University of Zurich[UZH]

Department of Informatics

Machine Learning-powered Iterative Combinatorial Auctions (Ph.D. Research Synopsis)

Gianluca Brero

Faculty of Business, Economics and Informatics

University of Zurich

Abstract

Over the last forty years, the continued increase in computational power has encouraged the spread of markets that use optimization software to find possible trades among participants. Combinatorial auctions (CAs) are certainly one of the key examples of these markets and have already found several applications in the real world, ranging from industrial procurement to spectrum sales. In a CA, a seller puts multiple indivisible items up for sale among several buyers who express their preferences via bids on arbitrary bundles of items. One of the major challenges when conducting CAs in practice is that the size of the bundle space grows exponentially in the number of items, which may allow bidders to only bid on a small fraction of their bundles of interest. Thus, CAs are often deployed in iterative forms, integrating price-based preference elicitation algorithms that provide bidding guidance. However, the currently adopted solutions are still unsatisfactory, leading to suboptimal bidding behavior and market inefficiencies.

In this thesis, I introduce iterative CAs that use preference elicitation algorithms based on statistical machine learning (ML) algorithms. Statistical ML algorithms allow the auctioneer to exploit prior data about bidders' valuations to reduce the costs of elicitation. My first research question focuses on how to use statistical ML algorithms to decrease the number of rounds in price-based auction designs. I introduce a Bayesian iterative CA design that alternates between refining beliefs over bidders' valuations and quoting the prices that are most likely to clear the auction under these beliefs. This Bayesian auction performs remarkably well against standard baselines under reasonable priors over bidders' valuations. Recent experimental studies have highlighted that, in price-based iterative CAs, bidders focus their bidding on a small set of bundles selected before the auction; this suboptimal bidding was one of the leading causes of inefficiencies in these experiments. My second research question focuses on how to use statistical ML algorithms to select which bundle bids bidders should submit to let the auction implement efficient allocations. I develop an elicitation algorithm based on *value queries* that, under reasonable prior beliefs over bidders' valuations, reveals highly efficient allocations even when only a small fraction of bundle bids are submitted. My third research question focuses on how to design value query-based auctions that integrate this elicitation algorithm. I introduce the Pseudo-VCG Machine Learning-based (PVML) auction, which employs my elicitation algorithm to determine the final allocation and charge payments that motivate truthful value reports. I compare PVML with the commonly used Combinatorial Clock Auction (CCA) design and show that PVML is competitive with the CCA in terms of allocative efficiency.

Acknowledgements

I think the word "advisor" is quite reductive when it comes to Sven Seuken. Over the last five years, Sven hasn't just advised me, but he has constantly encouraged me to wish for more, work more, become more. The word "leader" should probably be more appropriate. But leaders walk in front of you. Sven has always walked next to me, questioning my ideas, stimulating new insights, making me an independent researcher. Should I then call him "friend"? This would mean that this thesis is meant to obtain a Ph.D. degree from a friend—too dangerous, given my spaghetti accent! I'll just follow the suggestion he gave me in one of our first email exchanges: "just call me Sven". Thank you, Sven.

Most of my countless research discussions with Sven have been stimulated by an outstanding collaborator: Ben Lubin. Ben's sharpness and positive attitude are powerful engines constantly generating promising ideas to explore. Ben's humbleness completes the picture of a fantastic person.

There is a famous allegory where a dwarf, standing upon the shoulders of a giant, sees further than the others. During my Ph.D., I have identified a small group of giants upon which I was standing while doing my research. I am truly honored that one of these giants has agreed to serve as an external reviewer for this thesis. Thank you, David Parkes, for your extremely valuable feedback on my Ph.D. proposal and for having been a constant inspiration for my research.

During my Ph.D., I had the invaluable opportunity to work with Sébastien Lahaie. Sébastien has been an outstanding mentor both on research and personal level. His ability to spot subtle connections between diverse research fields is one of the aspects of research that I enjoy the most. I always had great fun while doing research with him, whether at Microsoft Research or while hiking at Dagstuhl.

Other mentors I want to thank include Giacomo Como and Paolo Tilli. Giacomo Como has guided my initial steps into research. His sharpness was fundamental for the success of my first research project, which has generated in me a positive attitude towards doing research. Paolo Tilli's beautiful mind constantly inspires my way of thinking.

The research of this thesis has been shaped and refined by countless discussions with incredibly smart people I met around the world. Among them, the current and former members of my fantastic research group: Vitor Bosshard, Ludwig Dierks, Stefania Ionescu, Timo Mennle, Dmitry Moor, Nils Olberg, Steffen Schuldenzucker, Mike Shann, and Jakob Weissteiner. Juan Camilo Castillo and Rupert Freeman, who gave me constructive feedback during our internship at Microsoft Research. Justus Winkelmann, who I was fortunate to meet in Jerusalem. Andrea Muolo, with whom I had a countless number of beers. Fabio Isler, who provided outstanding support on SATS. Manuel Beyeler, Nicolas Kuechler, and Andreas Perschak, the wonderful students I was fortunate to supervise.

After various attempts, I have just realized that there is no chance for me to write anything satisfactory on a more personal level and submit this thesis on time. Therefore, I limit myself to acknowledging my mother, Isabella, who has always encouraged my critical thinking, making sure I lived the life I wanted. My sister, Francesca, whose critical thinking has turned her into a hippie. My younger sister, Alessandra, whose critical thinking is turning her into a spiritual leader. And my father, Giovanni, the main victim of all this critical thinking. I finally thank Silvia, who left her Ithaca (and her Ippo) to go on this long journey with me. Sei il mio porto sicuro.

Ai miei nonni, Mariella, Teresina, Antonio e Francesco.

Contents

1 Introduction[1]

Over the last forty years, the continued increase in computational power has encouraged the spread of markets that use optimization software to find possible trades among participants. Combinatorial auctions have certainly emerged as a reference mechanism for these markets, finding applications in logistics markets (Caplice, 2007), industrial procurement (Sandholm, 2013), energy exchanges (Meeus et al., 2009), spectrum sales (Cramton, 2013), and markets for wind rights (Ausubel et al., 2011).

In a combinatorial auction (CA), a seller puts multiple indivisible items up for sale among several buyers who place bids on arbitrary bundles of items. By placing these *bundle bids*, a buyer can express complex preferences where items are complements, substitutes, or both. This prevents the *exposure problem* of simultaneous auctions for multiple heterogeneous items, where a bidder is exposed to paying more than her value for her final allocation because either she wins too few complementary items to realize their synergies or too many substitute items.

CAs are employed in increasingly large markets. Recently, in the 2014 Canadian spectrum auction, 98 licenses were sold simultaneously, providing bidders with an enormous range of bidding options (Industry Canada, 2013). One of the major challenges when conducting CAs in large markets is that bidders may only be able to bid on a small fraction of their bundles of interest. Furthermore, to properly address exposure problems, the auctioneer needs to treat bundle bids as "all-or-nothing" bids, which, in practice, do not assign any value for those bundles that did not receive a bid. The resulting "missing bids problem" has motivated the design of *iterative CAs*. Iterative CAs integrate preference elicitation algorithms to support bidders in selecting their bundle bids so that an efficient allocation is implemented. However, despite a great deal of effort dedicated to designing iterative CAs, the solutions adopted in the real world are still unsatisfactory, causing broader debates on how to best design large markets with many heterogeneous items (Goetzendorf et al., 2015). This debate was recently fostered by laboratory experiments performed by Bichler et al. (2013), who showed that

[1]This chapter freely borrows from my own prior work (Brero et al., 2017, Brero and Lahaie, 2018, Brero et al., 2018, 2019a,b).

because of the "missing bids problem" iterative CAs may even result in lower allocative efficiency than simultaneous single-item auctions. The consequences of this debate can be seen in spectrum sales: while many countries have already adopted CAs to sell their spectrum (see, e.g., Ausubel and Baranov, 2017), other countries still prefer standard non-combinatorial designs (e.g., Cramton and Ockenfels, 2017).

In this thesis, I introduce the new concept of *machine learning-powered iterative CAs*, which use preference elicitation algorithms based on statistical machine learning (ML) algorithms. These ML algorithms allow my iterative CAs to integrate prior data about bidders' preferences and use them to reduce the costs of elicitation. Given that spectrum auctions are probably the most prominent domain of application for CAs, I will use them as the leading example. In these auctions, prior data about bidders' preferences is often available, and it is commonly used to predict revenues of forthcoming auctions (Campbell, 2018), assess investments (Cusick et al., 2012), or test new auction designs (Weiss et al., 2017). In the standard CA designs used in spectrum auctions, prior data is only used to determine some design features such as the initial clock prices or the rate of price increase. In contrast, machine learning-powered CAs allow for a more principled integration of this data, better exploiting the information provided. Despite being motivated by spectrum sales, the auction designs introduced can be directly used in any domain of application for CAs; most of their features can also be used to design mechanisms for more general combinatorial exchanges where market participants can be buyers, sellers, or both. My goal is to provide concrete design solutions that, when reasonable priors are available, overcome the tradeoff between preventing exposure problems and avoiding efficiency losses due to missing bids.

1.1 Research Goals and Methodology

1.1.1 Research Goals

The goal of this thesis is to design machine learning-powered iterative CAs that reduce the cost of elicitation when prior data is available. While designing such CAs, I adopt the perspective of a market designer who wants to maximize allocative efficiency without making any explicit structural assumption over bidders' valuations.

I first consider the iterative auction designs currently used in practice. These auctions mainly interact with bidders via *demand queries*, i.e., they quote ask prices and ask bidders to report their profit-maximizing bundles at these prices. The goal of these auctions is to quote prices at which the bidders' demand meets the seller's supply. These prices are commonly called *clearing prices*, and, when an auction reveals clearing prices, the auctioneer can allocate items efficiently by assigning each bidder the bundle she demands. Unfortunately, price-based CAs often need an impractical number of rounds to reveal clearing prices: we only need to consider that the 2014 Canadian spectrum auction took more than a hundred rounds only to determine reasonable prices at which no item was over-demanded (Industry Canada, 2014). This challenge motivates my first research question:

Question 1. How can we use statistical machine learning algorithms to update ask prices with the goal of finding clearing prices in the lowest possible number of rounds?

In recent laboratory experiments, Scheffel et al. (2011) tested different price-based auction designs. The results highlighted that the primary source of inefficiency in all these auctions was that bidders only considered a small set of bundles when responding to prices. Motivated by these experiments, I investigate how statistical machine learning algorithms can be used to design auctions that identify relevant bundles on the bidders' behalf. Specifically, these auction designs are based on *value queries*, where the mechanism quotes some bundles to bidders and asks them to report their values for these bundles. In general, it is not clear whether it is easier for bidders to report their exact values for some bundles or to determine their favorite bundles at some quoted prices. However, the main goal of the value query-based iterative auctions of this thesis is to show how statistical machine learning algorithms can be used to identify bundles on which bidders should focus. Thus, value queries should be interpreted as proxies for queries that focus bidders on some bundles selected by the auction mechanism. In Section 1.5, I comment on how our value query-based auction designs can be modified to allow for reports consisting of upper and lower bounds on values.

Note that value query-based auction designs cannot provide efficiency guarantees over allocations unless bidders report their entire valuations (or some assumptions over these valuations can be made). Thus, I design my auctions with the goal of maximizing the *empirical efficiency* of the final allocation, i.e., the social welfare realized in this allocation normalized by the social welfare realized in an efficient allocation. This leads to my second research question:

Question 2. How can we use statistical machine learning algorithms to design value query-based elicitation algorithms with the goal of maximizing empirical efficiency in an iterative CA?

Once the elicitation algorithms are defined, it is important to integrate them in auctions that charge payments so that bidders are motivated to report their true values. When elicitation algorithms reveal efficient allocations, it is possible to use these algorithms to derive *VCG payments*, which, combined with some specific design features, motivate bidders to behave truthfully (see, e.g., Mishra and Parkes, 2007). As value query-based elicitation algorithms cannot reveal efficient allocations, it is necessary to take more careful considerations when integrating them into auction designs. This motivates my last-and-final research question:

Question 3. How can we use the elicitation algorithms identified in Question 2 to design value query-based auctions that motivate bidders to report their values truthfully?

1.1.2 Methodology

To answer the three aforementioned research questions, I introduce different ML-powered iterative CA designs and analyze them in the framework of *mechanism design* (see, e.g., Jackson, 2014). As these designs are meant to be practical, I assess them via empirical evaluations on synthetic settings. These settings are generated using two test suites:

- The Combinatorial Auction Test Suite (CATS) introduced by Leyton-Brown et al. (2000), which generates auction settings by sampling them from distributions based on different domains of application. CATS has been widely used over the last twenty years to test CA designs. In the settings generated by CATS, bidders have specific interests for a small number of bundles.

- The Spectrum Auction Test Suite (SATS) recently introduced by Weiss et al. (2017), which generates auction settings by sampling them from distributions inspired by spectrum sales applications. In the settings created by SATS, bidders are interested in a large number of bundles. Among the generative distributions provided by Weiss et al. (2017), we have the Multi-Region Value Model, which is based on the 2014 Canadian spectrum auction data and provides large settings where 98 items are auctioned off to 10 bidders.

Synthetic settings allow simulating rational bidding behaviors, leading to more objective evaluations that are not affected by cognitive biases. They also provide practical advantages, giving us access "ground-truth" information, such as bidders' real values or efficient allocations. Furthermore, sampling from the generative distributions provided by CATS and SATS allows us to generate reasonable prior data that can be integrated in our auctions.

In our simulations, bidders answer demand and value queries truthfully. We motivate this behavior using game theoretical models: We first assume that each bidder has a *quasi-linear utility function* over the auction outcomes, which is given by the difference between her value for the bundle of items she obtains and the price she has to pay. We then design our auctions so that each bidder trying to maximize her quasi-linear utility function is generally motivated to report truthfully.

1.2 Related Work

1.2.1 Iterative Combinatorial Auctions

Iterative CA designs are often preferred to sealed-bid ones in practical applications. One of the main reasons for preferring iterative CAs is that they address the problem of costly elicitation by providing adaptive bidding guidance across rounds (Parkes, 2006).

Commonly, iterative auctions provide this guidance by quoting ask prices at each round with the goal of revealing clearing prices. We can classify price-based iterative CAs along several dimensions, which include the price structure they use and the way they update prices if demand and supply do not match. Price structures usually range from "compact" item prices, which only quote a price for each item, to "expressive" bundle prices, which quote a price for each bundle. Furthermore, prices can be anonymous, i.e., the same prices are quoted to each bidder, or personalized.

de Vries et al. (2007) interpreted price update algorithms for CAs as optimization algorithms that solve a linear programming formulation of the efficient allocation problem. This interpretation allowed them to define two main classes for these algorithms. In the first class, we have *sub-gradient algorithms*, which use simple price updates based on excess demand. In the second class, we have *primal-dual algorithms*: these algorithms are believed to have faster convergence properties than the sub-gradient ones (de Vries et al., 2007), but they require bidders to report *all* of their profit-maximizing bundles at the quoted prices. This requirement is often considered impractical, which is why primal-dual iterative designs are not commonly used in real-world auctions.

There is an important class of price-based iterative auction designs based on personalized bundle prices. This price structure is appealing because it always allows to reveal clearing prices (Bikhchandani and Ostroy, 2002) and, consequently, to design auctions with provable efficiency guarantees and incentive properties. Among these auctions, we have iBundle (Parkes, 1999), with sub-gradient price updates, and the primal-dual auction designed by de Vries et al. (2007).[2] Building on the efficiency guarantees provided by these auctions, Mishra and Parkes (2007) showed how to modify these designs to charge bidders VCG payments for the final allocation, thus motivating them to report their demand at each round truthfully.

Item prices are often preferred to personalized bundle prices in price-based auction designs. An important argument in favor of item prices is that they are more informative for bidders, simplifying the search for profit-maximizing bundles. For this reason, many auction designs adopt linear price structures (e.g., Porter et al., 2003, Kwasnica et al., 2005, Bichler et al., 2009). Recent laboratory experiments performed by Scheffel et al. (2011) have confirmed the practical advantages of these auctions, showing that they achieve similar allocative efficiency to iBundle with much fewer rounds. However, in contrast to auctions using personalized bundle prices, item prices-based auctions are not guaranteed to reveal clearing prices, and, consequently, do not provide the good incentive properties induced by VCG prices.

The Combinatorial Clock Auction (CCA) introduced by Ausubel et al. (2006) has emerged as the standard CA design for spectrum sales (Ausubel and Baranov, 2017). The CCA attempts to resolve the tension between item and bundle prices by combining an item price-based elicitation with a supplementary round that allows bidders to submit additional bundle bids. As an alternative to the CCA, Lahaie (2011) introduced a primal-dual iterative auction design that employs price structures that are halfway between item and bundle prices. Recently, Lahaie and Lubin (2019) built on this idea and designed a sub-gradient auction with an adaptive price structure that led to promising results.

Unlike prior work, the price-based auctions I introduce in this thesis update prices using Bayesian learning models. These models allow designing price-update algorithms that also take into account prior beliefs over bidders' valuations when identifying new ask prices to quote. Thus, these algorithms do not belong to the classification provided by de Vries et al. (2007). Furthermore—even though this investigation is not part of this thesis—the Bayesian framework for price updates provides principled guidance on which price structure to employ. Indeed, the auctioneer can use her beliefs to estimate the benefits (measured in terms of clearing probabilities) of employing richer price structures.

One of the critical issues of iterative CAs is that, when no useful assumptions on valuations are available, it is not possible to derive a practical elicitation algorithm that

[2] Another appealing feature of these auctions is that they also use *ascending-price trajectories*. Even though in this thesis I do not investigate how machine learning-based algorithms can be used to design ascending price-based auctions, ascending auctions have many desirable properties for practical applications (see Cramton, 1998, for an extensive discussion).

is guaranteed to reveal an (even approximate) efficient allocation (Nisan and Segal, 2006). Thus, a strand of research has focused on designing practical elicitation algorithms for CAs when valuations have useful structures. In the context of price-based iterative CAs, much interest has been devoted to designing auctions for valuations where items are *substitutes* (i.e., increasing the price on one item does not decrease the demand on any other item). The reason behind this interest is that, under these valuations, item-based clearing prices are guaranteed to exist. Among the designs of price-based auctions for substitute valuations, we have Kelso Jr and Crawford (1982), Gul and Stacchetti (2000), Ausubel (2004). In particular, Ausubel (2004) introduced the idea of using multiple price-based elicitation threads simultaneously, with bidders responding to multiple ask prices at each auction round. We draw on this idea to design our auctions.

There are several other elicitation approaches for specific classes of valuations that are not based on prices (see Sandholm and Boutilier, 2006, for an overview). Most relevant to this thesis are the two works by Lahaie and Parkes (2004) and Blum et al. (2004). These works introduced the idea of using learning algorithms for preference elicitation in CAs. However, their goal was to build tractable elicitation algorithms for certain classes of valuations. Despite this difference in intentions, Lahaie and Parkes (2004) introduced an elicitation algorithm where learning algorithms are modular. As these learning algorithms can also have a statistical nature, the elicitation algorithm of Lahaie and Parkes (2004) can be interpreted as a precursor of the one I introduce in Brero et al. (2017). Yet, the algorithm introduced by Lahaie and Parkes (2004) critically requires demand queries. The advantage of using these queries is that they allow formulating termination conditions based on efficiency guarantees. But, due to the result by Nisan and Segal (2006), if their approach was applied in a general setting, it may require communicating exponentially-sized prices to the bidders in every round.

1.2.2 Machine Learning-based Mechanism Design

Another strand of research related to this thesis has focused on developing statistical machine learning approaches for *automated mechanism design*. The automated mechanism design (AMD) research agenda (Conitzer and Sandholm, 2002, 2004) seeks to use optimization algorithms to design incentive-compatible economic mechanisms. This agenda has only considered *direct-revelation mechanisms*, where agents are first asked to report all of their preferences, and then an outcome is determined.

Duetting et al. (2015) were the first to apply machine learning algorithms in the context of automated mechanism design. Specifically, they used support vector machines to design payment rules that allow integrating manually-selected allocation rules in incentive-compatible mechanisms. Narasimhan et al. (2016) generalized this idea to mechanisms that are not allowed to charge payments and designed incentive compatible social choice and matching mechanisms. More recently, Duetting et al. (2018) used deep learning methods to advance the design of auctions that maximize the expected revenue when bidders' valuations are sampled from some underlying probability distribution. Similar deep learning methods were also applied by Golowich et al. (2018) to automatically design matching mechanisms.

A more theoretical strand of research has studied machine learning-based AMD from a different perspective, providing sample-complexity results for revenue-maximizing auctions (see Balcan et al. (2018) for an overview) and general direct revelation mechanisms that are not allowed to charge payments (Narasimhan et al., 2016).

Unlike the work on AMD, my work focuses on iterative combinatorial auctions, which do not belong to the class of direct revelation mechanisms. Compared to direct revelation mechanisms, iterative combinatorial auctions introduce several design challenges related to their repeated interactions with bidders. For this reason, instead of using machine learning to automatically design entire mechanisms (which are then defined by the set of learned parameters), we design our iterative auctions manually and integrate machine learning algorithms to drive the elicitation process. Note that, given their iterative nature, our auctions are intuitively more robust to the quality of prior data as they can refine their inference across rounds.

1.3 Publications Contained in this Thesis

This thesis consists of five papers that answer the three research questions presented in Section 1.1. In this section, I restate the research questions and provide the list of papers that address each research question.

Question 1. How can we use statistical machine learning algorithms to update ask prices with the goal of finding clearing prices in the lowest possible number of rounds?

Publications:

1. A Bayesian Clearing Mechanism for Combinatorial Auctions. Gianluca Brero and Sébastien Lahaie. In *Proceedings of the Thirty-second AAAI Conference of Artificial Intelligence (AAAI-18)*, New Orleans, USA, February 2018.

2. Fast Iterative Combinatorial Auctions via Bayesian Learning. Gianluca Brero, Sébastien Lahaie, and Sven Seuken. In *Proceedings of the Thirty-third AAAI Conference of Artificial Intelligence (AAAI-19)*, Honolulu, USA, January 2019.

Question 2. How can we use statistical machine learning algorithms to design value query-based elicitation algorithms with the goal of maximizing empirical efficiency in an iterative CA?

Publications:

3. Probably Approximately Efficient Combinatorial Auctions via Machine Learning. Gianluca Brero, Benjamin Lubin, and Sven Seuken. In *Proceedings of the Thirty-first AAAI Conference of Artificial Intelligence (AAAI-17)*, San Francisco, USA, February 2017.

Question 3. How can we use the elicitation algorithms identified in Question 2 to design value query-based auctions that motivate bidders to report their values truthfully?

Publications:

4. Combinatorial Auctions via Machine Learning-based Preference Elicitation. Gianluca Brero, Benjamin Lubin, and Sven Seuken. In *Proceedings of the Twenty-seventh International Joint Conference on Artificial Intelligence and the Twenty-third European Conference on Artificial Intelligence (IJCAI-ECAI-18)*, Stockholm, Sweden, July 2018.

5. Machine Learning-powered Iterative Combinatorial Auctions. Gianluca Brero, Benjamin Lubin, and Sven Seuken. Working paper, November 2019.

1.4 Summary of Contributions

In this section, I provide a summary of all five papers and explain how they answer the three research questions.

1.4.1 A Bayesian Clearing Mechanism for Combinatorial Auctions

This paper provides the first answer to my first research question by introducing a Bayesian price update algorithm for item prices in settings where bidders are *single-*

minded.[3] The Bayesian update algorithm allows the auctioneer to quote ask prices that are likely to clear the auction under her current beliefs on bidders' valuations. These beliefs reflect the prior beliefs of the auctioneer and the information revealed by the bidders during the auction. Intuitively, the more accurate the prior beliefs, the lower the number of rounds to determine clearing prices.

The Bayesian price update algorithm is based on a joint probability distribution over bidders' valuations and market prices. This distribution has two important properties: first, it captures the auctioneer beliefs on valuations, and second, it assigns to market prices a probability density proportional to their clearing probability under these beliefs. After each demand observation, the algorithm refines the beliefs over bidders' valuations and then computes candidate clearing prices as maximum a posteriori (MAP) estimates. We capture the auctioneer's beliefs with Gaussian distributions over bidders' values for their bundle of interest. Gaussian distributions allow us to update beliefs after each demand observation using a technique called *assumed density filtering*, which is a special case of expectation propagation (Cowell et al., 1996). Linear prices and Gaussian beliefs on values allow us to use an expectation-maximization (EM) algorithm (Dempster et al., 1977) to determine the new candidate clearing prices as MAP estimates.

We evaluate our auction design with two experiments: a small one, that illustrates the behavior of our auction under biased and unbiased prior information, and a large-scale experiment that compares our auction against a competitive baseline.

The small experiment is performed on an instance of the *Local-Local-Global* (LLG) setting (Ausubel and Milgrom, 2006), which has been considered several times in the CA literature. This experiment shows that, even when the auctioneer has biased estimates over bidders' values, integrating these estimates with reasonably high variance allows the auction to determine clearing prices within a few rounds.

The larger-scale experiments are performed on settings with 12 items and 10 bidders. The valuations in these settings are generated using four distributions provided by CATS, each based on a different domain of application for CAs (Leyton-Brown et al., 2000). We generate prior beliefs for our Bayesian auction by fitting multivariate Gaussian distributions over items to bundle-value pairs sampled from the generative distribution of the setting at hand. As a baseline, we implemented a standard item-price auction where prices are updated according to excess demand, and the magnitude of the price update is proportional to a step-size parameter. Our Bayesian auction design performs remarkably well compared to this baseline, even when the baseline uses the step-size

[3]Single-minded bidders are interested in a single bundle of items and have a specified value for obtaining this bundle or any of its supersets, valuing zero any other bundle.

parameter at which the number of rounds is at its minimum.

This paper provides the first answer on how to use statistical learning models to design price update algorithms. Furthermore, it shows that relying on reasonable prior beliefs for valuations can be much more effective than carefully tuning price increments. However, this first design has three important limitations: 1. it only supports item prices, 2. it can only be used in settings with single-minded bidders, and 3. it only allows the auctioneer to exploit Gaussian beliefs over bidders' values.

1.4.2 Fast Iterative Combinatorial Auctions via Bayesian Learning

This paper completes the answer to the first research question by introducing a Bayesian price update algorithm (with a corresponding auction design) that is not subject to any of the three limitations described above.

At the core of our new auction is a practical Monte Carlo expectation-maximization (EM) algorithm that operates on the joint probability distribution of valuations and market prices. This algorithm determines approximate modal prices interpreting valuations as latent variables.

To draw a better comparison with Brero and Lahaie (2018), we implement our Bayesian auction using again Gaussian beliefs and item prices. We run two experiments. The first one is based on the single-minded CATS settings used in Brero and Lahaie (2018). The goal of this experiment is to compare the performance of the new Bayesian auction and the old one. The results show us that the two designs are competitive, even though the new Bayesian design is not specially designed for single-minded settings.

In the second experiment, we again use CATS to generate auction settings, but we properly interpret the generated bidders as multi-minded. The goal of this experiment is to compare the performance of our new Bayesian auction with the non-Bayesian baselines introduced in Brero and Lahaie (2018). The results again show that our auction design outperforms the non-Bayesian baseline, even when the baseline uses the step-size parameter at which the number of rounds is at its minimum.

This paper completes the answer to the first research question. We can conclude that, also in general settings, relying on reasonable priors for valuations is much more effective in decreasing the number of rounds than carefully tuning price increments.

1.4.3 Probably Approximately Efficient Combinatorial Auctions via Machine Learning

This paper provides the first answer to our second research question by showing how one can use statistical machine learning algorithms to design value query-based elicitation algorithms that achieve high empirical efficiency. Note that, when the only information available to the auctioneer consists of a set of bundle-value pairs for each bidder, there is no principled way to identify the ‘‘next most-useful query’’ to ask. Indeed, this task requires reasoning about the missing information from bidders' valuations, which requires making general statements about valuations given what has been observed so far.

In this paper, we perform this generalization by associating each bidder with a machine learning (ML) algorithm that is trained on her reported bundle-value pairs and infers her entire valuation. When all inferred valuations are determined, the auction computes an *inferred efficient allocation*, i.e., a feasible allocation that maximizes the inferred overall value. The ‘‘next most-useful query’’ assigned to each bidder corresponds to the bundle she obtains under this allocation. If at least one bidder is assigned a bundle she has not yet evaluated, all the new queries are submitted to bidders, and new value reports are obtained. This elicitation step is then iterated by including the new value reports in the training sets of the ML algorithms. If each bidder already evaluated her assigned bundle, the elicitation stops.

We integrate this algorithm in a ‘‘preliminary’’ auction design that allocates items according to the last inferred efficient allocation identified during the elicitation. The payment rule of this auction is left unspecified. To practically implement this auction, one needs to use ML algorithms that allow the computation of inferred efficient allocations without enumerating and evaluating all possible allocations. At the same time, the ML algorithms need to capture all reported values, not to reduce the expressivity of the auction. In this paper, we use support vector regression algorithms (SVRs) (Smola and Schölkopf, 2004) as our ML algorithms. The expressivity of these algorithms depends on the *kernel function* they employ. We then consider a class of expressive kernel functions that allow the auctioneer to determine inferred efficient allocations via integer programs.

We evaluate this auction with two experiments run on spectrum auction settings with 18 items. Specifically, we consider settings sampled from the Global Synergy Value Model (GSVM), introduced by Goeree and Holt (2010), and from the more complex Local Synergy Value Model (LSVM), introduced by Scheffel et al. (2012).

Similarly to the price-based designs, prior data is sampled from the generative distribution of the setting at hand. Importantly, to keep our mechanism computationally practical, this data is not integrated in the training set of the ML algorithms but only used to tune the hyper-parameters of the kernel functions of the SVRs.[4] The initial training set of the SVRs is obtained by asking bidders to evaluate a certain number of bundles selected by the mechanism uniformly at random in the bundle space.

In the first experiment, we test the efficiency of our preliminary auction, assuming that bidders report their values truthfully. We compare this efficiency with a baseline that computes the optimal allocation based on values reported for bundles selected uniformly at random, without using machine learning algorithms to identify new queries. We see that the machine learning algorithms used in our elicitation significantly increase empirical efficiency.

Similarly to the first experiment run in Brero and Lahaie (2018), in the second experiment, we test how robust our elicitation algorithm is to prior data. Specifically, we use data sampled from the LSVM domain to tune the kernel hyper-parameters in the GSVM domain, and vice-versa. We see that the results are generally stable and that our machine learning-powered auction outperforms the non-machine learning baselines even when prior data is inaccurate.

This paper provides the first answer to the second research question identified in Section 1.1, showing us how statistical learning algorithms can be used to design practical elicitation algorithms based on value queries. However, we only compared these algorithms with relatively weak baselines, and it is not clear whether they could provide concrete alternatives to price-based auctions.

1.4.4 Combinatorial Auctions via Machine Learning-based Preference Elicitation

This paper provides the first answer to my third research question introducing an auction design that integrates the elicitation algorithm presented in Brero et al. (2017) and provides bidders with proper incentives to report their values truthfully. These incentives are provided by carefully using our elicitation algorithm to derive VCG-style payments. We then call our auction mechanism *Pseudo-VCG Mechanism* (PVM).

Unlike the allocation rule presented in Brero et al. (2017), PVM determines its allocation by interpreting all the bundle-value pairs reported during the auction using an XOR bidding language. This allows keeping the mechanism expressive even when it

[4]From a Bayesian perspective, this usage of prior data can be interpreted as if data was only used to fit the covariance function of a multivariate Gaussian prior and not the mean function (Williams and Rasmussen, 2006).

employs machine learning algorithms that do not accurately capture bidders' reported values.

PVM emulates the VCG mechanism by additionally running our value query-based elicitation algorithm in each setting obtained by excluding one bidder from the main setting. With these additional elicitations, we can determine the externality that each bidder is imposing on the other bidders and use it to compute VCG-style payments. Our resulting mechanism is *social-welfare aligned*: in social welfare-aligned mechanisms, any beneficial manipulation a bidder finds when the other bidders are truthful must increase the social welfare by the same amount. We argue that, when the mechanism allows bidders to "push" self-selected bundle-value pairs, this property generally motivates bidders to report their values truthfully.[5]

As PVM requires running an additional elicitation for each bidder, it may be impractical in settings with many bidders. This paper also introduces a modified version of PVM where bidders are partitioned into different groups, and the mechanism runs only one elicitation per group that excludes all its bidders. The allocation identified using the bundles elicited during this elicitation is used to compute the externality that each bidder in this group imposes on the others. This modified version maintains all the incentive properties of PVM, even though it is more likely to run at a deficit.

We evaluate PVM with two experiments. These experiments are again performed on settings sampled from GSVM and LSVM. Furthermore, we also test PVM on the large settings with 98 items and 10 bidders generated via the Multi-Region Value Model (MRVM) distribution (Weiss et al., 2017), which is based on the Canadian 2014 spectrum auction results. To keep the computational requirements practical, in these experiments, we only consider support vector regression algorithms with quadratic kernels: these algorithms allow the auctioneer to determine inferred efficient allocations via succinct integer programs.

First, we test the empirical efficiency obtained in the allocation computed by interpreting the bundle-value pairs reported in our elicitation algorithm using an XOR bidding language. The empirical efficiency of the allocation identified using machine learning significantly outperforms the efficiency of the non-machine learning baselines introduced in Brero et al. (2017). Then, we test our PVM mechanism in its original version, and in the version were bidders are partitioned into groups. The mechanism achieves very high empirical efficiency in all its versions.

Besides introducing PVM, this paper also shows how the machine learning-based elicitation algorithm introduced in Brero et al. (2017) can be extended to support relaxed

[5]The extended discussion around this idea is provided in Brero et al. (2019b).

reports consisting of upper and lower bounds on values. This idea can be beneficial to simplify the bidding process: in many real-world applications, determining the exact value for a bundle can be a costly exercise for bidders (see, e.g., Parkes, 2006). From simulation results, we see that the machine learning algorithm is still able to identify allocations with high empirical efficiency even under these weaker reports. This new elicitation algorithm points to interesting directions for future work to design mechanisms that only ask bidders to report upper and lower bounds on their values.

This paper provides a first answer to the third research question by introducing a practical mechanism based on value queries that motivates bidders to report their values truthfully. However, there is still an important missing step to provide a proper answer to Question 2 and Question 3: how does our machine learning-powered auction perform when compared to appropriate baselines?

1.4.5 Machine Learning-powered Iterative Combinatorial Auctions

This paper is meant to provide a final answer to Question 2 and Question 3 by performing a comparison between PVM and CCA. In the process of providing this answer, it also significantly improves the design of the value query-based elicitation algorithm and the auction mechanism based on this algorithm.

The elicitation algorithm introduced in Brero et al. (2017) has two significant problems: 1. it is forced to stop whenever each bidder has already evaluated the bundle she is assigned in the inferred efficient allocation, and 2. it may ask different numbers of queries among bidders, thus being perceived as unfair. In this paper, we introduce a new elicitation algorithm that always assigns a new query to each bidder in every iteration. The elicitation algorithm identifies this query considering the set of all feasible allocations that do not assign this bidder a bundle she has already evaluated. It then determines the allocation in this set that maximizes the inferred overall value to bidders and uses this allocation to assign a new query to this bidder. This elicitation algorithm is integrated into a new mechanism we call Pseudo VCG Machine Learning-based (PVML) mechanism that asks the same number of queries to each bidder at each round.

We compare PVML against three different implementations of the CCA that vary depending on the heuristic bidders use to select their bids in the supplementary round. The first heuristic we test is one where bidders do not report any extra bids in the supplementary round. The second one assumes that bidders submit their true values for all bundles they demanded during the clock phase. In the third heuristic, bidders report their true values for n profit-maximizing bundles at the final clock prices; we test this design for $n \in \{100, 200, 500\}$, where $n = 500$ is motivated by the query cap used in the 2014 Canadian Spectrum Auction (Industry Canada, 2013). We adopt the same query caps in PVML.

As in Brero et al. (2017) and Brero et al. (2018), we use prior samples from the distribution used to generate the setting at hand to determine which kernel to use in the SVRs. We considered Linear, Quadratic, Gaussian, and Exponential kernels, with corresponding hyperparameters. To keep the computations practical, we set a time limit of one minute to the integer programs used to determine new queries. We found that the Quadratic kernel is the best performing one in all three domains as it both allows the auction to capture non-additive preferences between items and determine new queries via succinct integer programs.

In our experiments, we see that PVML outperforms the CCA in terms of efficiency, in particular in the large MRVM domains. Furthermore, while the CCA can have very low revenue depending on the heuristic used in the supplementary round, PVML always achieves similar revenue to VCG.

1.5 Conclusion and Future Work

In this thesis, I have introduced several iterative combinatorial auction (CA) designs using preference elicitation algorithms based on statistical learning models. These algorithms allow the auctioneer to exploit prior data on bidders' valuations to reduce the costs of elicitation.

I have first introduced a price-based iterative CA that exploits prior data to reduce the number of rounds to reveal clearing prices. This auction shows that exploiting reasonable priors over bidders' valuations is much more effective than carefully tuning price increments in standard designs. Motivated by laboratory experiments highlighting that price-based CA designs may lead to suboptimal bidding, I have shifted my focus to value query-based designs. I have first designed an elicitation algorithm that uses statistical machine learning algorithms to determine which value queries to ask at each round. I have then integrated this algorithm into an auction design called Pseudo VCG Machine Learning-based (PVML) mechanism, which provides bidders with good incentives to report their values truthfully. I have compared PVML with the Combinatorial Clock Auction (CCA) in realistic spectrum auction settings featuring up to 98 spectrum licenses. I have tested the most reasonable bidding heuristics for the CCA. The results have shown that, when reasonable priors are available, PVML achieves higher allocative efficiency than the CCA, even when the bidding heuristics in the CCA are optimized for efficiency.

It is not clear whether value queries simplify the bidding process compared to demand queries. Indeed, to answer a demand query, bidders only need to state their favorite bundle at some ask prices without having to report exact values. To address this critique, I am currently designing a new PVML version where, instead of exact values, bidders are only asked to report upper and lower bounds on bundles. In Brero et al. (2018), we have shown that, even under these weaker reports, our machine learning-based elicitation algorithm identifies highly efficient allocations. The new PVML version I am designing is based on combining the elicitation algorithm introduced in Brero et al. (2018) with the bound-refinement algorithm introduced in Lubin et al. (2008).

Another critique of the value query-based auction designs is that their performance is assessed via empirical efficiency measures. In general, worst-case efficiency guarantees can be more useful in practice, as they provide the auctioneer with principled beliefs that, when the elicitation stops, the most relevant aspects of bidders' valuations have been investigated. Given that efficiency is measured considering bidders' true valuations, which are assumed to be their private knowledge, efficiency guarantees can only be expressed based on assumed bidding behaviors. Commonly, bidders are assumed to be rational and truthful. Unfortunately, even under these assumptions, Nisan and Segal (2006) have shown that it is not possible to design practical CAs for large settings that always terminate with useful efficiency guarantees. However, under rational, truthful bidding, price-based designs are intuitively "safer" than value query-based ones, as they require bidders to explore their valuations several times during the auction when selecting their bids. It is then less likely that relevant aspects of bidders' valuations do not get revealed during the auction. In this thesis, I have opted for designs that are based on only one query type (i.e., demand or value queries) to make my analysis more principled. However, future work can investigate hybrid approaches based on both demand and value queries that achieve high empirical efficiency while providing better guarantees on worst-case efficiency when bidders are fully rational and truthful. At the same time, one could also derive efficiency guarantees based on different bidding behaviors, like the one highlighted by Bichler et al. (2013), where bidders preselect which bundles to bid on during the auction. In this scenario, value query-based auctions may provide even stronger efficiency guarantees than demand query-based ones, as they require bidders to evaluate bundles they would not have considered otherwise.

There are also some limitations related to the settings used to evaluate the new auction designs: The price-based approaches introduced were tested on CATS settings with 12 items and 10 bidders, and CATS only generates bidders interested in a small number of bundles. These settings are much smaller and qualitatively very different from the spectrum auction settings used for the value query-based approaches. Thus, it may not be computationally feasible to run my price-based auctions in realistic settings. Further work would be needed to scale the price-based auctions to larger settings with different kinds of valuations. Different prior distributions can be considered depending on the auction performance in these new settings. At the same time, the settings where I have tested the value query-based auctions are only related to spectrum sales. It would be interesting to test these designs in new kinds of settings related to other domains of applications (e.g., industrial procurement).

It is also important to note that I have evaluated my new auctions on synthetic settings where I have simulated truthful bidding behaviors justified by game-theoretic arguments based on quasi-linear utilities (although these arguments are not based on formal guarantees but on conditions that generally hold in practice). In practice, this model may be not realistic, and bidders may be interested in raising prices for their competitors or adopt collusive strategies. There are different works—e.g., by Knapek and Wambach (2012) or Janssen et al. (2017)—that have shown how bidders willing to increase rivals' costs or facing budget constraints can benefit from several strategic behaviors in the CCA. As similar behaviors may likely arise in our new auctions, testing these auctions with laboratory experiments is an interesting direction for future work.

Bibliography

Lawrence M. Ausubel. An efficient ascending-bid auction for multiple objects. *American Economic Review*, 94(5):1452--1475, 2004.

Lawrence M. Ausubel and Oleg Baranov. A Practical Guide to the Combinatorial Clock Auction. *The Economic Journal*, 127(605), 2017.

Lawrence M. Ausubel and Paul Milgrom. The Lovely but Lonely Vickrey Auction. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial auctions*. The MIT Press, 2006.

Lawrence M. Ausubel, Peter Cramton, Paul Milgrom, et al. The Clock-proxy Auction: A Practical Combinatorial Auction Design. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, chapter 5. MIT Press, 2006.

Lawrence M. Ausubel, Peter Cramton, et al. Auction Design for Wind Rights. *Report to Bureau of Ocean Energy Management, Regulation and Enforcement*, 2011.

Maria-Florina Balcan, Tuomas Sandholm, and Ellen Vitercik. A General Theory of Sample Complexity for Multi-item Profit Maximization. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 173--174. ACM, 2018.

Martin Bichler, Pasha Shabalin, and Alexander Pikovsky. A Computational Analysis of Linear Price Iterative Combinatorial Auction Formats. *Information Systems Research*, 20(1): 33--59, 2009.

Martin Bichler, Alok Gupta, and Wolfgang Ketter. Research commentary---designing smart markets. *Information Systems Research*, 21(4):688--699, 2010.

Martin Bichler, Pasha Shabalin, and Jürgen Wolf. Do Core-selecting Combinatorial Clock Auctions Always Lead to High Efficiency? An Experimental Analysis of Spectrum Auction Designs. *Experimental Economics*, 16(4):511--545, 2013.

Sushil Bikhchandani and Joseph M. Ostroy. The Package Assignment Model. *Journal of Economic theory*, 107(2):377--406, 2002.

Avrim Blum, Jeffrey Jackson, Tuomas Sandholm, and Martin Zinkevich. Preference Elicitation and Query Learning. *Journal of Machine Learning Research*, 5:649--667, 2004.

Gianluca Brero and Sébastien Lahaie. A Bayesian Clearing Mechanism for Combinatorial Auctions. In *Proceedings of the 32nd AAAI Conference of Artificial Intelligence*, pages 941--948, 2018.

Gianluca Brero, Benjamin Lubin, and Sven Seuken. Probably Approximately Efficient Combinatorial Auctions via Machine Learning. In *Proceedings of the 31st AAAI Conference of Artificial Intelligence*, pages 397--405, 2017.

Gianluca Brero, Benjamin Lubin, and Sven Seuken. Combinatorial Auctions via Machine Learning-based Preference Elicitation. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 128--136, 2018.

Gianluca Brero, Sébastien Lahaie, and Sven Seuken. Fast Iterative Combinatorial Auctions via Bayesian Learning. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence*, 2019a.

Gianluca Brero, Benjamin Lubin, and Sven Seuken. Machine Learning-powered Iterative Combinatorial Auctions. Working Paper, 2019b.

Donald E. Campbell. *Incentives: Motivation and the Economics of Information.* Cambridge University Press, 2018.

Chris Caplice. Electronic Markets for Truckload Transportation. *Production and Operations Management*, 16(4):423--436, 2007.

Vincent Conitzer and Tuomas Sandholm. Complexity of Mechanism Design. In *Proceedings of the 18th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*, Edmonton, Canada, 2002.

Vincent Conitzer and Tuomas Sandholm. Self-interested Automated Mechanism Design and Implications for Optimal Combinatorial Auctions. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 132--141. ACM, 2004.

Robert G. Cowell, A. Philip Dawid, and Paola Sebastiani. A Comparison of Sequential Learning Methods for Incomplete Data. *Bayesian Statistics*, 5:533--542, 1996.

Peter Cramton. Ascending Auctions. *European Economic Review*, 42(3-5):745--756, 1998.

Peter Cramton. Spectrum Auction Design. *Review of Industrial Organization*, 42(2):161--190, 2013.

Peter Cramton and Axel Ockenfels. The German 4G Spectrum Auction: Design and Behaviour, 2017.

Philip Cusick, Eric Pan, Richard Choe, and Derya Erdemli. Spectrum Overview and Valuation Matrix - Carrier by Carrier Spectrum Value Across the Wireless Industry. Technical report, J.P. Morgan, 2012.

Sven de Vries, James Schummer, and Rakesh V Vohra. On Ascending Vickrey Auctions for Heterogeneous Objects. *Journal of Economic Theory*, 132(1):95--118, 2007.

Arthur P. Dempster, Nan M. Laird, and Donald B. Rubin. Maximum Likelihood from Incomplete Data via the EM Algorithm. *Journal of the Royal Statistical Society, Series B*, pages 1--38, 1977.

Paul Duetting, Felix Fischer, Pichayut Jirapinyo, John K Lai, Benjamin Lubin, and David C. Parkes. Payment Rules through Discriminant-based Classifiers. *ACM Transactions on Economics and Computation*, 3(1):5, 2015.

Paul Duetting, Zhe Feng, Harikrishna Narasimhan, and David C. Parkes. Optimal Auctions through Deep Learning. arXiv:1706.0345, 2018.

Jacob K Goeree and Charles A Holt. Hierarchical Package Bidding: A Paper and Pencil Combinatorial Auction. *Games and Economic Behavior*, 70(1):146--169, 2010.

Andor Goetzendorf, Martin Bichler, Pasha Shabalin, and Robert W. Day. Compact Bid Languages and Core Pricing in Large Multi-item Auctions. *Management Science*, 61(7): 1684--1703, 2015.

Noah Golowich, Harikrishna Narasimhan, and David C Parkes. Deep Learning for Multi-Facility Location Mechanism Design. In *IJCAI*, pages 261--267, 2018.

Faruk Gul and Ennio Stacchetti. The English Auction with Differentiated Commodities. *Journal of Economic theory*, 92(1):66--95, 2000.

Industry Canada. Responses to clarification questions on the licensing framework for mobile broadband services (mbs) --- 700 mhz band. Technical report, Government of Canada, 2013. URL `http://www.ic.gc.ca/eic/site/smt-gst.nsf/eng/sf10651.html`.

Industry Canada. 700 mhz auction (2014) --- auction results. Technical report, Government of Canada, 2014. URL `http://www.ic.gc.ca/eic/site/smt-gst.nsf/eng/h_sf10598.html`.

Matthew O. Jackson. Mechanism Theory. *Available at SSRN 2542983*, 2014.

Maarten Janssen, Vladimir Karamychev, and Bernhard Kasberger. *Handbook of Spectrum Auction Design*, chapter Budget Constraints in Combinatorial Clock Auctions, pages 320--337. Cambridge University Press, 2017.

Alexander S. Kelso Jr and Vincent P. Crawford. Job Matching, Coalition Formation, and Gross Substitutes. *Econometrica: Journal of the Econometric Society*, pages 1483--1504, 1982.

Stephan Knapek and Achim Wambach. Strategic Complexities in the Combinatorial Clock Auction. CESifo Working Paper 3983, 2012.

Anthony M. Kwasnica, John O. Ledyard, Dave Porter, and Christine DeMartini. A New and Improved Design for Multiobject Iterative Auctions. *Management Science*, 51(3):419--434, 2005.

Sébastien Lahaie. A Kernel-Based Iterative Combinatorial Auction. In *Proceedings of the 25th Conference on Artificial Intelligence (AAAI)*, 2011.

Sébastien Lahaie and Benjamin Lubin. Adaptive-price Combinatorial Auctions. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 749--750. ACM, 2019.

Sébastien Lahaie and David C. Parkes. Applying Learning Algorithms to Preference Elicitation. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 180--188, 2004.

Kevin Leyton-Brown, Mark Pearson, and Yoav Shoham. Towards a Universal Test Suite for Combinatorial Auction Algorithms. In *Proceedings of the 2nd ACM Conference on Electronic Commerce*, pages 66--76. ACM, 2000.

Benjamin Lubin, Adam Juda, Ruggiero Cavallo, Sébastien Lahaie, Jeffrey Shneidman, and David C. Parkes. ICE: An Expressive Iterative Combinatorial Exchange. *Journal of Artificial Intelligence Research*, 33:33--77, 2008.

Leonardo Meeus, Karolien Verhaegen, and Ronnie Belmans. Block Order Restrictions in Combinatorial Electric Energy Auctions. *European journal of operational research*, 196 (3):1202--1206, 2009.

Debasis Mishra and David C. Parkes. Ascending Price Vickrey Auctions for General Valuations. *Journal of Economic Theory*, 132(1):335--366, 2007.

Harikrishna Narasimhan, Shivani Brinda Agarwal, and David C. Parkes. Automated Mechanism Design without Money via Machine Learning. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence*, 2016.

Noam Nisan and Ilya Segal. The Communication Requirements of Efficient Allocations and Supporting Prices. *Journal of Economic Theory*, 129(1):192--224, 2006.

David C. Parkes. iBundle: an Efficient Ascending Price Bundle Auction. In *Proceedings of the 1st ACM Conference on Electronic Commerce*, pages 148--157. ACM, 1999.

David C. Parkes. Iterative Combinatorial Auctions. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, chapter 2. MIT Press, 2006.

David Porter, Stephen Rassenti, Anil Roopnarine, and Vernon Smith. Combinatorial Auction Design. *Proceedings of the National Academy of Sciences*, 100(19):11153--11157, 2003.

Tuomas Sandholm. Very-Large-Scale Generalized Combinatorial Multi-Attribute Auctions: Lessons from Conducting $60 Billion of Sourcing. In Nir Vulkan, Alvin E. Roth, and Zvika Neeman, editors, *The Handbook of Market Design*, chapter 1. Oxford University Press, 2013.

Tuomas Sandholm and Craig Boutilier. Preference Elicitation in Combinatorial Auctions. In Peter Cramton, Yoav Shoham, and Richard Steinberg, editors, *Combinatorial Auctions*, chapter 10. The MIT Press, 2006.

Tobias Scheffel, Alexander Pikovsky, Martin Bichler, and Kemal Guler. An Experimental Comparison of Linear and Non-linear Price Combinatorial Auctions. *Information Systems Research*, 22(2):346--368, 2011.

Tobias Scheffel, Georg Ziegler, and Martin Bichler. On the Impact of Cognitive Limits in

Combinatorial Auctions: An Experimental Study in the Context of Spectrum Auction Design. *Experimental Economics*, 15:667--692, 2012.

Alex J. Smola and Bernhard Schölkopf. A Tutorial on Support Vector Regression. *Statistics and computing*, 14(3):199--222, 2004.

Michael Weiss, Benjamin Lubin, and Sven Seuken. SATS: A Universal Spectrum Auction Test Suite. In *Proceedings of the Sixteenth Conference on Autonomous Agents and MultiAgent Systems*, pages 51--59, 2017.

Christopher K. Williams and Carl E. Rasmussen. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.