TECHNISCHE UNIVERSITÄT MÜNCHEN

TUM School of Computation, Information and Technology

Designing Large-Scale Auction Markets

Gregor Schwarz

Vollständiger Abdruck der von der TUM School of Computation, Information and Technology der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitzender: Prof. Dr. Stefanie Rinderle-Ma

Prüfer der Dissertation: 1. Prof. Dr. Martin Bichler
2. Prof. Dr. Rainer Kolisch

Die Dissertation wurde am 15.12.2021 bei der Technischen Universität München eingereicht und durch die TUM School of Computation, Information and Technology am 29.10.2022 angenommen.

Abstract

While classical economic theory uses formal mathematical models to analyze existing markets, market design studies how to set rules in new markets such that goods are allocated efficiently and the room for strategic manipulation is minimized. At the core of designing large-scale auction markets is the development of suitable allocation and pricing rules which determine how goods are distributed among agents and what prices bidders have to pay for the allocated set of objects.

When agents bid truthfully, an ideal auction mechanism provably terminates in a Walrasian equilibrium, i.e., it produces an efficient allocation that clears the market and assigns all bidders a set of goods they desire the most at the given prices. The first contribution of this dissertation surveys conditions on the bidders' value functions that allow for the existence of Walrasian equilibria and shows how clock auction formats leading to such equilibria can be interpreted in the framework of primal-dual algorithms.

Due to complementarities in the bidders' preference relations, Walrasian equilibria generally do not exist in real-world auction markets. Designing suitable allocation and pricing schemes in these settings becomes particularly challenging when the number of bidders and items in the market is large. The exponential number of packages available in combinatorial auctions makes the communication of bids and the computation of the final allocation highly complex. Therefore, devising a compact, domain-specific bid language which allows for both a succinct preference elicitation and the computational tractability of the market plays a central role in market design.

In the second and third research project of this dissertation, we carefully design and analyze bid languages for two large-scale auction markets. First, we consider a wholesale market for road capacity that aims to implement dynamic congestion pricing to alleviate traffic jams in urban areas. As there are tens of thousands of road segments in large cities, the number of goods in the wholesale market is unparalleled. Our novel compact bid language does not only allow bidders to submit their preferences for multiple substitutable routes in a succinct manner but also keeps the market tractable even for a major city like Berlin as we show with a set of extensive numeric experiments. For the computation of market-clearing prices, we draw on pricing techniques commonly used in electricity markets today.

The second large-scale market considered in this dissertation is an auction for allocating electromagnetic spectrum in the United States. In this market, there are 14 license blocks for sale in each of the 406 distinct geographic areas, making it impossible for the 1,000 bidders to state their valuations for all available packages. The *Flexible Use and Efficient Licensing (FUEL)* bid language effectively reduces the communication and computation complexity in such markets. By simulating the market through numerical experiments, we analyze which features of the language allow for a fast computation of the final allocation and compare the market's efficiency to an auction design based on a standard XOR bid language.

Acknowledgements

This dissertation would not have been possible without all the people who supported me along the way. First and foremost, I am very grateful to my advisor Martin Bichler for his continuous support and guidance over the past few years. Thank you, Martin, for all the discussions, helpful comments, and ongoing encouragement! I also consider myself very fortunate for having had to opportunity to work together with Paul Milgrom. Thank you, Paul, for the insightful discussions on the FUEL bid language design and the fruitful collaboration! I also want to express my gratitude to Rainer Kolisch and Stefanie Rinderle-Ma for reviewing this dissertation and completing my examination committee.

Furthermore, I would like to thank my mentor Stefan Waldherr, my co-author and office mate Maximilian Fichtl, and my colleague Johannes Knörr for their valuable comments after reviewing parts of this dissertation. I am also grateful to my former and current colleagues at the Chair of Decision Sciences & Systems for the lively discussions on auction theory and all the other topics we covered along the way. It was a true pleasure working together with you and having you as my colleagues!

I would also like to use this opportunity to thank all members of the ICPC team at TUM, in particular Christian Müller, Stefan Toman, Kirill Martynov, and Gregor Matl. Thanks for getting me hooked up on competitive programming, reviewing my contest problems, supporting me in organizing several contest sites, and competing together as a team in various competitions.

A great thanks also goes to my entire family without whom I would have certainly not gotten to this point. Thank you for your unconditional love and support during the entire course of my education! Last but not least, I want to thank my wonderful fiancée Julia for all the moral support and for being the best partner in life I can possibly imagine.

Contents

1 Introduction **1**

2 Theoretical Background **7**
2.1 Market Model 7
2.2 Bid Languages 9
2.2.1 XOR and OR Bid Language 9
2.3 Design Desiderata for Auctions 11
2.4 Winner Determination Problem 14
2.5 Competitive Equilibria 16
2.5.1 Walrasian Equilibria 17
2.6 Payment Rules 19
2.6.1 Pay-as-Bid 20
2.6.2 VCG Prices 20
2.6.3 Core Pricing 22
2.6.4 IP Pricing 24
2.7 Large-Scale Auction Market Design in Practice 25
2.7.1 Congestion Pricing 26
2.7.2 Spectrum Auctions 30

3 Publication 1: Walrasian Equilibria from an Optimization Perspective **35**

4 Publication 2: A Wholesale Market Design for Road Capacity **55**

5 Publication 3: The FUEL Bid Language **71**

6 Discussion **93**

7 Conclusion **99**

References **100**

1 Introduction

Classical economic theory describes and analyses existing markets using formal mathematical models. Given the set of rules of a market, it aims to predict the behavior of agents as well as the allocation of goods. Market design, on the other hand, is a microeconomic engineering discipline that uses insights from game theory, mechanism design, mathematical optimization, and experimental studies to define rules for new markets. The goal of a market designer is to put regulations into place that allow bidders to state their preferences accurately, prevent strategic manipulations, and lead to welfare-maximizing allocations. According to Vulkan et al. (2013), market design can be understood as both a science and an art. It is a science in the sense that it relies on formal mathematical results from mechanism design, but it is also an art because agents in the field often behave differently than modeled in theory so that practical design decisions have to be made which go beyond theoretical findings.

Two prominent examples which sparked modern market design are the invention of the *Simultaneous Multi Round Auction (SMRA)* for radio spectrum sales by Paul Milgrom and Robert B. Wilson in 1994 (Milgrom, 2021) and the redesign of the *National Resident Matching Program (NRMP)* by Alvin E. Roth in 1995 (Roth and Peranson, 1999). Later incidences where market design played a distinctive role include the organization of kidney exchanges (Roth et al., 2004) and the design of electricity markets (Cramton, 2017). From 1994 until today, the SMRA and its modified versions have been used worldwide for selling electromagnetic spectrum and have raised hundreds of billions of dollars of revenue (Milgrom, 2021). The theory and application of matching mechanisms to kidney exchanges has saved numerous people's lives, while the redesign of the NRMP matches over 20,000 doctors including couples each year (Economic Sciences Prize Committee of the Royal Swedish Academy of Sciences, 2012). The significance of market design is not only reflected by these numbers but has also been recognized by the committee for the prize in economic sciences in memory of Alfred Nobel. In 2012, Alvin E. Roth and Lloyd S. Shapely were awarded the Nobel Memorial Prize in Economics "for the theory of stable allocations and the practice of market design" (Economic Sciences Prize Committee of the Royal Swedish Academy of Sciences, 2012). Paul Milgrom and Robert B. Wilson received the same prize in 2020 "for improvements to auction theory and inventions of new auction formats" (Committee for the Prize in Economic Sciences in Memory of Alfred Nobel, 2020).

This dissertation focuses on the design of large-scale auction markets. A fundamental part of designing such markets is to devise a suitable allocation and pricing rule. Given the bids of the participating bidders, the allocation rule determines which bids are accepted and how goods are distributed among bidders. The prices that bidders have to pay for their allocated set of items are determined by the pricing rule. The ultimate goal for a market designer is to develop an auction mechanism that always terminates in an outcome constituting a competitive equilibrium, i.e., the market is cleared and all bidders receive bundles they desire the most at the given prices. In this case, the first fundamental welfare theorem guarantees the efficiency of such an allocation, meaning that social welfare is maximized. If the respective competitive equilibrium prices are anonymous and linear, then the outcome is called a *Walrasian equilibrium*.

In the first contribution of this dissertation (see Chapter 3) we survey different conditions on the bidders' value functions that admit the existence of Walrasian equilibria. A prominent example of such a condition for multi-item single-unit markets is the *gross substitutes* condition which was originally proposed by Kelso and Crawford (1982). This concept was extended to multi-item multi-unit markets by Milgrom and Strulovici (2009) and is referred to as *strong substitutes* condition in the literature. If all bidders' preferences are strong substitutes, then it can be shown that a particular version of an ascending clock auction by Ausubel (2006) terminates in a Walrasian equilibrium. In our paper, we highlight the connection between this auction design and duality theory by interpreting the clock rounds as iterations of a primal-dual algorithm.

Unfortunately, the restrictions posed on the bidders' value functions which are necessary to guarantee the existence of Walrasian equilibria are so rigid that they often do not hold for practical applications of large-scale combinatorial auction markets. In particular, market designers face the following two challenges.

First, in real-world auction markets, bidders often exhibit synergies between different items. With such complementary value functions, the strong substitutes condition can no longer be fulfilled so that Walrasian equilibrium prices are unlikely to exist in these settings. Much research in market design focuses on the development of methods for computing prices that roughly have the same properties as equilibrium prices, thereby "approximating" equilibrium prices (O'Neill et al., 2005; Gribik et al., 2007; Schiro et al., 2015; O'Neill et al., 2020). Some of these pricing schemes are actively used in electricity markets today (Hytowitz et al., 2020).

Secondly, if the strong substitutes condition can no longer be satisfied, the winner determination problem degenerates to an NP-hard optimization problem (Lehmann et al., 2006). Especially for large-scale auction markets with hundreds of bidders and tens of thousands of goods, the allocation problem quickly becomes intractable, leading to a high computation complexity for the auctioneer. In addition to that, the synergies in the bidders' preference relations make it necessary for bidders to state their valuations not only over individual goods but entire packages of items. In order for the auctioneer to be able to identify the welfare-maximizing allocation, bidders have to submit bids for all bundles that could potentially be part of the winning allocation. As it is impossible for bidders to predict which bids will be accepted by the auctioneer, a naive approach is making bidders state their preferences for all available packages. However, due to the exponential number of available bundles, it is often infeasible for bidders to determine values for all of them. Moreover, even for a smaller subset of bundles, eliciting the bidders' preferences succinctly is difficult and often causes a high communication complexity (Nisan and Segal, 2006).

Both the communication and computation complexity of large-scale combinatorial auctions can be tackled by designing compact bid languages that are restrictive enough to allow for the tractability of the underlying optimization problem but that also use insights from the domain so that bidders are not limited too much in expressing their true valuations. Domain-specific bid languages have been devised for numerous applications, including procurement (Bichler et al., 2011), TV ads (Goetzendorff et al., 2015), and electricity markets (Cramton, 2017). The engineering of domain-specific bid languages is at the core of market design for combinatorial auctions as it is vital for a high efficiency of the market (Bichler, 2017; Milgrom, 2021).

In this dissertation, we consider two examples of large-scale auction markets, a wholesale market for road capacity and an auction design for selling 5G spectrum licenses in the United States, for which we effectively address the two design challenges mentioned above. By devising and analyzing two novel compact bid languages, we significantly tame the communication and computation complexity in these markets. To compute approximate equilibrium prices in the road capacity market, we rely on a pricing technique which has been implemented successfully for electricity markets. In the following, we briefly introduce the two large-scale auction markets considered in this dissertation.

In the second research project of this dissertation (see Chapter 4), we consider a wholesale market for road capacity. The goal of this market is to take action against the mispricing of the traffic's main resource, the road capacity, which has been identified as the primary reason for traffic congestion in urban areas (Cramton et al., 2018; de Palma and Lindsey, 2011). Cities like Singapore, London, and Stockholm have already adopted congestion pricing schemes that charge drivers for entering or driving within the city center. All of these pricing schemes set tolls with respect to a predefined schedule and do not take the actual traffic on the road network into account (Lehe, 2019). However, researchers agree that only dynamic real-time pricing of roads is suitable for charging drivers the actual externalities they impose on others and the environment (Cramton et al., 2019a; Cheng et al., 2017).

Beheshtian et al. (2020) propose a wholesale market for road capacity in analogy to electricity markets where dynamic real-time pricing has already been used for several years. In such markets, an *Independent System Operator (ISO)* first sells road capacity to multiple *Service Providers (SPs)* on a wholesale market consisting of multiple forward and a real-time market. The service providers then resell this capacity to end consumers on a separate retail market. We model the wholesale market as an auction in which the goods are defined as licenses for using specific road segments within particular time slots. Due to complementarities in the bidders' value functions, one cannot hope to find Walrasian equilibrium prices in this market. Therefore, we rely on a pricing technique called *IP Pricing* which produces linear, anonymous market-clearing prices and is commonly used in electricity markets today (O'Neill et al., 2005). A downside of this pricing method is that an outside party may have to pay bidders side-payments in order to prevent them from incurring a loss when winning a bid. While these side-payments can be substantial in electricity markets, we show with an extensive set of numerical experiments for the city of Berlin that these side-payments are negligible in our congestion pricing market.

While the number of bidders in our wholesale market is small, the quantity of goods is unparalleled. In a major city like Berlin, there are more than 34,000 non-residential road segments, each of which is a good on the wholesale market and can potentially be priced. As the service providers need to satisfy a customer demand between origin-destination pairs, they favor a bid language that allows them to specify bids on entire routes. While

Beheshtian et al. (2020) propose bids to be based on individual road segments, we devise a novel compact bid language that allows service providers to state their preferences over multiple substitutable routes between the same origin-destination pair, thereby effectively mitigating the exposure problem of the design by Beheshtian et al. (2020). Our extensive numeric experiments show that with our new bid language the underlying allocation problem is solvable within 15 minutes for a major city like Berlin even though it contains more than 1 million variables and 250 thousand constraints.

The analysis of domain-specific bid languages for large-scale auction markets is also the objective of our third research project of this dissertation (see Chapter 5). In mid 2019, a consortium of satellite providers offered the US Federal Communications Commission to conduct a private auction to sell electromagnetic spectrum in the C-band, which is currently used for broadcasting commercial television, to telecommunication providers. The sold spectrum could then be used by the telecommunication providers to establish a 5G wireless network for end consumers. There were 14 spectrum licenses in 406 different geographical areas up for sale, leading to 15^{406} distinct packages for which more than 1,000 bidders could submit bids in the auction. Nobel laureate Paul Milgrom proposed the *Flexible Use and Efficient Licensing (FUEL)* bid language that aimed to mitigate the missing bids and exposure problem of bidders in this auction (Milgrom, 2019). With a series of extensive numerical experiments, we analyze whether the FUEL bid language leads to a tractable allocation problem and which features of the bid language are critical for finding an allocation quickly. We also compare the efficiency of the FUEL auction design to a classical auction format that features a standard XOR bid language. Even though the US Federal Communications Commission rejected the private auction proposal in February 2020 and instead sold the C-band spectrum in a public auction using its usual standard simultaneous clock auction format (Federal Communications Commission, 2020), the FUEL bid language remains a potential candidate for other markets due to its generic design.

Outline

The remainder of this dissertation is structured as follows: Chapter 2 introduces notation and definitions of standard auction markets. It discusses different design desiderata for auctions and presents several payment rules which are commonly used in auction markets today. Finally, it puts congestion pricing markets and auctions for spectrum sales in their historic and scientific context. Chapter 3 includes the work on the existence of Walrasian equilibria. The second project on designing a wholesale market for road capacity is discussed in Chapter 4. The last project on evaluating the FUEL bid language for a large-scale spectrum auction in the US is presented in Chapter 5. Chapter 6 discusses the results and highlights possible directions for future research and Chapter 7 concludes this dissertation.

2 Theoretical Background

This chapter gives a short introduction to combinatorial auction markets. After introducing the basic notation and defining the economic environment, the main theorems and solution concepts of combinatorial auctions are discussed. This serves as a basis for the subsequent publications listed in this dissertation. The chapter is concluded by highlighting the connections to real-world applications of large-scale combinatorial auction markets which are analyzed in detail in Chapter 4 and 5. The notation as well as the definitions presented in this chapter mainly follow Bichler (2017) and Blumrosen and Nisan (2007).

2.1 Market Model

A combinatorial auction market consists of n *bidders* $i \in \mathcal{I} = \{1, \ldots, n\}$, a single *seller* (or *auctioneer*) indexed with $i = 0$, and m indivisible and heterogeneous *items* (or *goods*) $k \in \mathcal{K} = \{1, \ldots, m\}$. A *bundle* (or *package*) of items is denoted $S \in 2^{\mathcal{K}}$, with $2^{\mathcal{K}}$ representing the power set of $\mathcal{K}$. If there are multiple units available of each item, bundles have to be represented as vectors $x \in \mathbb{Z}^m_{\geq 0}$ or multisets. Unless stated otherwise, we will consider the single-unit case throughout this dissertation, i.e., the auctioneer offers exactly one unit of each item.

Each bidder i in the auction defines a *valuation function*

$$v_i : 2^{\mathcal{K}} \rightarrow \mathbb{R}_{\geq 0} \tag{2.1}$$

which gives the bidder's value for a bundle $S \subseteq \mathcal{K}$. If two bundles $S, T \subseteq \mathcal{K}$ with $S \cap T = \emptyset$ are *complements* for a bidder $i \in \mathcal{I}$, then the bidder's valuation for receiving both bundles is strictly larger than the sum of valuations for being allocated only one of the bundles, i.e., $v_i(S) + v_i(T) < v_i(S \cup T)$. Conversely, two bundles $S, T \subseteq \mathcal{K}$ with $S \cap T = \emptyset$ are *substitutes* of one another if $v_i(S) + v_i(T) > v_i(S \cup T)$.

Throughout this dissertation, we restrict the preferences of all bidders to valuation functions that fulfill the following three standard assumptions:

- *Monotonicity* (or *free disposal*): The valuation function v_i of bidder $i \in \mathcal{I}$ is weakly increasing, i.e., $v_i(\emptyset) = 0$ and $v_i(S) \leq v_i(S')$ for $S \subseteq S'$. In other words, bidders can always dispose additional items for free.
- *Independent private values*: A bidder's valuation function is independent of the values of other bidders, i.e., a bidder's value for a bundle does not change if the bidder learns the values of other bidders or the bundles allocated to them.
- *Quasilinearity*: The utility of bidder i for bundle $S \subseteq \mathcal{K}$ is given by $u_i(S) = v_i(S) - p_i(S)$ where $p_i : 2^{\mathcal{K}} \rightarrow \mathbb{R}_{\geq 0}$ maps a bundle S to the price $p_i(S)$ that bidder $i \in \mathcal{I}$ has to pay for it.

The *outcome* $o = (x, p) \in O$ of an auction consists of an *allocation* $x = (S_1, \ldots, S_n)$ and price functions $p = (p_1, \ldots, p_n)$ so that bidder i is allocated bundle S_i and has to pay $p_i(S_i)$. As bidders may not bid their true valuations, we define a bid function $b_i : 2^{\mathcal{K}} \rightarrow \mathbb{R}_{\geq 0}$ for each bidder $i \in \mathcal{I}$ that maps each bundle $S \subseteq \mathcal{K}$ to the amount that bidder i bids for it in the auction. If bidders bid *truthfully*, then their bids exactly represent their true valuations, i.e., $v_i(S) = b_i(S)$ for all $S \subseteq \mathcal{K}$ and $i \in \mathcal{I}$. Bidders are assumed to behave rationally in the auction, meaning that they aim to maximize their utility. They bid *straightforwardly*, if they submit bids for those bundles that maximize their utility at the given prices.

In some practical applications of combinatorial auctions (e.g., spectrum auctions), auctioneers only want to sell their items if a certain price is exceeded. These so-called *reserve prices* for items can be implemented by letting a dummy bidder submit bids on behalf of the auctioneer which match the auctioneer's reserve prices for the items. Whenever the dummy bidder is allocated an item in the auction, this item remains unsold. Unless stated otherwise, we assume that the auctioneer's value for each item is zero. Under this assumption, the reserve prices of all items are zero so that the dummy bidder is not needed.

2.2 Bid Languages

By placing a *bid* in an auction, bidders can communicate their willingness to pay for an object or even entire packages of objects. The *bid language* specifies the format to which the bids must adhere and also sets logical rules that define which bids may win simultaneously.

If there are m items up for sale, there exists an exponential number of $2^m - 1$ distinct packages for which bidders can submit bids. Already for a small number of items, bidders are unable to determine the value of each package (Parkes, 2006). Even if they could, the *communication complexity* for sharing their bids with the auctioneer would be too high (Nisan and Segal, 2006). In addition to that, the *computation complexity* quickly gets so severe for the auctioneer that the allocation problem becomes intractable. To address the communication and computation complexity, bidders either choose to submit bids for only a fraction of bundles or may be limited by the auctioneer in the number of admissible bids (Kroemer et al., 2016). As bidders are assumed to have a valuation of zero for bundles for which they did not submit a bid, the auction may suffer from substantial efficiency losses as Bichler et al. (2014) show in lab experiments. In the literature, this effect is known as the *missing bids problem*.

To tame the communication and computation complexity, market designers often exploit the domain knowledge on valuation and cost structures of bidders when designing domain-specific compact bid languages. Prominent examples are the design of bid languages for procurement auctions (Bichler et al., 2011; Olivares et al., 2012), TV ad auctions (Goetzendorff et al., 2015), or electricity markets (Cramton, 2017). This dissertation contributes to the line of research on parametric bid languages. In Chapter 4, a novel compact bid language for a congestion pricing wholesale market is proposed, while the communication and computation complexity of a parametric bid language for a spectrum auction in the US is analyzed thoroughly in Chapter 5.

2.2.1 XOR and OR Bid Language

In auction markets with only a single item up for sale, the design of a bid language is straightforward. Each bid defines a bidder's willingness to pay for the object to be sold, and a bidder is allocated the item if her bid is the highest among all submitted bids in the auction. When multiple goods are offered, the same bid language can be used if items are sold one after another in separate single-item auctions. However, such sequential auction procedures can lead to substantial strategic problems for bidders with complementary valuations. This is often the case in large combinatorial auctions such as spectrum auctions where bidders want to win spectrum licenses for multiple geographic areas to benefit from economies of scope (see Chapter 5). If they only win licenses for a subset of the desired regions, their market area may become scattered so that they cannot profit from network effects. In the literature, this situation is known as the *exposure problem* (Ausubel and Milgrom, 2002; Porter et al., 2003). It occurs when bidders end up winning only a subset of items of their desired bundle and the price they have to pay exceeds their valuation for these items. The exposure problem demonstrates the need for more complex bid languages that allow bidders to submit bids for entire packages of items. Auction markets that feature such package bids are called *combinatorial auctions*. Two elementary bid languages for such auctions are the *exclusive-OR (XOR)* and *additive-OR (OR)* bid language.

With an XOR bid language, bidders can submit multiple atomic bids for entire bundles of items but the auctioneer accepts at most one of them per bidder. If a bid becomes winning, the respective bidder is allocated all items specified in the bundle. If it is losing, the bidder is not allocated any of the items. The XOR bid language is *fully expressive* in the sense that it allows bidders to express every possible valuation (Nisan, 2000).

An additive-OR bid language also features atomic package bids but, in contrast to XOR, any non-intersecting combination of a bidder's bids may become winning simultaneously. While the OR bid language might allow bidders to express their bids more succinctly in some cases, this bid language is not fully expressive as it cannot represent valuations with substitutes (Nisan, 2006). A bidder who is interested in winning exactly one of two fully substitutable, non-intersecting bundles cannot express these preferences with an OR bid language. Fujishima et al. (1999) and Nisan (2000) suggest the OR^* bid language that supports such XOR constraints with the help of phantom items. Nisan (2000) shows that there are z OR^* bids and z^2 phantom items needed to represent z bids of an OR or XOR bid language.

Domain-specific compact bid languages for combinatorial auctions usually combine both the XOR and additive-OR bid language. By only allowing certain combinations of OR and XOR bids, the number of potential packages for which bidders can submit bids is reduced substantially. In addition to that, these bid languages often enforce a special bid structure which can be exploited to keep the allocation problem tractable.

2.3 Design Desiderata for Auctions

Auctions can be understood as mechanisms that map the bids of bidders to an outcome. An outcome consists of an allocation of items to bidders as well as the prices that bidders have to pay for the allocated set of goods. In general, one cannot expect bidders to state their true valuations in an auction. Instead, bidders may shade their bids or choose not to disclose their valuations for certain bundles altogether in order to alter the final allocation or minimize their payments. There exists a set of desirable properties that auction mechanisms aim to fulfill. These are briefly reviewed in the following.

An auction mechanism $\mathcal{M}$ maps the bids of the bidders to a respective outcome $o = (x, p)$, i.e., it implements a function $f : B^{\mathcal{I}} \to O$ where B is the set of all bid functions $b : 2^{\mathcal{K}} \to \mathbb{R}_{\geq 0}$, $\mathcal{I}$ is the set of bidders, and O denotes all possible outcomes. In order for the outcome to be realizable, the auction mechanism has to determine a *feasible* allocation, i.e., the auctioneer's supply is sufficient to cover the aggregated demand of all bidders.

Besides that, auctioneers often strive for allocations that maximize *social welfare* which is defined as the sum of utilities of all bidders and the auctioneer. As the auctioneer's value for all items is assumed to be zero, the auctioneer's utility is simply the sum of the prices which the bidders pay for their allocated bundles. Due to the quasi-linearity of the bidders' valuations, the prices occurring in the auctioneer's and bidders' utility functions cancel so that social welfare can be defined equivalently as the sum of all bidders' values for the allocated set of goods, i.e., $\sum_{i \in \mathcal{I}} v_i(S_i)$. Note that bidders may choose not to disclose their true valuations for strategic purposes. In this case, the auctioneer can only maximize the total bid value (i.e., $\sum_{i \in \mathcal{I}} b_i(S_i)$) in the auction. If bidders bid truthfully and disclose their true valuations, this value is equivalent to the social welfare.

Allocations are *efficient* if items are distributed in a way so that social welfare is maximized. Whenever this is not the case, the *(allocative) efficiency* is often considered as a metric for the quality of the allocation.

Definition 2.1 (Allocative Efficiency). *Given the valuations v_i of all bidders $i \in \mathcal{I}$, let $x^* = (S_1^*, \dots, S_n^*)$ be a welfare-maximizing allocation and $x = (S_1, \dots, S_n)$ the allocation produced by the auction mechanism. Then, the allocative efficiency of the outcome is defined as*

$$\frac{\sum_{i \in \mathcal{I}} v_i(S_i)}{\sum_{i \in \mathcal{I}} v_i(S_i^*)}. \tag{2.2}$$

As bidders cannot be forced to reveal their true preferences in an auction, they may bid strategically to influence the resulting allocation and prices. Such strategic manipulations reach from *bid shading* (i.e., submitting lower bids than the own valuation (Bichler, 2017)) over *shill bidding* (i.e., bidding with multiple identities (Sakurai et al., 1999; Yokoo et al., 2004)) to *spiteful bidding* (i.e., placing additional bids that are not intended to become winning but that drive up the payments of competitors (Morgan et al., 2003)) and other collusive behavior (Ausubel et al., 2006b). If an auction mechanism allows for such strategic behavior, this complicates the bidding process significantly as bidders have to spend a large amount of time and effort to learn the bidding strategies of their competitors and adapt their own bidding behavior accordingly. Market designers try to develop auction formats in which strategic bidding is limited or in the best case even impossible. The concept of *strategyproofness* (or *dominant-strategy incentive compatibility*) describes mechanisms where truth-telling is a dominant strategy.

Definition 2.2 (Strategyproofness)**.** *An auction mechanism is strategyproof (or dominant-strategy incentive compatible) if reporting true valuations always leads to a weakly higher utility regardless of the bids of the other bidders, i.e., for all bidders* $i \in \mathcal{I}$ *and bid functions* $(b_1, \ldots, b_i, \ldots, b_n) \in B^{\mathcal{I}}$*,*

$$u_i(f(v_i, b_{-i})) \geq u_i(f(b_i, b_{-i})) \tag{2.3}$$

with b_{-i} *denoting the bid functions of all bidders except for bidder* i *and* $u_i(o)$ *representing the utility of bidder* i *for the outcome* $o \in O$ *computed by the mechanism's function* f*.*

A slightly weaker notion of strategyproofness is *Bayesian-Nash incentive compatibility* which only states that truth-telling leads to a Bayesian Nash equilibrium, i.e., if the other bidders reveal their preferences truthfully, then a bidder can maximize her expected utility by also bidding her true valuations. Stronger notions of strategyproofness such as *group-strategyproofness* consider deviations of entire coalitions of bidders. For formal definitions of these properties we point the interested reader to (Bichler, 2017).

Mechanisms should terminate in an outcome in which the total payments of buyers match the amount that sellers receive. Otherwise, a third party outside the market needs to settle losses by paying market participants side-payments. On the other hand, if bidders pay more money than sellers actually receive, buyers and sellers may refuse to participate in such a market from the beginning. The described property is captured in the notion of *budget-balance*.

Definition 2.3 (Budget-Balance). *A mechanism is budget-balanced if the total payments of all buyers and sellers sum up to zero.*

A typical requirement of bidders for participating in an auction is that they are guaranteed not to incur a loss. This implies that the auction mechanism never terminates in an outcome where a bidder has to pay a higher price for an allocated set of items than originally bid in the auction. If this condition is met, then the mechanism fulfills *individual rationality*.

Definition 2.4 (Individual Rationality). *An auction mechanism is individual rational if for all bid functions $b = (b_1, \ldots, b_n)$ the mechanism produces an outcome $o = (x, p)$ with allocation $x = (S_1, \ldots, S_n)$ and price functions $p = (p_1, \ldots, p_n)$ such that no bidder $i \in \mathcal{I}$ has to pay more for the allocated set of goods than she bid in the auction, i.e., $b_i(S_i) \geq p_i(S_i)$ for all $i \in \mathcal{I}$.*

If bidders have to pay different amounts for the same set of items, then prices are *personalized*. This is often perceived as unfair among bidders as it violates the *law of one price*.

Definition 2.5 (Law of one price). *The law of one price is fulfilled if bidders are charged the same price for the same set of items, i.e., prices are anonymous: for each pair $i, j \in \mathcal{I}$ and for all bundles $S \subseteq \mathcal{K}$ it holds that $p_i(S) = p_j(S)$.*

Besides anonymity, mechanism designers also strive for *linear* (or *item-level*) prices. With a linear pricing scheme, each item $k \in \mathcal{K}$ is assigned an individual price $p(k)$, while the price of a bundle S is given as the sum of the prices of its components, i.e., $p(S) = \sum_{k \in S} p(k)$. Unfortunately, efficient outcomes can generally only be realized with linear, anonymous prices if the bidders' value functions fulfill very restrictive conditions. This will be discussed thoroughly in Chapter 3.

While all of these design desiderata are highly desirable for auctions, they cannot be fulfilled simultaneously in general. For example, Myerson and Satterthwaite (1983) show that for a bilateral trade problem with one buyer and one seller there does not exist a mechanism that simultaneously satisfies budget-balance, efficiency, Bayesian-Nash incentive compatibility, and individual rationality. A broad stream in literature analyzes the conditions under which certain design desiderata are guaranteed to be satisfied (d'Aspremont and Gérard-Varet, 1979; McAfee, 1992; Myerson and Satterthwaite, 1983). In practice, market designers have to consider carefully which of the properties are indispensable for their applications and where compromises can be made.

2.4 Winner Determination Problem

The auctioneer's task of finding a feasible, welfare-maximizing allocation in a single-unit environment with an XOR bid language can be stated as a binary program which is commonly referred to as *Winner Determination Problem (WDP)* in the literature (Bichler, 2017; Blumrosen and Nisan, 2007). For the WDP below, we assume that bidders disclose their true valuations in the auctions so that $v_i(S) = b_i(S)$ for all bidders $i \in \mathcal{I}$ and bundles $S \subseteq \mathcal{K}$.

$$\max \sum_{i \in \mathcal{I}} \sum_{S \subseteq \mathcal{K}} x_i(S)\, v_i(S) \tag{WDP}$$

$$\text{s.t.} \quad \sum_{S \subseteq \mathcal{K}} x_i(S) \leq 1 \qquad \forall i \in \mathcal{I} \tag{2.4}$$

$$\sum_{i \in \mathcal{I}} \sum_{S \subseteq \mathcal{K}: k \in S} x_i(S) \leq 1 \quad \forall k \in \mathcal{K} \tag{2.5}$$

$$x_i(S) \in \{0, 1\} \qquad \forall i \in \mathcal{I}, \forall S \subseteq \mathcal{K} \tag{2.6}$$

The first set of Constraints 2.4 restricts bidders to be allocated at most one bundle, thus enforcing an XOR bid language. Omitting these constraints gives the WDP with an OR bid language. Constraints 2.5 ensure that each item $k \in \mathcal{K}$ is allocated at most once. The binary variables defined in 2.6 guarantee that bidders cannot win bundles partially.

Unfortunately, there exists an equivalency between the WDP with an OR bid language and the *weighted set packing* optimization problem (Rothkopf et al., 1998; Sandholm, 2002), which was shown to be NP-complete (Karp, 1972). Lehmann et al. (2006) show that the WDP with an XOR bid language is equivalently NP-complete even if bidders are restricted to a single bid, each item occurs in exactly two bids, and all bids are valued the same.

An optimization problem is *convex* if its objective function is convex and the feasible area represents a convex set. As the items $k \in \mathcal{K}$ are assumed to be indivisible, the decision variables $x_i(S)$ are defined as binary decision variables in the WDP, making the WDP ultimately non-convex. If items and bundles can also be partially allocated, then the decision variables in the WDP are relaxed to continuous variables. This relaxed version of the WDP is called the *relaxed winner determination problem (RWDP)*, which is a convex optimization problem and is therefore solvable in polynomial time (Karmarkar, 1984).

$$
\begin{aligned}
\max \quad & \sum_{i \in \mathcal{I}} \sum_{S \subseteq \mathcal{K}} x_i(S)\, v_i(S) & & & \text{(RWDP)} \\
\text{s.t.} \quad & \sum_{S \subseteq \mathcal{K}} x_i(S) \leq 1 & \forall i \in \mathcal{I} & \quad (\pi_i) & (2.7) \\
& \sum_{i \in \mathcal{I}} \sum_{S \subseteq \mathcal{K} : k \in S} x_i(S) \leq 1 & \forall k \in \mathcal{K} & \quad (p(k)) & (2.8) \\
& x_i(S) \in \mathbb{R}_{\geq 0} & \forall i \in \mathcal{I}, \forall S \subseteq \mathcal{K} & & (2.9)
\end{aligned}
$$

As the RWDP is a linear program, one can also state its dual, the *dual of the relaxed winner determination problem (DRWDP)*.

$$
\begin{aligned}
\min \quad & \sum_{i \in \mathcal{I}} \pi_i + \sum_{k \in \mathcal{K}} p(k) & & & \text{(DRWDP)} \\
\text{s.t.} \quad & \pi_i + \sum_{k \in S} p(k) \geq v_i(S) & \forall i \in \mathcal{I}, \forall S \subseteq \mathcal{K} & \quad (x_i(S)) & (2.10) \\
& \pi_i \in \mathbb{R}_{\geq 0} & \forall i \in \mathcal{I} & & (2.11) \\
& p(k) \in \mathbb{R}_{\geq 0} & \forall k \in \mathcal{K} & & (2.12)
\end{aligned}
$$

The dual variables corresponding to constraints are denoted π_i and $p(k)$ intentionally. Given the (possibly fractional) allocation of the optimal solution of the RWDP, the values of the dual variables π_i and $p(k)$ exactly match the bidders' payoffs and item prices, respectively.

Under certain conditions, it can be shown that the RWDP has an optimal integral solution. Whenever this is the case, the WDP can be solved in polynomial time as the optimal solution to the RWDP is also an optimal solution to the WDP. Unfortunately, such cases are rather rare in practice. For example, an optimal solution of the RWDP exists, if the constraint matrix is totally unimodular. This is the case for assignment markets where bidders bid on multiple goods but win at most one of them (Bikhchandani and Ostroy, 2002). Another example are combinatorial auctions where the bidders' valuation functions fulfill the *gross substitutes* condition. This will be discussed more thoroughly in the context of Walrasian equilibria in Section 2.5.1.

2.5 Competitive Equilibria

A fundamental solution concept in economics is the notion of competitive equilibria. The latter model a state where supply matches demand and all bidders receive a bundle of items that they desire the most at the given prices. Arrow and Debreu (1954) as well as McKenzie (1954) could show that with perfect competition, convex preferences, and demand independence there always exists a set of competitive equilibrium prices supporting a competitive equilibrium. Unfortunately, the conditions for the existence of competitive equilibria are often not met in practice as in many cases items are indivisible and bidder preferences are non-convex. Before elaborating on the rich results on the existence of competitive equilibria, we want to define the concept of competitive equilibria more formally.

The *demand set* of a bidder contains the bundles that maximize the bidder's utility given the price function p_i of the bidder, i.e.,

$$D_i(p_i) = \arg\max_{S \subseteq \mathcal{K}} \{v_i(S) - p_i(S)\} \tag{2.13}$$

An outcome $o = (x, p)$ is called *envy-free* if it assigns each bidder a bundle from her demand set. An envy-free outcome corresponding to a feasible allocation always exists, e.g., prices could be raised so high that all bidders maximize their utility by requesting the empty bundle. If the auctioneer's supply matches the allocated set of items, then the outcome constitutes a *competitive equilibrium*.

Definition 2.6 (Competitive Equilibrium)**.** *A feasible allocation $x = (S_1, \ldots, S_n)$ and prices $p = (p_1, \ldots, p_n)$ are a competitive equilibrium if $S_i \in D_i(p_i)$ and $\bigcup_{i \in \mathcal{I}} S_i = \mathcal{K}$.*

So far, we assumed that each bidder $i \in \mathcal{I}$ has her own price function $p_i : 2^{\mathcal{K}} \rightarrow \mathbb{R}_{\geq 0}$ that maps arbitrary bundles $S \subseteq \mathcal{K}$ to a respective price that the bidder has to pay. This means that prices neither have to be anonymous nor linear. In fact, for personalized and non-linear prices it can be shown that a competitive equilibrium always exists when there is only a single seller in the market (Bikhchandani and Ostroy, 2002).

2.5.1 Walrasian Equilibria

If prices are *anonymous* and *linear*, then they can simply be denoted by a vector $p \in \mathbb{R}^m_{\geq 0}$ where the entry at index k represents the price for item $k \in \mathcal{K}$. The price for a bundle $S \subseteq \mathcal{K}$ is then given by $p(S) = \sum_{k \in S} p(k)$. In the presence of linear and anonymous prices, a competitive equilibrium is also called a *Walrasian equilibrium*.

Definition 2.7 (Walrasian Equilibrium). *A feasible allocation $x = (S_1, \dots, S_n)$ and prices $p \in \mathbb{R}^m_{\geq 0}$ constitute a Walrasian equilibrium if $S_i \in D_i(p) = \arg\max_{S \subseteq \mathcal{K}} \{v_i(S) - p(S)\}$ and $\bigcup_{i \in \mathcal{I}} S_i = \mathcal{K}$.*

The two famous fundamental theorems of welfare economics were originally shown for markets with divisible items (Arrow and Debreu, 1954; McKenzie, 1954). However, they also hold for economic environments where items are indivisible (Blumrosen and Nisan, 2007).

Theorem 2.1 (First Welfare Theorem). *If the allocation $x = (S_1, \dots, S_n)$ and prices $p \in \mathbb{R}^m_{\geq 0}$ constitute a Walrasian equilibrium, then x maximizes social welfare.*

Theorem 2.2 (Second Welfare Theorem). *If $x = (S_1, \dots, S_n)$ is a Pareto efficient allocation, then it can be supported by Walrasian equilibrium prices $p \in \mathbb{R}^m_{\geq 0}$ so that x and p represent a Walrasian equilibrium.*

Existence of Walrasian Equilibria

Walrasian equilibria are a very desirable outcome of combinatorial auctions. Not only does each bidder receive one of her most desired bundles so that the resulting allocation is welfare-maximizing, but also the prices are perceived as fair because they are linear and anonymous. It is therefore not surprising that economic researchers have spent a lot of effort to characterize the conditions which admit the existence of Walrasian equilibria (Bikhchandani and Mamer, 1997; Kelso and Crawford, 1982; Fujishige and Yang, 2003; Baldwin and Klemperer, 2019; Leme, 2017; Shioura and Tamura, 2015).

A central result of these efforts marks the contribution of Bikhchandani and Mamer (1997) who show via strong duality that Walrasian equilibrium prices always exist when the relaxed winner determination problem (RWDP) yields an optimal integer solution. They prove their theorem for a multi-item single-unit market. However, as already noted in their publication, this result can be transferred straightforwardly to multi-item multi-unit environments by treating each homogeneous unit of an item as a distinct good. Kelso and Crawford (1982) introduce the *gross substitutes* condition and find that if all bidders' valuation functions are gross substitutes, then a Walrasian equilibrium always exists.

Definition 2.8 (Gross Substitutes (GS))**.** *Let $p \in \mathbb{R}^m_{\geq 0}$ denote the prices for items $k \in \mathcal{K}$. An item k is demanded by bidder $i \in \mathcal{I}$ if it is part of a bundle $S' \subseteq \mathcal{K}$ being in the demand set of bidder i at prices p, i.e., $k \in S'$ and $S' \in D_i(p) = \arg\max_{S \subseteq \mathcal{K}}\{v_i(S) - p(S)\}$. The gross substitutes condition is fulfilled if for any prices $p' \geq p$ with $p'(k) = p(k)$ it holds that k is demanded at prices p' whenever it is demanded at prices p.*

Gul and Stacchetti (1999) introduce the *single improvement* and *no complementarities* condition which are both equivalent to the gross substitutes condition. They also propose an ascending auction for the multi-item single-unit setting that terminates in a Walrasian equilibrium if the bidders' valuation functions satisfy the GS condition (Gul and Stacchetti, 2000). Murota and Shioura (1999) introduce the notion of *$M^\natural$-concavity* for valuation function and Fujishige and Yang (2003) show that a bidder's valuation function satisfies the gross substitutes condition if and only if it is $M^\natural$-concave. Sun and Yang (2006) generalize the gross substitutes condition to the *gross substitutes and complements (GSC)* condition and show that Walrasian equilibria also exist is this generalized environment that allows complementarities across two different classes of goods.

While the gross substitutes condition and its equivalent formulations concern the existence of Walrasian equilibria in the multi-item single-unit economic environment, there exist also extensions of these concepts to the multi-item multi-unit setting. Milgrom and Strulovici (2009) generalize the notion of gross substitutes and propose the *strong substitutes* condition. A multi-unit valuation function can be reduced to the single-unit case by interpreting each homogeneous unit of a good as a distinct item. If this single-unit valuation function satisfies the gross substitutes condition, then the corresponding multi-unit valuations are strong substitutes. Milgrom and Strulovici (2009) also show that the single-improvement property can be extended to the multi-unit case, becoming the *binary single improvement* property. Shioura and Tamura (2015) note that the strong substitutes condition is equivalent to a multi-unit version of $M^\natural$-concavity. Also the GSC condition has a multi-unit extension which is the *generalized gross substitutes and complements (GGSC)* condition (Shioura and Yang, 2015). Last but not least, Baldwin and Klemperer (2019) consider so-called *demand types* which are a list of vectors describing the possible individual or aggregate demand changes with respect to a generic price change. They show that if the valuation functions of all bidders are concave and belong to a demand type composed of a unimodular set of vectors, then a Walrasian equilibrium exists.

Similar to the auction proposed by Gul and Stacchetti (2000) for the single-unit scenario, Ausubel (2005, 2006) proposes a tâtonnement process for the multi-unit setting that terminates in a Walrasian equilibrium if all bidders have strong substitutes valuations. The algorithm is based on a greedy minimization of the Lyapunov function whose minimum coincides with the lowest Walrasian equilibrium price vector. In the first publication contained in this dissertation (see Chapter 3), we connect this result to optimization algorithms by showing that a price vector is a minimizer of the Lyapunov function if and only if it minimizes the DRWDP. This result allows us to interpret the tâtonnement process presented in Ausubel (2005, 2006) as a primal-dual algorithm.

For a thorough summary of auction formats terminating in Walrasian equilibria as well as an analysis of conditions on individual and aggregate valuation functions that are sufficient for the existence of Walrasian equilibria, we point the reader to the publication listed in Chapter 3.

2.6 Payment Rules

The existence of linear and anonymous Walrasian equilibrium prices can only be guaranteed for rather restricted settings. Unfortunately, in many practical applications, the bidders' preferences contain complementarities so that linear and anonymous equilibrium prices generally do not exist. This makes it not obvious how to set prices that are acceptable to all market participants. This section provides an overview of the pricing rules that have received much attention over the past decades and are commonly used for practical applications today.

2.6.1 Pay-as-Bid

The classic *English auction* is undoubtedly by far the most known auction format. This is probably due to large auction houses like Sotheby's and Christie's using the English Auction for auctioning off their goods. In an English auction, starting from an item's reserve price, bidders submit successively higher bids than their competitors. If none of the bidders wants to raise their bid, the bidder with the highest bid wins the auction. The payment of the bidder equals the standing bid. This pricing rule is often referred to as *pay-as-bid* (or *first-price*) payment rule.

Because of its simplistic design, the pay-as-bid payment rule is often perceived as the natural choice among bidders, especially in smaller auctions where only one item is auctioned off. Apart from English auctions, the pay-as-bid pricing rule is also implemented in *first-price sealed-bid auctions*. In the latter auction format, bidders place their bids in a sealed form. After the bidding phase ends, the auctioneer opens all bids simultaneously and determines the bidder with the highest bid who is announced the winner of the auction and has to pay an amount equal to her bid.

In order to maximize their utility in first-price sealed-bid auctions, bidders should not bid up to their true valuation but need to shade their bid. While equilibrium bidding strategies can be computed for simple auctions with only a single good (Bichler, 2017), it is much more complex to derive them for other scenarios such as split-award auctions (Kokott et al., 2019). For first-price sealed-bid combinatorial auctions, no closed-form equilibrium bidding strategies can be derived (Bichler, 2017). In order to maximize their utility in these settings, bidders often need to invest a lot of time and effort to learn about the bidding behavior of their competitors so that they can estimate by how much they should shade their own bids. As a result, bidding becomes very complex for bidders in these types of auctions.

2.6.2 VCG Prices

In 1961, Vickrey published a groundbreaking paper in which he proposes a *second-price sealed-bid auction* which is also known as *Vickrey auction* today (Vickrey, 1961). In markets with only a single item up for sale, all bidders bid simultaneously in a sealed bid format. The bidder with the highest bid wins the auction but only has to pay the bid of the second highest bidder. Thus, the bidder's own bid does not influence the resulting price. It can be shown that the dominant strategy of all bidders is to bid their true valuations and that the resulting outcome is efficient, both being very desirable properties in auction theory (Ausubel et al., 2006b).

Vickrey's initial auction design and the later contributions of Clarke (1971) and Groves (1973) led to the celebrated *Vickrey-Clarke-Groves (VCG)* mechanism which is applicable also to multi-item multi-unit combinatorial auctions. Similar to Vickrey's original design, bidders have to pay the opportunity cost of the goods they are allocated in the auction. A more formal definition of the VCG price of bidder i is given in Definition 2.9.

Definition 2.9 (VCG Price)**.** *Let $x^* = (S_1^*, \ldots, S_n^*)$ be the optimal allocation when solving the WDP with the bids of all bidders $\mathcal{I}$ and let $y^* = (T_1^*, \ldots, T_{i-1}^*, T_{i+1}^*, \ldots, T_n^*)$ denote the optimal allocation considering only the bids of bidders $\mathcal{I}_{-i} = \mathcal{I} \setminus \{i\}$. Furthermore, let $w(\mathcal{I})$ be the optimal objective value of the WDP with respect to the stated valuations of all bidders $\mathcal{I}$. Then, the VCG price of bidder i is*

$$\begin{aligned} p_i^{VCG} &= \sum_{j \in \mathcal{I}_{-i}} v_j(T_j^*) - \sum_{j \in \mathcal{I}_{-i}} v_j(S_j^*) \qquad (2.14) \\ &= v_i(S_i^*) - \left(\sum_{j \in \mathcal{I}} v_j(S_j^*) - \sum_{j \in \mathcal{I}_{-i}} v_j(T_j^*) \right) \\ &= v_i(S_i^*) - \left(w(\mathcal{I}) - w(\mathcal{I}_{-i}) \right). \end{aligned}$$

Similar to the Vickrey auction, it is a dominant strategy for bidders to report their valuations truthfully in the VCG mechanism, i.e., the mechanism is strategyproof. This makes bidding substantially easier for bidders than in the first-price sealed-bid auction as they do not have to spend efforts to estimate the bids of their competitors when submitting their own bid. In addition to that, Green and Laffont (1979) and Holmström (1979) show that the VCG mechanism is the unique strategyproof and efficient mechanism in the independent private values setting.

Despite these positive theoretical results, the VCG mechanism is barely used in practice. Ausubel et al. (2006b) and Rothkopf (2007) list several properties of the VCG mechanism that make it impractical for real-word applications: The VCG mechanism breaks the law of one price as bidders potentially pay different amounts for the same set of items. It may result in low and sometimes even zero revenue, and the revenue is not monotonic in the number of bidders present in the auction or the amount they bid. The mechanism is also not group-strategyproof so that a group of bidders may collude and alter their bids to change the allocation.

For large-scale combinatorial auctions with many items and available bundles, the VCG mechanism puts a computational burden on the auctioneer as the computationally hard winner determination problem does not only have to be solved to determine the welfare-maximizing allocation but also once again for each winning bidder in the auction in order to compute the VCG payments (Rothkopf, 2007). Moreover, due to the large number of available packages, bidders are often unable to state their valuation for all possible bundles. Even if the sheer number of packages was not of concern, bidders might be reluctant to report their true preferences as this secret information might be exploited by their competitors if bids get public. Moreover, the VCG mechanism is no longer strategyproof if bidders have a limited budget, the independent private value model is violated, or bidders do not report their valuations for all available bundles. Last but not least, it can be shown that the VCG outcome is not necessarily in the core (Ausubel et al., 2006b), a solution concept that is discussed more thoroughly in the following section.

2.6.3 Core Pricing

The *core* can be seen as a measure of stability of an allocation. If an outcome $o = (x, p)$ consisting of an allocation $x = (S_1, \dots, S_n)$ and price functions $p = (p_1, \dots, p_n)$ is a core outcome, then no coalition of bidders can negotiate a side-deal with the auctioneer in which the accumulative payoff of everyone involved is strictly larger than in the original allocation. To define the core more formally, we have to introduce the notion of coalitional values.

Let $\mathcal{I}_0 = \{0, \dots, n\}$ denote the set of bidders including the auctioneer who is indexed with $i = 0$ and let $\mathcal{X}_C$ denote the set of all possible feasible allocations $x = (S_0, \dots, S_{|C|})$ involving coalition $C \subseteq \mathcal{I}_0$. Then, the *coalitional value* of a coalition $C \subseteq \mathcal{I}_0$ is given by

$$w(C) = \begin{cases} \max_{x \in \mathcal{X}_C} \sum_{i \in C} v_i(S_i), & \text{if } 0 \in C \\ 0, & \text{otherwise.} \end{cases} \tag{2.15}$$

As the auctioneer is necessary for trade, the coalitional value of coalitions not including the auctioneer is zero. This leads to the definition of the core:

Definition 2.10 (Core)**.** *A payoff vector* $\pi \in \mathbb{R}^n_{\geq 0}$ *is in the core if it is part of*

$$Core(\mathcal{I}_0, w) = \left\{ \pi \geq 0 : \sum_{i \in \mathcal{I}_0} \pi(i) = w(\mathcal{I}_0),\ \ w(C) \leq \sum_{i \in C} \pi(i) \quad \forall C \subseteq \mathcal{I}_0 \right\}. \tag{2.16}$$

An outcome is called a *core outcome* if the vector describing the bidders' payoff is in the core. The allocation and prices associated with the core outcome are referred to as *core allocation* and *core prices*, respectively. If an outcome is not in the core (i.e., the outcome is *unstable*), then the payments of winning bidders are so low that their losing competitors might be willing to pay a higher amount for the same set of goods. Allocations outside of the core are often perceived as unfair and are therefore impractical for auctioning off public goods such as spectrum (Day and Raghavan, 2007; Day and Milgrom, 2008). Therefore, ascending combinatorial auctions which are commonly used for spectrum sales worldwide are designed to terminate in a core outcome (Ausubel and Milgrom, 2002; Ausubel et al., 2006a).

While the core in a combinatorial auction setting is non-empty as the pay-as-bid price vector is always included (Day and Cramton, 2012), the core may contain several outcomes, making it unclear which one to choose. If the VCG payoff vector is in the core, then this would be a natural choice as it can be shown that the Vickrey payoff vector is *bidder dominant*, meaning that all bidders unanimously prefer it over any other element of the core (Ausubel et al., 2006b).

Unfortunately, the VCG mechanism generally does not lead to a core outcome (Ausubel et al., 2006b). Only in very restricted settings where the *bidders-are-substitutes condition* is fulfilled, VCG prices are guaranteed to support a core allocation (Bikhchandani and Ostroy, 2002). However, for practical applications this condition is often violated. If the VCG payoff vector is not in the core, then the associated VCG prices are always less than or equal to any core prices (Ausubel et al., 2006b). Researchers have therefore tried to develop pricing schemes that result in core prices which are as close as possible to VCG prices. Day and Raghavan (2007) propose a method based on constraint generation that terminates in *bidder-Pareto-optimal core prices*, i.e., prices that minimize the total payments of all bidders. Similar to Parkes et al. (2001), they argue that these prices minimize the possibility of bidders to collude and manipulate the outcome. Day and Cramton (2012) refine this method in the sense that the price vector among all bidder-Pareto-optimal core prices is selected which minimizes the Euclidean distance to the VCG prices.

Over the past years, the Vickrey-nearest bidder-optimal core-pricing rule based on Day and Cramton (2012) has gained much popularity in spectrum auctions around the world. As a pricing scheme for the *Combinatorial Clock Auction (CCA)* design, it has been used in Denmark, Austria, Canada, Switzerland, Australia, the United Kingdom, and others (Cramton, 2013; Ausubel and Baranov, 2017). The Vickrey-nearest core pricing rule was also proposed for the *Flexible Use and Efficient Licensing (FUEL)* auction design (Milgrom, 2019), which is analyzed in depth in Chapter 5.

2.6.4 IP Pricing

Electricity spot markets are two-sided markets where the load (i.e., the energy demand) is matched with the output of generators (i.e., the supply). These markets are highly non-convex due to start-up and shut-down costs of generators, minimum output requirements, and piece-wise linear cost functions. In such non-convex markets, the dual variables of the corresponding allocation problem no longer have a similar economic interpretation as in convex settings. Linear and anonymous equilibrium prices generally do not exist in electricity markets. Nevertheless, many electricity market operators around the world try to "convexify" the underlying optimization problem and mimic equilibrium prices by adopting linear, anonymous pricing schemes in combination with side-payments and penalties (Liberopoulos and Andrianesis, 2016).

A commonly used pricing scheme in electricity markets is *IP pricing* which was originally proposed by O'Neill et al. (2005). Under this pricing rule, a non-convex mixed-integer linear optimization program is solved to determine the optimal dispatch. Prices are then computed by solving a modified version of the allocation problem, the so-called *pricing problem*. In the latter, all integer variables are fixed to the value they attain in the optimal solution of the allocation problem. This eliminates the non-convexities so that the dual variables of the pricing problem can be interpreted as linear and anonymous prices.

Unfortunately, such prices are neither budget-balanced nor do they fulfill individual rationality. In fact, substantial side-payments (so-called *make-whole payments*) are necessary in energy markets to prevent market participants from incurring a loss. While energy prices are linear and anonymous, the side-payments are personalized and private, thereby reducing the transparency of the market (O'Neill et al., 2020; Hytowitz et al., 2020). Moreover, the stability of the allocation can only be enforced by setting suitable penalties for bidders that make deviation unattractive. High make-whole payments and penalties indicate a high degree of non-convexities in the market.

While IP pricing is used extensively in electricity markets today, make-whole payments remain a major concern as they are not incorporated in the public prices and therefore lead to wrong investment signals. There is ongoing research to develop other pricing mechanisms where make-whole payments and penalties are minimized. *Convex Hull Pricing (CHP)* as proposed by (Gribik et al., 2007) minimizes penalties but is computationally so challenging that it is not used for practical applications (Schiro et al., 2015). *Extended Locational Marginal Pricing (ELMP)* attempts to approximate convex hull pricing but its economic properties are not well understood yet (Schiro et al., 2015). Another direction is followed by O'Neill et al. (2020) who propose *Average Incremental Cost (AIC)* pricing which aims to eliminate make-whole payments. In the pricing run of this method, a generator's minimum output level is relaxed to zero and its marginal cost is replaced by the AIC price, i.e., its original marginal cost plus the fixed cost averaged over the generator's assigned dispatch. The price at which market participants trade is then defined as the highest AIC of supply dispatched.

In Chapter 4, we propose a wholesale market design for a congestion pricing market. To determine prices for individual road segments, we apply the IP pricing method presented above. Interestingly, unlike electricity markets, our numerical experiments show that the make-whole payments in our congestion pricing market are negligible compared to the overall payments made by the bidders.

2.7 Large-Scale Auction Market Design in Practice

Combinatorial auctions are the means of choice when the submission of package bids is essential for bidders to state their preferences accurately. This is often the case when bidders have complementary valuations. Combinatorial auctions are widely used in the field, prominent examples include spectrum auctions (Bichler and Goeree, 2017), allocation of airport time slots (Rassenti et al., 1982), procurement of freight transportation services (Caplice et al., 2006), industrial procurement (Bichler et al., 2006), and procurement of public transport services (Cantillon et al., 2006).

In Chapter 4 and 5 we consider two large-scale combinatorial auction markets, a wholesale market for road capacity and a spectrum auction in the United States. Before analyzing our bid language designs for these markets in detail, we would like to provide the reader with a short overview of the scientific and historic context of these markets. First, in Section 2.7.1, we consider different static congestion pricing schemes and their implementation in various cities around the world. The wholesale market design we analyze in Chapter 4 admits dynamic congestion pricing which, in contrast to static schemes, allows for setting tolls dynamically on individual road segments. Secondly, in Section 2.7.2, we give a brief overview of the evolution of spectrum auctions over the past decades, eventually leading to the large-scale 5G spectrum auction in the US which is subject to Chapter 5.

2.7.1 Congestion Pricing

Traffic congestion leads to substantial economic and ecological damage in large cities around the world. People sitting in traffic do not only lose time and waste fuel but also cause higher air pollution and carbon dioxide emissions, both being a threat to the health of a city's population (Levy et al., 2010). According to an article by *The Economist*, the total cost of traffic jams in Germany, Britain, and the US summed up to approximately $461 billion.[1] Due to an increase in population and a trend towards urbanization, traffic is predicted to worsen in the coming decades. According to a study by the German Federal Ministry of Transport and Digital Infrastructure, the car traffic in Germany will increase by approximately 10% between 2010 and 2030 (Schubert et al., 2014). The same study also predicts that the number of vehicles per capita raises by around 10%.

A widespread belief is that traffic congestion can be reduced in the future with a combination of ride-hailing services, self-driving cars, additional public transportation, and large investments in infrastructure. However, this belief cannot be supported by recent research (Duranton and Turner, 2011; Schaller, 2017; Simoni et al., 2019). Unlike other essential utilities such as electricity, water, gas, and communications for which consumers

[1] `https://www.economist.com/graphic-detail/2018/02/28/the-hidden-cost-of-congestion`, published: 28.02.2018, accessed: 09.12.2021

have to pay according to their level of usage, the road network in major cities is most often accessible free of charge. Researchers see this mispricing of road capacity as the main reason for traffic congestion (Cramton et al., 2018; Beheshtian et al., 2020).

Urban congestion pricing is designed to alleviate traffic jams by making drivers pay for the negative externalities they impose on other drivers and the environment. Road prices also give objective market signals that support policy makers in making the most effective investment choices (Cramton et al., 2019a). Charging drivers the marginal social cost of their trip has already been proposed by Pigou (1920). In the 1950s, Nobel laureate William Vickrey who is sometimes referred to as the "father of congestion pricing" proposed an increase of New York City's subway system fares at peak hours in order to reduce congestion (Vickrey, 1955). At the end of the 1950s, he also suggested to equip cars with electronic transponders and automatically impose charges whenever the car passes an intersection (Vickrey, 1959). Singapore was the first city to adopt downtown congestion pricing in 1975. London, Stockholm, Milan, and Gothenburg followed suit in the 2000s (Lehe, 2019). While concrete implementations of congestion pricing schemes are subject to public debate in Jakarta, New York City, and Vancouver (Lehe, 2019), various studies discuss implementing congestion pricing also in Germany, e.g., in Berlin (Berliner Senatsverwaltung für Umwelt, Verkehr und Klimaschutz, 2020) and Munich (Falck et al., 2020).

Classification of Congestion Pricing Schemes

Literature classifies congestion pricing schemes along different dimensions, among which are their type (*facility-bases*, *area-based*, *cordon-based*, *distance-based*), fluctuations of tolls over time (*static*, *dynamic*, *predictive*), and the technology used to implement tolls (de Palma and Lindsey, 2011; Cheng et al., 2017). The most commonly used road pricing type today is facility-based pricing that charges drivers for accessing certain facilities such as roads, bridges, and tunnels. Most often facility-based pricing serves the generation of revenue rather than reducing traffic congestion (de Palma and Lindsey, 2011).

Cordon-based pricing schemes charge drivers whenever they cross a cordon that is typically spanned around the center of a city. Prominent examples include Singapore's *Area License Scheme (ALS)* that was in practice from 1975 to 1998, its successor the *Electronic Road Pricing (ERP)*, the *Stockholm Congestion Tax (SCT)* established in 2006, as well as the pricing schemes in Milan and Gothenburg which were launched in 2008 and 2013, respectively (Lehe, 2019). When drivers must also pay for trips taking place within the toll area (and not only for those crossing its border), then this is referred to as area-based or zonal pricing. The *London Congestion Charge (LCC)* which was launched in 2003 uses this type of pricing scheme (Lehe, 2019). Cordon- and area-based congestion pricing schemes impose tolls on drivers for entering or driving inside a toll area but do not consider the actual distance traveled. This is the objective of distance-based pricing schemes which produce prices that can generally better reflect a driver's real road capacity consumption. National distance-based heavy goods vehicle tolls are imposed in various European countries including Germany but have not yet been used to impose congestion prices in urban areas (de Palma and Lindsey, 2011).

Congestion pricing schemes can also be categorized by how tolls fluctuate over time, distinguishing *static* from *dynamic* and *predictive* pricing policies (Cottingham et al., 2007; de Palma and Lindsey, 2011; Cheng et al., 2017). With static pricing, toll levels are announced long in advance and do not respond to the actual congestion on the road. Both flat tolls and those that vary by time of day with respect to a previously announced schedule are examples of static pricing methods. In dynamic pricing schemes, on the other hand, prices are responsive to the observed level of congestion and vary in real-time or near real-time. Instead of reacting to existing congestion levels, predictive schemes try to anticipate future congestion levels and set prices accordingly in advance. However, predictive schemes are still in their infancy (de Palma and Lindsey, 2011). They have been analyzed for toll lanes on highways (Dong et al., 2011) but have only recently been considered also for metropolitan areas (Vosough et al., 2020).

Theory suggests that dynamic pricing schemes better capture the inherent uncertainty and time-dependent nature of traffic flows (Beheshtian et al., 2020; Cheng et al., 2017; Wie and Tobin, 1998; Do Chung et al., 2012). Setting tolls dynamically allows to charge drivers the actual social cost of their trip which is a necessity for maximizing social welfare through efficient pricing. Dynamically adapting prices of individual road segments also guides drivers in making their route choice so that traffic is distributed effectively over the road network. Static pricing schemes do not provide the same benefits as they usually price drivers for using roads within an entire area but do not pose prices for individual road segments.

Despite these advantages of dynamic pricing schemes, the downtown congestion pricing models currently used in practice all implement static tolls (Cramton et al., 2019b;

Lehe, 2019). Technological advances in satellite positioning technology have recently sparked a discussion on the implementation of distance-based, dynamic road pricing in Singapore (Cramton et al., 2019a; Lehe, 2019). Implementing dynamic congestion pricing schemes is subject to ongoing research (Cheng et al., 2017). In the following, we consider a market-based approach to dynamic congestion pricing which is tractable also for a major city like Berlin.

A Wholesale Market Design for Congestion Pricing

In analogy to electricity markets, Cramton et al. (2019a) and Beheshtian et al. (2020) propose a two-stage wholesale market in which an *Independent System Operator (ISO)* sells road capacity to *Service Providers (SPs)*. Whenever end consumers want to travel across the road network, they need to purchase the respective road capacity for their trip from one of the service providers on a separate retail market. The wholesale market features several forward markets and a real-time market. Forward markets are run monthly, weekly, and daily before the traded road capacity licenses become valid. These markets are purely financial and allow service providers to take and adjust forward positions with respect to the predicted demand of their customers. In the physical real-time market, the ISO determines clearing prices for each road segment matching the supply with the actual demand of end consumers who travel through the network.

In both the forward and real-time market, service providers define piecewise-linear decreasing demand curves for individual road segments. However, their customers typically demand road capacity for an entire trip between an origin-destination pair and are often satisfied to be allocated licenses for one of several substitutable routes. As service providers are unable to express their customer demand in terms of origin-destination pairs in the auction, there is a risk of exposure for service providers if they only win a route fractionally. While this exposure problem is mitigated by the recurrent structure of forward markets and the large number of road capacity that service providers purchase within a time slot, it certainly makes the bidding process much more complex.

In the second research project of this dissertation (see Chapter 4), we propose and evaluate a compact bid language that aims to simplify the bidding process for service providers in the wholesale market. This novel bid language allows service providers to express their demand for multiple substitutable routes between the same origin-destination pair. Besides being able to specify a different willingness-to-pay for each individual route, bidders can set lower and upper bounds on the total number of trips to be allocated between a particular origin-destination pair. We model the allocation problem as a mixed-linear integer program and show its tractability for realistic problem instances by a series of numerical experiments. For this purpose, we rely on synthetic but calibrated traffic data of the MATSim Open Berlin Scenario which is designed to model the traffic in the city of Berlin accurately (Ziemke et al., 2019). Approximate competitive equilibrium prices are computed by applying the IP pricing method which is commonly used in electricity markets today (see Section 2.6.4). Unlike electricity markets, the non-convexities in our congestion pricing market are small so that the amount of side-payments is negligible. We find that the initial forward market which is the largest in terms of bids and trading volume can be solved within 15 minutes, while only a small fraction of road segments have to be priced. For a detailed discussion of the bid language and our results, we point the reader to Chapter 4.

2.7.2 Spectrum Auctions

As part of his Nobel prize lecture, the winner of the 2020 Nobel Memorial Prize in Economic Sciences, Paul Milgrom, published an excellent survey on the history of spectrum auctions on which the next few paragraphs are mainly based (Milgrom, 2021).

Competitive Hearings and Lotteries

Up to the mid-1980s, the demand for radio spectrum in the US was small so that the *Federal Communications Commission (FCC)* could conduct competitive hearings in order to decide how to allocate the available spectrum. In such “beauty contests”, competing companies had to lay out why the public benefited most if their firm was granted spectrum. With the rise of mobile phone technology, the interest in spectrum grew rapidly, making the time-consuming competitive hearings infeasible. Instead, the FCC decided to use lotteries to distribute spectrum but failed to pose any requirements on the participants who entered the lottery. This caused companies to win spectrum licenses that were unable to establish a viable network. Whenever these lottery winners did not resell their licenses to actual telecommunication companies within a reasonable time frame, the roll-out of mobile phone services was slowed down significantly.

Simultaneous Multiple Round Auction

In 1993, the US Congress intervened and ordered the FCC to use auctions for distributing spectrum. While the main goal was an efficient allocation of spectrum, earning a substantial amount of money was a secondary objective. As scientific literature on designing such large multi-item auctions did not exist at that point in time, the FCC sought guidance from different experts in the field. The 2020 Nobel laureates Paul Milgrom and Robert B. Wilson proposed the *Simultaneous Multiple Round Auction (SMRA)* which is run in a series of discrete rounds. Starting with a reserve price for each item, bidders can raise their bid on any good in each of the rounds. To ease the price discovery, bids are made public after each round. In order to prevent bid sniping, i.e., the submission of bids just before the auction closes, the process only terminates when no more bids were submitted in a round. The design of Milgrom and Wilson also features an activity rule which demands that bidders cannot bid for more spectrum in any round than they had bid in the previous round. This forces bidders to bid actively from the beginning and thereby reduces the number of required auction rounds.

The SMRA and its modified versions proved to be very successful not only in the US spectrum auction in 1994 but also in other parts of the world, leading to an accumulated transaction of more than $300 billion. Despite its success, the SMRA design is not without problems as it is susceptible to demand reduction (Ausubel et al., 2014), signaling through jump bids (Grimm et al., 2003), retaliating bids (Cramton and Schwartz, 2000), budget bluffing (Porter and Smith, 2006), and the exposure problem (Cramton, 1997; Brunner et al., 2010). Moreover, Milgrom (2000) shows that the auction only terminates in a Walrasian equilibrium if the licenses are mutual substitutes to bidders and bidders bid straightforwardly, which are rather strict assumptions that are often not met in practice.

Combinatorial Clock Auction

One way to tackle the exposure problem of the SMRA design is by allowing bidders to submit all-or-nothing package bids. This led to the development of *Combinatorial Clock Auctions (CCA)* (Porter et al., 2003; Ausubel et al., 2006a; Ausubel and Baranov, 2014; Cramton, 2013). Similar to the SMRA, such clock auction designs proceed in rounds. Given the prices for individual items in each round, bidders respond with a bundle of items they desire the most at the current prices. The prices of overdemanded items are raised between rounds and the auction terminates when aggregated demand of all bidders can be served with the available supply of items. It can be shown that if the bidders' valuations fulfill the gross substitutes condition and bidders bid straightforwardly, then the clock auction terminates in an efficient outcome (Ausubel and Milgrom, 2002). Without substitute valuations, however, the linear and anonymous prices do not suffice for the auction to be efficient even if bidders bid straightforwardly (Ausubel et al., 2006a).

To promote efficiency, Ausubel et al. (2006a) suggest to add a second auction phase after the clock stage terminates. This *supplementary stage* is sealed-bid, allows for package bids, and features an XOR bid language. Activity rules in the clock and supplementary phase of the auction ensure that if bidders bid straightforwardly in the clock phase and truthful in the supplementary phase, then the outcome is efficient (Ausubel et al., 2006a). A Vickrey-nearest core pricing rule is often applied to promote truthful bidding (Cramton, 2013). The two-stage CCA in combination with the Vickrey-closest core pricing rule has been used extensively for spectrum auctions around the world (Ausubel and Baranov, 2017). However, researchers have identified a series of weaknesses of this design including safe supplementary bids that allow bidders to safely win bundles with bids lower than their true valuation (Bichler et al., 2013a), spiteful bids to drive up the payments of the competing bidders (Janssen and Karamychev, 2013), and the absence of equilibrium bidding strategies (Bichler et al., 2013b). Despite that, due to the combinatorial package bidding with an XOR bid language in the supplementary phase, the auction is susceptible to the missing bids problem (described in Section 2.2).

C-band Auction

In mid 2019, a consortium of four of the world's largest satellite operators proposed to sell a portion of their C-band spectrum in a private auction to telecommunication companies.[1] While the satellite companies provided satellite downlink for television and radio stations on the C-band frequencies, telecommunication companies could use the spectrum for establishing terrestrial 5G services. The satellite providers proposed a sealed-bid package auction based on the novel *Flexible Use and Efficient Licensing (FUEL)* bid language for the sale of 14 licenses in each of the 406 different geographic areas (Milgrom, 2019).

[1]SpaceNews article on the formation of C-Band Alliance: `https://spacenews.com/telesat-changes-tune-joins-c-band-spectrum-group/`, published: 01.10.2018, accessed: 09.12.2021

Due to the high number of available goods and the need to allocate licenses quickly in order to facilitate a fast deployment of 5G technology, existing auction formats like the SMRA or CCA are unsuitable for the C-band auction. Multi-round auction formats such as the SMRA or CCA often result in a time-consuming bidding process that needs much strategic preparation in advance and is susceptible to bidding errors (Milgrom, 2019). Moreover, as spectrum licenses across neighboring geographic areas are highly complementary goods for bidders, the SMRA and related clock auction formats suffer from the exposure problem. In addition to that, the CCA does not scale up to the large number of 15^{406} distinct packages available in the auction. It is susceptible to the missing bids problem as we show in Chapter 5 by comparing FUEL with an XOR bid language which is commonly used in the supplementary stage of CCA designs. The FUEL bid language, on the other hand, features a novel concept of bid groups which allows bidders to specify valuations for a large number of packages in a succinct manner, making the allocation problem solvable within a few minutes even for large auctions like the one suggested for the C-band auction.

After initial positive replies from the FCC regarding the proposed private spectrum auction, the FCC decided against using the FUEL design and conducted the C-band spectrum sale as a public auction with a simultaneous clock auction format instead.[1,2] Due to its generic design, the FUEL bid language is well suited for other application domains such as procurement auctions (Bichler et al., 2006), fishery access rights (Iftekhar and Tisdell, 2012), or TV ads (Goetzendorff et al., 2015). In addition to that, there remains the intriguing idea to use FUEL instead of a classic XOR bid language in the supplementary stage of a CCA in order to tame the missing bids problem.

[1] Public statement of FCC's former chairman Ajit Pay on the C-band proceeding: `https://docs.fcc.gov/public/attachments/DOC-360855A8.pdf`, published: 18.11.2019, accessed: 09.12.2021

[2] FCC public notice establishing the procedures for Auction 107: `https://docs.fcc.gov/public/attachments/FCC-20-110A1.pdf`, published: 07.08.2020, accessed: 09.12.2021

3 Publication 1: Walrasian Equilibria from an Optimization Perspective

Peer-Reviewed Journal Paper

Title: Walrasian equilibria from an optimization perspective: A guide to the literature

Authors: Martin Bichler, Maximilian Fichtl, Gregor Schwarz

In: Naval Research Logistics

Abstract: An ideal market mechanism allocates resources efficiently such that welfare is maximized and sets prices in a way so that the outcome is in a competitive equilibrium and no participant wants to deviate. An important part of the literature discusses Walrasian equilibria and conditions for their existence. We use duality theory to investigate existence of Walrasian equilibria and optimization algorithms to describe auction designs for different market environments in a consistent mathematical framework that allows us to classify the key contributions in the literature and open problems. We focus on auctions with indivisible goods and prove that the relaxed dual winner determination problem is equivalent to the minimization of the Lyapunov function. This allows us to describe central auction designs from the literature in the framework of primal-dual algorithms. We cover important properties for existence of Walrasian equilibria derived from discrete convex analysis, and provide open research questions.

Contribution of dissertation author: Methodology, formal analysis, visualization, joint paper management

Reference: Bichler et al. (2021)

Received: 2 December 2019 | Revised: 4 September 2020 | Accepted: 14 October 2020 | Published on: 20 November 2020

DOI: 10.1002/nav.21963

RESEARCH ARTICLE

WILEY

Walrasian equilibria from an optimization perspective: A guide to the literature

Martin Bichler | **Maximilian Fichtl** | **Gregor Schwarz**

Department of Informatics, Technical University of Munich, Munich, Germany

Correspondence
Martin Bichler, Department of Informatics, Technical University of Munich, Munich, Germany.
Email: bichler@in.tum.de

Funding information
Deutsche Forschungsgemeinschaft (DFG), Grant/Award Number: BI 1057/1-8.

History
Accepted by Saša Pekeč, algorithms, computation, and economics.

Abstract

An ideal market mechanism allocates resources efficiently such that welfare is maximized and sets prices in a way so that the outcome is in a competitive equilibrium and no participant wants to deviate. An important part of the literature discusses Walrasian equilibria and conditions for their existence. We use duality theory to investigate existence of Walrasian equilibria and optimization algorithms to describe auction designs for different market environments in a consistent mathematical framework that allows us to classify the key contributions in the literature and open problems. We focus on auctions with indivisible goods and prove that the relaxed dual winner determination problem is equivalent to the minimization of the Lyapunov function. This allows us to describe central auction designs from the literature in the framework of primal-dual algorithms. We cover important properties for existence of Walrasian equilibria derived from discrete convex analysis, and provide open research questions.

KEYWORDS

duality, primal–dual algorithms, Walrasian equilibrium

1 | INTRODUCTION

Many markets match supply and demand for multiple goods or services (which we also refer to as items) via optimization. Typically, the auctioneer computes an allocation and linear (i.e., item-level), anonymous prices. Linear and anonymous competitive equilibrium prices are often referred to as Walrasian prices in honor of Léon Walras, a French mathematical economist, who pioneered the development of general equilibrium theory. Prominent examples include financial markets (Klemperer, 2010), day-ahead electricity markets (Meeus et al., 2009; Triki et al., 2005), environmental markets (Bichler et al., 2019), logistics (Caplice & Sheffi, 2003; Bichler et al., 2006; Ağralı et al., 2008) or spectrum auctions (Bichler & Goeree, 2017). In some of these markets the auctioneer computes prices that are in a competitive equilibrium with linear and anonymous prices (aka. a Walrasian equilibrium),[1] in others Walrasian prices even lead to efficiency losses (Özer & Özturan, 2009; Lessan & Karabatı, 2018; Bichler et al., 2018; Meeus et al., 2009; Madani & Van Vyve, 2015). This raises the question, which market characteristics admit Walrasian equilibria.

While this is an established and central question in the economic sciences, there have been a number of significant contributions in computer science, economics, and operations research in recent years. The literature on auction algorithms initiated by Bertsekas (1988) is one of the early examples of the fruitful interplay between optimization and equilibrium

[1]There are also competitive equilibria with nonlinear prices (Bikhchandani & Ostroy, 2002). However, some authors only use competitive equilibrium to refer to one with linear and anonymous prices.

theory. In this paper, we survey the literature and describe established and more recent results. We primarily draw on convex analysis and linear programming duality, and provide a consistent mathematical optimization framework to position and explain the key results of this broad literature.

1.1 | Competitive equilibrium

Early in the study of markets, general equilibrium theory was used to understand how markets could be explained through the demand, supply, and prices of multiple commodities or objects. The Arrow–Debreu model shows that under convex preferences, perfect competition, and demand independence there must be a set of competitive equilibrium prices (Arrow & Debreu, 1954; McKenzie, 1959; Gale, 1963; Kaneko, 1976). Market participants are price-takers, and they sell or buy goods in order to maximize their value subject to their budget or initial wealth in this model. The results derived from the Arrow–Debreu model led to the well-known welfare theorems, important arguments for markets as efficient or welfare-maximizing ways to allocate resources. Stability in the form of competitive equilibria where each participant maximizes his utility at the prices is central to this theory. More specifically, the theory focuses on Walrasian equilibria where there is one equilibrium price per good (aka. linear prices) and the price is the same for all bidders (aka. anonymous prices). The first theorem states that any Walrasian equilibrium leads to a Pareto efficient allocation of resources. The second theorem states that any efficient allocation can be attained by a Walrasian equilibrium under the Arrow–Debreu model assumptions.

However, general equilibrium theory assumes divisible goods and convex preferences, and the results do not carry over to markets with indivisible goods and complex (nonconvex) preferences and constraints. Also, in general equilibrium models money does not have outside value and bidders maximize value subject to a budget constraint (Cole et al., 2016). More importantly, bidders are assumed to be nonstrategic price-takers. Based on the work by Vickrey (1961), attention in economics shifted to auction theory, which focuses on small and imperfectly competitive markets, where strategic players can influence prices. These bidders have a quasilinear utility function, that is, they aim to maximize payoff (i.e., value minus price) (Krishna, 2009). Bayesian Nash equilibria (rather than competitive equilibria) are the central equilibrium solution concept in the auction literature, a branch of noncooperative and incomplete information game theory which led to remarkable results. Most importantly, the Vickrey–Clarke–Groves (VCG) mechanism was shown to be incentive-compatible, and truthful bidding to be a dominant strategy for bidders (Vickrey, 1961).

Many markets that have been implemented for trading financial products, electricity, or environmental access rights as discussed earlier are large markets involving many items and many market participants. Participants want to maximize payoff, but they might not be able to influence prices on such markets. As a consequence, much of the literature is based on a complete-information game-theoretical analysis where bidders are price-takers rather than an incomplete-information game (Baldwin & Klemperer, 2019). Competitive equilibria are the main design desideratum. Unfortunately, it is well known that in many of these markets linear (i.e., block-level) prices might not allow for a welfare-maximizing trade and that there might not be competitive equilibria (Meeus et al., 2009; Madani & Van Vyve, 2015b).

Such new markets have led to a renewed interest in the question of existence and computation of competitive equilibria (Kim, 1986; Bikhchandani & Mamer, 1997; Bikhchandani & Ostroy, 2002; Baldwin & Klemperer, 2019; Leme, 2017). The problem is fundamentally rooted in mathematical optimization, as we will show. In this survey, we will focus on central and recent results in competitive equilibrium theory and multiobject auction design and reformulate them in the language of optimization, specifically duality theory and primal-dual algorithms.

1.2 | Outline

There are various ways how surveys are written. Some articles collect and categorize a larger number of papers in a new and emerging field (Herroelen & Leus, 2005; Galindo & Batta, 2013; Olafsson et al., 2008), others provide a guide to a larger literature and introduce important concepts in a unified framework. Examples include a survey on bilevel programming by Colson et al. (2005) or a survey on the gross substitutes condition in economics by Leme (2017). We follow the latter path and discuss competitive equilibrium theory using duality theory and linear programming as a framework. While most of the literature on this subject is published in economics journals, key insights of this literature can be introduced conveniently using the mathematical framework of optimization. Fundamentally, auctions are algorithms for optimal resource allocation and there are plenty of questions where the OR community can contribute as we discuss in the last section.

The survey starts with markets for divisible goods and shows that the concave conjugate to the aggregate value function of all bidders yields prices, and that the minimizer of the Lyapunov function results in Walrasian prices if the aggregate value function is concave. A condition for concavity of the aggregate value function is concavity of the individual value functions, which is equivalent to diminishing marginal returns. The Lyapunov function is convex so that a simple subgradient algorithm finds the minimum efficiently. This algorithm has an interpretation as an auction.

We will next show that the same principles from duality theory carry over to markets with indivisible objects. For this, we describe the allocation problem as a binary program.

Whenever the linear programming relaxation of this binary program has integer solutions, then the dual variables of the capacity constraints have an interpretation as Walrasian prices for the respective resources. We prove that the dual of the linear programming relaxation of this binary program is equivalent to the Lyapunov function. Economic literature discusses conditions on individual value functions that allow for Walrasian equilibria. This is the case if the convolution of these individual functions results in a discrete concave aggregate value function.

As in the continuous case with divisible goods, we can use a steepest descent algorithm to find the minimizer of the Lyapunov function, which is equivalent to determining Walrasian prices for the market. This is exactly what the auction mechanism by Ausubel (2005) does, a central contribution to auction design. Primal-dual algorithms are well-known algorithms to solve linear programs, and they have a nice interpretation as a market with an auctioneer and the bidders optimizing alternatively. The steepest descent algorithm that minimizes the Lyapunov function is equivalent and we show the connections.

We contribute the equivalence of the Lyapunov function and the dual linear programming relaxation of the allocation problem in markets with indivisible goods, as well as the equivalence of primal-dual algorithms with central auction designs for selling multiple indivisible goods. These two results allow us to organize the material and use duality theory to discuss the literature on existence of Walrasian equilibria, and linear programming algorithms to discuss auction designs leading to Walrasian equilibria if it exists. The survey helps scholars with a background in mathematical optimization to understand central results in competitive equilibrium theory and draws important connections between competitive equilibrium theory, mathematical optimization, and discrete convexity.

In Section 2 we introduce the notation and standard assumptions in the economic literature for readers from operations research. Then we introduce important concepts for the understanding of Walrasian equilibria such as the Lyapunov function for markets with divisible goods in Section 3. The same concepts play a role for markets with indivisible goods and discrete value functions in Section 4. In Section 5 we use primal-dual algorithms and show that these are equivalent to important auction designs discussed in economics. Finally, we provide a research agenda and discuss open research problems for the operations research community.

2 | NOTATION AND ECONOMIC ENVIRONMENT

In the auction market, there are m types of items or goods, denoted by $k \in \mathcal{K} = \{1, \dots, m\}$, and n bidders $i \in \mathcal{I} = \{1, \dots, n\}$. In the multi-unit case, we have $s \in \mathbb{Z}^m_{\geq 0}$ units available, that is, $s(k)$ *homogeneous*units for each of the *heterogeneous* m items $k \in \mathcal{K}$. A bundle for bidder i is described by a vector $x_i \in \mathbb{Z}^m_{\geq 0}$. In case of single-unit supply the vector is binary, that is, $x_i \in \{0, 1\}^m$. We will sometimes omit the subscript i for convenience. Each bidder i has a value function $v_i : \mathbb{Z}^m_{\geq 0} \to \mathbb{Z}_{\geq 0}$ over bundles of items or objects x_i. We assume integer-valued functions v_i as it will be more convenient to analyze the optimality of auction algorithms. Moreover, integer-valued functions v_i allow to use integral prices in ascending auctions without losing efficiency.

Unless stated otherwise this paper we assume that bidders have *preferences* described via a valuation function with the following properties:

- Pure private values: Bidder i's value $v_i(x_i)$ does not change when she learns other bidder's information.
- Quasilinearity: Bidder i's (direct) utility from bundle x_i is given by $\pi_i(x_i, p) = v_i(x_i) - \langle p, x_i \rangle$, where $\langle \cdot, \cdot \rangle$ is the dot product.
- Monotonicity: The function $v_i : \mathbb{Z}^m_{\geq 0} \to \mathbb{Z}_{\geq 0}$ is weakly increasing with $v_i(0) = 0$ and, if $x_i \geq x_i'$, then $v_i(x_i) \geq v_i(x_i')$.

An auctioneer wants to find an allocation of items to bidders. Such an allocation is *feasible* when the supply suffices to serve the aggregate demand of the bidders. Furthermore, the auctioneer aims for *allocative efficiency*. This means the auctioneer wants to maximize *social welfare* which is the sum of the utilities of all participants (the bidders and the auctioneer). Maximization of welfare is also referred to as a *utilitarian* welfare function. In case of quasilinear utility functions, prices cancel and the social welfare is defined as $\sum_{i \in \mathcal{I}} v_i(x_i)$.

For the remainder of this survey we assume that the auctioneer's valuation for all items is zero. As a consequence, the auctioneer would sell items to bidders for a price of zero. In some auction scenarios, however, the auctioneer may want to set *reserve prices* which are the minimum prices at which the auctioneer would be willing to sell the goods. Often these reserve prices can be implemented by introducing a dummy bidder who simply bids the reserve prices on behalf of the auctioneer in the auction. In case the dummy bidder wins any items in the auction, these items remain unsold.

The goal of the auctioneer is to find an efficient allocation that yields linear (i.e., item-level) and anonymous market clearing prices $p = \{p(k)\}_{k \in \mathcal{K}} \in \mathbb{R}^m$. The *linearity of prices* refers to the property that individual prices are set for each item $k \in \mathcal{K}$; the price for a bundle x is then simply the sum of the prices of its components, that is, it is given by the dot product $\langle p, x \rangle$. *Anonymity* means that the resulting prices p are the same for all bidders and there is no price differentiation. Furthermore, prices p are *market clearing* when the aggregate demand of all bidders at the given prices p meets the supply s.

With linear and anonymous prices $p = (p(1), \dots, p(k), \dots, p(m))$, the bidder's *indirect utility function* is defined as

$$u_i(p) = \max_{x \in \mathbb{Z}^m_{\geq 0}} \{v_i(x) - \langle p, x \rangle\}.$$

The indirect utility function is widely used in economics and returns the maximal utility that bidder i can obtain for any bundle at prices p. The *demand correspondence* $D_i(p)$ is the set of bundles that maximize the indirect utility function at prices p, that is,

$$D_i(p) = \arg \max_{x \in \mathbb{Z}^m_{\geq 0}} \{v_i(x) - \langle p, x \rangle\}.$$

If in an outcome (consisting of an allocation and prices) all bidders are allocated a bundle from their demand set, then the outcome is *envy-free*. No bidder would want to get another bundle, as a bidder cannot increase her utility at these prices. Envy-free prices always exist. For example, if prices were higher than the valuations, then every bidder would only want the empty set. If in addition to envy-freeness all items are allocated, $\sum_{i \in \mathcal{I}} x_i = s$, then the outcome is a competitive equilibrium.

Definition 1 (Competitive equilibrium, CE). A price vector p^* and a feasible allocation $(x_1, \ldots, x_n)$ form a *competitive equilibrium* if $\sum_{i \in \mathcal{I}} x_i = s$ and $x_i \in D_i(p^*)$ for every bidder $i \in \mathcal{I}$.

If there were unsold items, an auctioneer could always add unsold units to the allocation of a bidder without decreasing welfare as bidders are assumed to have monotone value functions v_i.

In our setting with linear and anonymous prices, a competitive equilibrium is also called a *Walrasian equilibrium*. If there exists a Walrasian price vector p^* such that $p^* \leq p'$ for any other Walrasian price vector p', then p^* is called the *bidder-optimal* Walrasian price vector. For Walrasian equilibria the well-known welfare theorems hold:

Theorem 1 *First and second welfare theorem (following* Blumrosen and Nisan *(2007)) Let $x = (x_1, \ldots, x_n)$ be an equilibrium allocation induced by a Walrasian equilibrium price vector p, then x yields the optimal social welfare. Conversely, if x is a Pareto efficient allocation, then it can be supported by a Walrasian price vector p so that the pair (p, x) forms a Walrasian equilibrium.*

3 | WALRASIAN EQUILIBRIA WITH DIVISIBLE GOODS AND CONJUGACY

In this article, we focus on markets with indivisible goods. However, for instructive purposes, we briefly consider the case of divisible goods to introduce relevant concepts. These can then be transferred to the indivisible case. Our aim is to give an intuitive graphical and analytical interpretation of how the aggregate valuation function is connected to the indirect utility function, the Lyapunov function and the market prices.

We consider a market with multiple bidders $i \in \mathcal{I}$ and multiple divisible goods $k \in \mathcal{K}$ with $|\mathcal{I}| = n$ and $|\mathcal{K}| = m$. The aggregate value function $v_{\mathcal{I}}$ is defined as the supremum convolution of concave functions $v_i : \mathbb{R}^m_{\geq 0} \to \mathbb{R}$ where v_i is the value function of the ith bidder.

$$v_{\mathcal{I}}(s) = \max_{\{x_i\}_{i \in \mathcal{I}}} \left\{ \sum_{i \in \mathcal{I}} v_i(x_i) \mid x_i \in \mathbb{R}^m_{\geq 0} \text{ and } \sum_{i \in \mathcal{I}} x_i = s \right\}.$$

By compactness and continuity, the maximum exists. Concavity implies that $v_i((1-\alpha)x + \alpha y) \geq (1-\alpha)\ v_i(x) + \alpha\ v_i(y)$ with $x, y \in \mathbb{R}_{\geq 0}$ and $\alpha \in (0, 1)$. The economic interpretation of a concave valuation function is that it exhibits decreasing marginal valuations. Since every function v_i is concave, also their convolution $v_{\mathcal{I}}$ is concave.

The aggregate indirect utility is defined as $u_{\mathcal{I}}(p) = \sum_i u_i(p)$ and the aggregate demand set is given by the Minkowski sum $D_{\mathcal{I}}(p) = \sum_i D_i(p)$.

For the sake of simplicity of the following graphical interpretation of indirect utility and the concept of conjugacy, we consider a market with multiple bidders but only a single divisible good $x \in \mathbb{R}_{\geq 0}$. However, our explanations carry over directly to markets with multiple goods. It is also worth mentioning that in the presence of only a single bidder i the aggregate valuation function $v_{\mathcal{I}}$ becomes the individual valuation function v_i of the single bidder. Thus, even though the following example illustrates the aggregate valuation and indirect utility function of multiple bidders, it similarly applies to the valuation and indirect utility function of an individual bidder.

In our example, we assume $v_{\mathcal{I}}(x) = \ln(x+1)$. It is well known that for concave functions $v_{\mathcal{I}}$ local optimality implies global optimality and this yields efficient optimization algorithms.

At a given price, every rational bidder $i \in \mathcal{I}$ only demands a quantity of good x which maximizes her utility at this price. The utility of such a quantity is described by the indirect utility function $u_i(p) = \max_x\{v(x) - \langle p, x \rangle\}$, which is convex as it is the maximum of affine linear functions. As the aggregate indirect utility function $u_{\mathcal{I}}(p)$ is a sum of convex functions, it must also be convex.

A quantity x^* is demanded at prices p if and only if $v_{\mathcal{I}}(x^*) - \langle p, x^* \rangle \geq v_{\mathcal{I}}(x) - \langle p, x \rangle$ for all $x \in \mathbb{R}$. When rearranging terms to $v_{\mathcal{I}}(x^*) + \langle p,\ x - x^* \rangle \geq v_{\mathcal{I}}(x)$, it becomes clear that the left-hand side of the inequality describes the tangent at $v_{\mathcal{I}}(x^*)$ (see Figure 1). In other words, a quantity x^* is demanded at prices p whenever the slope of the tangent at $v_{\mathcal{I}}(x^*)$ equals the price p. The aggregate utility of quantity x^* is given by $\pi_{\mathcal{I}}(x^*, p) = v_{\mathcal{I}}(x^*) - \langle p,\ x^* \rangle$. As $x^* \in D_{\mathcal{I}}(p)$, the aggregate utility $\pi_{\mathcal{I}}(x^*, p)$ equals the aggregate indirect utility $u_{\mathcal{I}}(p)$. The graphical interpretation of the aggregate indirect utility function $u_{\mathcal{I}}(p)$ is the intercept of the tangent at $v_{\mathcal{I}}(x^*)$ with the ordinate.

We can now compute the quantity of good x that generates maximum utility at prices p. In our illustrative example with $v_{\mathcal{I}}(x) = \ln(x+1)$, the aggregate utility $\pi_{\mathcal{I}}(x, p) = \ln(x+1) - \langle p,\ x \rangle$ at given prices p is maximized when

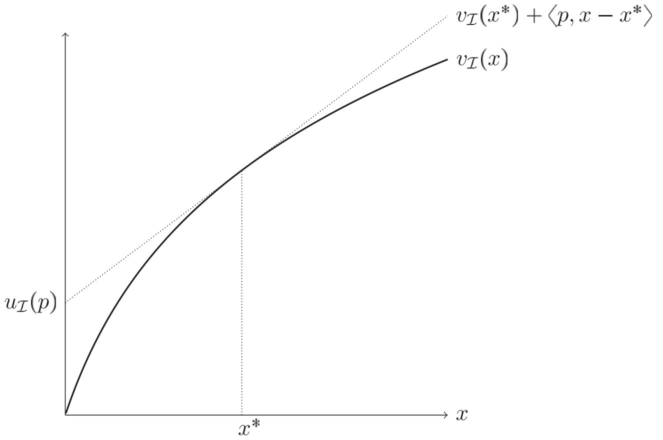

FIGURE 1 Graphical representation of $v_{\mathcal{I}}(x) = \ln(x+1)$ with tangent at $v_{\mathcal{I}}(x^*)$

$\partial\pi_{\mathcal{I}}/\partial x = 1/(x+1) - p = 0$. This means, at a price of $p = 1/3$ for example, the total utility $\pi_{\mathcal{I}}$ is maximized for a demand of $x^* = 2$. Thus, the aggregate indirect utility function at prices $p = 1/3$ equals $u_{\mathcal{I}}(1/3) = \pi_{\mathcal{I}}(2,\ 1/3) = \ln(3) - 2/3$. The concave conjugate (or Legendre transformation) of $v_{\mathcal{I}}$ is defined as $v^\bullet_{\mathcal{I}}(p) = \min_x\{\langle p,x\rangle - v_{\mathcal{I}}(x)\}$, which is the aggregate indirect utility function multiplied by -1. We also note that convex and concave conjugates are connected via $v^\bullet_{\mathcal{I}}(p) = -(-v_{\mathcal{I}})^*(-p)$, so $u_{\mathcal{I}}(p) = (-v_{\mathcal{I}})^*(-p)$. From these results, we can make the following connection: In order to construct the concave conjugate $v^\bullet_{\mathcal{I}}(p)$ of $v_{\mathcal{I}}(x) = \ln(x+1)$ for a fixed p, we must calculate the minimum of $\langle p, x\rangle - \ln(x+1)$. Taking the derivative, we see that a minimizing x must solve $x = 1/p - 1$, so we get $v^\bullet_{\mathcal{I}}(p) = 1 - p + \ln(p)$ and consequently $u_{\mathcal{I}}(p) = -v^\bullet_{\mathcal{I}}(p) = p - \ln(p) - 1$. For a given price of $p = 1/3$ the reader may verify that the bidders' aggregate indirect utility equals $u_{\mathcal{I}}(1/3) = 1/3 - \ln(1/3) - 1 = \ln(3) - 2/3$, which is in line with our calculations above.

Unlike in this single-item example, the price p is not known in an auction setting. Instead, the auctioneer tries to find a price vector p^* for which the supply s is a maximizer of the aggregate utility function $\pi_{\mathcal{I}}(x,\ p^*)$. Note that such a p^* is a Walrasian equilibrium price vector, because s maximizes $\pi_{\mathcal{I}}(s,\ p^*) = v_{\mathcal{I}}(s) - \langle p^*, s\rangle$ and the aggregate demand of the bidders equals the supply s.

We will now return to a market with multiple divisible goods $k \in \mathcal{K}$. First, we introduce important notions from convex analysis.

Definition 2 Let $f: \mathbb{R}^d \to \mathbb{R} \cup \{+\infty\}$ be a convex function. The *subdifferential* of f at x is the set of all tangents of f at x:

$$\partial f(x) = \{y \in \mathbb{R}^d | f(x') \geq f(x) + \langle y, x' - x\rangle \forall x' \in \mathbb{R}^d\}.$$

Any element of $\partial f(x)$ is called a *subgradient*. The *convex conjugate* or *Legendre transform* of f is the convex function

$$f^*(y) = \sup_{x\in\mathbb{R}^d} \langle y, x\rangle - f(x).$$

Under additional mild assumptions on the convex function f, the conjugate of the conjugate is again f, $f^{**} = f$, and subdifferentials of f and f^* are connected in the following way: $y \in \partial f(x) \Leftrightarrow x \in \partial f^*(y)$. For more details, we refer to Rockafellar (2015). The concave conjugate defined above and the convex conjugate are related as follows: If g is concave, then $g^\bullet(y) = -(-g)^*(-y)$. In particular, we have for the indirect utility function $u_{\mathcal{I}}(p) = (-v_{\mathcal{I}})^*(-p)$. We make the following important observation: The bundle x is in the demand set $D_{\mathcal{I}}(p)$, if and only if $v_{\mathcal{I}}(x) - \langle p,\ x\rangle \geq v_{\mathcal{I}}(x') - \langle p,\ x'\rangle$ for all $x' \in \mathbb{R}^{|\mathcal{K}|}$. By rearranging terms we see that this is equivalent to $-v(x') \geq -v(x) + \langle -p,\ x' - x\rangle$ and thus to $-p \in \partial(-v_{\mathcal{I}})(x)$. Convex analysis tells us that this is equivalent to $x \in \partial(-v_{\mathcal{I}})^*(-p) = -\partial u_{\mathcal{I}}(p)$. Consequently, demand sets are equal to subdifferentials of the indirect utility function—a fact that allows us to interpret auctions as descent algorithms.

The Lyapunov function was a central concept already in the early literature on general equilibrium theory (Arrow & Hahn, 1971). The same function plays a central role in more recent auction designs for markets with indivisible goods (Ausubel, 2006). Since this function plays such a central role, we introduce it in detail for the continuous case.

Definition 3 (Lyapunov function). The *Lyapunov function* is defined as $L(p) = \sum_{i\in\mathcal{I}} u_i(p) + \langle p, s\rangle$, where s is the supply and $u_i(p)$ is the indirect utility function of bidder $i \in \mathcal{I}$ at prices p.

The Lyapunov function has its roots in the dynamical systems literature (La Salle & Lefschetz, 2012). Since the indirect utility $u_i(p)$ is convex in p, also the Lyapunov function is convex, because it is the sum of convex functions. For convex functions such as $L(p)$ the vector p^* minimizes L iff 0 is a subgradient at p^*. The first-order condition for $L(p)$ yields $-\sum_{i\in\mathcal{I}} x_i + s = 0$, where $x_i \in D_i(p)$.
$\forall i \in \mathcal{I}$. In words, the prices are minimized when supply equals demand:

Proposition 1 *A vector $p^* \in \mathbb{R}^m$ is a Walrasian equilibrium price vector for supply s if and only if it is a minimizer of the Lyapunov function $L(p) = u_{\mathcal{I}}(p) + \langle p,\ s\rangle$.*

Proof If there is a Walrasian equilibrium, then $\sum_{i\in\mathcal{I}} x_i = s$ and $x_i \in D_i(p^*)$ need to hold. The minimizer p^* of $L(p)$ requires that $\partial L(p) = s - \sum_{i\in\mathcal{I}} x_i = 0$, which is equivalent to the first condition of a Walrasian equilibrium. Also, when $L(p) = \sum_{i\in\mathcal{I}} \max_{x_i}\{v_i(x_i) - \langle p, x_i\rangle\} + \langle p, s\rangle$ attains the minimum, then each bidder is assigned a bundle x_i that maximizes her utility $v_i(x_i) - \langle p,\ x_i\rangle$. This implies $x_i \in D_i(p^*)$ for all i, so that the second condition of a Walrasian equilibrium is fulfilled. Thus, if $L(p)$ is minimized then both conditions of a Walrasian equilibrium are satisfied. By reversing the argument it becomes evident that any price vector p^* supporting s in a Walrasian equilibrium is also a minimizer for $L(p)$. ■

Similar results can be found in Ausubel and Milgrom (2006) or later in Murota (2016). One way to find Walrasian equilibria is now to minimize the Lyapunov function. Since we can interpret the subdifferential of u_i at price p as the demand set at this price—for an auction setting it is natural to utilize standard subgradient methods for (approximately) minimizing $L(p)$—computing subgradients is then equivalent to asking bidders for their demand sets at a given price. Note that it is in general not possible to compute *exact* minimizers to general convex functions—algorithms for minimizing a convex function f can in general only provide complexity bounds for finding an ε-approximate solution x', in the sense that

$$f(x') \leq \varepsilon + \min_x f(x).$$

Note that in general x' does not even have to be close to the true minimizer x without additional assumption on f. Since the aim of our treatment of divisible economies is mainly to motivate the ideas in the indivisible case, we will not go into more detail here. If no additional regularity assumptions on L are imposed, it can be shown that finding ε-approximate solutions has a worst-case running time of $\Theta(1/\varepsilon^2)$ (Nesterov, 2018). Interestingly, for markets with indivisible goods where Walrasian equilibria exist, we will show that the Lyapunov function equals the dual of the allocation problem.

Central results of convex economic theory with divisible goods are reasonable approximations to large economies where nonconvexities vanish in the aggregate (Starr, 1969). However, most markets are such that indivisibilities and nonconvexities matter. As one would assume, the analysis of markets with indivisible items has proven much harder.

4 | EXISTENCE OF WALRASIAN EQUILIBRIA WITH INDIVISIBLE GOODS

In this section, we discuss sufficient and necessary conditions for the individual value functions of bidders such that Walrasian equilibria exist in markets with indivisible goods.

4.1 | Conditions on aggregate value functions

A simple multi-item market with remarkable properties is the assignment market by Shapley and Shubik (1971). In assignment markets each bidder can bid on multiple items but wants to win at most one (aka. *unit-demand*). As a consequence, the allocation problem reduces to an assignment problem, that is, the problem of finding a maximum weight matching in a weighted bipartite graph. On an aggregate level, the LP relaxation of the assignment problem is always integral. This is a consequence of the unit demand on an individual level and the resulting total unimodularity of the constraint matrix, and this is a sufficient condition for the existence of Walrasian prices. The environment of assignment markets allows for incentive-compatible auctions. Besides, simple ascending clock auctions yield bidder-optimal Walrasian prices (Demange et al., 1986).

4.1.1 | The allocation problem

Let us first extend the assignment market to a more general multi-item, multi-unit market which allows for package bids. Let $\mathcal{X}_i \subseteq \mathbb{Z}^m_{\geq 0}$ denote all bundles for which bidder i submitted a bid. For simplicity, we make the natural assumption that every bidder submits a bid with value 0 for the empty bundle. Let $z_i(x) \in \{0, 1\}$ be a binary decision variable denoting whether bidder i wins bundle $x \in \mathcal{X}_i$. The allocation or winner determination problem WDP can then be written as an integer program as follows:

$$\begin{aligned} \max \quad & \sum_{i\in\mathcal{I}}\sum_{x\in\mathcal{X}_i} v_i(x)z_i(x) && && \text{(WDP)} \\ \text{s.t.} \quad & \sum_{x\in\mathcal{X}_i} z_i(x) \leq 1 && \forall i \in \mathcal{I} && (\pi_i) \\ & \sum_{i\in\mathcal{I}}\sum_{x\in\mathcal{X}_i} x(k)z_i(x) \leq s(k) && \forall k \in \mathcal{K} && (p(k)) \\ & z_i(x) \in \{0,1\} && \forall i \in \mathcal{I}, \forall x \in \mathcal{X}_i \end{aligned}$$

For a given supply s the WDP determines an allocation of bundles to bidders maximizing social welfare. The LP relaxation RWDP in standard form replaces $z_i(x) \in \{0, 1\}$ by $z_i(x) \geq 0$ and introduces additional slack variables. We use the standard form with slack variables (a_i, b_k) because it will be helpful in our algorithmic treatment of the subject.

$$\begin{aligned} \max \quad & \sum_{i\in\mathcal{I}}\sum_{x\in\mathcal{X}_i} v_i(x)z_i(x) && && \text{(RWDP)} \\ \text{s.t.} \quad & \sum_{x\in\mathcal{X}_i} z_i(x) + a_i = 1 && \forall i \in \mathcal{I} && (\pi_i) \\ & \sum_{i\in\mathcal{I}}\sum_{x\in\mathcal{X}_i} x(k)z_i(x) + b_k = s(k) && \forall k \in \mathcal{K} && (p(k)) \\ & z_i(x), a_i, b_k \geq 0 && \forall i \in \mathcal{I}, \forall x \in \mathcal{X}_i, \forall k \in \mathcal{K} \end{aligned}$$

In contrast to the assignment problem where bidders have unit demand, the RWDP does not yield integer solutions in general.

Example 1 Consider a market with three items $\mathcal{K} = \{A, B, C\}$ and two bidders with valuations v_1 and v_2

	$x_\emptyset$	x_A	x_B	x_C	x_{AB}	x_{AC}	x_{BC}	x_{ABC}
x	(0, 0, 0)	(1, 0, 0)	(0, 1, 0)	(0, 0, 1)	(1, 1, 0)	(1, 0, 1)	(0, 1, 1)	(1, 1, 1)
$v_1(x)$	0	1	2	1	2	2	2	2
$v_2(x)$	0	1	2	2	3	2	3	3

The optimal solution of the RWDP given these valuations is fractional:

$z_1(x_B) = z_1(x_{AC}) = z_2(x_C) = z_2(x_{AB}) = 0.5$ with all other decision variables set to 0. The optimal value of the RWDP with respect to this fractional solution is 4.5. An optimal integral solution (e.g., assigning bundle x_{AC} to the first and x_B to the second bidder) only leads to a social welfare of 4.

Let us also introduce the dual DRWDP of the RWDP.

$$\begin{aligned} \min \quad & \sum_{i\in\mathcal{I}} \pi_i + \sum_{k\in\mathcal{K}} s(k)p(k) && && \text{(DRWDP)} \\ \text{s.t.} \quad & \pi_i + \sum_{k\in\mathcal{K}} x(k)p(k) \geq v_i(x) && \forall i \in \mathcal{I}, \forall x \in \mathcal{X}_i && (z_i(x)) \\ & \pi_i \geq 0 && \forall i \in \mathcal{I} && (a_i) \\ & p(k) \geq 0 && \forall k \in \mathcal{K} && (b_k) \end{aligned}$$

We will draw on these models in the subsequent sections.

4.1.2 | Integrality of the linear program

Bikhchandani and Mamer (1997) describe a multi-item, single-unit market. Their central theorem shows that there exist clearing prices for the indivisible single-unit problem if and only if the RWDP has an integer solution. In this case, the set of equilibrium prices is the set of solutions to the dual LP projected to the price coordinates. The result can be proven via the strong duality theorem in linear programming (Blumrosen & Nisan, 2007). As was already noted by Bikhchandani and Mamer (1997), the result for multi-item, *multi*-unit markets also directly follows from their result, by considering each of the multiple units as separate items. As the proof is a particularly nice application of duality theory, we provide a direct proof in the Appendix. Note that this theorem proves the welfare theorems from general equilibrium theory (see Theorem 1).

Theorem 2 *Walrasian prices exist for the supply s if and only if the RWDP has an optimal integral solution.*

The proof can be found in Appendix A.

As indicated, the RWDP typically does not yield an integral solution, and there can be a significant integrality gap between the objective function value of the RWDP and that of the optimal integer program WDP. In the next sections, we will discuss conditions on the individual value functions, which yield integral solutions of the RWDP and Walrasian prices.

Before we do this, let us return to the Lyapunov function that has proven so helpful in our analysis of markets with divisible goods. A minimizer to this function yielded the Walrasian prices in Section 3, where we analyzed markets with divisible goods. It turns out that the Lyapunov function is actually equivalent to the DRWDP, as we show in the following proposition.

Proposition 2 *A vector $p^* \in \mathbb{R}^m$ minimizes the DRWDP if and only if it is a minimizer of the Lyapunov function $L(p) = u_{\mathcal{I}}(p) + \langle p, s\rangle$.*

Proof We can substitute the utilities π_i in the dual objective function $\min \sum_{i\in\mathcal{I}} \pi_i + \sum_{k\in\mathcal{K}} s(k)p(k)$ by the tight dual constraints $\pi_i = v_i(x) - \sum_{k\in\mathcal{K}} x(k)p(k)$ of the optimal DRWDP and get the following convex function:

$$\min_p \sum_{i\in\mathcal{I}} \max_{x\in\mathbb{Z}^m_{\geq 0}} \left[v_i(x) - \sum_{k\in\mathcal{K}} x(k)p(k) \right] + \sum_{k\in\mathcal{K}} s(k)p(k). \tag{4.1}$$

Note that this is equivalent to minimizing the Lyapunov function $L(p) = \sum_{i\in\mathcal{I}} u_i(p) + \langle p, s\rangle$. Obviously, $\langle p, s\rangle$ in $L(p)$ is equal to $\sum_{k\in\mathcal{K}} s(k)p(k)$, and $u_i(p)$ equals $\max_{x\in\mathbb{Z}^m_{\geq 0}} \left[v_i(x) - \sum_{k\in\mathcal{K}} x(k)p(k) \right]$ for every bidder i. Since the equivalence of the Lyapunov function and the DRWDP holds for any price vector p, minimizing prices of the Lyapunov function also constitutes a minimal solution to the DRWDP and vice versa. ■

In summary, both the Lyapunov function and the LP approach yield equilibrium prices, and such prices are minimizers of both problems. We will leverage this insight, when we analyze auction algorithms to solve the RWDP in Section 5.

4.2 | Conditions for individual value functions

In practical applications a market designer often wants to understand which assumptions on the individual value functions v_i allow for integer solutions of the LP relaxation and Walrasian prices. Discrete convex analysis identifies classes of convex functions defined on a subset of the discrete lattice $\mathbb{Z}^m$, which allow for integrality and efficient optimization algorithms.

First, we discuss single-unit, multi-item auctions. There are several classes of integrally convex functions such as separable-convex functions on $\mathbb{Z}^m$ or gross substitutes set functions on $\{0, 1\}^m$, which yield a discrete concave aggregate value function $v_{\mathcal{I}}$ and integral solutions of the RWDP, such that Walrasian equilibria exist.

4.2.1 | Single-unit multi-item auctions

Let us first define monotonicity and submodularity, two well-known properties of set functions that allow for efficient function minimization.

Definition 4 For a finite set $\mathcal{K}$ of items, the set function $v : 2^{\mathcal{K}} \rightarrow \mathbb{R}$ is

- *monotone* if $v(S) \leq v(T)$ for all $S, T \subseteq \mathcal{K}$ with $S \subseteq T$,

- *submodular* if $v(S \cup \{k\}) - v(S) \geq v(T \cup \{k\}) - v(T)$ for all $S, T \subseteq \mathcal{K}$ with $S \subseteq T$ and for all $k \notin T$.

In the above definition, submodularity can be understood as diminishing marginal values. Alternatively, submodularity can be defined as $v(S) + v(T) \geq v(S \cup T) + v(S \cap T)$ for all S, T. The vector notation $v : \{0, 1\}^m \rightarrow \mathbb{R}$ in the single-unit case maps a set S to a vector $x \in \{0, 1\}^m$ by setting $x(k) = 1$ whenever $k \in S$ and $x(k) = 0$ otherwise.

It is well-known that the minimization of unconstrained submodular functions can be done in polynomial time, for example via the ellipsoid method (Grötschel et al., 1981). The ellipsoid method is notoriously slow in practice. However, there are also more effective algorithms such as the Fujishige-Wolfe algorithm (Chakrabarty et al., 2014) and specialized subgradient methods (Chakrabarty et al., 2017). Unfortunately, even when submodularity and monotonicity are satisfied, this does not guarantee the integrality of a welfare maximization problem such as the RWDP.

Example 2 The reader may verify that the valuation functions of both bidders in example 1 satisfy monotonicity and submodularity. However, the optimal solution of the RWDP is not integral.

The subset of submodular valuations called gross substitutes valuations, however, has this desirable property. Gross substitutes roughly means that a bidder regards the items as substitute goods or independent goods but not complementary goods.

Definition 5 (Gross substitutes, GS). Let p denote the prices on all items, with item k demanded by bidder i if there is some bundle S, with $k \in S$, for which S maximizes the utility $v_i(S') - \sum_{j \in S'} p(j)$ across all bundles $S' \subseteq \mathcal{K}$. The gross substitutes condition requires that, for any prices $p' \geq p$ with $p'(k) = p(k)$, if item $k \in \mathcal{K}$ is demanded at the prices p then it is still demanded at p'.

The definition includes both substitute goods and independent goods, but rules out complementary goods.[2]

Example 3 Consider a market with three items $\mathcal{K} = \{A, B, C\}$ and a single bidder with a valuation function v fulfilling the gross substitutes condition

	$x_\emptyset$	x_A	x_B	x_C	x_{AB}	x_{AC}	x_{BC}	x_{ABC}
x	(0, 0, 0)	(1, 0, 0)	(0, 1, 0)	(0, 0, 1)	(1, 1, 0)	(1, 0, 1)	(0, 1, 1)	(1, 1, 1)
$v(x)$	0	1	2	3	3	3	5	5

At prices $p = (0, 1, 2)$ the bidder's indirect utility is $u(p) = 2$ and the bidder's demand set is given by $D(p) = \{x_{AB}, x_{BC}, x_{ABC}\}$, that is, items A, B, and C are demanded as for each item there exists at least one bundle in the demand set containing the item. If the price for item A is raised to 1 but stays constant for items B and C, then the gross substitutes condition implies that items B and C must still be demanded at the new prices $p' = (1, 1, 2)$. This is obviously true as the demand set at the new prices p' is given by $D(p') = \{x_{BC}\}$. Note that price vectors p and p' were only chosen for illustrative purposes. In fact, valuation function v satisfies the gross substitutes condition for any price vectors $p, p' \in \mathbb{R}^3_{\geq 0}$ with $p' \geq p$.

Kelso and Crawford (1982) show that if all agents have GS valuations, then a Walrasian equilibrium always exists, which implies that the RWDP has an optimal integral solution. Ausubel and Milgrom (2002) prove that a bidder has GS valuations if and only if the indirect utility function u is submodular. Gross substitutes appear to be a rather restricted type of valuations, but it contains important subclasses such as unit-demand valuations (Shapley & Shubik, 1971) and additive valuations. Gul and Stacchetti (1999) show that GS excludes complementarity between goods and show equivalence with the so called single improvement property. The latter property states that whenever a bundle is not optimal at the given prices, then a better bundle can be found which is derived from the original one by performing any of the following operations: removing an item, adding an item, or doing both. Leme (2017) provides a survey of the extensive literature on the gross substitutes condition and its alternative definitions for multi-item, single-unit markets, and show that additive valuations ⊂ GS ⊂ submodular valuations ⊂ subadditive valuations. We also refer to Shioura and Tamura (2015) for an extensive survey of GS.

Sun and Yang (2006) identify the gross substitutes and complements (GSC) condition, which also guarantees for Walrasian equilibria in single-unit, multi-item markets. It describes an exchange economy with two classes of goods, where each class only contains substitutes, but there are complements across these classes of goods. Teytelboym (2014) generalizes the GSC condition in the sense that goods are partitioned into more than two classes. His generalized version of the GSC condition is satisfied if it is possible to partition goods into several classes so that whenever considering the bidders' valuations for items contained in only two of these classes in isolation, there exist some bidders for which these valuations satisfy the GSC condition.

4.2.2 | Multi-unit ulti-item auctions

Let us now concentrate on more general conditions for $x \in \mathbb{Z}^m_{\geq 0}$ rather than $x \in \{0, 1\}^m$. $A \subset \mathbb{Z}^m$ is *integrally convex* if

[2]Sometimes the word "gross" used by Kelso and Crawford (1982) is omitted, but it is useful to distinguish the single-unit case from substitutes valuations in other environments, such as the strong substitutes definition that we will introduce later.

$A = (conv\,A) \cap \mathbb{Z}^m$. First, we define the *convex closure* $\bar{f}$ of f as

$$\bar{f}(x) = \sup_{p\in\mathbb{R}^m,\alpha\in\mathbb{R}} \{\langle p, x\rangle + \alpha \mid \langle p, y\rangle + \alpha \le f(y)\ \forall y \in \mathbb{Z}^m\}.$$

Geometrically, the epigraph of $\bar{f}$ is the convex hull of the epigraph of f. If the convex closure coincides with f on the set of integer vectors, that is, if $f(x) = \bar{f}(x)$ for all $x \in \mathbb{Z}^m$, f is called *convex-extensible*. In the same way, we can define the concave closure of f by $-\overline{(-f)}$. The definition can be restricted to the integral neighborhood of a bundle $x \in \mathbb{R}^m$ and is then referred to as a *local convex extension* $\tilde{f}$ of f (Murota, 2003, Chap. 3). Formally, set $N(x) = \{y \in \mathbb{Z}^m \mid \lfloor x(k)\rfloor \le y(k) \le \lceil x(k)\rceil \forall k = 1, \ldots, m\}$. Then the local convex extension is given by

$$\tilde{f}(x) = \sup_{p\in\mathbb{R}^m,\alpha\in\mathbb{R}} \{\langle p, x\rangle + \alpha \mid \langle p, y\rangle + \alpha \le f(y)\ \forall y \in N(x)\}.$$

Definition 6 A function $f : \mathbb{Z}^m \to \mathbb{R}$ is called *integrally convex* if the local convex extension of f is convex, or *integrally concave* if the function $-f$ is integrally convex.

Integrally convex functions share with convex functions the property that local minima are also global minima (Murota, 2016). We have already seen in the divisible case that concavity of the valuation functions is necessary for equilibrium prices to exist. We also want to make this connection here in the indivisible case, by explaining how convexity is related to integrality of the WDP—which is necessary and sufficient for the existence of equilibrium prices. To start with, consider the aggregate valuation function $v_{\mathcal{I}}(s)$, given by the value of the WDP for the supply s, and the "relaxed" aggregate valuation function $\tilde{v}_{\mathcal{I}}(s)$, given by the value of the RWDP at s. Note $\tilde{v}_{\mathcal{I}}$ is well-defined for all *real* supply vectors $s \ge 0$ and attains finite values at each such s. A central observation is the following: $\tilde{v}_{\mathcal{I}}$ is the concave extension of $v_{\mathcal{I}}$. This shows that $v_{\mathcal{I}}$ is concave-extensible, and thus $v_I = \tilde{v}_{\mathcal{I}}$ if and only if for every integral supply vector s, the RWDP has an integral solution, which—as we have seen—is equivalent to the existence of equilibrium prices. While the stronger assumption of integral concavity is not necessary for the existence of equilibrium prices, it is not hard to imagine, that this property is of importance for the algorithmic problem of computing equilibrium prices. Loosely speaking, since the value of the concave extension can then be evaluated at any point s by considering an easy to characterize neighborhood of s, the computation of subgradients of $v_{\mathcal{I}}$ gets much simpler. Unfortunately, concave extensibility, and even integral concavity of the individual valuation functions does not suffice to guarantee concave extensibility of the aggregate valuation function, or equivalently, existence of equilibrium prices. It is thus of central importance to identify conditions on the individual valuations that imply concave extensibility of the aggregate valuation, or equivalently integrality of the RWDP.

Definition 7 A function $f : \mathbb{Z}^m \to \mathbb{R} \cup \{\infty\}$ is said to be $M^\natural$-convex if for $x, y \in \text{dom}f$ and $j \in supp^+(x - y)$

(i) $f(x) + f(y) \ge f(x - \mathbb{1}_j) + f(y + \mathbb{1}_j)$ or
(ii) $f(x) + f(y) \ge f(x - \mathbb{1}_j + \mathbb{1}_k) + f(y + \mathbb{1}_j - \mathbb{1}_k)$ for some $k \in supp^-(x - y)$.

A function f is $M^\natural$-concave if the function $-f$ is $M^\natural$-convex. A set $X \subseteq \mathbb{Z}^m$ is an $M^\natural$-convex set if its indicator function δ_X is $M^\natural$-convex.

Here $\mathbb{1}_j$ denotes the jth unit vector, whereas the positive and negative support are defined as $supp^+(x) = \{k \in \mathcal{K} | x(k) > 0\}$ and $supp^-(x) = \{k \in \mathcal{K} | x(k) < 0\}$, respectively. The effective domain is $\text{dom}f = \{z \in \mathbb{Z}^m | f(z) \ne \infty\}$. An $M^\natural$-convex function is integrally convex, and thus convex-extensible (Murota, 2003, Theorem 6.42). Since the exchange property (ii) is closely related to the exchange axiom of a matroid, the M stands for "matroid". It means that if we add the jth unit-vector to one point x and exchange it with the ith unit vector of another point y, then the function value decreases or stays the same. Fujishige & Yang, 2003 showed that for the single-unit case the GS condition is equivalent to $M^\natural$-concavity.

Theorem 3 (Fujishige and Yang (2003)). A value function $v : \{0, 1\}^m \to \mathbb{R}$ satisfies the GS condition if and only if it is an $M^\natural$-concave function.

This equivalence extends to multi-unit extensions of the gross substitutes property. Milgrom and Strulovici (2009) distinguish between *weak* and *strong* substitutes. The weak substitutes condition can be seen as the natural extension of the original gross substitutes property to the multi-unit setting by simply quantifying the demand for every item. Note however, that weak substitutes do not correspond to $M^\natural$ functions anymore (Shioura & Tamura, 2015). The strong substitutes condition, on the other hand, transforms a multi-unit to a single-unit valuation function by treating each copy of a good as an individual item. Whenever the corresponding single-unit valuation function fulfills the original gross substitutes property (as defined by Kelso and Crawford (1982)), the multi-unit valuation function satisfies the strong substitutes condition.

Definition 8 (Strong substitutes, SS). Let $\mathcal{K} = \{k_1, k_2, \ldots, k_m\}$ be the set of items with $d_i \in \mathbb{N}$ denoting the number of units available of item k_i. Treating each copy of a good as an individual item leads to the definition of a set $\mathcal{K}_s = \{(k_i, z) | k_i \in \mathcal{K}, 1 \le z \le d_i\}$. A multi-unit valuation function $v : \mathbb{N}_0^m \to \mathbb{R}$ can then be transformed to a single-unit valuation function $v_s : \{0,1\}^{\mathcal{K}_s} \to \mathbb{R}$ by setting $\mathrm{v_s}(x_s) = v(x)$

for $x_s \in \{0,1\}^{\mathcal{K}_s}$ where $x(i) = \sum_{z=1}^{d_i} x_s(k_i, z)$. The valuation v fulfills the strong substitutes condition if v_s is a gross substitutes valuation function.

There exist many equivalent definitions of the strong substitutes condition, among them the binary single-improvement property as shown by Milgrom and Strulovici (2009).

Danilov et al. (2001) and Milgrom and Strulovici (2009) show that a Walrasian equilibrium exists for every finite set of strong substitutes valuations. Ausubel (2006) shows that in case of strong substitutes valuations the Lyapunov function is submodular which ensures the existence of a bidder-optimal Walrasian price vector. While the strong substitutes property is a sufficient condition for the existence of Walrasian equilibria, it is not a necessary one and alternatives exist.

Shioura and Yang (2015) extend the gross substitutes and complements (GSC) condition to a multi-unit and multi-item economy with two classes of items, where units of the same type are substitutable, whereas goods across two classes are complementary. When there is only one class of indivisible goods, their generalized gross substitutes and complements (GGSC) condition becomes identical to the strong-substitute valuation of Milgrom and Strulovici (2009). Further, if each type of good has only one unit, it becomes the gross substitute condition of Kelso and Crawford (1982).

Baldwin and Klemperer (2019) provide an innovative approach characterizing preferences where Walrasian equilibria exist. Instead of working with the value functions, their framework is based on properties of the geometric structure of the regions in the price space where a bidder demands different bundles. A demand type is defined by a list of vectors that give the possible ways in which the individual or aggregate demand can change in response to a small price change. Intuitively, given a valuation v_i, consider the set $\mathcal{L}_i = \{p | |D_i(p)| > 1\}$ of all prices at which more than one bundle is in the bidder's demand set. $\mathcal{L}_i$ can be shown to form a so-called *polyhedral complex*, and in particular is a union of hyperplanes, which splits price space into multiple full-dimensional regions where a unique bundle is demanded, which are called *unique demand regions (UDRs)*. Now given a set $\mathcal{D}$ of integral vectors, v_i is of the demand type defined by $\mathcal{D}$ if all normals of all hyperplanes in $\mathcal{L}_i$ are integral multiples of vectors in $\mathcal{D}$.[3] We say that the demand type defined by $\mathcal{D}$ is *unimodular* if any linear independent subset of vectors in $\mathcal{D}$ can be extended with integral vectors to a basis with determinant in $\{-1, 1\}$. It can be shown, that if participants' valuations are concave and all have the same unimodular demand type $\mathcal{D}$, then a Walrasian equilibrium exists. There are several proofs for the unimodularity theorem, see Baldwin and Klemperer (2019); Danilov et al. (2001); Tran and Yu (2015). The authors further show that an equilibrium is guaranteed for more classes of pure complements than of pure substitutes preferences. Note that while all agents being drawn from an equal certain valuation type (SS, GGSC, pure complements) allows for Walrasian equilibria, agent valuations drawn from a mixture of these types in general do not allow for one. Unimodularity of the demand types is a sufficient condition for the existence of Walrasian equilibria. Remarkably, it is also *necessary*: Given valuations of the agents, there exist equilibrium prices for *every* given supply if and only if the agents' demand types are unimodular. Again, whenever the unimodularity condition holds, the optimal solution to the RWDP is integral.

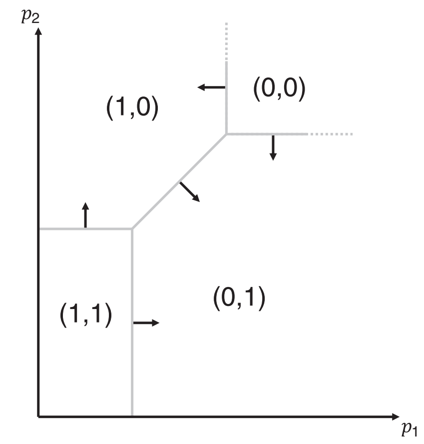

FIGURE 2 Illustration of $\mathcal{L}_i$ (gray). For each indifference hyperplane, we indicate one of the two normal vectors associated with this hyperplane. We can directly see that these normals all lie in $\mathcal{D}$ as defined in Example 4. The tuples (x_1, x_2) indicate the bundles that are demanded in the respective UDRs

Example 4 Consider a market with two items $\mathcal{K} = \{A, B\}$ and a single bidder with a valuation function v, given by the following table

	$x_\emptyset$	x_A	x_B	x_{AB}
x	(0, 0)	(1, 0)	(0, 1)	(1, 1)
$v(x)$	0	2	3	4

The set $\mathcal{L}$ is shown in Figure 2. We can see that v is of the demand type given by $\mathcal{D} = \{\pm(1,0), \pm(0,1), \pm(1,-1)\}$. It can be checked that $\mathcal{D}$ is actually unimodular.

4.2.3 | From individual to aggregate value functions

We now want to understand when we can expect individual value functions v_i to yield aggregate value functions $v_{\mathcal{I}}$ that are integrally concave. The aggregation of value functions is referred to as *convolution* (see Section 3).

[3]The normals of these hyperplanes have the following economic meaning: Consider a path in price space starting in some UDR. Each time the path crosses an indifference hyperplane, and thus entering another UDR, the demanded bundle changes by the normal vector of the crossed hyperplane, which points into the opposite direction of the price path. In Figure 2 for example, if the price path goes from the UDR (0, 1) to the UDR (1, 0) in a straight line, we cross the hyperplane with normal (1, −1), and of course (1, 0) = (0, 1) + (1, −1).

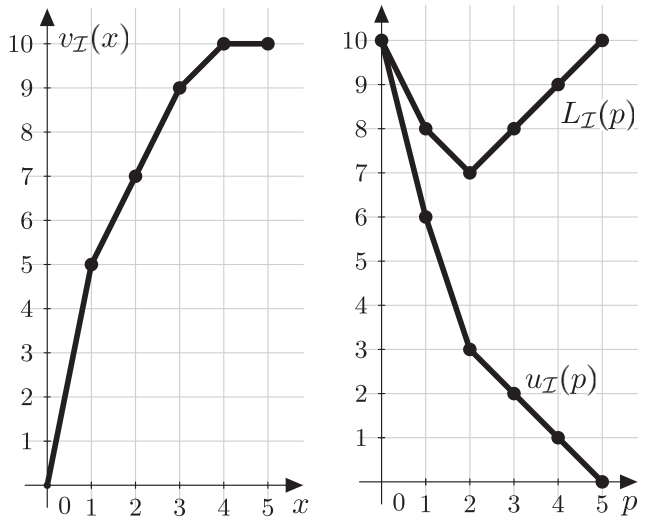

FIGURE 3 For a market with two units of a single indivisible item x, the figure shows the aggregate valuation function $v_{\mathcal{I}}(x)$, the aggregate utility function $u_{\mathcal{I}}$, and the Lyapunov function $L_{\mathcal{I}}(p)$. The Lyapunov function is minimized at $p = 2$, denoting the Walrasian equilibrium prices. Note that $p = 2$ is also the supergradient of $v_{\mathcal{I}}(x)$ at $x = 2$

Murota (2016)[p. 196] shows that if the individual value functions v_i of all bidders $i \in \mathcal{I}$ are $M^\natural$-concave, also their convolution is $M^\natural$-concave. Similarly, one can define the aggregate demand correspondence $D_{\mathcal{I}}(p)$, which is equal to the Minkowski sum $\sum_{i\in\mathcal{I}} D_i(p)$.

For $M^\natural$-concave functions there is a supergradient at any point that determines a Walrasian price p. To show this, let us consider an arbitrary bounded, integrally convex set $A \subset \mathbb{Z}^m_{\geq 0}$. Let $v_{\mathcal{I}}: A \to \mathbb{Z}$ be an $M^\natural$-concave valuation on this set. A bundle $x \in A$ is demanded at price $p \in \mathbb{R}^m$ iff $v_{\mathcal{I}}(x) - \langle p, x\rangle \geq v_{\mathcal{I}}(x') - \langle p, x'\rangle \ \forall x' \in A$, which is equivalent to $v_{\mathcal{I}}(x) + \langle p, x' - x\rangle \geq v_{\mathcal{I}}(x') \ \forall x' \in A$ (as for divisible goods in Section 3). Figure 3 now illustrates an integrally concave value function on the left and the resulting indirect utility function $u_{\mathcal{I}}(p)$ as well as the Lyapunov function $L_{\mathcal{I}}(p)$ for a single item on the right.

With indivisible items and an integrally concave aggregate value function $v_{\mathcal{I}}$, bundle x is demanded at p if and only if p is a supergradient of $v_{\mathcal{I}}$ at x. The superdifferential $\partial v_{\mathcal{I}}(x)$ of an integrally concave function $v_{\mathcal{I}} : \mathbb{Z}^m_{\geq 0} \to \mathbb{R} \cup \{-\infty\}$ at $x \in \text{dom } v_{\mathcal{I}}$ is defined as

$$\partial v_{\mathcal{I}}(x) = \{p \in \mathbb{R}^m_{\geq 0} \mid v_{\mathcal{I}}(y) - v_{\mathcal{I}}(x) \leq \langle p, y - x\rangle \ \forall y \in \mathbb{Z}^m_{\geq 0}\}.$$

The individual and aggregate value functions are nondecreasing such that the gradient p^* of the superdifferential is $p^* \geq 0$. With an integrally concave value function $v_{\mathcal{I}}$ there exists an integral equilibrium price vector p^* (Murota et al., 2016). The integrality of the prices follows from the fact that an integer-valued $M^\natural$-concave function $v_{\mathcal{I}}$ on $\mathbb{Z}^m_{\geq 0}$ has an integral subgradient at every point x in dom $v_{\mathcal{I}}$. As both $v_{\mathcal{I}}(x)$ and the subgradient at x are integral, the tangent at $v_{\mathcal{I}}(x)$ has an integral slope p, which can be verified in Figure 3.

An underlying assumption in the study of competitive equilibria is that agents are price-takers, that is, agents honestly report their true demand in response to prices in each round of an auction. Mechanism design, a line of research initiated by Hurwicz (1972), wants to understand how such markets perform when agents are strategic about their demands. Unfortunately, Gul and Stacchetti (1999) showed that even if goods are substitutes, Walrasian markets are not incentive-compatible. The assignment market, where bidders have unit-demand is an exception where straightforward bidding is actually an ex post equilibrium (Shapley & Shubik, 1971; Demange et al., 1986).

5 | ALGORITHMIC AUCTION MODELS

Auctions can be understood as algorithms to solve a welfare maximization problem. Some algorithms provide models that allow us to understand when an auction can be expected to be efficient and when it yields a Walrasian equilibrium.

The auction proposed by Ausubel (2005) for strong substitutes valuations follows a greedy steepest descent algorithm to minimize the (integrally convex) Lyapunov function (Murota & Tamura, 2003). This algorithm has an intuitive interpretation as an ascending auction: subgradients of the Lyapunov function at p are oversupplies at this price: $\partial L(p) = s - D_{\mathcal{I}}(p)$.[4] Knowing that the Lyapunov function is equivalent to the DRWDP (see Proposition 1), the overall auction can now be described as a primal-dual algorithm to solve the RWDP. For the price minimization, both algorithms require all subgradients at each point, that is, the entire demand set needs to be revealed. A specific version of a primal-dual algorithm yields the same steps.

We focus on primal-dual algorithms as a consistent algorithmic framework to model Walrasian auction mechanisms. Let us first describe the auction by Ausubel (2005) as a steepest descent algorithm before we introduce the overall primal-dual auction framework.

5.1 | The auction by Ausubel (2005)

The auction algorithm starts with an arbitrary price vector p below the bidder-optimal Walrasian prices, possibly $p(k) = 0$ for all $k \in \mathcal{K}$. The algorithm then searches iteratively in each round $t \in T$ for a subset of goods $S \subseteq \mathcal{K}$ such that $L(p^t) - L(p^t + \mathbb{1}_S)$ is maximized. Here, p^t denotes the prices in round t. This is equivalent to determining the direction of steepest descent to find the global minimum of this function:

[4]Note that subgradient and steepest descent algorithms for convex minimization are equivalent for differential functions but not for the minimization of discrete functions as in the case of markets with indivisible goods. The difference between the two algorithms is that the steepest descent algorithm evaluates all subgradients at a point, while subgradient algorithms use only a single subgradient. This is equivalent to eliciting the entire demand demand correspondence or only a single bundle from the demand correspondence. As a result, the primal-dual algorithm needs fewer iterations to converge to the exact solution (de Vries et al., 2007).

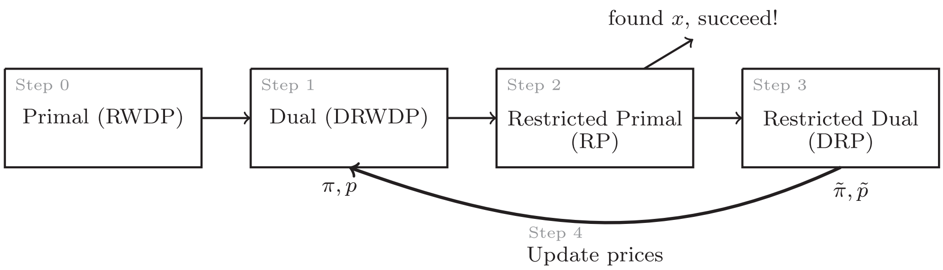

FIGURE 4 A primal-dual algorithm following Papadimitriou and Steiglitz (1998)

(i) At p^t the auctioneer asks each bidder $i \in \mathcal{I}$ for her entire demand set $D_i(p^t)$.

(ii) For all potential price update vectors $\widetilde{p} \in \{\mathbb{1}_S : S \subseteq \mathcal{K}\}$ the auctioneer determines each bidder's decrease of the indirect utility. The auctioneer chooses the price update $\widetilde{p} \in \{\mathbb{1}_S : S \subseteq \mathcal{K}\}$ such that the Lyapunov function is decreased the most, that is, $L(p^t) - L(p^t + \widetilde{p})$ is maximized. In case there are multiple such minimizers, the $\widetilde{p}$ with the smallest number of positive entries is selected. This price vector is referred to as the minimal minimizer and is guaranteed to be unique.

(iii) If no nonempty subset S can be found satisfying $L(p^t) - L(p^t + \mathbb{1}_S) > 0$, then the submodularity of the Lyapunov function guarantees that p^t is the bidder-optimal Walrasian price vector and the algorithm terminates. Otherwise the price p^{t+1} is set to $p^t + \widetilde{p}$ and the algorithm continues.

With integer valuations, $L(p)$ decreases by at least 1 in each iteration and therefore converges after finitely many steps. Murota et al. (2016) analyze the convergence and number of iterations of this steepest descent algorithm. In particular, if the auction algorithm is initialized with $p(k) = 0$ for all $k \in \mathcal{K}$ and p^* is the minimal equilibrium price, the algorithm terminates in exactly $\|p^*\|_\infty = \max_{k\in\mathcal{K}} |p^*(k)|$ iterations. The price update step described in this subsection can now be interpreted as an operation in a primal-dual algorithm to solve the WDP, as we show next.

5.2 | The primal-dual auction framework

Let us now describe the auction by Ausubel (2005) in the context of the more general primal-dual framework. Primal-dual algorithms (Papadimitriou & Steiglitz, 1998) can be used to compute solutions of the RWDP and DRWDP (see Section 4.1.1). Based on a feasible solution of the DRWDP, one derives a restricted primal RP that determines whether supply equals demand at these prices or not. If this is not the case, the dual restricted primal DRP determines the price increment, which is then added to the current price vector of the dual DRWDP, before a new restricted primal is computed. The overall process is illustrated in Figure 4. There is some flexibility in choosing each iteration's direction of price adjustment. In this primal-dual auction framework, we compute the price update that yields the steepest descent of the DRWDP.

Instead of solving the RWDP and the DRWDP directly, the primal-dual algorithm replaces these linear programs by a series of other linear programs known as the restricted primal RP and the dual of the restricted primal DRP. As the primal dual algorithm follows the same price trajectory as Ausubel's auction as we will show below, exactly $\|p^*\|_\infty$ iterations must be executed where p^* is the minimal equilibrium price vector (Murota et al., 2016). In each iteration two linear programs (the RP and DRP) must be solved which both are of exponential size in the number of goods. Clearly, the primal dual algorithm does not give any runtime benefits over solving the RWDP and DRWDP directly. However, executing the primal-dual algorithm instead of solving the RWDP and DRWDP directly allows to interpret the auction by Ausubel (2005) in terms of a primal-dual framework. Moreover, unlike the solution obtained by solving the RWDP and DRWDP directly, the allocation and prices computed by the primal-dual algorithm are guaranteed to constitute the Walrasian equilibrium with bidder-optimal prices.

Let us discuss the algorithm in more detail. In an ascending auction the components of the initial price vector are set to $p(k) = 0$ for all $k \in \mathcal{K}$. To obtain an initial feasible dual solution, the dual is solved with these prices to find initial values for the indirect utility π_i of every bidder i.

With a feasible dual solution, one can exploit the complementary slackness conditions to derive an optimal primal solution which defines a welfare-maximizing allocation of bundles to bidders. Naturally, not every feasible dual solution allows for an optimal primal solution. To check this, one solves an optimization problem known as the restricted primal RP problem.

$$
\begin{aligned}
\max\; & -\sum_{i\in\mathcal{I}} \lambda_i c_i - \sum_{k\in\mathcal{K}} \mu_k d_k && \text{(RP)}\\
\text{s.t.}\; & \sum_{x\in\mathcal{X}_i} z_i(x) + a_i + c_i = 1 \quad \forall i \in \mathcal{I}(\widetilde{\pi}_i) \\
& \sum_{i\in\mathcal{I}} \sum_{x\in\mathcal{X}_i} x(k) z_i(x) + b_k + d_k = s(k) \quad \forall k \in \mathcal{K}(\widetilde{p}(k)) \\
& z_i(x), a_i, b_k \geq 0 \quad \forall z_i(x) \in \mathcal{J}_z, \forall a_i \in \mathcal{J}_a, \forall b_k \in \mathcal{J}_b \\
& z_i(x) = 0, a_i = 0, b_k = 0 \quad \forall z_i(x) \notin \mathcal{J}_z, \forall a_i \notin \mathcal{J}_a,
\end{aligned}
$$

$$\forall b_k \notin \mathcal{J}_b$$
$$c_i, d_k \geq 0 \quad \forall i \in \mathcal{I}, \forall k \in \mathcal{K}$$

Given a feasible dual solution for the DRWDP, any tight dual constraint $\pi_i \geq v_i(x) - \sum_{k\in\mathcal{K}} x(k)p(k)$ corresponds to a bundle x that maximizes the utility of bidder i at prices p. Thus, the set of tight dual constraints $\mathcal{J}_z$ corresponds to the bidders' demand sets. In case the given dual solution is optimal, the complementary slackness conditions mandate that whenever the dual constraint has slack, that is, $\pi_i > v_i(x) - \sum_{k\in\mathcal{K}} x(k)p(k)$, the corresponding primal variable $z_i(x)$ defining whether bidder i is allocated bundle x equals zero. The interpretation of this is that a bidder is never allocated a bundle not being part of her demand set. Of course, if the given dual solution is not optimal, there might not exist an allocation such that each bidder receives a bundle from her demand set. Therefore, additional slack variables c_i and d_k are introduced to the RP that measure by how much the complementary slackness conditions are violated. A violation may either occur due to bidder i not being allocated a bundle from her demand set ($c_i > 0$) or an item k remaining (partially) unsold ($d_k > 0$). The restricted primal problem tries to find an allocation in which these violations are minimized. In fact, when the optimal solution of the RP equals 0, the complementary slackness conditions are fulfilled so that the current primal and dual solution constitute a Walrasian equilibrium. Otherwise, the price of some items needs to be raised.

Complementary slackness conditions must also hold for the dual constraints $\pi_i \geq 0$ and $p(k) \geq 0$. We denote the set of tight dual constraints by $\mathcal{J}_a$ and $\mathcal{J}_b$ respectively. Due to complementary slackness, the primal variable a_i must equal zero whenever the corresponding dual constraint $\pi_i \geq 0$ has slack. In other words this means that whenever a bidder's indirect utility is positive, she must be allocated a nonempty bundle from her demand set. Similarly, complementary slackness implies that when a price of an item $k \in \mathcal{K}$ is greater than zero, then slack variable b_k must equal zero, which guarantees that all units of item k are allocated in an optimal solution.

In the primal-dual framework of Papadimitriou and Steiglitz (1998) all coefficients λ_i and μ_k in the objective function of the restricted primal RP equal 1. Note that as long as λ_i and μ_k are chosen to be strictly positive, their specific values do not influence the termination criterion of the primal-dual algorithm as one only checks whether the objective equals zero. However, the particular choice of λ_i and μ_k affects the constraints in the dual of the restricted primal DRP, and we will take advantage of this to find a particular price update vector when solving the DRP.

In case the RP objective does not equal zero, the current dual solution of the DRWDP is updated using the solution to the dual of the restricted primal DRP. Solving the DRP essentially means computing a direction $\tilde{\pi}$, $\tilde{p}$ in which the dual objective function can be improved the most. We set $\tilde{\pi}, \tilde{p}$ such that it minimizes the function $\sum_{i\in\mathcal{I}}(\pi_i + \tilde{\pi}_i) + \sum_{k\in\mathcal{K}} s(k)(p(k) + \tilde{p}(k))$. This is equivalent to finding a subgradient to the Lyapunov function as we will show below.

As there may exist multiple potential directions $(\tilde{\pi}, \tilde{p})$ that minimize the Lyapunov function, we need to make small adaptions to the DRP such that the gradient found by the DRP is equivalent to the minimal minimizer in Ausubel's auction. For this purpose we introduce additional constraints $0 \leq \tilde{p}(k) \leq 1$ for all $k \in \mathcal{K}$. As proven in Ausubel (2005), the Lyapunov function restricted to the unit $|\mathcal{K}|$-dimensional cube $\{p + \tilde{p} : 0 \leq \tilde{p}(k) \leq 1 \; \forall k \in \mathcal{K}\}$ is minimized on the vertices of this cube. Thus, limiting price updates $\tilde{p}(k)$ to the interval [0, 1] for all $k \in \mathcal{K}$ ensures that the same potential price updates as in Ausubel's auction (i.e., $\{\mathbb{1}_S : S \subseteq \mathcal{K}\}$) are considered. Note that this also implies that in each iteration of our primal-dual auction framework the respective prices and price updates are integer valued.

Another adaption to be made is to chose λ_i suitably large for all $i \in \mathcal{I}$ so that the decrease of utility for each bidder i is unrestricted when raising prices. To guarantee that the gradient found by the DRP is not only a minimizer of the Lyapunov function but a minimal minimizer, price penalties $\tau_k > 0$ are added to the objective function that are small enough so that their impact on the objective value is negligible.

$$\begin{aligned}
\min \quad & \sum_{i\in\mathcal{I}} \tilde{\pi}_i + \sum_{k\in\mathcal{K}} (s(k) + \tau_k)\tilde{p}(k) && \text{(DRP)} \\
\text{s.t.} \quad & \tilde{\pi}_i + \sum_{k\in\mathcal{K}} x(k)\tilde{p}(k) \geq 0 \qquad \forall i, x : z_i(x) \in \mathcal{J}_z && (z_i(x)) \\
& \tilde{\pi}_i \geq 0 \qquad \forall i : a_i \in \mathcal{J}_a && (a_i) \\
& \tilde{\pi}_i \geq -\lambda_i \qquad \forall i : a_i \notin \mathcal{J}_a && (c_i) \\
& \tilde{p}(k) \geq 0 \qquad \forall k : b_k \in \mathcal{J}_b && (b_k) \\
& \tilde{p}(k) \geq -\mu_k \qquad \forall k : b_k \notin \mathcal{J}_b && (d_k) \\
& 0 \leq \tilde{p}(k) \leq 1 \qquad \forall k \in \mathcal{K}
\end{aligned}$$

In the following we make the connection between the DRP and the price update step of Ausubel's ascending auction explicit by demonstrating how to transform one approach into the other. Recall that in Ausubel (2006) the goal is to find a $\tilde{p} \in \{\mathbb{1}_S : S \subseteq \mathcal{K}\}$ leaving all entries of $p + \tilde{p}$ nonnegative and minimizing

$$L(p + \tilde{p}) - L(p).$$

Ausubel (2006) shows that for a fixed $\tilde{p}$ it holds that

$$\begin{aligned}
L(p + \tilde{p}) - L(p) = \sum_{i\in\mathcal{I}} \max_{x\in D_i(p)} \left\{ -\sum_{k\in\mathcal{K}} x(k)\tilde{p}(k) \right\} \\
+ \sum_{k\in\mathcal{K}} s(k)\tilde{p}(k).
\end{aligned}$$

The term $\max_{x\in D_i(p)} \left\{-\sum_{k\in\mathcal{K}} x(k)\tilde{p}(k)\right\}$ is clearly equal to

$$\begin{aligned}
\min \quad & \tilde{\pi}_i \\
\text{s.t.} \quad & \tilde{\pi}_i \geq -\sum_{k\in\mathcal{K}} x(k)\tilde{p}(k) \qquad \forall x \in D_i(p)
\end{aligned}$$

Consequently, by adjusting notation and noting that $\mathcal{J}_z$ represents the demand set $D_i(p)$, we can rewrite the problem of minimizing $L(p+\widetilde{p}) - L(p)$:

$$\begin{aligned} \min \quad & \sum_{i\in\mathcal{I}} \widetilde{\pi}_i + \sum_{k\in\mathcal{K}} s(k)\widetilde{p}(k) \\ \text{s.t.} \quad & \widetilde{\pi}_i + \sum_{k\in\mathcal{K}} x(k)\widetilde{p}(k) \geq 0 \qquad \forall i, x : z_i(x) \in \mathcal{J}_z \\ & p(k) + \widetilde{p}(k) \geq 0 \qquad \forall k \in \mathcal{K} \\ & 0 \leq \widetilde{p}(k) \leq 1 \qquad \forall k \in \mathcal{K} \end{aligned}$$

As argued above, all price updates and consequently also the prices are integral in each step of our primal-dual auction framework. Hence, the second last set of inequalities can be replaced by

$$\widetilde{p}(k) \geq 0 \qquad \forall k : b_k \in \mathcal{J}_b$$

since $\mathcal{J}_b$ represents all indices where $p(k)$ equals 0.

The only remaining difference to the DRP is that we are missing the inequalities $\widetilde{\pi}_i \geq 0$ for $a_i \in \mathcal{J}_a$. From the definition we see, however, that $a_i \in \mathcal{J}_a$ if and only if the utility of bidder i at price p is 0. But this means that the empty bundle is in her demand set. Hence, $\widetilde{\pi}_i \geq 0$ is actually one of the constraints $\widetilde{\pi}_i + \sum_{k\in\mathcal{K}} x(k)\widetilde{p}(k) \geq 0$. As a result we get that one step of the Lyapunov minimization approach is exactly the same as one step of the primal-dual algorithm.

We restricted our attention so far on explaining the relationship between the primal-dual algorithm and the ascending version of the tâtonnement process described by Ausubel (2005). However, similar observations can also be made for the descending version. The only adaptions to be made in our argument concern the formulation of the DRP. Instead of applying positive price penalties τ_k in the objective function, negative ones have to be used to ensure that a maximal minimizer is found in each iteration. Furthermore, the price updates $\widetilde{p}(k)$ need to be bounded to the interval $[-1, 0]$ instead of [0, 1]. Of course, this also implies that μ_k must be chosen suitably large, that is, $\mu_k \geq 1$, in order to allow for price updates of -1.

While the auction described by Ausubel (2005) requires the bidders' valuations to satisfy the strong substitutes condition, the primal-dual algorithm also works for other environments, in particular for economies where the preferences of the bidders fulfill the more general GGSC condition. Sun and Yang (2006) propose the *dynamic double-track auction* (DDT) that terminates in a Walrasian equilibrium if bidders bid straightforwardly and have GSC valuations. Given two sets S_1 and S_2 describing two classes of goods, the auctioneer announces start prices of zero for items in S_1 and suitable high start prices in S_2 such that items in S_1 are overdemanded while items in S_2 are underdemanded. In the course of the auction the auctioneer simultaneously adjusts prices of items S_1 upwards but those of items in S_2 downwards.

Shioura and Yang (2015) introduce the *generalized double-track auction* which is an extension of the DDT to multi-item multi-unit economies where bidders' valuations satisfy the GGSC condition. Their auction starts with an arbitrary integral price vector and then proceeds in two phases. While in the first phase the auctioneer adjusts prices of items in S_1 upwards and prices in S_2 downwards, the price update directions are reversed in the second phase.

Similar to the auction proposed by Ausubel (2005), the price updates in the generalized double-track auction correspond to the steepest descent direction of the Lyapunov function, which can be embedded into a primal-dual algorithm. Essentially, the primal-dual algorithm for the generalized double-track auction combines the DRP adaptions for the ascending and descending version of the auction by Ausubel (2005) as described above. Let the set S_1 and S_2 denote the set of items with an upward and downward moving price trajectory, respectively. While price updates for items in S_1 are bounded to the interval [0, 1], they are restricted to interval $[-1, 0]$ for items in S_2. Similarly, the price penalties in the objective of the DRP are positive for items in S_1 and negative for items in S_2. Once the generalized double-track auction moves from the first to the second phase, the price trajectories of items in S_1 and S_2 are inverted so that the adaptions made to the DRP for items in S_1 now apply for items in S_2 and vice versa.

5.3 | Allocation of items

While our paper focuses on the process of determining equilibrium prices, of course, the auctioneer must determine an equilibrium allocation as well. That is, given a target supply s and an equilibrium price vector p^*, we must find allocations $x_i \in D_i(p^*)$ for every bidder, such that $\sum_{i\in\mathcal{I}} x_i = s$. Since we assume access to demand oracles, that is, each bidder i reports her whole demand set $D_i(p^*)$ in each iteration of the auction, and as demand sets only contain integer points, we could just try every of the finitely many combinations of allocations $x_i \in D_i(p^*)$ in order to match the target supply. This approach is however not very efficient: the number of combinations we possibly have to check is $\Pi_{i\in\mathcal{I}} |D_i(p^*)|$, which can clearly be exponential.

The allocation problem can also be interpreted as a flow problem: Consider the directed graph $G = (V, A)$ consisting of $|\mathcal{I}| \cdot |\mathcal{K}|$ vertices $b_i(k)$, describing bidder i's demand of good k, and $|\mathcal{K}|$ vertices $t(k)$, describing the total supply of good k. For each $i \in \mathcal{I}$ and $k \in \mathcal{K}$, there is an arc pointing from $t(k)$ to $b_i(k)$. Now consider a flow x on this graph, where $x_i(k)$ denotes the amount of flow from vertex $b_i(k)$ to vertex $t(k)$. We interpret $x_i(k)$ as the number of units of good k bidder i receives. As usually, given a flow x, and a node v in the graph, the *excess* at node v is the difference of the flow entering the node and the flow leaving the node:

$$\partial x(v) = \sum_{(w,v)\in A} x(w,v) - \sum_{(v,w)\in A} x(v,w).$$

We call the vector ∂x the *boundary* of x. In our above defined graph, we have $\partial x(b_i(k)) = x_i(k)$ and $\partial x(t(k)) = -\sum_{i\in\mathcal{I}} x_i(k)$. The total number of goods of type k should be equal to the supply of good k. Hence, we have the constraint $\partial x(t(k)) = -s(k)$.

Also, each bidder should receive an allocation in her demand set $D_i(p^*)$, so $(\partial x(b_i(1)), \ldots, \partial x(b_i(|\mathcal{K}|)) \in D_i(p^*)$ should hold. Thus, the allocation problem can be interpreted as finding a feasible flow with respect to these constraints on the boundary. In the case of strong-substitutes valuations, the demand sets $D_i(p^*)$ are all $M^\natural$-convex, so this is an instance of the *M-convex submodular flow problem*. Polynomial-time algorithms have been developed for this problem, many of them are based on well-known algorithms for min-cost flows. For an overview, see for example (Murota, 2003, Ch. 10).

6 | SUMMARY AND RESEARCH AGENDA

A number of assumptions are crucial for the existence of Walrasian equilibria. Apart from (a) *integral concavity of the aggregate value function*, (b) the *bidders' valuations need to be independent* of each other, and all bidders need to be pure payoff maximizers, that is, have a (c) *quasilinear utility function*. Also, we assume that (d) the bidders are *price-takers and truthfully reveal their demand correspondence* in each round. With these assumptions we can guarantee Walrasian equilibria. However, these are strong assumptions, which might not hold in the field.

(i) Bidder valuations in real-world auctions include complements and substitutes such that Walrasian equilibria might not even exist. Competitive equilibria with nonlinear and personalized prices always exist in ascending auctions under the assumptions above.[5]

(ii) Quasilinearity is not always given as there might exist budget constraints, spitefulness, or market-power effects. For example, if bidders have financial constraints, quasilinearity is violated, and ascending auctions with budget constrained bidders have only been analyzed recently (Gerard van der Laan, 2016; Yang et al., 2018). Even if one tries to set budget constraints endogenously for bidders, it might not always be possible to implement an efficient outcome via an auction (Bichler & Paulsen, 2018).

(iii) Finally, bidders might not bid straightforward in a simple clock auction and behave strategically. A number of papers discusses variations or extensions of simple clock auctions, which yield incentive compatibility (Ausubel, 2006). These are, however, quite different from the simple clock auctions we see in the field.

The assumptions (i)–(iii) above also lead to corresponding research challenges for the operations research community.

1. Most resource allocation problems analyzed in operations research (e.g., scheduling or packing problems) do not satisfy the assumptions that allow for Walrasian equilibria. Duality breaks for nonconvex integer programming problems and new concepts for competitive equilibrium prices need to be derived. The literature on integer programming duality can provide useful insights and guidance how to derive equilibrium prices for such nonconvex allocation problems (Wolsey, 1981).
2. Budget constraints play a major role in many markets. We need to understand equilibria in markets where bidders maximize payoff, but are financially restricted. Very recent results suggest that budget constraints have a substantial impact on the computational complexity of the allocation and pricing problem and require bilevel integer programs which are known to be Σ_2^p-hard (Bichler & Waldherr, 2019). Overall, it will be useful to analyze utility models different from the standard quasi-linear utility function as they have been observed in advertising and other domains where bidders might not maximize payoff but their net present value or return on investment (Fadaei & Bichler, 2017; Baisa, 2017; Baldwin et al., 2020). Effective ways to compute market equilibria in such an environment still need to be developed.
3. Finally, incentive-compatibility plays an important role in small markets where participants can influence the price. Recent research tries to design simple ascending auction and pricing rules that are incentive-compatible (Baranov, 2018). Incentive-compatibility is very restrictive in most environments. For example, in markets with purely quasilinear utilities, the Vickrey–Clarke–Goves mechanism is unique (Green & Laffont, 1979). For larger markets it can also be useful to understand weaker notions of robustness against strategic manipulation (Azevedo & Budish, 2018).

[5]For example, Sun and Yang (2014) introduces an ascending and incentive-compatible auction in markets with only complements using non-linear and anonymous prices. Ausubel and Milgrom (2002), Parkes and Ungar (2000) and de Vries et al. (2007) discuss ascending auctions for markets where bidders have substitutes and complements and allow for discriminatory and non-linear prices. These auctions are incentive-compatible if the bidders' valuations were gross substitutes.

Overall, competitive equilibrium theory is closely related to mathematical optimization and it provides a rich field for operations research to contribute.

ACKNOWLEDGMENT

The financial support from the Deutsche Forschungsgemeinschaft (DFG) (BI 1057/1-8) is gratefully acknowledged. Open access funding enabled and organized by Projekt DEAL.

ORCID

Martin Bichler https://orcid.org/0000-0001-5491-2935

REFERENCES

Ağralı, S., Tan, B., & Karaesmen, F. (2008). Modeling and analysis of an auction-based logistics market. *European Journal of Operational Research*, *191*(1), 272–294.

Arrow, K. J., & Debreu, G. (1954). Existence of an equilibrium for a competitive economy. *Econometrica: Journal of the Econometric Society*, *22*, 265–290.

Arrow, K. J., & Hahn, F. H. (1971). General competitive equilibrium. *Advanced Textbooks in Economics*, *12*.

Ausubel, L. (2005). *Walrasian tatonnement for discrete goods* (Unpublished paper).

Ausubel, L., & Milgrom, P. (2006). *The lovely but lonely vickrey auction*. In P. Cramton, Y. Shoham, & R. Steinberg (Eds.), Combinatorial auctions (pp. 17–40). Cambridge, MA, Ch. 1: MIT Press.

Ausubel, L. M. (2006). An efficient dynamic auction for heterogeneous commodities. *American Economic Review*, *96*(3), 602–629.

Ausubel, L. M., & Milgrom, P. R. (2002). Ascending auctions with package bidding. *Advances in Theoretical Economics*, *1*(1).

Azevedo, E. M., & Budish, E. (2018). Strategy-proofness in the large. *The Review of Economic Studies*, *86*(1), 81–116.

Baisa, B. (2017). Auction design without quasilinear preferences. *Theoretical Economics*, *12*(1), 53–78.

Baldwin, E., Edhan, O., Jagadeesan, R., Klemperer, P., & Teytelboym, A. (2020). *The equilibrium existence duality: Equilibrium with indivisibilities & income effects*. arXiv preprint arXiv:2006.16939.

Baldwin, E., & Klemperer, P. (2019). Understanding preferences: "Demand types," and the existence of equilibrium with indivisibilities. *Econometrica*, *87*(3), 867–932.

Baranov, O. (2018). An efficient ascending auction for private valuations. *Journal of Economic Theory*, *177*, 495–517.

Bertsekas, D. P. (1988). The auction algorithm: A distributed relaxation method for the assignment problem. *Annals of Operations Research*, *14*(1), 105–123.

Bichler, M., Davenport, A., Hohner, G., & Kalagnanam, J. (2006). Industrial procurement auctions. *Combinatorial Auctions*, *1*, 593–612.

Bichler, M., Fux, V., & Goeree, J. (2018). A matter of equality: Linear pricing in combinatorial exchanges. *Information Systems Research*, *29*(4), 1024–1043.

Bichler, M., Fux, V., & Goeree, J. K. (2019). Designing combinatorial exchanges for the reallocation of resource rights. *Proceedings of the National Academy of Sciences of the United States of America*, *116*(3), 786–791.

Bichler, M., & Goeree, J. K. (2017). Handbook of spectrum auction design. Cambridge, MA: Cambridge University Press.

Bichler, M., & Paulsen, P. (2018). A principal-agent model of bidding firms in multi-unit auctions. *Games and Economic Behavior*, *111*, 20–40.

Bichler, M., & Waldherr, S. (2019). *Computing core-stable outcomes in combinatorial exchanges with financially constrained bidders*. In Proceedings of the 2019 ACM Conference on Economics and Computation (pp. 747–747). New York, NY: ACM.

Bikhchandani, S., & Mamer, J. W. (1997). Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*, *74*(2), 385–413.

Bikhchandani, S., & Ostroy, J. M. (2002). The package assignment model. *Journal of Economic Theory*, *107*(2), 377–406.

Blumrosen, L., & Nisan, N. (2007). Combinatorial auctions. *Algorithmic Game Theory*, *267*, 300.

Caplice, C., & Sheffi, Y. (2003). Optimization-based procurement for transportation services. *Journal of Business Logistics*, *24*(2), 109–128.

Chakrabarty, D., Jain, P., & Kothari, P. (2014). *Provable submodular minimization using Wolfe's algorithm*. In Proceedings of the 27th International Conference on Neural Information Processing Systems—Volume 1. NIPS'14 (pp. 802–809). Cambridge, MA: MIT Press.

Chakrabarty, D., Lee, Y. T., Sidford, A., & Wong, S. C.-w. (2017). *Subquadratic submodular function minimization*. In Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing. STOC 2017 (pp. 1220–1231). New York, NY: Association for Computing Machinery. https://doi.org/10.1145/3055399.3055419.

Cole, R., Devanur, N. R., Gkatzelis, V., Jain, K., Mai, T., Vazirani, V. V., & Yazdanbod, S. (2016). *Convex program duality, fisher markets, and nash social welfare*. arXiv preprint arXiv:1609.06654.

Colson, B., Marcotte, P., & Savard, G. (2005). Bilevel programming: A survey. *4OR*, *3*(2), 87–107.

Danilov, V., Koshevoy, G., & Murota, K. (2001). Discrete convexity and equilibria in economies with indivisible goods and money. *Mathematical Social Sciences*, *41*(3), 251–273.

de Vries, S., Schummer, J., & Vohra, R. V. (2007). On ascending vickrey auctions for heterogeneous objects. *Journal of Economic Theory*, *132*(1), 95–118.

Demange, G., Gale, D., & Sotomayor, M. (1986). Multi-item auctions. *Journal of Political Economy*, *94*(4), 863–872.

Fadaei, S., & Bichler, M. (2017). Truthfulness with value-maximizing bidders: On the limits of approximation in combinatorial markets. *European Journal of Operational Research*, *260*(2), 767–777.

Fujishige, S., & Yang, Z. (2003). A note on Kelso and Crawford's gross substitutes condition. *Mathematics of Operations Research*, *28*(3), 463–469.

Gale, D. (1963). A note on global instability of competitive equilibrium. *Naval Research Logistics Quarterly*, *10*(1), 81–87.

Galindo, G., & Batta, R. (2013). Review of recent developments in or/ms research in disaster operations management. *European Journal of Operational Research*, *230*(2), 201–211.

Gerard van der Laan, Z. Y. (2016). An ascending multi-item auction with financially constrained bidders. *Journal of Mechanism and Institution Design*, *1*, 109–149.

Green, J., & Laffont, J. (1979). Incentives in public decision making. Amsterdam, Netherlands: North Holland.

Grötschel, M., Lovász, L., & Schrijver, A. (1981). The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, *1*(2), 169–197.

Gul, F., & Stacchetti, E. (1999). Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, *87*, 95–124.

Herroelen, W., & Leus, R. (2005). Project scheduling under uncertainty: Survey and research potentials. *European Journal of Operational Research*, *165*(2), 289–306.

Hurwicz, L. (1972). On informationally decentralized systems. *Decision and Organization*, *12*.

Kaneko, M. (1976). On the core and competitive equilibria of a market with indivisible goods. *Naval Research Logistics Quarterly*, *23*(2), 321–337.

Kelso, A. S., & Crawford, V. P. (1982). Job matching, coalition formation, and gross substitute. *Econometrica*, *50*, 1483–1504.

Kim, S. (1986). Computation of a large-scale competitive equilibrium through optimization. *Computers and Operations Research*, *13*(4), 507–515.

Klemperer, P. (2010). The product-mix auction: A new auction design for differentiated goods. *Journal of the European Economic Association*, *8*(2–3), 526–536.

Krishna, V. (2009). Auction theory. Boston, MA: Academic Press.

La Salle, J., & Lefschetz, S. (2012). Stability by Liapunov's direct method with applications by Joseph L Salle and Solomon Lefschetz (Vol. *4*). Amsterdam, Netherlands: Elsevier.

Leme, R. P. (2017). Gross substitutability: An algorithmic survey. *Games and Economic Behavior*, *106*, 294–316.

Lessan, J., & Karabatı, S. (2018). A preference-based, multi-unit auction for pricing and capacity allocation. *Computers and Operations Research*, *91*, 237–246.

Madani, M., & Van Vyve, M. (2015). Computationally efficient mip formulation and algorithms for european day-ahead electricity market auctions. *European Journal of Operational Research*, *242*(2), 580–593.

McKenzie, L. W. (1959). On the existence of general equilibrium for a competitive market. *Econometrica: Journal of the Econometric Society*, *27*, 54–71.

Meeus, L., Verhaegen, K., & Belmans, R. (2009). Block order restrictions in combinatorial electric energy auctions. *European Journal of Operational Research*, *196*, 1202–1206.

Milgrom, P., & Strulovici, B. (2009). Substitute goods, auctions, and equilibrium. *Journal of Economic Theory*, *144*(1), 212–247.

Murota, K. (2003). *Discrete convex analysis*. Society for Industrial and Applied Mathematics. Retrieved from https://epubs.siam.org/doi/abs/10.1137/1.9780898718508.

Murota, K. (2016). Discrete convex analysis: A tool for economics and game theory. *Journal of Mechanism and Institution Design*, *1*(1), 151–273.

Murota, K., Shioura, A., & Yang, Z. (2016). Time bounds for iterative auctions: A unified approach by discrete convex analysis. *Discrete Optimization*, *19*, 36–62.

Murota, K., & Tamura, A. (2003). New characterizations of m-convex functions and their applications to economic equilibrium models with indivisibilities. *Discrete Applied Mathematics*, *131*(2), 495–512.

Nesterov, Y. (2018). Lectures on convex optimization. Berlin, Heidelberg: Springer.

Olafsson, S., Li, X., & Wu, S. (2008). Operations research and data mining. *European Journal of Operational Research*, *187*(3), 1429–1448.

Özer, A. H., & Özturan, C. (2009). A model and heuristic algorithms for multi-unit nondiscriminatory combinatorial auction. *Computers and Operations Research*, *36*(1), 196–208 Part Special Issue: Operations Research Approaches for Disaster Recovery Planning.

Papadimitriou, C. H., & Steiglitz, K. (1998). Combinatorial optimization: Algorithms and complexity. Chelmsford, MA: Courier Corporation.

Parkes, D. C., & Ungar, L. H. (2000). Iterative combinatorial auctions: Theory and practice (p. 7481). AAAI/IAAI.

Rockafellar, R. T. (2015). Convex analysis. Princeton, NJ: Princeton University Press.

Shapley, L. S., & Shubik, M. (1971). The assignment game I: The core. *International Journal of Game Theory*, *1*(1), 111–130.

Shioura, A., & Tamura, A. (2015). Gross substitutes condition and discrete concavity for multi-unit valuations: A survey. *Journal of the Operations Research Society of Japan*, *58*(1), 61–103.

Shioura, A., & Yang, Z. (2015). Equilibrium, auction, and generalized gross substitutes and complements. *Journal of the Operations Research Society of Japan*, *58*(4), 410–435.

Starr, R. M. (1969). Quasi-equilibria in markets with non-convex preferences. *Econometrica: Journal of the Econometric Society*, *37*, 25–38.

Sun, N., & Yang, Z. (2006). Equilibria and indivisibilities: Gross substitutes and complements. *Econometrica*, *74*(5), 1385–1402.

Sun, N., & Yang, Z. (2014). An efficient and incentive compatible dynamic auction for multiple complements. *Journal of Political Economy*, *122*(2), 422–466.

Teytelboym, A. (2014). Gross substitutes and complements: A simple generalization. *Economics Letters*, *123*(2), 135–138.

Tran, N. M. & Yu, J. (2015). *Product-mix auctions and tropical geometry*.

Triki, C., Beraldi, P., & Gross, G. (2005). Optimal capacity allocation in multi-auction electricity markets under uncertainty. *Computers and Operations Research*, *32*(2), 201–217.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, *16*(1), 8–37.

Wolsey, L. A. (1981). Integer programming duality: Price functions and sensitivity analysis. *Mathematical Programming*, *20*(1), 173–195.

Yang, Z., Yu, J., et al. (2018). *Efficient ascending menu auctions with budget constrained bidders* (Tech. rep).

How to cite this article: Bichler M, Fichtl M, Schwarz G. Walrasian equilibria from an optimization perspective: A guide to the literature. *Naval Research Logistics* 2021;68:496–513. https://doi.org/10.1002/nav.21963

APPENDIXA.

Proof of Theorem 2. First, let $\{z_i^*(x)\}_{i\in\mathcal{I},x\in\mathcal{X}_i}$ be an optimal solution to the RWDP and $(\{\pi_i^*\}_{i\in\mathcal{I}}, \{p^*(k)\}_{k\in\mathcal{K}})$ be an optimal solution to the DRWDP. By assumption, the optimal value of the WDP is equal to the one of the RWDP, so we may assume that all $z_i^*(x)$ are in $\{0, 1\}$. We may further assume without loss of generality that for each bidder i, there exists exactly one x with $z_i^*(x) = 1$: If $z_i^*(x) = 0$ for all $x \in \mathcal{X}_i$, we can just set $z_i^*(\mathbf{0}) = 1$, where $\mathbf{0}$ is the empty bundle, without altering the value of the WDP, since $v_i(\mathbf{0}) = 0$. Similarly, if for some $k \in \mathcal{K}$, $\sum_{i\in\mathcal{I}}\sum_{x\in\mathcal{X}_i} x(k)z_i^*(x) < s(k)$, we may distribute the remaining items of type k arbitrarily among the agents. This does not decrease the value of the WDP because of monotonicity of the agents' valuations. The (possibly altered) variables $z_i^*(x)$ thus constitute an allocation where the whole supply is distributed among the agents—so the first criterion of a Walrasian equilibrium is satisfied. Let us now check that every bidder receives a bundle in her demand set: If $z_i^*(\bar{x}) = 1$, that is, bidder i receives bundle $\bar{x}$, we have by complementary slackness $\pi_i = v_i(\bar{x}) - \sum_{k\in\mathcal{K}} \bar{x}(k)p^*(k)$. Since π_i^* is part of an optimal solution,

$$\pi_i^* = \max_{x\in\mathcal{X}_i} v_i(x) - \sum_{k\in\mathcal{K}} x(k)p^*(k).$$

Otherwise, we could decrease π_i^*, making the value of the DRWDP smaller. Consequently, $v_i(\bar{x}) - \sum_{k\in\mathcal{K}} \bar{x}(k)p^*(k) = \max_{x\in\mathcal{X}_i} v_i(x) - \sum_{k\in\mathcal{K}} x(k)p^*(k)$, so $\bar{x}$ is in the demand set of bidder i at prices $\{p^*(k)\}_{k\in\mathcal{K}}$. The second condition of a Walrasian equilibrium is thus satisfied, and $\{p^*(k)\}_{k\in\mathcal{K}}$ are equilibrium prices.

For the other direction, let $\{p^*(k)\}_{k\in\mathcal{K}}$ be equilibrium prices together with an allocation, described by binary variables $\{z_i^*(x)\}_{i\in\mathcal{I},x\in\mathcal{X}_i}$. Let $\bar{x}$ be the bundle with $z_i^*(\bar{x}) = 1$. Set $\pi_i^* = v_i(\bar{x}) - \sum_{k\in\mathcal{K}} \bar{x}(k)p(k)$. Since $\bar{x}$ is in the demand set of bidder i, $\pi_i^* \geq v_i(x) - \sum_{k\in\mathcal{K}} x(k)p(k)$ for all bundles x, so $(\{p^*(k)\}, \{\pi_i^*\})$ is feasible for the DRWDP ($\pi_i^* \geq 0$ follows from choosing $x = \mathbf{0}$ in the above inequality). By definition of the Walrasian equilibrium, $\{z_i^*(x)\}$ is also feasible for the (R)WDP. All inequalities in the WDP actually hold with equality—so complementary slackness of the primal problem is trivially fulfilled. From the choice of π_i^* we also directly see, that complementary slackness is satisfied for the dual problem. It follows that the optimal value of the WDP equals the optimal value of the DRWDP.

4 Publication 2: A Wholesale Market Design for Road Capacity

Peer-Reviewed Conference Paper

Title: Designing a Large-Scale Wholesale Market for Urban Congestion Pricing

Authors: Martin Bichler, Gregor Schwarz

In: 17th International Conference on Wirtschaftsinformatik (WI22)

Abstract: The mispricing of scarce road capacity is one of the main reasons for traffic congestion in major cities, leading to substantial economic losses and environmental damage. In analogy to electricity markets, a two-stage market for road capacity was recently proposed in which an Independent System Operator (ISO) sells road capacity to multiple Service Providers (SPs) who then resell it to end consumers. As urban road networks consist of tens of thousands of road segments, the number of products in the market is unparalleled, making the preference elicitation and market tractability challenging. We propose a compact bid language based on origin-destination pairs and apply mixed-integer programming techniques to maximize welfare in the market. Utilizing calibrated traffic data from the MATSim Open Berlin Scenario, our numerical tests show that for a major city like Berlin our market can be solved to optimality within 15 minutes.

Contribution of dissertation author: Methodology, experimental design, software, formal analysis, investigation, visualization, joint paper management

Reference: Schwarz and Bichler (2022)

Comment: The long version of this paper was submitted to the academic journal *Transportation Research Part A: Policy and Practice.*

Designing a Large-Scale Wholesale Market for Urban Congestion Pricing

Gregor Schwarz and Martin Bichler

Technical University of Munich, Department of Informatics, Munich, Germany
{gregor.schwarz,bichler}@in.tum.de

Abstract. *The mispricing of scarce road capacity is one of the main reasons for traffic congestion in major cities, leading to substantial economic losses and environmental damage. In analogy to electricity markets, a two-stage market for road capacity was recently proposed in which an Independent System Operator (ISO) sells road capacity to multiple Service Providers (SPs) who then resell it to end consumers [1]. As urban road networks consist of tens of thousands of road segments, the number of products in the market is unparalleled, making the preference elicitation and market tractability challenging. We propose a compact bid language based on origin-destination pairs and apply mixed-integer programming techniques to maximize welfare in the market. Utilizing calibrated traffic data from the MATSim Open Berlin Scenario, our numerical tests show that for a major city like Berlin our market can be solved to optimality within 15 minutes.*

Keywords: congestion pricing, wholesale markets, auctions, MATSim

1 Introduction

Due to population growth and urbanization, traffic congestion has increased substantially in major cities around the world causing significant economic losses and environmental damage. A central challenge is the mispricing of the core resource for mobility, the road capacity. In most cities, road capacity is not priced at all leading to severe traffic congestion. In 2017, the total costs due to traffic jams in Britain, Germany, and the US were estimated \$461 billion.[1] Those numbers come from the lost productivity of workers sitting in traffic, the increased cost of transporting goods through congested areas and wasted fuel, among other factors.

While Vickrey proposed urban congestion pricing already in the 1960s [2], its practical implementation has lagged behind for decades. Singapore was the first city to implement congestion pricing in 1975, others such as London, Milan, and Stockholm followed suit only in the 2000s [3]. Almost all existing models are based on static traffic assignment theory, setting fixed prices based on prior estimates of the demand at certain times of the day [4, 5]. Maximizing social welfare requires efficient pricing, i.e., the cost of a trip must equal its marginal social cost, which makes it necessary to price individual road segments [6]. Static pricing schemes are not suitable for this as they either only consider pure travel distances or entire areas but not individual road segments.

[1] `https://www.economist.com/graphic-detail/2018/02/28/the-hidden-cost-of-congestion`

Dynamic congestion pricing is the determination of time-varying prices to control traffic congestion. By pricing roads dynamically, drivers can select their route through the network based on these prices which effectively redirects traffic through the network. More than that, prices also guide drivers when to start their journey, thereby not only distributing traffic over the network but also over time. With the help of these dynamic prices, the roads' capacity constraints are respected so that harmful traffic congestion can be avoided [7]. Prior literature agrees that dynamic congestion pricing offers the most efficient approach to manage and operate roads [8, 9]. Even though nowaday's privacy-preserving mobile applications for GPS-based tracking of vehicles allow for real-time dynamic congestion pricing, it has not been implemented in practice [10].

In other major utilities markets (e.g., electricity) real-time dynamic pricing is applied regularly already today. These markets follow a two-tiered structure with an *Independent System Operator (ISO)* managing the market and multiple *Service Providers (SPs)* who serve end consumers. A market-based congestion pricing model similar to electricity markets has received much scholarly attention over the past few years [1, 7, 11, 12]. Despite some fundamental differences, both electricity and road pricing markets aim to allocate scarce network resources in the presence of locational and temporal supply-demand variations, making electricity markets a suitable template for the design of road pricing markets (see [1] for a detailed discussion).

The authors of [1] envision a wholesale market which is managed by an ISO and allows SPs to buy and trade road capacity. Their wholesale market consists of several forward markets and a real-time market. Forward markets let participants plan and hedge risk by taking positions consistent with their underlying demands, while deviations from forward positions can be settled on the real time market. Forward markets could be run monthly, weekly, and daily to adapt to changing information about the demand on a particular day. SPs resell their capacity bought on the wholesale market to end consumers on a separate retail market. SPs compete in this retail market to offer customers attractive prices for individual trips but also package deals for recurrent work trips. This competition is vital as it enhances the effective redirection of traffic and ultimately brings down consumer prices. Without SPs, the ISO would act as a monopolist in the market.

Market design studies the rules for allocating and pricing goods on a market with multiple decision makers. Designing a wholesale market for road capacity is challenging due to the unparalleled number of products available in the market. In a city like Berlin, there are more than 34 thousand road segments that can potentially be priced. Thus, the number of products in our wholesale market is significantly higher than in other utilities markets, e.g., spectrum sales where there are typically not more than a few hundred products [13]. The decision makers in our wholesale market, the service providers, have preferences over entire routes consisting of many individual road segments. Eliciting these preferences in a succinct manner and pricing the road segments with respect to the disclosed valuations of the market participants is far from obvious.

1.1 Contributions

We propose a novel, compact bid language that allows SPs to state their preferences not only for individual road segments as in [1] but for multiple substitutable routes. Each bid defines an SP's demand for road capacity between two locations in the network (i.e., an *origin-destination (OD)* pair). SPs may associate each bid with a set of substitutable routes connecting these locations and specify an individual willingness-to-pay for each of them. We determine an allocation by solving a mixed-integer linear optimization program. To approximate competitive equilibrium prices, we draw on ideas from electricity market design and rely on a pricing technique called *IP Pricing* which is commonly used in electricity markets today [14]. Unlike to electricity markets, almost no side-payments are required in our market to prevent SPs from losses.

In order to evaluate the computational tractability as well as the volatility of road prices throughout the day for realistic problem instances, we rely on traffic data taken from the *MATSim Open Berlin Scenario*, a calibrated transport simulation scenario that aims to model the traffic in Berlin's urban area accurately [15]. Our tests show that the underlying optimization problems are tractable for multi-million cities like Berlin and can be solved within 15 minutes to optimality. The resulting allocation problem has more than $\sim$90k binary variables, $\sim$950k continuous variables, and 250k constraints. It is far from obvious that such problems can be solved with modern optimization software. Interestingly, our tests indicate that only a small fraction of road segments have to be priced to reach an efficient allocation. The solver makes effective use of substitutable routes between OD pairs and effectively distributes the available capacity.

The paper contributes to a growing stream of research on market design in Information Systems [16]. In particular, pricing in non-convex markets has played a considerable role in this literature, with combinatorial auctions as leading examples [17–21]. By proposing a novel design for a road capacity wholesale market, the paper also advances recent research on congestion pricing in urban areas [1, 3, 5, 6].

2 Market Design

The products in our wholesale market are licenses for road capacity (measured in vehicles/hour). Each license is associated with a single road segment and a one-hour time slot during a specific day. Licenses for different time slots are traded on separate markets. Initially, all road capacity is owned by the ISO. The wholesale market for licenses of a particular time slot consists of multiple forward markets (taking place monthly, weekly, and daily before the time slot) and a real-time market. SPs can buy road capacity from the ISO or trade capacity among each other on both the forward and real-time market.

SPs resell the road capacity bought on the wholesale market to end consumers on a separate retail market. Products in the retail market may include licenses for unlimited road use, monthly packages for daily work trips, or capacity for one-time trips. SPs compete against one another to offer end consumers affordable route alternatives with few detours within the desired time slot. This competition incentivizes innovation and brings end consumer prices down. In case of high demand for a road segment, SPs will raise prices in the retail market, thereby encouraging drivers to switch to other modes of transportation or shift their trips to other time slots.

The end consumers' trips may stretch over several time intervals. However, as they purchase road capacity for their entire trips from SPs, it is the obligation of the SPs to ensure that no more vehicles drive across a road segment than their capacity allows.

The ISO sets suitable penalties to prevent SPs from overbooking. Due to the temporal discretization of the wholesale market, SPs may need to combine licenses from different time slots to serve a customer, making bidding in the wholesale market more complex. However, the subsequent structure of the forward markets which allows to settle demand fluctuations and the fact that SPs serve tens of thousands of customers within a time slot mitigates this problem. In the following, we exclusively focus on the wholesale market and leave the design of the retail market for future research.

Bidding for tens of thousands of individual road segments in the wholesale market is impractical for SPs. The number of OD pairs is much more manageable and is what traffic planners work with today. Bidding on entire routes vastly simplifies the bidding process of SPs who need to satisfy a certain customer demand between an OD pair. It also prevents SPs from winning much capacity on one segment of a route but little on the next, which would render the capacity on the first worthless.

In our bid language each buy bid is associated with a set of substitutable routes for an OD pair. Bidders can specify an individual willingness-to-pay for each of the routes (see Example 1). Buy bids also have a lower and upper capacity bound, restricting the total allocated capacity for all routes to this interval if the bid is accepted. Specifying lower bounds allows SPs to run marketing campaigns that require them to offer at least a certain amount of capacity for particular routes. Sell bids only correspond to single road segments and must have a lower bound of zero. Limiting sell bids to individual road segments prevents potential fitting problems that may occur in practice: SP1 sells capacity for route $\{k_1, k_2, k_3, k_4\}$, while SP2 wants to buy capacity only for k_1. The trade will not take place because a route must either be sold entirely or not at all. Limiting sell bids to single road segments is a mild restriction for SPs, but it significantly increases the probability of finding matches with the package buy bids.

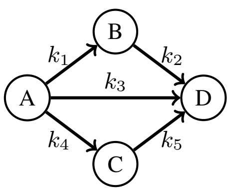

SP1 7-8am BUY [100, 500]	
Route	Price
$S_1 = \{k_1, k_2\}$	0.19
$S_2 = \{k_3\}$	0.22
$S_3 = \{k_4, k_5\}$	0.18

SP2 7-8am SELL [0, 200]	
Segment	Price
k_1	0.08

SP2 7-8am SELL [0, 200]	
Segment	Price
k_2	0.07

SP2 7-8am SELL [0, 200]	
Segment	Price
k_3	0.15

Example 1. SP1 wants to buy capacity for at least 100 and at most 500 vehicles traveling from node A to D in the 7-8 am time slot. SP1 specifies three alternative routes: $S_1 = \{k_1, k_2\}$, $S_2 = \{k_3\}$, $S_3 = \{k_4, k_5\}$. SP2 wants to sell up to 200 licenses each for road segments k_1, k_2, and k_3. SP1 may win in total 400 licenses for routes from A to D: 200 for S_1 and 200 for S_2.

The bid language can be extended in several ways, e.g., by guaranteeing bidders that they either win road capacity in one area of the city or another. Such XOR-constraints add further non-convexities to the market, making the price determination more challenging and requiring higher side-payments for SPs. This will be discussed thoroughly in the long version of this paper.

3 Allocation Problem

As the *allocation problem (AP)* is solved hourly and these hourly auctions are independent, no index for time is needed. Let I be the set of bidders (ISO and SPs), K the set of road segments, $S \subseteq K$ a route, B_i the set of buy bids of bidder i, and $v_i(k)$ and $v_i(S)$ the valuation of bidder i for road segment k and route S, respectively. Each buy bid $b \in B_i$ represents a set of alternative routes and is associated with an upper bound u_i^b and lower bound ℓ_i^b. Bidders also specify an upper bound u_i^k for each road segment k they intend to sell. The binary decision variable $z_i(b)$ denotes whether buy bid b is accepted, the continuous decision variables $x_i(S)$ and $y_i(k)$ define the amount of licenses a bidder wins for route S or sells for segment k, respectively. The ISO is modeled as a bidder who only sells capacity and whose value for all road segments is zero.

$$\max_{x,y} \sum_{i \in I} \sum_{S \subseteq K} v_i(S) \cdot x_i(S) - \sum_{i \in I} \sum_{k \in K} v_i(k) \cdot y_i(k) \tag{AP}$$

$$\text{s.t.} \sum_{i \in I} \sum_{S \subseteq K : k \in S} x_i(S) - \sum_{i \in I} y_i(k) = 0 \qquad \forall k \in K \quad (p(k)) \tag{1a}$$

$$\ell_i^b \cdot z_i(b) \leq \sum_{S \in b} x_i(S) \leq u_i^b \cdot z_i(b) \qquad \forall i \in I, b \in B_i \tag{1b}$$

$$y_i(k) \leq u_i^k \qquad \forall i \in I, k \in K \tag{1c}$$

$$x_i(S) \in \mathbb{R}_{\geq 0} \qquad \forall i \in I, S \subseteq K \tag{1d}$$

$$y_i(k) \in \mathbb{R}_{\geq 0} \qquad \forall i \in I, k \in K \tag{1e}$$

$$z_i(b) \in \{0, 1\} \qquad \forall i \in I, b \in B \tag{1f}$$

The objective function maximizes the gains of trade, constraints (1a) ensure that for each segment the bought and sold road capacity match, constraints (1b) fix a buy bid's lower and upper quantity bounds in case the bid is accepted, and constraints (1c) define the upper quantity bounds for the sell bids. As variables $x_i(S)$ and $y_i(k)$ are continuous to make real life problem instances tractable, it is possible that bidders end up buying or selling fractional capacities for certain road segments. These are of no value to SPs as they cannot be resold to end consumers. However, our experiments in Section 6.1 show that the number of fractionally allocated road segments is negligible.

4 Pricing Problem

Because of the binary variables, the (AP) is a non-convex integer program so that the dual variables corresponding to constraint (1a) cannot be interpreted as shadow prices. The problem is ultimately a non-convex mixed integer programming problem. There is an extensive literature on the existence of Walrasian equilibrium prices [22,23]. Such prices are linear (i.e., the route price is the sum of its components' prices) and anonymous (i.e., prices are the same for all bidders). In a Walrasian equilibrium no bidder wants to change the set of allocated road capacities at the given prices. Unfortunately, Walrasian equilibria generally do not exist for non-convex allocation problems as the (AP) [22].

Pricing in such non-convex markets has been a significant concern in the IS literature [16], but also for the design of electricity markets [24]. We follow an approach called *IP pricing* that is being used by several ISOs in US electricity markets [14]. To calculate IP prices, the binary variables $z_i(b)$ are fixed to their optimal value in the (AP), while all other variables and constraints remain unchanged. As $x_i(S)$ and $y_i(k)$ are continuous, the resulting pricing problem is a linear program which allows to interpret the dual variables of constraint (1a) as shadow prices. IP pricing may cause bidders to incur losses so that side-payments (so-called *make-whole payments*) are necessary as compensation. In our market, given the optimal solutions $x_i^*(S)$, $y_i^*(k)$ and dual prices $p(k)$, bidder i receives a make-whole payment $mwp_i(b) = \max\left\{0, \sum_{S \in b} x_i^*(S) \cdot (p(S) - v_i(S))\right\}$ for a buy bid $b \in B_i$ and $mwp_i(k) = \max\{0, y_i^*(k) \cdot (v_i(k) - p(k))\}$ for a sell bid, where $p(S) = \sum_{k \in S} p(k)$. In electricity markets these side-payments can be substantial and may lead to a significant bias in the market prices. However, in our road pricing market such side-payments are negligible as our tests indicate because the level of non-convexities due to the lower bounds on the demand is less pronounced.

5 Experimental Design

In our experiments we focus on the sale of road capacity for a single day consisting of 24 one-hour time slots. In the initial forward market the ISO holds all road capacity, while SPs can only buy but not sell licenses. The subsequent forward markets allow SPs to handle demand fluctuations from their forecasts by trading capacities or buying previously unsold licenses from the ISO. In our experiments we focus on the initial forward market as it is by far the largest in terms of trading volume and number of bids. As no road capacity market has been implemented in practice yet, the number of SPs in the market is somewhat speculative. We base our guess on other large utilities markets in Germany, e.g., electricity and telecommunication, and set to number of SPs to four.

5.1 MATSim Open Berlin Scenario

In order to analyze the road pricing market for realistic problem instances, we generate bids based on synthetic but calibrated traffic demand from the *MATSim Open Berlin Scenario*[2], which was published 2019 for the open-source traffic simulator MATSim. The scenario contains full day activity-travel patterns for a 10% sample of all adults living in the states of Berlin and Brandenburg and is designed to model the traffic in the city of Berlin realistically. The scenario is solely built upon openly available data and is calibrated with respect to traffic counts, choice of transport modes, and mode-specific trip duration and distance distributions [15]. To simulate our market for the entire population and also reflect the anticipated increase of traffic volume [25], we scale up this sample demand to values between 100% and 130%.

The underlying road network consists of all road categories within Berlin and all main roads in the surrounding state of Brandenburg, covering an area of roughly 150×250 km and consisting of $\sim$73k nodes and $\sim$159k one-way car-only links. For our market we

[2] `https://github.com/matsim-scenarios/matsim-berlin`

only consider private cars and freight delivery as transport modes since public transport moves along a completely separated network. In our experiments only non-residential roads within the urban center of Berlin can potentially be priced, which sum up to ∼34k segments. The latter are highlighted in Figure 1.

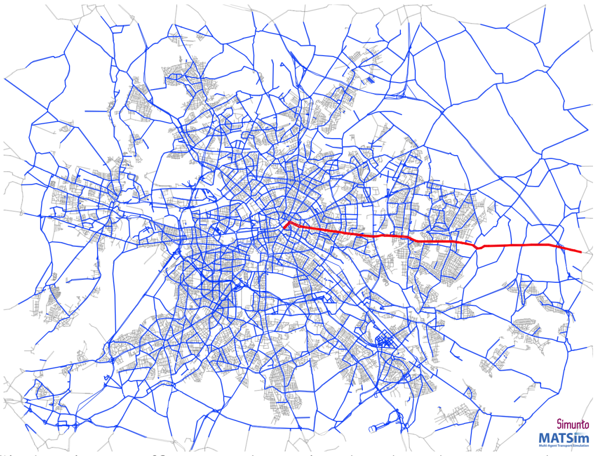

Figure 1. MATSim's private traffic network restricted to the urban area. Blue road segments can potentially be priced, gray segments are residential streets or lie outside of the city and cannot be priced, red segments represent the fastest route from Berlin Alexanderplatz to the eastern outskirts.

5.2 Value Model

The *flow rate* $f(k)$ of a road segment k is the number of vehicles passing it per hour. The *capacity* $c(k)$ of segment k is the maximum sustainable flow rate at which vehicles can reasonably be expected to traverse it. The ratio of both, i.e., the *volume to capacity ratio* $\lambda(k) = f(k)/c(k)$, is the utilization of a segment [26]. The number of licenses that the ISO sells for a segment equals the segment's capacity $c(k)$. This prevents congestion and allows cars to drive at *free flow speed*, i.e., they drive at the segment's speed limit. Given the length $d(k)$ and the free flow speed $s(k)$ of a road segment, one can calculate the *free flow travel time* $t(k) = d(k)/s(k)$. Summing up these numbers for the segments assembling a route, gives the free flow travel time for an entire route $t(S)$.

As Berlin's road network consists of ∼73k nodes, the number of OD pairs is too large for SPs to determine bids manually. Instead, SPs need to rely on a value model that determines bid values automatically. In our experiments, SPs have the value function

$$v(S) = \begin{cases} \mu \sum_{k \in S} d(k) \cdot \lambda(k) & \text{if } S = S^*, \\ \min \left\{ \mu \sum_{k \in S} d(k) \cdot \lambda(k),\ v(S^*) \cdot \dfrac{t(S^*)}{t(S)} \right\} & \text{otherwise,} \end{cases}$$

where $\mu = 0.00018$ euros$/m$ denotes the distance-based marginal cost of any driver [27], $R \subseteq 2^K$ the set of alternative routes, and $S^* \in R$ the fastest route. This value function ensures that SPs place higher bids for longer routes and those passing across highly demanded segments (reflected by $\mu \sum_{k \in S} d(k) \cdot \lambda(k)$). Moreover, SPs never bid more for alternative routes that require detours. Since S^* is the fastest route, it must hold that $t(S^*) \leq t(S)$ for all $S \neq S^*$. Therefore, $v(S^*) \cdot t(S^*)/t(S)$ scales down the value $v(S)$ of an alternative route with respect to the additional time for detours.

Assuming that all drivers take the fastest route in the Open Berlin Scenario, we can calculate exact values for $\lambda(k)$. In reality, SPs must forecast these numbers so that we assume that the SPs in our experiments deviate from these values by up to 20%. SPs must also predict their number of customers. Designing an accurate demand forecast for an SP is out of scope of this article so that we adopt a simplified model. Each trip between an OD pair represents the journey of an individual customer. For each time slot all available customers are distributed uniformly at random among the SPs, which submit their bids according to this demand. By doing so, the total demand of all SPs matches the one of the Open Berlin Scenario. This allows for a meaningful analysis of the allocation and pricing of road segments. Distributing customers between the same OD pair uniformly at random among the SPs causes the latter to submit buy bids for almost all requested OD pairs in the market, thus driving up the number of binary and continuous variables in the allocation problem. Creating these hard problem instances is done intentionally as it allows us to derive a strong upper bound on the tractability of real life problem instances.

For each OD pair of any customer in a time slot, we generate up to 10 alternative routes. SPs are assumed to include all these routes in their buy bids, as this increases their chance to win capacity even when a specific road segment on another route is overdemanded. A full description of the route set generation algorithm can be found in the long version of this paper.

Internally, the MATSim traffic simulator calculates the utility of agents based on a modified version of the Charypar-Nagel scoring function [28]. An agent's activity score is given by summing up the agent's marginal gain for performing the activity for a certain amount of time and subtracting the agent's disutilities incurred when traveling to activity's location [27]. Most of the parameters of the Charypar-Nagel utility function are agent-specific, making it impractical for our wholesale market as SPs need to purchase road capacity for thousands of agents which are unknown to them at the time of bidding.

6 Results

In this section we report the results of our computational experiments using the Gurobi Optimizer 9.1.0 to solve the allocation and pricing problem up to a "MIPGap" of 10^{-4}, i.e., the solution computed by Gurobi differs from the optimal solution by no more than 0.01%. Our test computer contains two Intel Xeon CPU E5-2620 @ 2.00GHz and 64GB of RAM. We analyze the tractability of the allocation problem as well as the emerging road prices for the initial forward market with one ISO and four service providers. The market is solved separately for all 24 time slots consisting of one hour each. Additionally, we also consider four different demand scalings (100%, 110%, 120%, 130%).

6.1 Allocation

To get a better understanding for the market size, let us consider time slot 17 with a demand scaling of 100%. In this specific time slot, service providers demand $\sim$244k trips between $\sim$23k distinct OD pairs. They submit $\sim$92k buy bids, while the ISO places one sell bid for each of the $\sim$34k road segments in the auction. This corresponds to an allocation problem with $\sim$92k binary and $\sim$951k continuous variables as well as $\sim$252k constraints. Despite this vast number of variables and constraints, the allocation problem can be solved surprisingly fast (see Figure 2), while the fraction of trips served with either an alternative or the fastest route is very large (see Figures 4 and 5). In total, the ISO sells more than 14.4 million road segment licenses (see Figure 3).

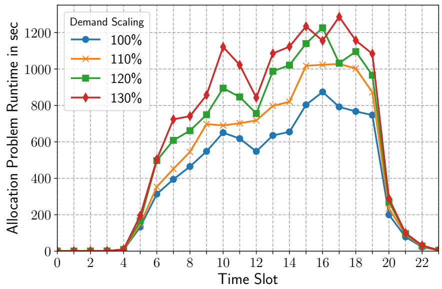

Figure 2. Runtime of the allocation problem

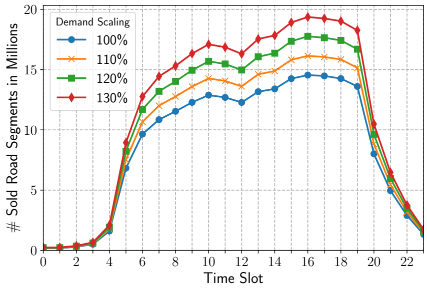

Figure 3. Number of sold road segments

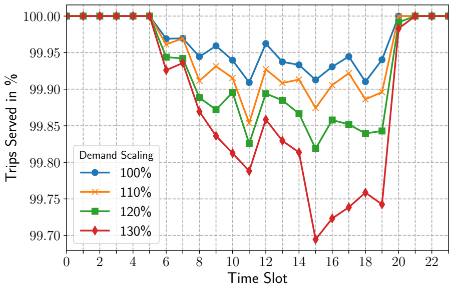

Figure 4. Percentage of demanded trips that can be served

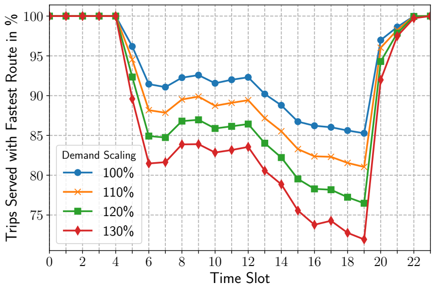

Figure 5. Percentage of demanded trips that are served with the fastest route

Result 1 *With a demand scaling of* 100% *the allocation and pricing problem can be solved in less than* 875 *and* 385 *seconds, respectively. In each time slot, more than* 85.2% *of the demanded trips can be served with the fastest route and less than only* 0.01% *of the trips cannot be satisfied. Even when demand is scaled to* 130%*, the allocation and pricing problem can be solved in less than* $1{,}300$ *and* $1{,}200$ *seconds, respectively. In each time slot, more than* 70% *of the trips can be served with the fastest route and less than* 0.031% *trips cannot be served. The optimization problem effectively distributes the traffic demand in substitutable routes on the available capacity.*

Increasing the demand above 100% causes more lower bounds of buy bids to become binding. This drives up the runtime of the allocation problem and reduces the number of trips being served with fastest routes.

For the sake of tractability, variables $x_i(S)$ and $y_i(k)$ are defined as continuous variables in the allocation problem (see Section 3). Whenever SPs are allocated a fractional capacity for a road segment, they incur a financial loss as they cannot resell this capacity to end consumers. However, our tests show that the number of fractionally allocated road segments is negligible compared to the overall number of segments traded.

Result 2 *The percentage of fractionally sold licenses is below* 0.029% *for the entire day when demand is scaled to* 100%*. For a demand scaling of* 130% *this value is* 0.073%*.*

6.2 Road Prices

Only a small fraction of the $\sim$34k road segments which can potentially be priced is actually priced in our experiments (see Figure 6). While some roads that are priced with varying amounts throughout the day, the vast majority of roads is never priced, leading to a low average road segment price (see Figure 7). The price fluctuations for roads in the city center are visualized in Figure 8.

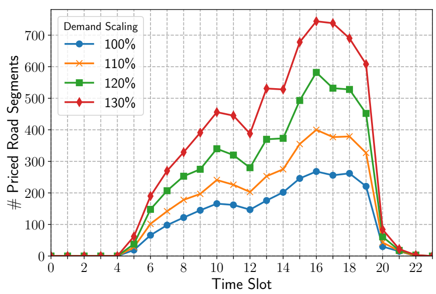

Figure 6. Number of priced road segments

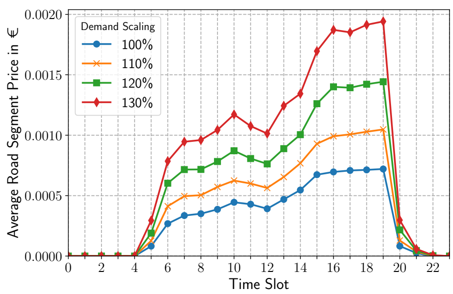

Figure 7. Average road segment price

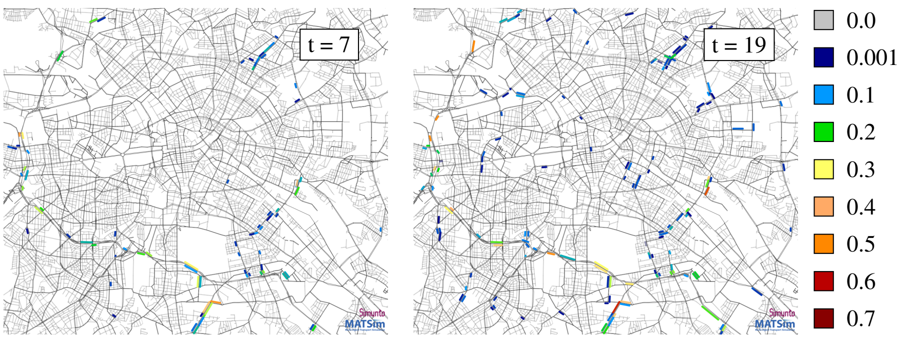

Figure 8. Prices (in cent) of city center segments at two time slots with demand scaled to 100%.

Result 3 *If the demand is scaled to* 100%*, only* 599 *of all* $34{,}179$ *road segments are priced at least once throughout the day and at most* 271 *road segments are priced simultaneously per time slot. Averaged over all time slots, the road segment price is*

0.031 cents and the route price is 13.6 cents. When scaling the demand to 130%, there are $1{,}607$ distinct road segments priced throughout the day, while at most 729 road segments are priced simultaneously per time slot. The average road segment and route price raise to 0.082 and 26.0 cents, respectively.

When calculating road prices with IP pricing, SPs may incur a loss when having to trade road capacities at these prices so that side-payments (so-called make-whole payments) become necessary (see Section 4). With the proposed bid language, our experiments show that these payments are negligible. Due to a high proportion of trips that can be served, the lower bounds specified in the buy bids are almost never binding, leading to a degeneration of the allocation problem to a linear program. However, adding further non-convex constraints to the bid language will lead to an increase of make-whole payments as we will show in the long version of this article.

Result 4 *For a demand scaling of 100%, the sum of all make-whole payments over all SPs and time slots is less than 3.1 cents, the social welfare (i.e., the sum of all bidder valuations for traded road capacities) is 2.4 million euros, and the total payment of all SPs is $471{,}000$ euros. Scaling the demand to 130% results in total make-whole payments of less than 7.5 cents, while the social welfare increases to 3.8 million euros and the total payment of all SPs to 1.1 million euros.*

To get a better understanding on how prices evolve for entire routes, we track the price for traveling between Berlin Alexanderplatz and the eastern outskirts (see Figure 1) over all time slots (see Figure 9). The route prices reflect the traffic waves during the course of a day. While in the morning it is more expensive to enter the city center from the eastern outskirts (see Figure 9a) than leaving the city center towards the East (see Figure 9b), it is the other way around in the afternoon.

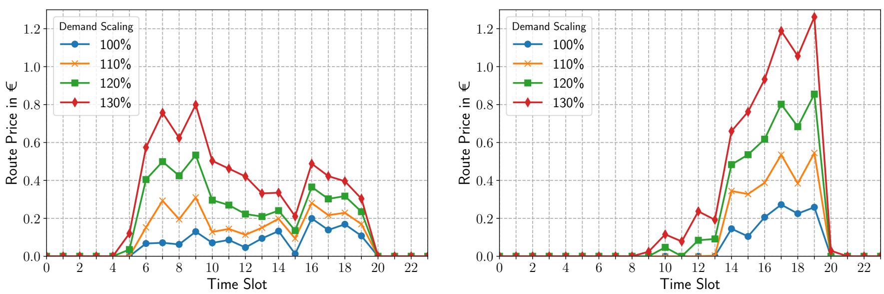

(a) Eastern outskirts → Alexanderplatz

(b) Alexanderplatz → Eastern outskirts

Figure 9. Price trajectory for traveling between Berlin Alexanderplatz and the eastern outskirts.

7 Discussion

The adoption of congestion pricing schemes has led to substantial revenues and significant traffic reductions in cities like London, Stockholm, and Singapore [3]. Due to these positive results, other cities such as Munich or Berlin currently discuss the implementation of similar policies [29, 30]. Often, congestion pricing is critically opposed upon its adoption, followed by a shift in public opinion in favor of such policies after the positive effects of traffic congestion become evident [3]. The skepticism towards congestion pricing are due to the fear of negative welfare effects (i.e., the high income class benefits at the expense of the poor) and privacy concerns (i.e., cars need to be tracked). Some authors [11, 12] provide some compelling arguments mitigating these concerns and we point the reader to these articles for a detailed discussion. The revenue redistribution has been identified as a crucial aspect for the public acceptance of road pricing [31].

The implementation of market-based dynamic congestion pricing is a gradual process, likely to start out with a simple static pricing scheme and then incrementally extending this policy towards dynamic congestion pricing. Singapore started out with a simple static cordon pricing scheme in 1975. In 2020, the city began to equip vehicles with GPS devices, which is the technological basis for distance-based and dynamic congestion charges in the future [3]. Researchers have already laid out a detailed plan on how the existing policy can gradually extended to reach such a futuristic pricing scheme [11].

For modeling a wholesale market for road capacity it is inevitable to make assumptions that may seem speculative, e.g., forecasting the number of SPs in the market. Certainly, also our value model does not capture an SP's true valuation in its entirety and resulting prices may change with another choice of parameters. However, these decisions do not weaken the core message of our research, which is to show that with our novel bid language a wholesale market for road capacity is feasible for a major city like Berlin.

8 Conclusion

Markets for utilities such as electricity are typically organized in a two-tiered structure. An independent system operator provides the basic infrastructure and a market on which multiple SPs compete. This sets incentives for innovation and the competition drives down prices for the end consumer. A similar market structure is compelling for congestion pricing as well. Wholesale markets are a central piece of the overall market design. Compared to other markets the sheer volume of products traded, the number of road segments per time slot, is far beyond what is traded on other markets. This leads to a huge welfare maximization problem and it is unclear whether such a market is tractable.

We propose a novel, compact bid language where SPs only need to specify prices for substitutable routes on origin-destination pairs. The demand for such pairs is readily available nowadays. We show that the welfare maximization problem is indeed huge with more than one million variables and 250 thousand constraints for a city like Berlin, but it can be solved with today's optimization technology. We can also derive effective prices in spite of a significant number of binary variables in the problem. Actually, only a small number of road segments has a positive price and the optimization does a good job in distributing the demand for origin-destination pairs on different routes throughout the network, thus keeping the city congestion-free. It is important to mention that the bid language is also decisive for tractability. Overall, the study shows that a wholesale market for road capacity is feasible and practical even for a city as large as Berlin.

References

1. Beheshtian, A., Geddes, R.R., Rouhani, O.M., Kockelman, K.M., Ockenfels, A., Cramton, P., Do, W.: Bringing the efficiency of electricity market mechanisms to multimodal mobility across congested transportation systems. Transportation Research Part A: Policy and Practice 131, 58–69 (2020)
2. Vickrey, W.S.: Pricing in urban and suburban transport. The American Economic Review 53(2), 452–465 (1963)
3. Lehe, L.: Downtown congestion pricing in practice. Transportation Research Part C: Emerging Technologies 100, 200–223 (2019)
4. Chiu, Y.C., Bottom, J., Mahut, M., Paz, A., Balakrishna, R., Waller, T., Hicks, J.: Dynamic traffic assignment: A primer. Dynamic Traffic Assignment: A Primer (2011)
5. de Palma, A., Lindsey, R.: Traffic congestion pricing methodologies and technologies. Transportation Research Part C: Emerging Technologies 19(6), 1377–1399 (2011)
6. Cheng, Q., Liu, Z., Liu, F., Jia, R.: Urban dynamic congestion pricing: an overview and emerging research needs. International Journal of Urban Sciences 21(sup1), 3–18 (2017)
7. Cramton, P., Geddes, R.R., Ockenfels, A.: Using technology to eliminate traffic congestion. Journal of Institutional and Theoretical Economics (JITE) 175(1), 126–139 (2019)
8. de Palma, A., Picard, N.: Route choice decision under travel time uncertainty. Transportation Research Part A: Policy and Practice 39(4), 295–324 (2005)
9. Do Chung, B., Yao, T., Friesz, T.L., Liu, H.: Dynamic congestion pricing with demand uncertainty: a robust optimization approach. Transportation Research Part B: Methodological 46(10), 1504–1518 (2012)
10. Hensher, D.A.: Tackling road congestion: What might it look like in the future under a collaborative and connected mobility model? Transport policy 66, A1–A8 (2018)
11. Cramton, P., Geddes, R.R., Ockenfels, A.: Markets for road use: Eliminating congestion through scheduling, routing, and real-time road pricing. Working paper (2019)
12. Cramton, P., Geddes, R.R., Ockenfels, A.: Set road charges in real time to ease traffic. Nature 560, 23–25 (2018)
13. Ausubel, L.M., Baranov, O.: A practical guide to the combinatorial clock auction (2017)
14. O'Neill, R.P., Sotkiewicz, P.M., Hobbs, B.F., Rothkopf, M.H., Stewart Jr, W.R.: Efficient market-clearing prices in markets with nonconvexities. European journal of operational research 164(1), 269–285 (2005)
15. Ziemke, D., Kaddoura, I., Nagel, K.: The matsim open berlin scenario: A multimodal agent-based transport simulation scenario based on synthetic demand modeling and open data. Procedia Computer Science 151, 870–877 (2019), https://www.sciencedirect.com/science/article/pii/S1877050919305848
16. Bichler, M., Gupta, A., Ketter, W.: Research commentary - designing smart markets. Information Systems Research 21(4), 688–699 (2010)
17. Bapna, R., Goes, P., Gupta, A., Jin, Y.: User heterogeneity and its impact on electronic auction market design: An empirical exploration. MIS quarterly pp. 21–43 (2004)
18. Adomavicius, G., Gupta, A.: Toward comprehensive real-time bidder support in iterative combinatorial auctions. Information Systems Research 16(2), 169–185 (2005)
19. Adomavicius, G., Curley, S.P., Gupta, A., Sanyal, P.: Impact of information feedback in continuous combinatorial auctions: An experimental study of economic performance. MIS Quarterly pp. 55–76 (2013)
20. Bichler, M., Hao, Z., Adomavicius, G.: Coalition-based pricing in ascending combinatorial auctions. Information Systems Research 28(1), 159–179 (2017)
21. Adomavicius, G., Curley, S.P., Gupta, A., Sanyal, P.: How decision complexity affects outcomes in combinatorial auctions. Production and Operations Management 29(11), 2579–2600 (2020)

22. Bikhchandani, S., Mamer, J.W.: Competitive equilibrium in an exchange economy with indivisibilities. Journal of economic theory 74(2), 385–413 (1997)
23. Baldwin, E., Klemperer, P.: Understanding preferences: demand types, and the existence of equilibrium with indivisibilities. Econometrica 87(3), 867–932 (2019)
24. Liberopoulos, G., Andrianesis, P.: Critical review of pricing schemes in markets with non-convex costs. Operations Research 64(1), 17–31 (2016)
25. Schubert, M., Kluth, T., Nebauer, G., Ratzenberger, R., Kotzagiorgis, S., Butz, B., Schneider, W., Leible, M.: Verkehrsverflechtungsprognose 2030 Los 3: Erstellung der Prognose der deutschlandweiten Verkehrsverflechtungen unter Berücksichtigung des Luftverkehrs. Tech. rep., Bundesministerium für Verkehr und digitale Infrastruktur (6 2014)
26. Transportation Research Board: Highway Capacity Manual. National Research Council, Washington, DC (2000)
27. Horni, A., Nagel, K., Axhausen, K. (eds.): Multi-Agent Transport Simulation MATSim. Ubiquity Press, London (Aug 2016)
28. Charypar, D., Nagel, K.: Generating complete all-day activity plans with genetic algorithms. Transportation 32(4), 369–397 (2005)
29. Berliner Senatsverwaltung für Umwelt, V.u.K.: Ergänzende instrumente zur finanzierung des berliner Öpnv: Machbarkeitsstudie. Tech. rep. (2020)
30. Falck, O., Fichtl, A., Janko, A., Kluth, T., Wölfl, A.: Verkehrliche wirkungen einer anti-stau-gebühr in münchen. Tech. Rep. 115, ifo Forschungsberichte (2020)
31. Creutzig, F., Javaid, A., Koch, N., Knopf, B., Mattioli, G., Edenhofer, O.: Adjust urban and rural road pricing for fair mobility. Nature climate change 10(7), 591–594 (2020)

5 Publication 3: The FUEL Bid Language

Peer-Reviewed Conference Paper

Title: Taming the Communication and Computation Complexity of Combinatorial Auctions: The FUEL Bid Language

Authors: Martin Bichler, Paul Milgrom, Gregor Schwarz

In: Conference on Information Systems and Technology 2020

Abstract: Combinatorial auctions have found widespread application for allocating multiple items in the presence of complex bidder preferences. The enumerative XOR bid language is the *de facto* standard bid language for spectrum auctions, despite the difficulties of enumerating all the relevant packages or solving the resulting NP-hard winner determination problem. We introduce the FUEL bid language, which was proposed for radio spectrum auctions to ease both communications and computations. We introduce a mathematical model of the resulting allocation problem and conduct extensive computational experiments, showing that the winner determination problem of the FUEL bid language can be solved reliably for large realistic-sized problem sizes in less than half an hour on average. In contrast, auctions with an XOR bid language quickly become intractable even for much smaller problem sizes. For the XOR bid language, the missing bids problem and computational hardness incur significant welfare losses compared to the FUEL bid language.

Contribution of dissertation author: Experimental design, software, formal analysis, investigation, visualization, joint paper management

Reference: Bichler et al. (2020)

Comment: The long version of this paper is currently in the 2nd round of revisions for the academic journal *Management Science*.

Taming the Communication and Computation Complexity of Combinatorial Auctions: The FUEL Bid Language

Martin Bichler[1], Paul Milgrom[2], Gregor Schwarz[1]
Technical University of Munich[1]
Stanford University and Auctionomics[2]

Abstract

Combinatorial auctions have found widespread application for allocating multiple items in the presence of complex bidder preferences. The enumerative XOR bid language is the *de facto* standard bid language for spectrum auctions, despite the difficulties of enumerating all the relevant packages or solving the resulting NP-hard winner determination problem. We introduce the FUEL bid language, which was proposed for radio spectrum auctions to ease both communications and computations. We introduce a mathematical model of the resulting allocation problem and conduct extensive computational experiments, showing that the winner determination problem of the FUEL bid language can be solved reliably for large realistic-sized problem sizes in less than half an hour on average. In contrast, auctions with an XOR bid language quickly become intractable even for much smaller problem sizes. For the XOR bid language, the missing bids problem and computational hardness incur significant welfare losses compared to the FUEL bid language.

Key words: Market design, spectrum auction, bid languages

1 Introduction

Many information systems nowadays are designed to coordinate activities or allocate scarce resources. The design of respective information systems has a number of challenges because incentives of the participants need to be considered, but also computational problems play a role. This has led to a fruitful line of research on the design of electronic markets (Banker and Kauffman 2004, Bichler et al. 2010, Adomavicius et al. 2012). The information systems literature has made many contributions to the design of electronic market mechanisms, specifically to allocation and payment rules (Xia et al. 2004, Adomavicius and Gupta 2005, Bichler et al. 2009), to the analysis of bidder behavior (Scheffel et al. 2011, Adomavicius et al. 2012), as well as to the design of markets for specific domains (Guo et al. 2007, Bapna et al. 2007). We extend this line of research and focus on auction mechanisms for large multi-object markets. Auctions for radio spectrum are the most prominent example of such markets as they are often involving hundreds of licenses to be sold to telecom companies by a regulator (Bichler and Goeree 2017). A common recommendation in such auctions world-wide is to use a combinatorial auction with XOR bidding, in which bidders simply enumerate values for all possible combinations of items. The auctioneer then solves a combinatorial optimization problem to find the allocation that maximizes the total bid.

As the number of items to be allocated becomes large, however, a full XOR-based approach to auctioning quickly becomes impractical for two reasons. The first is related to communication complexity (Nisan and Segal 2006). In combinatorial spectrum auctions in Canada using XOR bidding, there have sometimes been more than 100 spectrum licenses for sale, leading to more than 2^{100} different packages — far too many for any bidder to enumerate (Kroemer et al. 2017). The second reason is that computations at this scale can be impractical. To address that problem, the auctioneer in Canada limited the number of XOR bids that each bidder is permitted to submit to 2,000, treating the many missing packages as if they have received bids of zero. In lab experiments comparing an XOR design to alternatives, the resulting *missing bids problem* from XOR bidding leads to substantial efficiency losses, even with many fewer than 100 licenses (Bichler et al. 2014).

We focus on the design of large markets with hundreds of items and bidders and use the recent plans for a private C-band auction as a case for our analysis. In mid-2019, a consortium of companies providing satellite downlink for commercial television in the United States proposed to conduct a private sale of their C-band spectrum rights to support a fast 5G wireless deployment. They suggested a novel combinatorial auction design dubbed FUEL (*Flexible Use and Efficient Licensing*) (Milgrom 2019).[1] Spectrum licenses were to be offered in 406 geographical areas — the *Partial Economic Areas* (PEAs) with 14 licenses to use 20 MHz of bandwidth in each, so the number of possible combinations that any bidder might win in the proposed auction was 15^{406}. The FUEL bid language was introduced in an attempt to tame both the communication and computational complexity by requiring bidders to use a particular parameterized valuation model to describe values for all license combinations or for certain subsets of them.[2]

This paper explains the rationale for the language and tests and compares the computational tractability of the FUEL and XOR designs assuming a high level of participation in the auction by bidders with similar license values. High participation and similar values are thought to make the optimization more challenging by providing more near-optimal combinations of bids for the software to rule out. Our simulations of the auction show that even with a vastly reduced bid set, accurate computations with XOR bids require significantly more computation time than FUEL. The optimization problem coded using the FUEL bid language utilizes many binary variables, and just as for the XOR auction, computing the optimal solution is NP-hard. However, our computational experiments show that, in practice, even in auctions with more participating bidders than are expected for an actual spectrum auction, optimal solutions for the FUEL auction can be computed on a desktop computer in mere minutes using commercial off-the-shelf optimization software. Moreover, computation time tends to grow only linearly in the number of bid groups.

With the FUEL restrictions on bid groups in place, we are able, in practice, to solve large problems with more than 400 licenses and more than 1000 bidders using a state-of-the-art branch-

[1] In February 2020, the US Federal Communications Commission decided against using a private auction, so the FUEL design will not be adopted for the C-band auction.

[2] The design of parametric bid languages for various auction problems has previously received scholarly attention in Milgrom (2009), Bichler et al. (2011), Eilat and Milgrom (2011), Bichler et al. (2017).

and-cut algorithm in a few minutes of runtime. Similar problems are intractable when coded using the XOR bid language and beyond what one would expect to solve to optimality.[3]

The FUEL design may also serve as a template for other large combinatorial auctions in which there are several kinds of items with economies of scale and scope among them. By exploiting computationally tractable hierarchical valuation structures, some of the most important barriers to large-scale auctions may be overcome. Auctions with regional lots as in the procurement of school meals (Kim et al. 2014) and the sale of fishery access rights (Iftekhar and Tisdell 2012) could be appropriate candidates, but so could be TV ad auctions (Goetzendorff et al. 2015).[4]

2 Bid Languages

A bid in an auction expresses a bidder's willingness to pay money for various outcomes and depends both on the bidder's private preferences and its bidding strategy. A *bid language* defines the format used to communicate the bids. For combinatorial auctions some common bid languages are built from *bundles* (also known as *packages*), which are subsets of the item set, *atomic bids*, which associate a price with a bundle, and *logical rules*, which govern which bids can win simultaneously.

The two most popular and intuitive bid languages of this kind are *exclusive-OR* (XOR) and *additive-OR* (OR). In both bid languages bidders can specify multiple atomic bids and in case a bid becomes winning, the bidder gets all items contained in the bundle. While in an XOR bid language at most one of each bidder's atomic bids become winning, any non-intersecting combination of atomic bids can win in an OR bid language.

Spectrum auctions have so far resisted the design of particular bid languages and, until recently, only XOR bid languages have been used, leading to significant limitations in applications like the Canadian auctions in 2014 with 100 licenses (Kroemer et al. 2017). The even larger number of

[3]We use a time limit of 30 minutes for each problem instance in our experiments. If problems could not be solved to near-optimality within 30 minutes or if there was a large integrality gap after 30 minutes, they typically could also not be solved to optimality or near-optimality with several hours of computation time.

[4]The bidding language is just one element of a sealed-bid auction design. Two others are the winner determination rule and the payment rule, which together determine bidders' incentives. If a hierarchical language like FUEL makes optimization tractable, then that would allow selecting the allocation that maximizes the total bid and setting payments using the Vickrey-Clarke-Groves payment rule, which is incentive-compatible (Green and Laffont 1977).

licenses available in spectrum auctions in the United States are one reason why combinatorial auctions have not yet been adopted there. On the one hand, the missing bids problem in such auctions is huge and the auctioneer cannot confidently solve a large winner determination problem with hundreds of licenses and an XOR bid language. On the other hand, using simple auction formats such as the simultaneous multi-round auction and related clock auctions limits expressiveness of the bids and creates an *exposure problem* for bidders, in which they may win some but not all of the licenses they need for a viable network.

3 FUEL Auction Design

Similar to previous auctions designed by the Federal Communications Commission (FCC), the market area for the C-band auction is geographically subdivided into smaller entities, so-called *Economic Areas* (EAs). As some local market participants are expected to be only interested in spectrum for some part of an Economic Area, each EA is split again into *Partial Economic Areas* (PEAs), with the number of PEAs in an EA ranging from 1 to 12. In total, there are 170 EAs and 406 PEAs across the contiguous United States.

In each PEA, 280 MHz of spectrum is sold in the C-band auction. The spectrum in a PEA is split into 14 homogeneous blocks, each containing 20 MHz of the 280 MHz available per PEA. All license blocks can be made available within 36 months of the FCC's final order. Furthermore, in 46 of the 50 most populous PEAs the satellite companies are able to provide 100 MHz (5 blocks) of spectrum even earlier, within 18 months of the FCC's approval of the auction process. With respect to these temporal constraints, licenses for spectrum blocks that are available within 18 months are referred to as *early*, the remaining licenses are called *late*. It is possible to serve a bidder's demand for late blocks with available early blocks, but the reverse is not possible.

3.1 Bid Language

Assuming the C-band auction is executed with 406 PEAs and 14 spectrum blocks per PEA, the number of potential distinct packages equals 15^{406}: far too many to enumerate. The FUEL bid language circumvents this problem by allowing each bidder to submit a small number of *bid groups*. Each bid group is based on a single package bid, called the *base bid* consisting of a *base package* and a *base price*. Furthermore, a bid group incorporates *adjustments*, defining the price that applies to a package that *increments* or *decrements* the number of licenses to be purchased in a PEA. Each increment is associated with a markup to the base price and each decrement is associated with a discount (see Figure 1). Adjustments are intended to provide a natural and intuitive way for bidders to specify their demand for spectrum and at the same time avoid the missing-bids problem.

Bidder 1 LATE SMALL Base price: 177																
		#Licenses														
EA	**PEA**	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**
60	155		-90	Base	101	174										
60	354			Base												

Figure 1: Bidder 1 submits a bid group whose base package contains 2 licenses in PEA 155 and 2 licenses in PEA 354. She also defines adjustments (in $) for 1, 3, and 4 licenses in PEA 155. If the auctioneer accepts her bid group and assigns her 4 licenses in PEA 155 and 2 licenses in PEA 354 (highlighted blue), her bid for this set of licenses is $\$177 + \$174 + \$0 = \$351$.

In 46 of the PEAs, 5 of the 14 spectrum blocks are available early and bidders may wish to bid more for those blocks. The FUEL bid language addresses this issue by allowing each bidder to designate its bid groups as bids for *early/mixed* spectrum or for *late* spectrum and interpreting the base bids and adjustments differently for each specification.

A late bid group works exactly as described above, with the interpretation that its adjustments apply only to late spectrum. The seller, at its sole discretion, can convert early spectrum to late spectrum by delaying rights to use that spectrum. For an early/mixed bid group, the base bid is interpreted as follows. If the base package contains a PEA providing 5 early and 9 late licenses and a bidder states a demand for k licenses in her base bid, then this is interpreted as demand for $\min\{5,k\}$ early licenses and $\max\{0,k-5\}$ late licenses. Increments in the bid group constitute demand for additional late license blocks, while decrements from the base package first reduce the number of early licenses in the package and, if the reduction takes the number below zero, then apply to reduce the number of late licenses.

Bid groups are further classified with respect to the MHz-pop of their base package. The MHz-pop of a set of licenses for the same PEA is given by the product of the frequency bandwidth in MHz and the population of the respective PEA. Summing up the MHz-pop of all PEAs present in the base package gives the MHz-pop of the base package. If the MHz-pop of a base package is no less than the MHz-pop equivalent of two national licenses (i.e. two licenses in all 406 PEAs), then the corresponding bid group is considered to be a *nationwide* bid group and is labeled *large*, otherwise it is a *local* bid group and is classified *small*. While large bid groups may include any combination of PEAs, small bid groups may only contain PEAs from the same EA. Bidders are also restricted in the number of small and large bid groups they are allowed to submit. The exact numbers were to be determined according to the computational feasibility of the underlying allocation problem. Moreover, bidders may either win a single large bid group or multiple small bid groups, but never large and small bid groups at the same time. In addition to that, small bid groups can never become winning simultaneously if they contain bids on the same EA.

Reserve prices are commonly used in auctions to set the minimum price at which the auctioneer is willing to sell the products. The FUEL bid language implements reserve prices by placing a bid group on behalf of the auctioneer. In the course of the auction, the reserve bid is treated like any other bid group of a bidder. If the auctioneer's reserve bid is winning when solving the underlying allocation problem, the respective licenses remain unsold.

The auctioneer in the C-band auction sets reserve prices by submitting a single bid group which includes bids on each of the 406 PEAs. The number of demanded licenses in the base package as well as the base price is set to 0. Through suitable markups the auctioneer then specifies linear reserve price for multiple licenses of the PEA. It is sufficient for the auctioneer to specify reserve prices only for late licenses as every early license can be transformed to a late license.

3.2 Pricing

As presented in Appendix A, the allocation problem of the C-band auction can be represented as a binary program. The winning bid groups and also their associated winning adjustments are given by the optimal solution of the corresponding binary program. In accordance with other auctions conducted by the FCC, the determination of winning bids and the price that bidders are obliged to pay are separate calculations. While the binary program is used to determine the allocation of licenses to bidders, a Vickrey-nearest core-selecting rule is applied to compute the prices that the bidders will have to pay.

In the following we will exclusively focus on the most computationally challenging problem: the allocation problem that follows the main bidding round. Incentives and payment rules are not treated in detail here. Rather, we observe that VCG payment rules require the winner determination problem to be solved to optimality. So, among the advantages of the FUEL language is to enable such payment computations.[5]

4 XOR Bid Language

As we compare the empirical complexity of auctions with the FUEL bid language to auctions with a standard XOR bid language, we briefly introduce the XOR bid language for the application of the C-Band auction.

Similar to FUEL bid groups an XOR bid consists of a set of PEAs for which the bidder would like to acquire licenses. For each of these PEAs the bidder specifies two numbers which indicate the amount of early and late licenses that the bidder would like to purchase in the respective PEA. Every XOR bid is also associated with a price which expresses the bidder's willingness to pay for the set of licenses. In contrast to the FUEL bid language, the XOR bid language neither distinguishes between early and late nor between small and large bid groups. While local and national bidders may win at most one of their XOR bids, the auctioneer is exempt from this rule. In order to represent the auctioneer's reserve bid for a single PEA offering 14 licenses, only 4 XOR bids need to be generated: one XOR bid each for 1, 2, 4, and 8 licenses.

The corresponding winner determination problem can be found in Appendix B. It is well-known that the winner determination problem with an XOR bid language is strongly NP-hard and can be modeled as a weighted set packing problem (Lehmann et al. 2006).

[5]Core-selecting payment rules, which are widely used in spectrum auctions, also require the allocation problem to be solved to optimality (Day and Milgrom 2008, Goetzendorff et al. 2015).

5 Experimental Design

In our experiments we differentiate between *local* and *national* bidders. Local bidders are only active in up to a dozen PEAs, while national bidders are interested in licenses in almost all 406 PEAs. For our tests we assume that there are 10 national and 1,000 local bidders, which we chose to be larger than (but of the same order as) the actual numbers in any previous spectrum auction.

5.1 Value Model

A widespread international metric for comparing the prices of spectrum is the license price per MHz-pop. The value model of the FUEL bid generator is based on this convention using the PEA population data provided by the FCC.[6] The worth w_p of a single license in PEA $p \in P$ is chosen to be a constant fraction of the license's MHz-pop.

Naturally, bidders' valuations for licenses in a particular PEA differ depending on their financial strength and their current possession of frequencies. To generate idiosyncratic bidder valuations, we introduce value factors r_{ip} for each bidder $i \in I$ and PEA $p \in P$ which scale the worth of a PEA for a particular bidder. In general, local bidders are financially weaker than national bidders so that we choose r_{ip}^{local} and r_{ip}^{national} uniformly at random from the intervals $[1.0, 1.3]$ and $[1.1, 1.4]$, respectively. The auctioneer's idiosyncratic value factor is set to 1.0 for all PEAs, so that the auctioneer's reserve prices equal a constant fraction of the MHz-pop in each PEA.

To provide a functional 5G network, bidders need spectrum bandwidth of at least 40 MHz. Therefore, bidders have only little interest in being allocated less than 2 licenses (i.e. less than 40 MHz) but also their marginal valuation for more than 5 licenses is very small. As a consequence, the valuation of a bidder is represented best by a sigmoid function whose point of inflection Δ_i is a bidder specific value chosen uniformly at random from the interval $[2, 4]$. Scaling this sigmoid function by the idiosyncratic bidder valuations and shifting it so that is crosses the origin gives $d_{ip}(x) = r_{ip} w_p \left(\frac{1}{1+e^{-x+\Delta_i}} - \frac{1}{1+e^{\Delta_i}} \right)$, where x is the number of licenses demanded by bidder $i \in I$ in PEA $p \in P$, and $r_{ip} w_p$ describes the bidder's idiosyncratic valuations for the respective PEA.

[6] `https://transition.fcc.gov/bureaus/oet/info/maps/areas/data/FCC_PEA_website.xlsx`

In 46 of the PEAs, 5 out of the 14 offered licenses can be made available 18 months earlier than the remaining 9 licenses. Depending on the bidder's current stock of licenses, receiving early licenses might be essential but may also have little value to the bidder. This is modeled by factor t_i, drawn uniformly at random from the interval $[1.0, 1.5]$ for each bidder $i \in I$, that describes by which factor the bidder values early over late licenses. A bidder's valuation for winning x early and y late licenses in PEA $p \in P$ is then given by

$$v_{ip}(x,y) = \begin{cases} 0 & \text{if } x+y=0, \\ r_{ip}\, w_p \left(\frac{1}{1+e^{-x-y+\Delta_i}} - \frac{1}{1+e^{\Delta_i}} \right) \left(t_i \frac{x}{x+y} + \frac{y}{x+y} \right) & \text{otherwise.} \end{cases}$$

For exemplary values of w_p, r_{ip}, Δ_i, and t_i the valuation function $v_{ip}(x,y)$ is plotted in Figure 2.

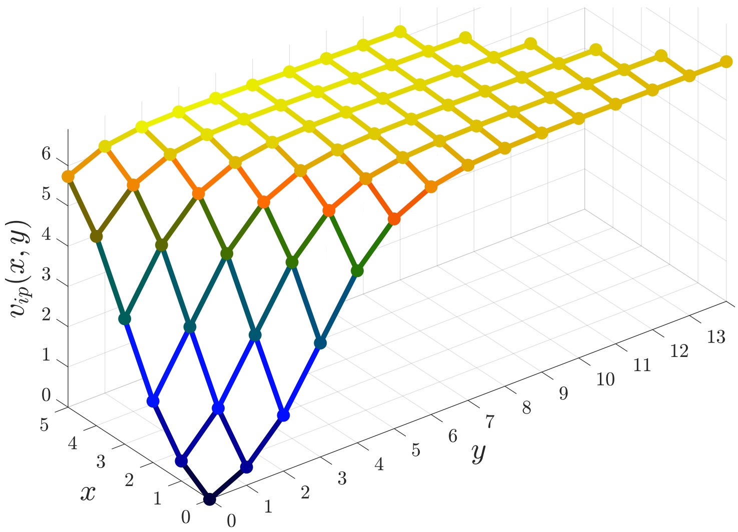

Figure 2: Valuation function $v_{ip}(x,y)$ with $w_p = 5$, $r_{ip} = 1.25$, $\Delta_i = 3$, and $t_i = 1.1$.

5.2 FUEL Bid Generation

We assume that every local bidder is active in only a single EA. The number of PEAs included in the bidder's small bid group is chosen uniformly at random between 1 and the number of PEAs available in the respective EA (at most 12). In practice, some local bidders may be active in several PEAs belonging to different EAs, but for the purposes of estimating runtime, these bidders can equivalently be represented as multiple independent local bidders. Unlike local bidders, we assume that national bidders are active throughout the whole United States: each of their bid groups covers at least 380 of the 406 PEAs. The PEAs not contained in a bid group are chosen uniformly at random among the 50% least populous PEAs, so that national bidders are always able to provide service in all of the most densely populated areas whenever one of their bid groups is winning.

When local and national bidders place a bid on licenses in a PEA, they have to state a base bid and may additionally specify several adjustments. According to the sigmoid value model (see Section 5.1), a bidder's largest marginal gain for a license is at Δ_i, the point of inflection of the bidder's valuation function v_{ip}. Therefore, the base package either contains $\lfloor \Delta_i \rfloor$ or $\lceil \Delta_i \rceil$ licenses in each PEA. The number of adjustments a bidder specifies for each PEA is chosen uniformly at random from the interval $[0,4]$. The selected adjustments always constitute a consecutive interval around the base bid as it is assumed that this models the bidders' valuations most accurately.

If a FUEL bid group contains PEAs for which early licenses are available, then the bid group is marked early/mixed with a probability of $5/14 \approx 40\%$ as this refers to the fraction of early licenses available in the respective PEAs. Otherwise, the bid group is considered late.

The number of bid groups that local and national bidders are allowed to submit is a parameter that was still undefined for the C-band auction at the time the design was proposed. For the following tests we assume that both local and national bidders are allowed to submit up to five bid groups. This number assures computational tractability within a time limit of 30 minutes and is expected to be a sufficiently large for bidders to express their valuations accurately.

5.3 Deriving XOR from FUEL Bids

In an XOR bid language bidders are unable to specify adjustments. If a bidder wants to state the same information as in a FUEL bid, she has to place one XOR bid for each possible combination of adjustments and adapt the price of the XOR bid according to the chosen markups and discounts.

Unfortunately, the number of XOR bids necessary to reproduce a FUEL bid group can become very large. If a national bidder submits a bid group containing all 406 PEAs and specifies two adjustments in each of them, then there would be $3^{406} \approx 5.14 \cdot 10^{193}$ XOR bids necessary to replicate the FUEL bid group. In order to guarantee computational tractability of the XOR allocation problem, it is indispensable to restrict bidders in the number of XOR bids to be submitted. Thus, bidders can only submit bids for a fraction of all possible adjustment combinations (see Figure 3).

EARLY SMALL Base price: 780				
		#Licenses		
EA	**PEA**	**3**	**4**	**5**
7	44	Base	418	741
7	271	-50	Base	

$\Longrightarrow$

XOR Price: 1,471			
		#Licenses	
EA	**PEA**	**2**	**3**
7	44 Early		Base
7	44 Late	Base	
7	271 Late		Base

Figure 3: A random XOR bid is generated from a FUEL bid by picking a random adjustment combination (highlighted blue).

5.4 XOR Bid Generation

In contrast to the FUEL bid language, bids in the XOR bid language are no longer subject to any EA restrictions. In particular this means that bidders can submit bids for any subset of PEAs even though they belong to different EAs. When generating XOR bids it is assumed that any local bidder's market area contains between 1 and 10 PEAs (potentially belonging to different EAs) which form a highly cohesive component. Each XOR bid of a local bidder contains between 1 and the maximum number of PEAs available in her market area, so that each local bidder's XOR bid encompasses 3.25 PEAs on average. Similar to a national bidder's FUEL bid groups, each XOR bid of a national bidder contains at least 380 of the 406 PEAs. The PEAs not contained in an XOR bid of a national bidder are chosen uniformly at random among the 50% least populous PEAs.

Both local and national bidders demand between 1 and 5 licenses in each PEA. The exact number of licenses is chosen uniformly at random from this interval for each XOR bid and each PEA. In case an XOR bid contains a PEA offering early licenses, then the number of early licenses demanded by a bidder is chosen uniformly at random between 1 and the number of licenses demanded by the bidder in this PEA.

6 Results

We conduct our computational experiments using the Gurobi Optimizer 8.1.1 to solve the winner determination problem up to a tolerance ("MIPGap") of 0.001, i.e., the solution computed by Gurobi differs from the optimal solution by no more than 0.1%. The time limit is set to 30 minutes for all test instances. Our test computer contains two Intel(R) Xeon(R) CPU E5-2620 @ 2.00GHz and 64GB of RAM. All test instances are available upon request.

6.1 FUEL Bid Groups

The original FUEL proposal did not specify the number of bid groups that local and national bidders would be allowed to submit, but proposed to choose those to ensure the computational tractability of the winner determination program.

Let z_L and z_N denote the number of bid groups that local and national bidders are allowed to submit, respectively. For each configuration of z_L and z_N, we generated 25 random instances with the FUEL bid generator. Table 1 summarizes the number of bid groups submitted by all bidders including the auctioneer (denoted $\sum z$), the average number of binary variables and constraints in the winner determination problem, the average runtime, the number of test instances that exceed the time limit of 30 minutes, the maximum MIPGap of all 25 test instances, and the average number of licenses that remain unsold out of 5,684 (406×14) licenses. Test instances that exceed the time limit of 30 minutes are weighted with 1,800 seconds when computing the average runtime.

z_L	z_N	$\sum z$	**Binary Variables**	**Constr.**	**Runtime in sec.**	**TLE**	**Max. MIPGap**	**Unsold Licenses**
1	1	1,011	27,935	6,613	5	0 of 25	0.0010	1.1
3	3	3,031	63,688	18,871	153	0 of 25	0.0010	2.5
5	5	5,051	99,351	30,121	356	0 of 25	0.0010	1.8
7	7	7,071	135,246	41,374	745	1 of 25	0.0014	1.3
10	10	10,101	188,709	58,193	1,212	3 of 25	0.0030	1.2
15	15	15,151	278,317	86,368	1,640	16 of 25	0.0101	1.3

Table 1: Average values of 25 test instances for different configurations of the number of small (large) bid groups that local (national) bidders are submitting.

Result 1. *If the 1,000 local bidders are allowed to submit 5 small bid groups and the 10 national bidders are allowed to submit 5 large bid groups, we can compute the winning allocation within 356 seconds and a MIPGap of only* 0.001 *on average.*

The number of small and large bid groups that bidders are allowed to submit has a direct impact on the number of binary variables and constraints present in the winner determination problem. Restricting the number of bid groups in the auction can therefore reduce the runtime substantially.

6.2 FUEL vs. XOR

To compare the efficiency of the standard XOR and FUEL bid language, we first generate the FUEL bids and then derive XOR bids from them as described in section 5.3. For our FUEL instances, we assume that local and national bidders submit 5 bid groups as our tests in section 6.1 imply that such instances can be solved within the time limit of 30 minutes.

In order to keep the XOR allocation problem tractable, we restrict local and national bidders in the maximal number of XOR bids they are allowed to submit and denote these upper bounds by z_L and z_N, respectively. The adjustment combinations for which XOR bids are generated are selected uniformly at random. As a consequence, bidders in the XOR auction have the same valuations as in the FUEL auction, but they are only able to state a fraction of the potential winning FUEL bid combinations.

Table 2 shows the efficiency (denotes Eff.) for different combinations of the maximum number of bids that local and national bidders may submit (denotes z_L and z_N). The interpretation of the remaining columns is the same as for table 1.

Type	z_L	z_N	$\sum z$	**Bin. Vars.**	**Constr.**	**Runtime in sec.**	**TLE**	**Max MIPGap**	**Unsold Licenses**	**Eff.**
FUEL	5	5	5,051	99,351	30,121	356	0 of 25	0.0010	1.8	1.000
XOR	1	1	2,634	2,634	1,462	32	0 of 25	0.0010	806.3	0.756
XOR	5	5	6,605	6,605	1,462	277	2 of 25	0.0022	299.3	0.903
XOR	10	10	10,358	10,358	1,462	654	8 of 25	0.0115	202.6	0.925
XOR	15	15	13,565	13,565	1,462	770	9 of 25	0.0052	164.0	0.936

Table 2: Average values of 25 test instances for different limitations on the number of XOR bids that local and national bidders are allowed to submit.

Result 2. *If bidders are only allowed to submit the same number of bids in the XOR as in the FUEL auction, more than* 5.2% *of all licenses remain unsold, the welfare loss compared to FUEL is almost* 10%, *and there are already 2 out of 25 test instances that are intractable. Even when bidders are allowed to submit three times as many XOR as FUEL bids, still more than* 2.8% *of the licenses remain unsold and the welfare loss is* 6.4%.

6.3 Unrestricted XOR

Deriving XOR bids from previously generated FUEL bids (as done in the tests of Section 6.2) implies that the XOR bids encompass the same EA restrictions as the original FUEL bids. In particular, this means that the XOR bids which are derived from a local bidder's small bid group contain bids for only a single EA. A fully combinatorial XOR bid language, however, does not pose such restrictions on the bids but allows the auction participants to bid on licenses for any subset of PEAs. Due to the additional interdependencies between bids containing PEAs of different EAs, the corresponding winner determination problem becomes more complex.

In the following test, we check whether an XOR bid language that only restricts bidders in the maximum number of admissible XOR bids without imposing any further restrictions is computational tractable for the C-band auction. Similar to our previous tests, we assume that there are 10 national and 1,000 local bidders. Table 3 shows the test results for different configurations of the maximum number of XOR bids that local and national bidders may submit in the auction. The interpretation of the columns is the same as for tables 1 and 2.

z_L	z_N	$\sum z$	**Binary Variables**	**Constr.**	**Runtime in sec.**	**TLE**	**Max MIPGap**	**Unsold Licenses**
1	1	2,634	2,634	1,462	34	0 of 25	0.0010	748.5
2	2	3,644	3,644	1,462	699	5 of 25	0.0037	444.9
3	3	4,654	4,654	1,462	1,800	25 of 25	0.0158	309.4
10	10	11,193	11,193	1,462	1,800	25 of 25	0.0123	146.4
50	50	45,545	45,545	1,462	1,800	25 of 25	0.0217	55.0

Table 3: Average values of 25 test instances for different restrictions on the number of XOR bids that both local and national bidders are allowed to submit.

Result 3. *Even if both local and national bidders are restricted to submit no more than two XOR bids without restrictions on the EAs, 5 out of 25 instances could not be solved within the time limit.*

Even if all bidders are restricted to submit no more than 3 XOR bids, none of the 25 instances can be solved within the time limit even though this number of XOR bids is far too small to give a reasonable account of national bidders' preferences. Such limited bids also result in many licenses remaining unsold.

7 Conclusions

The missing bids problem is one of the key problems in larger combinatorial auctions with the standard XOR bid language. Such design issues are at the core of auction design research in information systems. In this paper we investigate the empirical hardness of the FUEL bid language based on the case of the planned C-band auction for the US, which constitutes an important real-world case. Even though the winner determination problem of the FUEL bid language is NP-hard and contains roughly 100,000 binary variables and 30,000 constraints, our experiments indicate that this auction can consistently be solved in less than 30 minutes, and usually much less. We find evidence that the short solution times result predominantly from the hierarchical structure created by FUEL, which allows the Gurobi optimizer to decompose the binary program effectively.

In contrast to FUEL, we show that a fully enumerative XOR bid language quickly becomes computationally intractable. More importantly, bidders would need to specify an exponentially large set of XOR bids to express the same preferences as in a FUEL bid group with adjustments. Although the FUEL bid language is not fully expressive and limits the set of values that can be expressed relative to an XOR bid language, it is based on common spectrum valuation methods and may often be able to express values close to the bidders' actual ones. To the extent that FUEL bids fail to capture actual values, that must be weighed against its powerful mitigation of the missing bids problem that inevitably arises in large auctions using XOR bids. Our experiments show that both the missing bids problem and computation failures using an XOR bid language can lead to significant welfare losses.

In summary, by allowing bidders to use bid groups with adjustments to their base bids, the FUEL bid language gives bidders an intuitive and compact way to describe their valuations and effectively address the missing bids problem. The hierarchical structure of the bid groups makes it possible to solve very large problem instances exactly on a desktop computer in a matter of minutes. The specifics of the bid language allow for exact solutions in large-scale auctions with several hundred items, for which this would not have been considered feasible only recently.

Appendix A FUEL Winner Determination Problem

Let I_0, A, and P, denote the set of bidders including the auctioneer, the set of EAs and the set of PEAs, respectively. In each PEA $p \in P$ there are $e_p \in \{0,5\}$ early and $\ell_p \in \{9,14\}$ late licenses up for sale; the set $P^E \subseteq P$ denotes the set of PEAs with $e_p = 5$. Each bidder $i \in I_0$ submits a set G_i of bid groups. Large early and late bid groups of bidder i are denoted G_i^{TE} and G_i^{TL}, respectively. Similarly, small early and late bid groups of bidder i are denoted G_i^{SE} and G_i^{SL}. For both of these sets, we define subsets G_{ia}^{SE} and G_{ia}^{SL}, respectively, which include all small early and late bid groups of bidder i that include bids on EA $a \in A$. The PEAs contained in the base package of bid group $g \in G_i$ of bidder i are denoted P_i^g. For each PEA $p \in P_i^g$ bidder i specifies a number b_i^{pg}, which defines the number of licenses contained in the bidder's base bid and a set $K_i^{pg} \subseteq K = \{0,1,\dots,14\}$ of adjustments. For each of the adjustments $k \in K_i^{pg}$ bidder i specifies a markup or discount μ_i^{gpk} that defines how the bid group's base price ω_i^g is adjusted in case the adjustment k becomes winning. The number M denotes the maximum number of small bid groups that any bidder submits.

There exist three types of decision variables. Binary variables $x_i^g \in \{0,1\}$ denote whether bidder $i \in I_0$ wins bid group $g \in G_i$. The information which adjustment k is accepted for each PEA p contained in a winning bid group g is conveyed in binary variable $y_i^{gpk} \in \{0,1\}$. Finally, the decision variable $z_i \in \{0,1\}$ denotes whether bidder i wins multiple small ($z = 0$) or large ($z = 1$) bid groups.

The objective is the sum of base prices of winning bid groups and the respective base price markups/discounts of the winning adjustments. Constraint (2) ensures that exactly one adjustment must be accepted for each PEA contained in a winning bid group. Constraints (3) and (4) limit the supply of early and late licenses for each PEA. Recall that early licenses can be used to serve the demand for late licenses. Constraint (5) and (6) guarantee that a bidder wins at most one large bid group and never small and large bid groups simultaneously. Constraint (7) ensures that at most one small bid group of a bidder may become winning per EA.

$$\max \sum_{i\in I_0}\sum_{g\in G_i}(x_i^g\,\omega_i^g)+\sum_{i\in I_0}\sum_{g\in G_i}\sum_{p\in P_i^g}\sum_{k\in K_i^{gp}}(y_i^{gpk}\,\mu_i^{gpk}) \tag{1}$$

$$\text{s.t.} \sum_{k\in K_i^{gp}} y_i^{gpk}=x_i^g \qquad \forall i\in I_0,\forall g\in G_i,\forall p\in P_i^g \tag{2}$$

$$\sum_{i\in I_0}\sum_{g\in G_i^{SE}\cup G_i^{TE}}\sum_{k\in K_i^{gp}}\Big(y_i^{gpk}\max\Big\{0,\big(\min\{e_p,b_i^{gp}\}-\max\{0,b_i^{gp}-k\}\big)\Big\}\Big)\le e_p \quad \forall p\in P^E \tag{3}$$

$$\sum_{i\in I_0}\sum_{g\in G_i}\sum_{k\in K_i^{gp}}(y_i^{gpk}k)\le e_p+\ell_p \qquad \forall p\in P \tag{4}$$

$$\sum_{g\in G_i^{TE}\cup G_i^{TL}} x_i^g\le z_i \qquad \forall i\in I_0 \tag{5}$$

$$\sum_{g\in G_i^{SE}\cup G_i^{SL}} x_i^g\le M(1-z_i) \qquad \forall i\in I_0 \tag{6}$$

$$\sum_{g\in G_{ia}^{SE}\cup G_{ia}^{SL}} x_i^g\le 1 \qquad \forall i\in I_0,\forall a\in A \tag{7}$$

$$x_i^g\in\{0,1\} \qquad \forall i\in I_0,\forall g\in G_i \tag{8}$$

$$y_i^{gpk}\in\{0,1\} \qquad \forall i\in I_0,\forall g\in G_i,\forall p\in P_i^g,\forall k\in K_i^{gp} \tag{9}$$

$$z_i\in\{0,1\} \qquad \forall i\in I_0 \tag{10}$$

Appendix B XOR Winner Determination Problem

We reuse the notation introduced in Appendix A for sets I_0, A, P, P^E and Parameters e_p, ℓ_p, ω_i^g. We further introduce set I which includes all bidders except for the auctioneer. Furthermore, let c_i^{gp} and d_i^{gp} denote the number of early and late licenses, respectively, demanded in XOR bid $g \in G_i$ for PEA $p \in P$. There exists only a single type of decision variable, namely $x_i^g \in \{0,1\}$ which indicates whether bidder i wins XOR bid $g \in G_i$.

The objective function sums up prices of winning XOR bids. Constraint (2) and (3) are the supply constraints for early and late licenses, respectively. Constraint (4) ensures that at most one XOR bid becomes winning per bidder.

$$\max \sum_{i\in I_0} \sum_{g\in G_i} (x_i^g \, \omega_i^g) \tag{1}$$

$$\text{s.t.} \sum_{i\in I_0} \sum_{g\in G_i} (x_i^g \, c_i^{gp}) \leq e_p \qquad \forall p \in P^E \tag{2}$$

$$\sum_{i\in I_0} \sum_{g\in G_i} (x_i^g \, (c_i^{gp} + d_i^{gp})) \leq e_p + \ell_p \qquad \forall p \in P \tag{3}$$

$$\sum_{g\in G_i} x_i^g \leq 1 \qquad \forall i \in I \tag{4}$$

$$x_i^g \in \{0,1\} \qquad \forall i \in I_0, \forall g \in G_i \tag{5}$$

References

Adomavicius, G., Curley, S. P., Gupta, A., and Sanyal, P. (2012). Effect of information feedback on bidder behavior in continuous combinatorial auctions. *Management Science*, 58(4):811–830.

Adomavicius, G. and Gupta, A. (2005). Toward comprehensive real-time bidder support in iterative combinatorial auctions. *Information Systems Research*, 16(2):169–185.

Banker, R. D. and Kauffman, R. J. (2004). 50th anniversary article: The evolution of research on information systems: A fiftieth-year survey of the literature in management science. *Management Science*, 50(3):281–298.

Bapna, R., Das, S., Garfinkel, R., and Stallaert, J. (2007). A market design for grid computing. *INFORMS Journal on Computing*, 20(1):100–111.

Bichler, M., Goeree, J., Mayer, S., and Shabalin, P. (2014). Spectrum auction design: Simple auctions for complex sales. *Telecommunications Policy*, 38(7):613–622.

Bichler, M. and Goeree, J. K. (2017). *Handbook of spectrum auction design*. Cambridge University Press.

Bichler, M., Gupta, A., and Ketter, W. (2010). Research commentary–designing smart markets. *Information Systems Research*, 21(4):688–699.

Bichler, M., Hao, Z., and Adomavicius, G. (2017). Coalition-based pricing in ascending combinatorial auctions. *Information Systems Research*, 28(1):159–179.

Bichler, M., Schneider, S., Guler, K., and Sayal, M. (2011). Compact bidding languages and supplier selection for markets with economies of scale and scope. *European Journal of Operational Research*, 214(1):67–77.

Bichler, M., Shabalin, P., and Pikovsky, A. (2009). A computational analysis of linear-price iterative combinatorial auctions. *Information Systems Research*, 20(1):33–59.

Day, R. and Milgrom, P. (2008). Core-selecting package auctions. *international Journal of game Theory*, 36(3-4):393–407.

Eilat, A. and Milgrom, P. (2011). The caf auction: Design proposal. Technical report.

Goetzendorff, A., Bichler, M., Shabalin, P., and Day, R. W. (2015). Compact bid languages and core pricing in large multi-item auctions. *Management Science*, 61(7):1684–1703.

Green, J. and Laffont, J.-J. (1977). Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica: Journal of the Econometric Society*, pages 427–438.

Guo, Z., Koehler, G., and Whinston, A. B. (2007). A market-based optimization algorithm for distributed systems. *Management Science*, 53(8):1345–1358.

Iftekhar, M. S. and Tisdell, J. (2012). Comparison of simultaneous and combinatorial auction designs in fisheries quota market. *Marine Policy*, 36(2):446–453.

Kim, S. W., Olivares, M., and Weintraub, G. Y. (2014). Measuring the performance of large-scale combinatorial auctions: A structural estimation approach. *Management Science*, 60(5):1180–1201.

Kroemer, C., Bichler, M., and Goetzendorff, A. (2017). About inconsistent bidding and its impact on efficiency in the combinatorial clock auction. *Handbook of Spectrum Auction Design*, page 338.

Lehmann, D., Müller, R., and Sandholm, T. (2006). The winner determination problem. *Combinatorial auctions*, pages 297–318.

Milgrom, P. (2009). Assignment messages and exchanges. *AEJ-Micro*, 1(2):95–113.

Milgrom, P. (2019). FUEL for 5G: Flexible use and efficient licensing. Technical report, Auctionomics, C-Band Alliance.

Nisan, N. and Segal, I. (2006). The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129(1):192–224.

Scheffel, T., Pikovsky, A., Bichler, M., and Guler, K. (2011). An experimental comparison of linear and non-linear price combinatorial auctions. *Information Systems Research*, 22(2):346–368.

Xia, M., Koehler, G. J., and Whinston, A. B. (2004). Pricing combinatorial auctions. *European Journal of Operational Research*, 154(1):251–270.

6 Discussion

A central question in mechanism design is which market properties admit Walrasian equilibria. The latter are so desirable to attain because they describe a very favorable market situation: the market is cleared, the allocation of goods is efficient and envy-free, and the resulting prices are anonymous and linear. As outlined in our first contribution of this dissertation (see Chapter 3), the conditions to be satisfied for Walrasian equilibria to exist are very restrictive. One possibility to enforce their existence is to define the bid language in a way so that bidders can only express strong substitutes valuations. The product-mix auction which has successfully been used for financial markets in the UK follows such an approach (Klemperer, 2010).

While it can be proven that the bid language of the product mix auction is capable of expressing all possible strong substitutes valuations (Baldwin and Klemperer, 2021), it still restricts bidders' preferences to a very specific class of valuation functions. In particular, strong substitutes valuations do not allow the expression of any kind of complementarities. However, in most practical applications, stating synergistic preferences is essential for guaranteeing a high allocative efficiency (Cramton, 2013). Therefore, market designers have focused on developing domain-specific bid languages that admit complementarities; examples include TV ad auctions (Goetzendorff et al., 2015), procurement auctions (Bichler et al., 2011), and electricity markets (Cramton, 2017). While the possibility to state synergies eases the preference elicitation, it makes the design of a proper allocation and pricing rule more challenging because competitive equilibria generally do not exist in these settings.

Two projects in this dissertation focus on the design and analysis of combinatorial bid languages for real-world applications. In Chapter 4, we consider a market-based approach to dynamic congestion pricing. Our research is based on a proposal for a road capacity wholesale market by Beheshtian et al. (2020). Instead of letting bidders bid on individual road segments, we develop a bid language that allows them to submit bids for road capacity between OD pairs. Bids can include multiple substitutable routes between the same OD pair and bidders can specify a different willingness-to-pay for each of them. As one cannot hope for competitive equilibrium prices in this setting, we draw on ideas from electricity markets, namely the IP pricing technique (see Section 2.6.4) to "approximate" equilibrium prices. As opposed to electricity markets, our tests show that the make-whole payments are negligible in our wholesale market. Utilizing traffic data from the MATSim Open Berlin scenario (Ziemke et al., 2019), our numerical experiments further indicate that realistic problem instances of our wholesale market can be solved within 15 minutes for a major city like Berlin.

A possible direction for future research is to simulate not only the auction market but also the resulting traffic on the road when drivers choose their trips with respect to the actual allocated set of road capacity. As the road capacity licenses in our wholesale market are issued for 60 minute time intervals, the market operator does not have any control over the point in time when drivers actually use the respective road segment within this time slot. Especially during rush hours, one cannot expect a uniform traffic distribution over the 60 minute time interval so that congestion may still occur at peak times. Of course, this could be addressed by simply selling less road capacity from the beginning or using smaller time intervals during rush hour. However, while the first approach would unnecessarily reduce the throughput, the second one would not only increase the overhead due to the additional auctions necessary but would also make bidding more complex for service providers as end consumers may now need licenses from multiple time slots to complete their trips. To decide upon a reasonable time slot length, more simulations of the impact of congestion pricing on the actual traffic are necessary.

In addition to that, policy makers have to overcome some considerable administrative and regulatory obstacles before a market-based approach to congestion pricing can be implemented successfully: Legislation must pass new laws to set the legal framework for road pricing policies (Berliner Senatsverwaltung für Umwelt, Verkehr und Klimaschutz, 2020). Security experts must devise a virtual network that allows for secure registration and tracking of vehicles. Policy makers must address the citizen's skepticism towards road pricing that often stems from privacy and equity concerns (Cramton et al., 2018). While recent research provides compelling arguments in favor of dynamic congestion pricing (Cramton et al., 2019a), the discussion about the societal consequences of such policies is still an ongoing process (Creutzig et al., 2020).

While from a theoretical perspective setting road prices dynamically is the most effective way to charge drivers for externalities they impose on others (Cheng et al., 2017; Cramton et al., 2019a), the technical burden for setting tolls in response to the actual traffic on the road is much higher than implementing static pricing schemes. All cities that adopted congestion pricing in the past are using simple cordon- or area-based static pricing schemes (Lehe, 2019). Back in 1975, when the city of Singapore launched its road pricing scheme, technology simply did not allow for more complex pricing policies. However, Singapore has recently started to equip cars with GPS transponders, opening up the possibility for dynamic distance-based congestion pricing (Cramton et al., 2019a; Lehe, 2019). Singapore's approach of gradually refining its pricing policy whenever technology allows and social consensus is given serves as a template for other cities, e.g., policy makers in Berlin described such an incremental procedure to congestion pricing in a recent study (Berliner Senatsverwaltung für Umwelt, Verkehr und Klimaschutz, 2020).

The third project of this dissertation studies an auction for selling electromagnetic spectrum in the United States (see Chapter 5). The market we consider for evaluating the FUEL bid language consists of 406 geographic areas in each of which 14 license blocks are up for sale. This leads to 15^{406} distinct packages, far too many for bidders to evaluate. The high communication and computation complexity which are typical for large-scale combinatorial auctions is mitigated by the novel FUEL bid language that allows bidders to express their valuations succinctly. With an extensive set of numerical experiments, we can identify the hierarchical structure of bids to be the main reason for the tractability of the FUEL allocation problem even in the presence of more than 1,000 bidders. In contrast to a standard XOR bid language, FUEL effectively prevents the missing bids problem, leading to a substantially higher efficiency of the auction.

Unfortunately, the US Federal Communications Commission decided against running a private auction in which satellite companies would have sold their C-band spectrum to telecommunication providers. While this led to FUEL not having been implemented in practice yet, its generic bid language design makes FUEL also attractive for auctions in other domains. FUEL is particularly suitable whenever there are economies of scale in a product category and economies of scope between them. Possible application fields are TV ads (Goetzendorff et al., 2015), procuring school meals in Chile (Olivares et al., 2012), or selling fishery access rights (Iftekhar and Tisdell, 2012).

More than that, it is an intriguing idea to use FUEL instead of a standard XOR bid language in the supplementary phase of combinatorial clock auctions. While different variants of combinatorial clock auctions have been used worldwide for selling spectrum (Ausubel and Baranov, 2017; Cramton, 2013), these designs are not without problems. One drawback of some variants is using the XOR bid language in the supplementary phase. Due to the high number of available packages in spectrum auctions, bidders are often restricted in the number of bids they are allowed to submit in the supplementary phase in order to keep the allocation problem tractable. A case in point is the 2014 Canadian spectrum auction that featured 98 licenses. Bidders were restricted to a maximum of 2,000 package bids in the supplementary phase, clearly too few to state their preferences accurately.[1] For the spectrum auction we considered in Chapter 5, FUEL proved suitable to effectively mitigate this missing bids problem while still allowing a high tractability of the market.

An open question for future research is how prices can be computed when applying the FUEL bid language. In the FUEL whitepaper (Milgrom, 2019), it is suggested to use Vickrey-closest core payments (see Section 2.6.3). While core pricing is the de-facto standard for combinatorial clock auctions (Ausubel and Baranov, 2017; Cramton, 2013), this pricing technique may lead to prices that are perceived as unfair by bidders. In the 2012 Swiss spectrum auction, two telecommunication providers had to pay vastly different amounts of money (482 million CHF versus 360 million CHF) for almost the same set of licenses, clearly violating the law of one price (Bichler, 2017; Levin and Skrzypacz, 2016).

Besides that, the computation of Vickrey-closest core payments is non-trivial and computationally challenging (Cramton, 2013). To compute VCG prices, the winner determination problem has to be solved once for each winning bidder. This computation is followed by a method proposed by Day and Raghavan (2007) which iteratively finds the most violated core constraint and adds it to the partial core representation until there is no more coalition that could gain from deviating. In large markets with hundreds or thousands of bidders (e.g., electricity markets), computing Vickrey-closest core payments may no longer be feasible.

While core stability is an arguably fair and desirable property, it can only be fulfilled in practical settings when compromising on the linearity and anonymity of prices. There-

[1] Licensing Framework for Mobile Broadband Services (MBS) – 700 MHz Band: `https://www.ic.gc.ca/eic/site/smt-gst.nsf/eng/sf10572.html`; published: 07.03.2013; accessed: 09.12.2021

fore, an interesting direction for future research on the FUEL bid language is to calculate prices using techniques currently used in electricity markets. Depending on the degree of non-convexities available in the market, considerable side-payments might be necessary to prevent bidders from incurring a loss. However, this can be seen as a tradeoff for fulfilling the law of one price.

7 Conclusion

While various allocation and pricing rules have been proposed for auction markets over the past decades, the research in this field is still evolving. This is partly due to the advances in technology which allow solving large allocation problems today that were considered infeasible a couple of years ago. In addition to that, more and more electronic markets emerge every day which replace central market operators that allocated resources (often inefficiently) in the past. These new markets frequently require other rules than existing ones as allocation and pricing rules that were applied successfully in one market may not be suitable for another. For example, mitigating the communication and computation complexity drives the design decisions of large-scale auction markets but can be neglected in settings where only a single object is up for sale as it is commonly the case in auction houses like Sotheby's or Christie's.

Auction markets differ in both their size and design desiderata. While the law of one price is often sacrificed in spectrum sales to guarantee core stability (Cramton, 2013; Day and Cramton, 2012; Levin and Skrzypacz, 2016), electricity markets produce anonymous prices but can only fulfill individual rationality through substantial side-payments and satisfy market stability by threatening bidders with high penalties in case they deviate from the allocation (O'Neill et al., 2020; Hytowitz et al., 2020). In other markets such as financial markets in the UK for providing liquidity to banks, the bid language only admits the expression of strong substitutes valuations (Klemperer, 2010). While this guarantees for strong theoretic properties such as the existence of Walrasian equilibria, it is not suitable for most other practical applications where expressing complementary valuations is crucial for achieving high allocative efficiency.

The focus of this dissertation is the design of large-scale auction markets. At the core of this research is the design and analysis of combinatorial bid languages. While the bid languages considered in this dissertation have not been applied in the field yet, their generic design and our numerical tests provide valuable insights for both the research community and market designers in the field.

References

Arrow, K.J., Debreu, G., 1954. Existence of an equilibrium for a competitive economy. Econometrica: Journal of the Econometric Society , 265–290.

Ausubel, L.M., 2005. Walrasian tatonnement for discrete goods. Unpublished.

Ausubel, L.M., 2006. An efficient dynamic auction for heterogeneous commodities. American Economic Review 96, 602–629.

Ausubel, L.M., Baranov, O., 2017. A practical guide to the combinatorial clock auction. The Economic Journal 127, F334–F350. doi:10.1111/ecoj.12404.

Ausubel, L.M., Baranov, O.V., 2014. Market design and the evolution of the combinatorial clock auction. American Economic Review 104, 446–51.

Ausubel, L.M., Cramton, P., Milgrom, P., 2006a. The clock-proxy auction: A practical combinatorial auction design. Handbook of Spectrum Auction Design , 120–140.

Ausubel, L.M., Cramton, P., Pycia, M., Rostek, M., Weretka, M., 2014. Demand reduction and inefficiency in multi-unit auctions. The Review of Economic Studies 81, 1366–1400.

Ausubel, L.M., Milgrom, P., et al., 2006b. The lovely but lonely Vickrey auction. Combinatorial Auctions 17, 22–26.

Ausubel, L.M., Milgrom, P.R., 2002. Ascending auctions with package bidding. The B.E. Journal of Theoretical Economics 1, 20011001. URL: https://doi.org/10.2202/1534-5963.1019, doi:doi:10.2202/1534-5963.1019.

Baldwin, E., Klemperer, P., 2019. Understanding preferences: Demand types, and the existence of equilibrium with indivisibilities. Econometrica: Journal of the Econometric Society 87, 867–932.

Baldwin, E., Klemperer, P., 2021. Proof that the strong substitutes product-mix auction bidding language can represent any strong substitutes preferences. Working Paper. Oxford University. URL: http://elizabeth-baldwin.me.uk/papers/strongsubsproof.pdf.

Beheshtian, A., Geddes, R.R., Rouhani, O.M., Kockelman, K.M., Ockenfels, A., Cramton, P., Do, W., 2020. Bringing the efficiency of electricity market mechanisms to multimodal mobility across congested transportation systems. Transportation Research Part A: Policy and Practice 131, 58–69.

Berliner Senatsverwaltung für Umwelt, Verkehr und Klimaschutz, 2020. Ergänzende Instrumente zur Finanzierung des Berliner ÖPNV: Machbarkeitsstudie. Technical Report. Berliner Senatsverwaltung für Umwelt, Verkehr und Klimaschutz.

Bichler, M., 2017. Market design: A linear programming approach to auctions and matching. Cambridge University Press.

Bichler, M., Davenport, A., Hohner, G., Kalagnanam, J., 2006. Industrial procurement auctions. Combinatorial Auctions 1, 593–612.

Bichler, M., Fichtl, M., Schwarz, G., 2021. Walrasian equilibria from an optimization perspective: A guide to the literature. Naval Research Logistics (NRL) 68, 496–513.

Bichler, M., Goeree, J., Mayer, S., Shabalin, P., 2014. Spectrum auction design: Simple auctions for complex sales. Telecommunications Policy 38, 613–622.

Bichler, M., Goeree, J.K., 2017. Handbook of spectrum auction design. Cambridge University Press.

Bichler, M., Milgrom, P., Schwarz, G., 2020. Taming the communication and computation complexity of combinatorial auctions: The FUEL bid language, in: Conference on Information Systems and Technology (CIST2020).

Bichler, M., Schneider, S., Guler, K., Sayal, M., 2011. Compact bidding languages and supplier selection for markets with economies of scale and scope. European Journal of Operational Research 214, 67–77.

Bichler, M., Shabalin, P., Wolf, J., 2013a. Do core-selecting combinatorial clock auctions always lead to high efficiency? An experimental analysis of spectrum auction designs. Experimental Economics 16, 511–545.

Bichler, M., Shabalin, P., Ziegler, G., 2013b. Efficiency with linear prices? A game-theoretical and computational analysis of the combinatorial clock auction. Information Systems Research 24, 394–417.

Bikhchandani, S., Mamer, J.W., 1997. Competitive equilibrium in an exchange economy with indivisibilities. Journal of Economic Theory 74, 385–413.

Bikhchandani, S., Ostroy, J.M., 2002. The package assignment model. Journal of Economic Theory 107, 377–406.

Blumrosen, L., Nisan, N., 2007. Combinatorial auctions, in: Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.V. (Eds.), Algorithmic Game Theory. Cambridge University Press. chapter 11, pp. 267–300.

Brunner, C., Goeree, J.K., Holt, C.A., Ledyard, J.O., 2010. An experimental test of flexible combinatorial spectrum auction formats. American Economic Journal: Microeconomics 2, 39–57.

Cantillon, E., Pesendorfer, M., et al., 2006. Auctioning bus routes: The London experience. Combinatorial Auctions 22, 573–592.

Caplice, C., Sheffi, Y., et al., 2006. Combinatorial auctions for truckload transportation. Combinatorial Auctions 21, 539–571.

Cheng, Q., Liu, Z., Liu, F., Jia, R., 2017. Urban dynamic congestion pricing: An overview and emerging research needs. International Journal of Urban Sciences 21, 3–18.

Clarke, E.H., 1971. Multipart pricing of public goods. Public Choice , 17–33.

Committee for the Prize in Economic Sciences in Memory of Alfred Nobel, 2020. Improvements to auction theory and inventions of new auction formats. Technical Report. The Royal Swedish Academy of Sciences.

Cottingham, D.N., Beresford, A.R., Harle, R.K., 2007. Survey of technologies for the implementation of national-scale road user charging. Transport Reviews 27, 499–523.

Cramton, P., 1997. The FCC spectrum auctions: An early assessment. Journal of Economics & Management Strategy 6, 431–495.

Cramton, P., 2013. Spectrum auction design. Review of Industrial Organization 42, 161–190.

Cramton, P., 2017. Electricity market design. Oxford Review of Economic Policy 33, 589–612.

Cramton, P., Geddes, R.R., Ockenfels, A., 2018. Set road charges in real time to ease traffic. Nature 560, 23–25.

Cramton, P., Geddes, R.R., Ockenfels, A., 2019a. Markets for road use: Eliminating congestion through scheduling, routing, and real-time road pricing. Working Paper.

Cramton, P., Geddes, R.R., Ockenfels, A., 2019b. Using technology to eliminate traffic congestion. Journal of Institutional and Theoretical Economics (JITE) 175, 126–139.

Cramton, P., Schwartz, J.A., 2000. Collusive bidding: Lessons from the FCC spectrum auctions. Journal of Regulatory Economics 17, 229–252. doi:10.1023/A:1008174031940.

Creutzig, F., Javaid, A., Koch, N., Knopf, B., Mattioli, G., Edenhofer, O., 2020. Adjust urban and rural road pricing for fair mobility. Nature Climate Change 10, 591–594.

d'Aspremont, C., Gérard-Varet, L.A., 1979. Incentives and incomplete information. Journal of Public Economics 11, 25–45.

Day, R.W., Cramton, P., 2012. Quadratic core-selecting payment rules for combinatorial auctions. Operations Research 60, 588–603.

Day, R.W., Milgrom, P., 2008. Core-selecting package auctions. International Journal of Game Theory 36, 393–407.

Day, R.W., Raghavan, S., 2007. Fair payments for efficient allocations in public sector combinatorial auctions. Management Science 53, 1389–1406.

Do Chung, B., Yao, T., Friesz, T.L., Liu, H., 2012. Dynamic congestion pricing with demand uncertainty: A robust optimization approach. Transportation Research Part B: Methodological 46, 1504–1518.

Dong, J., Mahmassani, H.S., Erdoğan, S., Lu, C.C., 2011. State-dependent pricing for real-time freeway management: Anticipatory versus reactive strategies. Transportation Research Part C: Emerging Technologies 19, 644–657.

Duranton, G., Turner, M.A., 2011. The fundamental law of road congestion: Evidence from us cities. American Economic Review 101, 2616–52.

Economic Sciences Prize Committee of the Royal Swedish Academy of Sciences, 2012. Stable allocations and the practice of market design. Technical Report. The Royal Swedish Academy of Sciences.

Falck, O., Fichtl, A., Janko, A., Kluth, T., Wölfl, A., 2020. Verkehrliche Wirkungen einer Anti-Stau-Gebühr in München. Technical Report 115. ifo Forschungsberichte.

Federal Communications Commission, 2020. Auction 107 clock phase technical guide. Technical Report. Federal Communications Commission. URL: `https://www.fcc.gov/file/19165/download`.

Fujishige, S., Yang, Z., 2003. A note on Kelso and Crawford's gross substitutes condition. Mathematics of Operations Research 28, 463–469.

Fujishima, Y., Leyton-Brown, K., Shoham, Y., 1999. Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches, in: IJCAI, pp. 548–553.

Goetzendorff, A., Bichler, M., Shabalin, P., Day, R.W., 2015. Compact bid languages and core pricing in large multi-item auctions. Management Science 61, 1684–1703.

Green, J., Laffont, J.J., 1979. Incentives in public decision-making. Elsevier North-Holland.

Gribik, P.R., Hogan, W.W., Pope, S.L., 2007. Market-clearing electricity prices and energy uplift. Cambridge, MA , 1–46.

Grimm, V., Riedel, F., Wolfstetter, E., 2003. Low price equilibrium in multi-unit auctions: The GSM spectrum auction in Germany. International journal of Industrial Organization 21, 1557–1569.

Groves, T., 1973. Incentives in teams. Econometrica: Journal of the Econometric Society, 617–631 .

Gul, F., Stacchetti, E., 1999. Walrasian equilibrium with gross substitutes. Journal of Economic Theory 87, 95–124.

Gul, F., Stacchetti, E., 2000. The English auction with differentiated commodities. Journal of Economic Theory 92, 66–95.

Holmström, B., 1979. Groves' scheme on restricted domains. Econometrica: Journal of the Econometric Society , 1137–1144.

Hytowitz, R.B., Frew, B., Stephen, G., Ela, E., Singhal, N., Bloom, A., Lau, J., 2020. Impacts of price formation efforts considering high renewable penetration levels and system resource adequacy targets. Technical Report. National Renewable Energy Lab.(NREL).

Iftekhar, M.S., Tisdell, J., 2012. Comparison of simultaneous and combinatorial auction designs in fisheries quota market. Marine Policy 36, 446–453.

Janssen, M., Karamychev, V.A., 2013. Gaming in combinatorial clock auctions. Technical Report. Tinbergen Institute.

Karmarkar, N., 1984. A new polynomial-time algorithm for linear programming, in: Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing, pp. 302–311.

Karp, R.M., 1972. Reducibility among combinatorial problems, in: Complexity of computer computations. Springer, pp. 85–103.

Kelso, A.S., Crawford, V.P., 1982. Job matching, coalition formation, and gross substitutes. Econometrica: Journal of the Econometric Society , 1483–1504.

Klemperer, P., 2010. The product-mix auction: A new auction design for differentiated goods. Journal of the European Economic Association 8, 526–536.

Kokott, G.M., Bichler, M., Paulsen, P., 2019. The beauty of Dutch: Ex-post split-award auctions in procurement markets with diseconomies of scale. European Journal of Operational Research 278, 202–210.

Kroemer, C., Bichler, M., Goetzendorff, A., 2016. (Un)expected bidder behavior in spectrum auctions: About inconsistent bidding and its impact on efficiency in the combinatorial clock auction. Group Decision and Negotiation 25, 31–63.

Lehe, L., 2019. Downtown congestion pricing in practice. Transportation Research Part C: Emerging Technologies 100, 200–223.

Lehmann, D., Müller, R., Sandholm, T., 2006. The winner determination problem. Combinatorial Auctions , 297–318.

Leme, R.P., 2017. Gross substitutability: An algorithmic survey. Games and Economic Behavior 106, 294–316.

Levin, J., Skrzypacz, A., 2016. Properties of the combinatorial clock auction. American Economic Review 106, 2528–51.

Levy, J.I., Buonocore, J.J., Von Stackelberg, K., 2010. Evaluation of the public health impacts of traffic congestion: A health risk assessment. Environmental Health 9, 1–12.

Liberopoulos, G., Andrianesis, P., 2016. Critical review of pricing schemes in markets with non-convex costs. Operations Research 64, 17–31.

McAfee, R.P., 1992. A dominant strategy double auction. Journal of Economic Theory 56, 434–450.

McKenzie, L., 1954. On equilibrium in Graham's model of world trade and other competitive systems. Econometrica: Journal of the Econometric Society , 147–161.

Milgrom, P., 2000. Putting auction theory to work: The simultaneous ascending auction. Journal of Political Economy 108, 245–272.

Milgrom, P., 2019. FUEL for 5G: Flexible use and efficient licensing. Technical Report. Auctionomics, C-Band Alliance. URL: `https://milgrom.people.stanford.edu/sites/g/files/sbiybj4391/f/fuel_for_5g_white_paper_june_6_2019.pdf`.

Milgrom, P., 2021. Auction research evolving: Theorems and market designs. American Economic Review 111, 1383–1405.

Milgrom, P., Strulovici, B., 2009. Substitute goods, auctions, and equilibrium. Journal of Economic Theory 144, 212–247.

Morgan, J., Steiglitz, K., Reis, G., 2003. The spite motive and equilibrium behavior in auctions. Contributions in Economic Analysis & Policy 2, 1–25.

Murota, K., Shioura, A., 1999. M-convex function on generalized polymatroid. Mathematics of Operations Research 24, 95–105.

Myerson, R.B., Satterthwaite, M.A., 1983. Efficient mechanisms for bilateral trading. Journal of Economic Theory 29, 265–281.

Nisan, N., 2000. Bidding and allocation in combinatorial auctions, in: Proceedings of the 2nd ACM Conference on Electronic Commerce, pp. 1–12.

Nisan, N., 2006. Bidding languages, in: Cramton, P., Shoham, Y., Steinberg, R. (Eds.), Combinatorial auctions. MIT Press. chapter 9, pp. 400–420.

Nisan, N., Segal, I., 2006. The communication requirements of efficient allocations and supporting prices. Journal of Economic Theory 129, 192–224.

Olivares, M., Weintraub, G.Y., Epstein, R., Yung, D., 2012. Combinatorial auctions for procurement: An empirical study of the Chilean school meals auction. Management Science 58, 1458–1481.

O'Neill, R., Hytowitz, R.B., Whitman, P., Mead, D., Dautel, T., Chen, Y., Eldridge, B., Siskind, A., Kheloussi, D., Kolkmann, D., Smith, A., Castillo, A., Mays, J., 2020. Essays on average incremental cost pricing for independent system operators. Unpublished.

O'Neill, R.P., Sotkiewicz, P.M., Hobbs, B.F., Rothkopf, M.H., Stewart Jr, W.R., 2005. Efficient market-clearing prices in markets with nonconvexities. European Journal of Operational Research 164, 269–285.

de Palma, A., Lindsey, R., 2011. Traffic congestion pricing methodologies and technologies. Transportation Research Part C: Emerging Technologies 19, 1377–1399.

Parkes, D., 2006. Iterative combinatorial auctions, in: Cramton, P., Shoham, Y., Steinberg, R. (Eds.), Combinatorial Auctions. MIT Press. chapter 2, pp. 41–78.

Parkes, D.C., Kalagnanam, J.R., Eso, M., 2001. Achieving budget-balance with Vickrey-based payment schemes in exchanges, in: Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence, pp. 1161–1168.

Pigou, A.C., 1920. The economics of welfare. Macmillan and Company.

Porter, D., Rassenti, S., Roopnarine, A., Smith, V., 2003. Combinatorial auction design. Proceedings of the National Academy of Sciences 100, 11153–11157.

Porter, D., Smith, V., 2006. FCC license auction design: A 12-year experiment. Journal of Law, Economics & Policy 3, 63.

Rassenti, S.J., Smith, V.L., Bulfin, R.L., 1982. A combinatorial auction mechanism for airport time slot allocation. The Bell Journal of Economics , 402–417.

Roth, A.E., Peranson, E., 1999. The redesign of the matching market for American physicians: Some engineering aspects of economic design. American Economic Review 89, 748–780.

Roth, A.E., Sönmez, T., Ünver, M.U., 2004. Kidney exchange. The Quarterly Journal of Economics 119, 457–488.

Rothkopf, M.H., 2007. Thirteen reasons why the Vickrey-Clarke-Groves process is not practical. Operations Research 55, 191–197.

Rothkopf, M.H., Pekeč, A., Harstad, R.M., 1998. Computationally manageable combinational auctions. Management Science 44, 1131–1147.

Sakurai, Y., Yokoo, M., Matsubara, S., 1999. A limitation of the generalized Vickrey auction in electronic commerce: Robustness against false-name bids., in: AAAI/IAAI, pp. 86–92.

Sandholm, T., 2002. Algorithm for optimal winner determination in combinatorial auctions. Artificial Intelligence 135, 1–54.

Schaller, B., 2017. Unsustainable? The growth of app-based ride services and traffic, travel and the future of New York City. Schaller Consulting .

Schiro, D.A., Zheng, T., Zhao, F., Litvinov, E., 2015. Convex hull pricing in electricity markets: Formulation, analysis, and implementation challenges. IEEE Transactions on Power Systems 31, 4068–4075.

Schubert, M., Kluth, T., Nebauer, G., Ratzenberger, R., Kotzagiorgis, S., Butz, B., Schneider, W., Leible, M., 2014. Verkehrsverflechtungsprognose 2030 Los 3: Erstellung der Prognose der deutschlandweiten Verkehrsverflechtungen unter Berücksichtigung des Luftverkehrs. Technical Report. Bundesministerium für Verkehr und digitale Infrastruktur.

Schwarz, G., Bichler, M., 2022. Designing a large-scale wholesale market for urban congestion pricing, in: 17th International Conference on Wirtschaftsinformatik (WI22).

Shioura, A., Tamura, A., 2015. Gross substitutes condition and discrete concavity for multi-unit valuations: A survey. Journal of the Operations Research Society of Japan 58, 61–103.

Shioura, A., Yang, Z., 2015. Equilibrium, auction, and generalized gross substitutes and complements. Journal of the Operations Research Society of Japan 58, 410–435.

Simoni, M.D., Kockelman, K.M., Gurumurthy, K.M., Bischoff, J., 2019. Congestion pricing in a world of self-driving vehicles: An analysis of different strategies in alternative future scenarios. Transportation Research Part C: Emerging Technologies 98, 167–185.

Sun, N., Yang, Z., 2006. Equilibria and indivisibilities: Gross substitutes and complements. Econometrica: Journal of the Econometric Society 74, 1385–1402.

Vickrey, W.S., 1955. A proposal for revising New York's subway fare structure. Journal of the Operations Research Society of America 3, 38–68.

Vickrey, W.S., 1959. Statement on the pricing of urban street use. Hearing of US Joint Committee on Metropolitan Washington Problems 11, 42–65.

Vickrey, W.S., 1961. Counterspeculation, auctions, and competitive sealed tenders. The Journal of Finance 16, 8–37.

Vosough, S., Poorzahedy, H., Lindsey, R., 2020. Predictive cordon pricing to reduce air pollution. Transportation Research Part D: Transport and Environment 88, 102564.

Vulkan, N., Roth, A.E., Neeman, Z., 2013. Introduction, in: Vulkan, N., Roth, A.E., Neeman, Z. (Eds.), The Handbook of Market Design. Oxford University Press, pp. 1–6.

Wie, B.W., Tobin, R.L., 1998. Dynamic congestion pricing models for general traffic networks. Transportation Research Part B: Methodological 32, 313–327.

Yokoo, M., Sakurai, Y., Matsubara, S., 2004. The effect of false-name bids in combinatorial auctions: New fraud in internet auctions. Games and Economic Behavior 46, 174–188.

Ziemke, D., Kaddoura, I., Nagel, K., 2019. The MATSim Open Berlin Scenario: A multimodal agent-based transport simulation scenario based on synthetic demand modeling and open data. Procedia Computer Science 151, 870–877.