s10479-007-0176-2

Constantinides, P., Henfridsson, O., & Parker, G. G. (2018). Introduction—platforms and infrastructures in the digital age. *Information Systems Research*, *29*(2), 381–400.

Cramton, P., Geddes, R. R., & Ockenfels, A. (2018, 7). Set road charges in real time to ease traffic. *Nature*, *560*(7716), 23–25. Retrieved from https://www.nature.com/articles/d41586-018-05836-0 doi: 10.1038/d41586-018-05836-0

Crönert, T., & Minner, S. (2021). Location selection for hydrogen fuel stations under emerging provider competition. *Transportation Research Part C: Emerging Technologies*, *133*, 103426.

De Berg, M. (2000). *Computational geometry: Algorithms and applications*. Springer Science & Business Media.

Department for Transport. (2019). National Travel Survey: England 2018. *Transport*(July 2019), 1–34. Retrieved from https://assets.publishing.service.gov.uk/government/uploads/

system/uploads/attachment_data/file/823068/national-travel-survey-2018.pdfhttps://www.gov.uk/government/collections/national-travel-survey-statistics.

Drezner, Z., & Eiselt, H. (2024, July). Competitive location models: A review. *European Journal of Operational Research*, *316*(1), 5–18. Retrieved from http://dx.doi.org/10.1016/j.ejor.2023.10.030 doi: 10.1016/j.ejor.2023.10.030

Dvorkin, Y. (2019, 12). A Chance-Constrained Stochastic Electricity Market. *IEEE Transactions on Power Systems*, *35*(4), 2993–3003. doi: 10.1109/tpwrs.2019.2961231

Fan, M., Stallaert, J., & Whinston, A. B. (2003, 3). Decentralized Mechanism Design for Supply Chain Organizations Using an Auction Market. *Information Systems Research*, *14*(1), 1–22. Retrieved from http://pubsonline.informs.org/doi/abs/10.1287/isre.14.1.1.14763 doi: 10.1287/isre.14.1.1.14763

Flath, C., Ilg, J., & Weinhardt, C. (2012). *Decision Support for Electric Vehicle Charging* (Vol. 14; Tech. Rep.). Retrieved from http://aisel.aisnet.org/amcis2012/proceedings/GreenIS/14

Flath, C., Ilg, J. P., Gottwalt, S., Schmeck, H., & Weinhardt, C. (2014). Improving electric vehicle charging coordination through area pricing. *Transportation Science*, *48*(4), 619–634. Retrieved from http://pubsonline.informs.org doi: 10.1287/trsc.2013.0467

Flath, C., Ilg, J. P., & Weinha, C. (2012, 7). Decision support for electric vehicle charging. In *18th americas conference on information systems 2012, amcis 2012* (Vol. 3, pp. 1794–1803). Retrieved from https://aisel.aisnet.org/amcis2012/proceedings/GreenIS/14

Fudenberg, D., & Tirole, J. (1991). *Game Theory.* MIT press.

Funke, S. Á., Plötz, P., & Wietschel, M. (2019). Invest in fast-charging infrastructure or in longer battery ranges? a cost-efficiency comparison for germany. *Applied Energy*, *235*, 888–899.

Gabriel, S. A., Conejo, A. J., Fuller, J. D., Hobbs, B. F., & Ruiz, C. (2013). *Complementarity Modeling in Energy Markets* (Vol. 180). New York, NY: Springer New York. Retrieved from `http://link.springer.com/10.1007/978-1-4419-6123-5` doi: 10.1007/978-1-4419-6123-5

Gaur, V., & Seshadri, S. (2005, 4). Hedging Inventory Risk Through Market Instruments. *Manufacturing & Service Operations Management*, *7*(2), 103–120. Retrieved from `http://pubsonline.informs.org/doi/abs/10.1287/msom.1040.0061` doi: 10.1287/msom.1040.0061

Geng, Y., & Cassandras, C. G. (2013). A model for acceptance of a parking reservation system. *Transportation Research Part C: Emerging Technologies*, *26*, 167–180.

German National Centre for Charging Infrastructure. (2023). *The Deutschlandnetz.* `https://nationale-leitstelle.de/en/foerdern/deutschlandnetz/`. (Accessed: March 20, 2025)

German National Centre for Charging Infrastructure. (2024). *OBELISöffentlich: Stamm- und Betriebsdaten geförderter öffentlich zugänglicher Ladestationen für Elektrofahrzeuge.* `https://mobilithek.info/offers/714073450865197056`.

Gottwalt, S., Ketter, W., Block, C., Collins, J., & Weinhardt, C. (2011). Demand side management-A simulation of household behavior under variable prices. *Energy Policy*, *39*, 8163–8174. Retrieved from `www.bundesregierung.de/Webs/Breg/DE/Energiekonzept/energiekonzept.` doi: 10.1016/j.enpol.2011.10.016

Groves, T. (1973, 7). Incentives in Teams. *Econometrica*, *41*(4), 617. Retrieved from `https://www.jstor.org/stable/1914085?origin=crossref` doi: 10.2307/1914085

Guo, Z., Deride, J., & Fan, Y. (2016). Infrastructure planning for fast charging stations in a competitive market. *Transportation Research Part C: Emerging Technologies*, *68*, 215–227.

Guo, Z., Koehler, G. J., & Whinston, A. B. (2007, 8). A Market-Based Optimization Algorithm for Distributed Systems. *Management Science*, *53*(8), 1345–1358. Retrieved from `http://pubsonline.informs.org/doi/abs/10.1287/mnsc.1060.0690` doi: 10.1287/mnsc.1060.0690

Guo, Z., Koehler, G. J., & Whinston, A. B. (2012, 9). A Computational Analysis of Bundle Trading Markets Design for Distributed Resource Allocation. *Information Systems Research*, *23*(3-part-1), 823–843. Retrieved from `http://pubsonline.informs.org/doi/abs/10.1287/isre.1110.0366` doi: 10.1287/isre.1110.0366

Gupta, J. K., Egorov, M., & Kochenderfer, M. (2017). Cooperative Multi-Agent Control Using Deep Reinforcement Learning. In *International conference on autonomous agents and multiagent systems (aamas).* Retrieved from `http://ala2017.it.nuigalway.ie/papers/ALA2017_Gupta.pdf`

Hardman, S., Jenn, A., Tal, G., Axsen, J., Beard, G., Daina, N., … others (2018). A review of consumer preferences of and interactions with electric vehicle charging infrastructure. *Transportation Research Part D: Transport and Environment*, *62*, 508–523.

Hastig, G. M., & Sodhi, M. S. (2020). Blockchain for supply chain traceability: Business requirements and critical success factors. *Production and Operations Management*, *29*(4), 935–954.

He, L., Hu, Z., & Zhang, M. (2020). Robust repositioning for vehicle sharing. *Manufacturing & Service Operations Management*, *22*(2), 241–256.

He, L., Ma, G., Qi, W., & Wang, X. (2021). Charging an electric vehicle-sharing fleet. *Manufacturing & Service Operations Management*, *23*(2), 471–487.

Healy, P. J., Linardi, S., Lowery, J. R., & Ledyard, J. O. (2010, 11). Prediction Markets: Alternative Mechanisms for Complex Environments with Few

Traders. *Management Science*, *56*(11), 1977–1996. Retrieved from `http://pubsonline.informs.org/doi/abs/10.1287/mnsc.1100.1226` doi: 10.1287/mnsc.1100.1226

Hogan, W. W. (1992, 9). Contract networks for electric power transmission. *Journal of Regulatory Economics*, *4*(3), 211–242. doi: 10.1007/BF00133621

Hotelling, H. (1929). Stability in competition. *The Economic Journal*, *39*(153), 41–57.

Houde, J.-F. (2012). Spatial differentiation and vertical mergers in retail markets for gasoline. *American Economic Review*, *102*(5), 2147–2182.

Houghtalen, L., Ergun, O., & Sokol, J. (2011). Designing mechanisms for the management of carrier alliances. *Transportation Science*, *45*(4), 465–482. doi: 10.1287/trsc.1100.0358

Hu, J., & Wellman, M. P. (2003). Nash Q-Learning for General-Sum Stochastic Games. *Journal of Machine Learning Research*, *4*, 1039–1069.

Hu, X., & Ralph, D. (2007, 10). Using EPECs to model bilevel games in restructured electricity markets with locational prices. *Operations Research*, *55*(5), 809–827.

Huang, S., Wu, Q., Oren, S. S., Li, R., & Liu, Z. (2015). Distribution Locational Marginal Pricing Through Quadratic Programming for Congestion Management in Distribution Networks. *IEEE Transactions on Power Systems*, *30*(4), 2170–2178. doi: 10.1109/TPWRS.2014.2359977

Huang, Y., & Kockelman, K. M. (2020). Electric vehicle charging station locations: Elastic demand, station congestion, and network equilibrium. *Transportation Research Part D: Transport and Environment*, *78*, 102179.

International Energy Agency. (2024). *Global EV Outlook 2024* (Tech. Rep.). Paris: International Energy Agency. Retrieved from `https://www.iea.org/reports/global-ev-outlook-2024` (Accessed: 2024-10-02)

Kabra, A., Belavina, E., & Girotra, K. (2020). Bike-share systems: Accessibility and availability. *Management Science*, *66*(9), 3803–3824.

Kahlen, M., Ketter, W., & van Dalen, J. (2018). Electric vehicle virtual power plant dilemma: Grid balancing versus customer mobility. *Production and Operations Management*, *27*(11), 2054–2070.

Kahlen, M., Schroer, K., Ketter, W., & Gupta, A. (2024). Smart markets for real-time allocation of multiproduct resources: the case of shared electric vehicles. *Information Systems Research*, *35*(2), 871–889.

Karaca, O., & Kamgarpour, M. (2017, 12). Game theoretic analysis of electricity market auction mechanisms. In *2017 ieee 56th annual conference on decision and control (cdc)* (pp. 6211–6216). IEEE. Retrieved from `http://ieeexplore.ieee.org/document/8264596/` doi: 10.1109/CDC.2017.8264596

Karaca, O., & Kamgarpour, M. (2018, 11). *Core-Selecting Mechanisms in Electricity Markets.* Retrieved from `http://arxiv.org/abs/1811.09646`

Ketter, W., Collins, J., Saar-Tsechansky, M., & Marom, O. (2018, 11). Information Systems for a Smart Electricity Grid. *ACM Transactions on Management Information Systems*, *9*(3), 1–22. Retrieved from `https://dl.acm.org/doi/10.1145/3230712` doi: 10.1145/3230712

Ketter, W., Peters, M., Collins, J., & Gupta, A. (2016, 2). A Multiagent Competitive Gaming Platform to Address Societal Challenges. *MIS Quarterly*, *40*(2), 447–460. Retrieved from `https://misq.org/a-multiagent-competitive-gaming-platform-to-address-societal-challenges.html` doi: 10.25300/MISQ/2016/40.2.09

Ketter, W., Schroer, K., & Valogianni, K. (2023). Information systems research for smart sustainable mobility: A framework and call for action. *Information Systems Research*, *34*(3), 1045-1065. Retrieved from `https://doi.org/10.1287/isre.2022.1167` doi: 10.1287/isre.2022.1167

Kim, S., & Park, E. (2024). STAD-GCN: Spatial–temporal attention-based dynamic graph convolutional network for retail market price prediction. *Expert Systems with Applications*, *255*, 124553.

Kouvelis, P., & Li, R. (2019, 10). Integrated Risk Management for Newsvendors with Value-at-Risk Constraints. *Manufacturing & Service Operations Management*, *21*(4), 816–832. Retrieved from `http://pubsonline.informs.org/doi/10.1287/msom.2018.0726` doi: 10.1287/msom.2018.0726

Kraft, D. (1988). A software package for sequential quadratic programming. *Forschungsbericht- Deutsche Forschungs- und Versuchsanstalt fur Luft- und Raumfahrt*.

Kraft, D. (1994). Algorithm 733: Tomp–fortran modules for optimal control calculations. *ACM Transactions on Mathematical Software (TOMS)*, *20*(3), 262–281.

Kraftfahrt-Bundesamt. (2024). *Bestand an Kraftfahrzeugen und Kraftfahrzeuganhängern nach Zulassungsbezirken.*

Krishna, V. (2010). *Auction Theory.* Elsevier Inc. doi: 10.1016/C2009-0-22474-3

LaMonaca, S., & Ryan, L. (2022). The state of play in electric vehicle charging services–a review of infrastructure provision, players, and policies. *Renewable and Sustainable Energy Reviews*, *154*, 111733.

Lamontagne, S., Carvalho, M., Frejinger, E., Gendron, B., Anjos, M. F., & Atallah, R. (2023). Optimising electric vehicle charging station placement using advanced discrete choice models. *INFORMS Journal on Computing*, *35*(5), 1195–1213.

Lee, E. (2008). Cyber Physical Systems : Design Challenges University of California , Berkeley. *Distributed Computing*, 363–369. Retrieved from `http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=4519604` doi: 10.1109/ISORC.2008.25

Lee, H., & Clark, A. (2018). *Charging the future: Challenges and opportunities for electric vehicle adoption* (HKS Working Paper No. No. RWP18-026).

Li, Z., & Wang, G. (2025). Regulating powerful platforms: Evidence from commission fee caps. *Information Systems Research*, *36*(1), 126–140.

Lu, Y., Gupta, A., Ketter, W., & Van Heck, E. (2019). Dynamic decision making in sequential business-to-business auctions: A structural econometric approach. *Management Science*, *65*(8), 3853–3876.

Ma, G., Lim, M. K., Mak, H.-Y., & Wan, Z. (2019). Promoting clean technology adoption: To subsidize products or service infrastructure? *Service Science*, *11*(2), 75–95.

Ma, S., Zheng, Y., & Wolfson, O. (2013). Real-time city-scale taxi ridesharing. *IEEE Transactions on Knowledge and Data Engineering*, *25*(11), 2646–2659.

Mak, H. Y., Rong, Y., & Shen, Z. J. M. (2013). Infrastructure planning for electric vehicles with battery swapping. *Management Science*, *59*(7). doi: 10.1287/mnsc.1120.1672

Markus, M. L., Silver, M. S., Desanctis, G., Poole, M. S., & Cummings, J. (2008, 10). A Foundation for the Study of IT Effects: A New Look at DeSanctis and Poole's Concepts of Structural Features and Spirit. *Journal of the Association for Information Systems*, *9*(10), 5. Retrieved from `https://aisel.aisnet.org/jais/vol9/iss10/5` doi: 10.17705/1jais.00176

McKinsey. (2021). *The future of EV charging infrastructure: Executive perspectives.* `https://www.mckinsey.com/capabilities/operations/our-insights/the-future-of-ev-charging-infrastructure-executive-perspectives#/`. (Accessed: March 20, 2025)

McKinsey. (2023). *Can public EV fast-charging stations be profitable in the United States?* `https://www.mckinsey.com/features/mckinsey-center-for-future-mobility/our-insights/can-public`

-ev-fast-charging-stations-be-profitable-in-the-united-states. (Accessed: March 20, 2025)

Mengelkamp, E., Gärttner, J., Rock, K., Kessler, S., Orsini, L., & Weinhardt, C. (2017). Designing microgrid energy markets. A case study: The Brooklyn Microgrid. *Applied Energy*, *210*, 870–880. Retrieved from https://doi.org/10.1016/j.apenergy.2017.06.054 doi: 10.1016/j.apenergy.2017.06.054

Mengelkamp, E., Notheisen, B., Beer, C., Dauer, D., & Weinhardt, C. (2018). A blockchain-based smart grid: towards sustainable local energy markets. *Computer Science - Research and Development*, *33*(1-2), 207–214. doi: 10.1007/s00450-017-0360-9

Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., & Riedmiller, M. (2013). *Playing Atari with Deep Reinforcement Learning.* Retrieved from https://arxiv.org/pdf/1312.5602.pdf)http://arxiv.org/abs/1312.5602

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, *518*, 529–533. Retrieved from https://daiwk.github.io/assets/dqn.pdf doi: 10.1038/nature14236

Munizaga, M. A., & Palma, C. (2012). Estimation of a disaggregate multimodal public transport origin–destination matrix from passive smartcard data from santiago, chile. *Transportation Research Part C: Emerging Technologies*, *24*, 9–18.

Myerson, R. B. (2013). *Game Theory.* Harvard University Press. doi: 10.2307/j.ctvjsf522

Nadiminti, R., Mukhopadhyay, T., & Kriebel, C. H. (2002, 12). Research Report: Intrafirm Resource Allocation with Asymmetric Information and Negative Externalities. *Information Systems Research*, *13*(4), 428–434. Retrieved from http://pubsonline.informs.org/doi/abs/10.1287/isre.13.4.428.70 doi: 10.1287/isre.13.4.428.70

Narasimhan, R., & Talluri, S. (2009, 4). *Perspectives on risk management in supply chains* (Vol. 27) (No. 2). doi: 10.1016/j.jom.2009.02.001

Ning, J., & Sobel, M. J. (2018, 2). Production and Capacity Management with Internal Financing. *Manufacturing & Service Operations Management*, *20*(1), 147–160. Retrieved from `http://pubsonline.informs.org/doi/10.1287/msom.2017.0655` doi: 10.1287/msom.2017.0655

OpenStreetMap. (2024). *OpenStreetMap.* `https://wiki.openstreetmap.org/wiki/Key:amenity`.

Orlikowski, W. J., & Robey, D. (1991). Information technology and the structuring of organizations. *Information Systems Research*, *2*(2), 143–169.

Parag, Y., & Sovacool, B. K. (2016, 4). Electricity market design for the prosumer era. *Nature Energy*, *1*(4), 16032. Retrieved from `http://www.nature.com/articles/nenergy201632` doi: 10.1038/nenergy.2016.32

Parker, G. G., Tan, B., & Kazan, O. (2019). Electric power industry: Operational and public policy challenges and opportunities. *Production and Operations Management*, *28*(11), 2738–2777.

Pedersen, A. B., Risius, M., & Beck, R. (2019). Blockchain decision path: when to use blockchains? which blockchains do you mean? *MIS Quarterly Executive*, *18*(2), 3.

Peters, M., Saar-Tsechansky, M., Ketter, W., Williamson, S. A., Groot, P., & Heskes, T. (2018, 6). A scalable preference model for autonomous decision-making. *Machine Learning*, *107*(6), 1039–1068. Retrieved from `http://link.springer.com/10.1007/s10994-018-5705-5` doi: 10.1007/s10994-018-5705-5

Pollitt, M. (2008, 2). The arguments for and against ownership unbundling of energy transmission networks. *Energy Policy*, *36*(2), 704–713. doi: 10.1016/j.enpol.2007.10.011

Powell, M. J. (1970). A new algorithm for unconstrained optimization. In *Nonlinear programming* (pp. 31–65). Elsevier.

Qian, X., & Olsen, T. L. (2020, 5). Operational and Financial Decisions Within Proportional Investment Cooperatives. *Manufacturing & Service Operations Management*, *22*(3), 545–561. Retrieved from `http://pubsonline.informs.org/doi/10.1287/msom.2018.0762` doi: 10.1287/msom.2018.0762

Rajkumar, R., Lee, I., Sha, L., & Stankovic, J. (2010). Cyber-physical systems: The next computing revolution. *Proceedings - Design Automation Conference*, 731–736. doi: 10.1145/1837274.1837461

Ralph, D., & Smeers, Y. (2012). Pricing Risk Under Risk Measures: An Introduction to Stochastic-Endogenous Equilibria. *SSRN Electronic Journal*(August), 1–45. doi: 10.2139/ssrn.1903897

Ralph, D., & Smeers, Y. (2015, 1). Risk Trading and Endogenous Probabilities in Investment Equilibria. *SIAM Journal on Optimization*, *25*(4), 2589–2611. Retrieved from `http://epubs.siam.org/doi/10.1137/110851778` doi: 10.1137/110851778

Ramchurn, S. D., Vytelingum, P., Rogers, A., & Jennings, N. R. (2012). Putting the 'smarts' into the smart grid: A grand challenge for artificial intelligence. *Communications of the ACM*, *55*(4), 86–97. doi: 10.1145/2133806.2133825

Reddy, P., & Veloso, M. (2011). Learned Behaviors of Multiple Autonomous Agents in Smart Grid Markets. *Proceedings of the Twenty-Fifth AAAI Conference on Artificial Intelligence*, 1396–1401. Retrieved from `http://www.aaai.org/ocs/index.php/aaai/aaai11/paper/download/3773/4082`

Robu, V., Gerding, E. H., Stein, S., Parkes, D. C., Rogers, A., & 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. doi: 10.1613/jair.4064

Robu, V., Stein, S., Gerding, E. H., Parkes, D. C., Rogers, A., & Jennings, N. R. (2012). An online mechanism for multi-speed electric vehicle charging. In *Lecture notes of the institute for computer sciences, social-informatics and telecommunications engineering* (Vol. 80 LNICST, pp. 100–112). Retrieved from `http://www.pod-point.com/pod-point-home` doi: 10.1007/978-3-642-30913-7{_}22

Rogers, A., Ramchurn, S. D., & Jennings, N. R. (2012). Delivering the smart grid: Challenges for autonomous agents and multi-agent systems research. In *Proceedings of the national conference on artificial intelligence* (Vol. 3, pp. 2166–2172). Retrieved from `www.aaai.org`

Roth, A. E., & Sotomayor, M. A. O. (1990). *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis.* Cambridge University Press. Retrieved from `https://www.cambridge.org/core/books/twosided-matching/FF14E482078F3DBD8D82A50581C110C0` doi: 10.1017/CCOL052139015X

Sarker, S., Chatterjee, S., Xiao, X., & Elbanna, A. (2019). The sociotechnical axis of cohesion for the is discipline: Its historical legacy and its continued relevance. *MIS Quarterly*, *43*(3), 695–720.

Schneider, K. P., Mather, B. A., Pal, B. C., Ten, C. W., Shirek, G. J., Zhu, H., ... Kersting, W. (2018). Analytic Considerations and Design Basis for the IEEE Distribution Test Feeders. *IEEE Transactions on Power Systems*, *33*(3), 3181–3188. Retrieved from `https://ieeexplore.ieee.org/abstract/document/8063903/` doi: 10.1109/TPWRS.2017.2760011

Schroer, K., Ketter, W., Lee, T. Y., Gupta, A., & Kahlen, M. (2022). Data-driven competitor-aware positioning in on-demand vehicle rental networks. *Transportation Science*, *56*(1), 182–200.

Shann, M., Alan, A., Seuken, S., Costanza, E., & Ramchurn, S. D. (2017). Save money or feel Cozy? A field experiment evaluation of a smart thermostat that learns heating preferences. In *Proceedings of the international joint conference on autonomous agents and multiagent systems, aamas* (Vol. 2, pp. 1008–1016). Retrieved from https://nest.com/support/article/What-is-Rush-Hour-Rewards

Shen, B., Dong, C., & Minner, S. (2022). Combating copycats in the supply chain with permissioned blockchain technology. *Production and Operations Management*, *31*(1), 138–154.

Shi, L., Sethi, S. P., & Çakanyıldırım, M. (2022). Promoting electric vehicles: Reducing charging inconvenience and price via station and consumer subsidies. *Production and Operations Management*, *31*(12), 4333–4350.

Staudt, P., Schmidt, M., Gärttner, J., & Weinhardt, C. (2018, 11). A decentralized approach towards resolving transmission grid congestion in Germany using vehicle-to-grid technology. *Applied Energy*, *230*, 1435–1446. Retrieved from https://linkinghub.elsevier.com/retrieve/pii/S0306261918313540 doi: 10.1016/j.apenergy.2018.09.045

Stone, P., & Veloso, M. (2000, 6). Multiagent systems: a survey from a machine learning perspective. *Autonomous Robots*, *8*(3), 345–383. Retrieved from https://link.springer.com/article/10.1023/A:1008942012299 doi: 10.1023/A:1008942012299/METRICS

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (Second ed.). The MIT Press. Retrieved from http://incompleteideas.net/book/the-book-2nd.html

The Economist. (2021). *Lone Star scars - Why Snowmageddon won't change Texas.* Dallas. Retrieved from https://www.economist.com/united-states/2021/02/23/why-snowmageddon-wont-change-texas

The New York Times. (2023). *Slow Rollout of National Charging System Could Hinder E.V. Adoption.* https://www.nytimes.com/2023/12/23/us/politics/electric-vehicle-chargers-network.html. (Accessed: March 20, 2025)

The White House. (2023). *Full charge: The economics of building a national EV charging network.* https://www.whitehouse.gov/briefing-room/blog/

`2023/12/11/full-charge-the-economics-of-building-a-national-ev-charging-network/`. (Accessed: March 20, 2025)

Valogianni, K., Ketter, W., & Collins, J. (2015). A Multiagent Approach to Variable-Rate Electric Vehicle Charging Coordination. In *International conference on autonomous agents and multiagent systems (aamas)* (Vol. 2, pp. 1131–1139). Retrieved from `www.ifaamas.org`

Valogianni, K., Ketter, W., Collins, J., & Adomavicius, G. (2019). Heterogeneous Electric Vehicle Charging Coordination: A Variable Charging Speed Approach. In *Proceedings of the 52nd hawaii international conference on system sciences (hicss)* (Vol. 6, pp. 3679–3688). Retrieved from `https://hdl.handle.net/10125/59804` doi: 10.24251/hicss.2019.444

Valogianni, K., Ketter, W., Collins, J., & Zhdanov, D. (2020, 4). Sustainable Electric Vehicle Charging using Adaptive Pricing. *Production and Operations Management*, *29*(6), poms.13179. Retrieved from `https://onlinelibrary.wiley.com/doi/abs/10.1111/poms.13179` doi: 10.1111/poms.13179

van Velzen, A., Annema, J. A., van de Kaa, G., & van Wee, B. (2019, 6). Proposing a more comprehensive future total cost of ownership estimation framework for electric vehicles. *Energy Policy*, *129*, 1034–1046. Retrieved from `https://linkinghub.elsevier.com/retrieve/pii/S0301421519301612` doi: 10.1016/j.enpol.2019.02.071

Vázquez-Canteli, J. R., & Nagy, Z. (2019, 2). Reinforcement learning for demand response: A review of algorithms and modeling techniques. *Applied Energy*, *235*(October 2018), 1072–1089. Retrieved from `https://doi.org/10.1016/j.apenergy.2018.11.002` doi: 10.1016/j.apenergy.2018.11.002

Verhoef, E. T., Nijkamp, P., & Rietveld, P. (1996). The economics of regulating road transport. *Edward Elgar Publishing Limited*.

Vickrey, W. (1961, 3). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, *16*(1), 8–37. Retrieved from `http://doi.wiley.com/10.1111/j.1540-6261.1961.tb02789.x` doi: 10.1111/j.1540-6261.1961.tb02789.x

Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., ... Silver, D. (2019, 11). Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature*, *575*(7782), 350–354. Retrieved from `https://www.nature.com/articles/s41586-019-1724-z` doi: 10.1038/s41586-019-1724-z

Vukolić, M. (2016). The quest for scalable blockchain fabric: Proof-of-work vs. BFT replication. In *Lecture notes in computer science (including subseries lecture notes in artificial intelligence and lecture notes in bioinformatics)* (Vol. 9591, pp. 112–125). Springer Verlag. doi: 10.1007/978-3-319-39028-4{_}9

Vytelingum, P., Voice, T. D., Ramchurn, S. D., Rogers, A., & Jennings, N. R. (2010). Agent-based micro-storage management for the smart grid. In *Proceedings of the international joint conference on autonomous agents and multiagent systems, aamas* (Vol. 1, pp. 39–46). Retrieved from `www.ifaamas.org`

Wang, L., Geng, X., Ma, X., Zhang, D., & Yang, Q. (2019, 8). Ridesharing car detection by transfer learning. *Artificial Intelligence*, *273*, 1–18. Retrieved from `https://www.sciencedirect.com/science/article/pii/S0004370219300165` doi: 10.1016/j.artint.2018.12.008

Wang, T., & Wang, R. (2018). A network-city model of spatial competition. *Economics Letters*, *170*, 168–170.

Wang, Y., Currim, F., & Ram, S. (2022). Deep learning of spatiotemporal patterns for urban mobility prediction using big data. *Information Systems Research*, *33*(2), 579–598.

Williams, J. H., Debenedictis, A., Ghanadan, R., Mahone, A., Moore, J., Morrow Iii, W. R., ... Torn, M. S. (2012). The Technology Path to Deep Greenhouse Gas Emissions Cuts by 2050: The Pivotal Role of Electricity. *Science*, *335*, 53–59. Retrieved from `www.sciencemag.orgsciencevolhttp://science.sciencemag.org/`

Wired. (2024). *Tesla's charging network is now open to other EVs—and Ford is first in line.* `https://www.wired.com/story/public-ev-charging-tesla-ford/`. (Accessed: March 20, 2025)

Wood, G. (2014). Ethereum: A secure decentralised generalised transaction ledger. *Ethereum project yellow paper*, *151*(2014), 1–32.

Wörner, A., Tiefenbeck, V., Wortmann, F., Meeuw, A., Ableitner, L., Fleisch, E., & Azevedo, I. (2022, 6). Bidding on a Peer-to-Peer Energy Market: An Exploratory Field Study. *Information Systems Research*. doi: 10.1287/isre.2021.1098

Wu, F., Varaiya, P., Spiller, P., & Oren, S. (1996). Folk Theorems on Transmission Access: Proofs and Counterexamples. *Journal of Regulatory Economics*, *10*, 5–23.

Wu, O. Q., Yücel, Ş., & Zhou, Y. (2022). Smart charging of electric vehicles: An innovative business model for utility firms. *Manufacturing & Service Operations Management*, *24*(5), 2481–2499.

Xu, Z., Hu, Z., Song, Y., & Yuan, Y. (2020). Electric vehicle charging management: A review of the contributions of academia and industry. *IEEE Transactions on Smart Grid*, *11*(3), 2407–2421.

Yang, D., Sarma, N. J., Hyland, M. F., & Jayakrishnan, R. (2021). Dynamic modeling and real-time management of a system of ev fast-charging stations. *Transportation Research Part C: Emerging Technologies*, *128*, 103186.

Yang, H., & Lau, H.-Y. (2015). Optimizing taxi fleet operations using intelligent transportation systems. *IEEE Transactions on Vehicular Technology*, *64*(8), 3640–3649.

Yoo, Y., Boland, R. J., Lyytinen, K., & Majchrzak, A. (2012, 10). Organizing for Innovation in the Digitized World. *Organization Science*, *23*(5), 1398–1408. doi: 10.1287/orsc.1120.0771

YouGov. (2018). *Over nine in ten not working the usual 9-5 week | YouGov.* Retrieved from https://yougov.co.uk/topics/economy/articles-reports/2018/08/24/over-nine-ten-not-working-usual-9-5-week

Yu, J. J., Tang, C. S., Li, M. K., & Shen, Z.-J. M. (2022). Coordinating installation of electric vehicle charging stations between governments and automakers. *Production and Operations Management*, *31*(2), 681–696.

Yuan, W., Huang, J., & Zhang, Y. J. A. (2015). Competitive charging station pricing for plug-in electric vehicles. *IEEE Transactions on Smart Grid*, *8*(2), 627–639.

Zhang, W., & Dou, Y. (2022). Coping with spatial mismatch: Subsidy design for electric vehicle and charging markets. *Manufacturing & Service Operations Management*, *24*(3), 1595–1610.

Zhang, W., Zhang, J., & Li, X. (2011). A taxicab dispatch system based on current demands and real-time traffic conditions. *Transportation Research Part C: Emerging Technologies*, *19*(5), 975–988.

Zhang, Y., Li, B., & Qian, S. (2023). Ridesharing and digital resilience for urban anomalies: Evidence from the new york city taxi market. *Information Systems Research*, *34*(4), 1775–1790.

Zhao, X., Ai, P., Lai, F., Luo, X., & Benitez, J. (2022, 9). Task management in decentralized autonomous organization. *Journal of Operations Management*, *68*(6-7), 649–674. doi: 10.1002/joom.1179

Zheng, Z., Xie, S., Dai, H., Chen, X., & Wang, H. (2017). An Overview of Blockchain Technology: Architecture, Consensus, and Future Trends. In *Proceedings - 2017 ieee 6th international congress on big data, bigdata congress 2017* (pp. 557–564). Retrieved from https://www.researchgate.net/publication/318131748 doi: 10.1109/BigDataCongress.2017.85

Zhu, F., Li, X., Valavi, E., & Iansiti, M. (2021). Network interconnectivity and entry into platform markets. *Information Systems Research*, *32*(3), 1009–1024.

CURRICULUM VITAE

Philipp Artur Kienscherf

PERSONAL DATA

Date of Birth	6th March 1992
Place of Birth	Gütersloh
Nationality	German

EDUCATION

Since 04/2018	**Cologne Institute of Information Systems, University of Cologne** Doctoral Candidate in Information Systems
10/2014 - 09/2017	**RWTH Aachen University** Master of Science (M.Sc.), Business Administration and Engineering: Electrical Power Engineering
09/2016 - 03/2017	**Czech Technical University in Prague** Semester abroad
10/2011 - 09/2014	**RWTH Aachen University** Bachelor of Science (B.Sc.), Business Administration and Engineering: Electrical Power Engineering
07/2011	**Städtisches Gymnasium Harsewinkel** Abitur

WORKING EXPERIENCE

Since 03/2026	**EY-Parthenon** Manager
10/2019 - 02/2026	**Institute of Energy Economics at the University of Cologne (EWI)** Different positions, last as Head of Research Area
04/2018 - 10/2022	**University of Cologne** Research Associate

LANGUAGES

German	Mother tongue
English	Proficient

PUBLICATIONS

Articles in Peer-Reviewed Journals:

- Kaufmann, J., Kienscherf, P. A., & Ketter, W. (2020). Modeling and managing joint price and volumetric risk for volatile electricity portfolios. Energies, 13(14), 3578.
- Kienscherf, P. A., Collins, J., Joe-Wong, C., Ketter, W., & Sen, S. (2020). Time-dependent electricity pricing using variable announcement horizons. Energy Informatics, 3, 1-21.

Working Papers:

- Kienscherf, P. A., Lu, Y., He, L., Ketter, W. (2024). Spatial Competition in Fast-Charging Networks.
- Nussberger, M., Diers, H., Kienscherf, P. A. (2025). A Day Late and a Dollar Short: Intertemporal Revenue Cap Regulation Considering Stranded Assets. *EWI Working Paper* 02/25.

Selected Further Publications:

- Lilienkamp, A., Kienscherf, P. A., Schroer, K., Gierkink, M. (2020). Analyse zukünftiger Elektrofahrzeugnutzung auf Basis von App-Daten. *et – Energiewirtschaftliche Tagesfragen*, Vol. 70 (11), 2020, pp. 53-55.

PRESENTATIONS AND TALKS

- Kienscherf, P. A., Wörner, A., Tiefenbeck, V. & Ketter, W. (2019). "A Blockchain-Based Distributed Mechanism for Resource Allocation." SIG-AIAA Workshop. Munich, Germany, 2019.
- Kienscherf, P. A., Reinhold, N. & and Ketter, W. (2019). "A Blockchain-Based Distributed Mechanism for Resource Allocation." Workshop on Information Technology and Systems (WITS). Munich, Germany, 2019.
- Kienscherf, P. A., Collins, J., Joe-Wong, C., Ketter, W., & Sen, S. (2020). Time-dependent electricity pricing using variable announcement horizons. Energy Informatics 2020. Online.